Combining Voting and Abstract Argumentation to Understand Online Discussions

Michael Bernreiter
DBAI, TU Wien
Vienna, Austria
michael.bernreiter@tuwien.ac.at

Jan Maly
DBAI, TU Wien
Vienna, Austria
jan.maly@tuwien.ac.at

Oliviero Nardi
DBAI, TU Wien
Vienna, Austria
oliviero.nardi@tuwien.ac.at

Stefan Woltran
DBAI, TU Wien
Vienna, Austria
stefan.woltran@tuwien.ac.at

ABSTRACT

Online discussion platforms are a vital part of the public discourse in a deliberative democracy. However, how to interpret the outcomes of the discussions on these platforms is often unclear. In this paper, we propose a novel and explainable method for selecting a set of most representative, consistent points of view by combining methods from computational social choice and abstract argumentation. Specifically, we model online discussions as abstract argumentation frameworks combined with information regarding which arguments voters approve of. Based on ideas from approval-based multiwinner voting, we introduce several voting rules for selecting a set of preferred extensions that represents voters' points of view. We compare the proposed methods across several dimensions, theoretically and in numerical simulations, and give clear suggestions on which methods to use depending on the specific situation.

KEYWORDS

argumentation; voting; digital democracy; online discussions

ACM Reference Format:
Michael Bernreiter, Jan Maly, Oliviero Nardi, and Stefan Woltran. 2024. Combining Voting and Abstract Argumentation to Understand Online Discussions. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 10 pages.

1 INTRODUCTION

In recent years, a veritable "deliberative wave" has swept through many democratic societies [44], bringing with it many new discursive participation formats, from citizen assemblies to online discussion platforms. These formats allow citizens to discuss (often very divisive) political issues and thus can enable us to understand which opinions are held by well-informed citizens. In this paper, we focus on text-based online discussion platforms used to inform political decision making, such as *Polis* (pol.is), *Your Priorities* (yrpri.org) or *Decidim* (decidim.org) to name just a few of the rapidly expanding number of tools used around the world [32]. These platforms allow citizens to submit comments, but also to approve (and sometimes disapprove) comments others have posted. While such platforms are easy to set up and use, discussions can be unorderly and chaotic, which makes it hard to interpret and summarize them. One approach for solving this problem – used for example by *Polis* [48] – is to use machine learning and statistics to find clusters of voters with similar opinions, as well as a set of comments representing their joint opinion. However, with this method the comments used to represent a cluster of voters might be inconsistent. Additionally, the obtained results are generally not explainable, which is highly problematic for processes that inform political decision making.

In this paper, we propose an explainable method for picking consistent points of view that represent the opinions of voters using methods from computational social choice (ComSoC) combined with tools from abstract argumentation [23]. Picking a representative set of most popular comments based on citizens' approvals is an aggregation problem that is closely related to the well-studied setting of approval-based multiwinner voting [40], for which many voting rules guaranteeing fair representation have been introduced lately [3, 46]. However, as every comment is only one argument in a larger discussion, we would argue that it does not suffice to pick a single representative set of popular comments, but that we have to find *sets* of comments that represent citizens' points of view.

Moreover, these points of view should be *consistent* if we want them to influence political decision making. To guarantee this, we need additional semantic information about the comments. Depending on the structure of the discussion, different formalisms could be used to capture this information. We focus on abstract argumentation frameworks (AFs), one of the most well-studied formalisms for representing conflicting arguments. AFs are particularly well-suited for our setting due to their minimal and easy to understand syntax that only requires the specification of attacks between different arguments. These attack relations either have to be added by a moderator, crowd-sourced from the participants, or mined from natural language text using argument mining techniques [7, 13, 35, 45].

Our contributions are as follows: we formally introduce *Approval-Based Social Argumentation Frameworks* to model online discussions in which citizens can approve arguments. We study the problem of selecting a small but representative set of so-called preferred extensions, which are maximal consistent and defendable sets of arguments. We study this problem from two sides. First, in smaller

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand. © 2024 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

discussions, we might want to pick an, ideally small, set of extensions that represents every voter perfectly. Whether this is possible depends, as we shall show, on how we conceptualize perfect representation. Secondly, for large discussions, we might be more interested in picking a small number of extensions that represent the voters *as well as possible*. We propose several methods for doing so, based on well-known voting rules from the ComSoC literature, and compare them with respect to their computational complexity, their axiomatic properties and their performance in numerical simulations. Based on these results, we make precise recommendations for which methods are best suited for different applications.

The full version of this paper (containing full proofs) [9] and the code used in the experiments [10] are available online.

Related Work

Social choice theory has been used before to analyze discussions in participation platforms (like *Polis*), e.g., by Halpern et al. [36] and Fish et al. [31] However, these works do not use argumentation to capture the relationship between comments, nor are they concerned with choosing a consistent subset of them. Hence, they are technically quite distinct from ours. To our knowledge, our paper is the first to consider voter representation in abstract argumentation.

That said, there is already a significant amount of literature on combining argumentation and social choice [8]. One strain of literature [2, 14] applies ideas from judgment aggregation [30] to argumentation. Here, the objective is to aggregate individual judgments about the acceptability of arguments. Crucially, these works assume that voters submit ballots adhering to strict rationality constraints, which is not realistic in online discussions, where voters are often not even aware of all arguments at the time of voting.

Another related approach is that of *weighted* AFs [11, 26], where attacks between arguments have weights that can be established by letting agents vote on arguments, similar to our scenario. Similarly, in *social* AFs [41], voters can approve (or disapprove) arguments. The strength of an argument is then related to its social support and the support of its attacking arguments. Beyond abstract argumentation, in the *relational reasoning model* [33] voters can judge the acceptability of a set of arguments and their relationships (by assigning them weights). The goal here is to aggregate such judgments and collectively evaluate a set of target arguments. While there are conceptual similarities to our approach, all three formalisms redefine argumentation semantics in light of votes (e.g. by considering the strength of arguments or attacks). In contrast, we rely on standard semantics for AFs and use the additional approval information to select a representative subset of extensions. Most importantly, none of the papers mentioned above study the problem of voter representation, which is the main objective of our paper.

Finally, there exists a significant body of work focused on the merging of AFs [17, 19, 22, 27, 50] where each agent is endowed with a framework, and the goal is to merge these into a single one.

2 PRELIMINARIES

The basic problem considered in this paper, namely selecting representative comments based on the approvals of voters, is an aggregation problem with the following components: a finite set of voters N, a finite set of candidates C, and a vector of approval ballots $\bar{A} = (A(i))_{i \in N}$, where $A(i) \subseteq C$. Note that this is equivalent to the input of an approval-based single- or multiwinner election. The difference to these formalisms lies in the outcome W we try to select: we want to select a set W of subsets of C, i.e., $W \subseteq 2^C$ where 2^C denotes the power set of C. Often, we will study the problem of selecting exactly k subsets, i.e., we will stipulate $|W| = k$. However, in contrast to most multiwinner voting settings, we will not constrain the cardinality of the selected subsets in W. Instead, we impose consistency constraints using abstract argumentation.

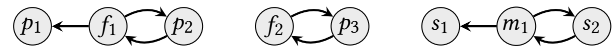

Figure 1: AF for the Canadian election reform discussion.

Argumentation. AFs [23] are a well-studied formalism in which discussions can be represented and reasoned about. Arguments (denoted by lower-case letters $a, b, c, \ldots$) in AFs are abstract entities, i.e., we are not concerned with their internal structure but rather with the relationship between them. Specifically, an argument x can attack another argument y, which implies that x and y are in conflict and cannot be jointly accepted. To accept y, it must be defended against x's attack, i.e., either y attacks x or there is another argument z which attacks x and can be accepted alongside y.

Definition 1. *An argumentation framework (AF) is a tuple $F = (Arg, Att)$ where Arg is a finite set of arguments and $Att \subseteq Arg \times Arg$ is an attack relation between arguments. Let $S \subseteq Arg$. S* attacks *b (in F) if $(a, b) \in Att$ for some $a \in S$; $S_F^+ = \{b \in Arg \mid \exists a \in S : (a, b) \in Att\}$ denotes the set of arguments attacked by S. An argument $a \in Arg$ is* defended *(in F) by S if $b \in S_F^+$ for each b with $(b, a) \in Att$.*

AFs are usually visualized as directed graphs, where a node is an argument and an edge an attack between arguments (see Figure 1).

AF-semantics are functions σ that assign a set $\sigma(F) \subseteq 2^{Arg}$ of extensions to each AF $F = (Arg, Att)$. Conflict-free ($\sigma = cf$) semantics select sets $S \subseteq Arg$ where no two arguments attack each other. Admissible semantics ($\sigma = adm$) select conflict-free sets that defend themselves. Preferred semantics ($\sigma = prf$) select subset-maximal admissible sets. Many alternative AF-semantics have been defined [5], but we focus on the well-established preferred semantics.

Definition 2. *Let $F = (Arg, Att)$ be an AF. For $S \subseteq Arg$ it holds that*

- *$S \in cf(F)$ iff there are no $a, b \in S$ such that $(a, b) \in Att$;*
- *$S \in adm(F)$ iff $S \in cf(F)$ and each $a \in S$ is defended by S;*
- *$S \in prf(F)$ iff $S \in adm(F)$ and $S \not\subset T$ for all $T \in adm(F)$.*

Example 1. *Let F be the AF from Figure 1. Note that $\{p_1, p_2\} \in cf(F)$ but $\{f_1, p_2\} \notin cf(F)$. For admissible semantics, $\{p_2\} \in adm(F)$ since p_2 defends itself against the attack (f_1, p_2). However, $\{p_1\} \notin adm(F)$ since p_1 does not defend itself against (f_1, p_1) while $\{p_1, p_2\} \in adm(F)$ since p_2 defends p_1. As for preferred extensions, we have that $prf(F) = \{\{p_1, p_2, p_3, s_1, s_2\}, \{p_1, p_2, p_3, m_1\}, \{f_2, p_1, p_2, s_1, s_2\}, \{f_2, p_1, p_2, m_1\}, \{f_1, p_3, s_1, s_2\}, \{f_1, p_3, m_1\}, \{f_1, f_2, s_1, s_2\}, \{f_1, f_2, m_1\}\}$.*

Complexity Theory. We assume familiarity with complexity classes P and NP. Moreover, Θ_2P is the class of decision problems solvable in polynomial time with access to $O(\log n)$ NP-oracle calls [51].

3 APPROVAL-BASED SOCIAL AFS

Let us now introduce our main object of study, *Approval-Based Social AFs* (ABSAFs), which model discussions where agents can approve arguments that they find convincing.

Definition 3. *An ABSAF $\mathcal{S} = (F, N, \bar{A})$ consists of an AF $F = (Arg, Att)$, a finite set N of agents (also called voters), and a vector $\bar{A} = (A(i))_{i \in N}$ of approval ballots where, for every agent $i \in N$, $A(i) \subseteq Arg$ with $A(i) \neq \emptyset$ is the set of arguments approved by i.*

Throughout the paper, we denote the number of voters by n and identify each voter by a natural number i, i.e., $N = \{1, \ldots, n\}$.

We do not stipulate any constraints for the submitted ballots – not even conflict-freeness – as ballots containing conflicting arguments appear in real-world examples (see Example 2). The goal is to select a (usually small) set of coherent viewpoints representing the agents.

Definition 4. *An outcome $\Omega \subseteq \sigma(F)$ of an ABSAF $\mathcal{S} = (F, N, \bar{A})$ is a set of σ-extensions. We call $\pi \in \Omega$ a point of view (or viewpoint).*

Example 2. *We consider a discussion on an election reform proposed in Canada, taken from the "Computational Democracy Project" (compdemocracy.org). For succinctness, we show two arguments in favor of the first-past-the-post (FPTP) electoral system (f_1, f_2), three arguments for proportional representation (PR) (p_1, p_2, p_3), two arguments for the single transferable vote (STV) system (s_1, s_2), and one argument for a mixed-member proportional (MMP) system (m_1). The "cid" code refers to the comment-id in the original dataset.*[1]

- f_1 *(cid 16): "In systems with PR, opinions are too strongly divided and nothing gets done."*
- f_2 *(cid 15): "FPTP results in more stable governance."*
- p_1 *(cid 21): "A party's share of seats in the House of Commons should reflect its share of the popular vote."*
- p_2 *(cid 42): "PR will reduce hyper-partisanship and promote cooperation between parties."*
- p_3 *(cid 43): "PR will reduce wild policy swings and result in more long-lasting policies."*
- s_1 *(cid 162): "STV's advantage over MMP is that it doesn't explicitly enshrine political parties in our electoral system."*
- s_2 *(cid 168): "I like the simplicity of stating my favorite candidate, as well as alternative choices."*
- m_1 *(cid 163): "MMP has a better chance of obtaining public support than STV due to its relative simplicity."*

We manually added attacks between arguments. The resulting AF $F = (Arg, Att)$ is shown in Figure 1 and is the same as in Example 1. Moreover, we extracted the ballots, along with how many agents voted for the given ballot, from the original dataset:

33$\times\{p_1\}$; ***31***$\times\{p_1, p_2, p_3\}$; ***16***$\times\{p_2\}, \{p_3\}$; ***11***$\times\{f_2\}$; ***10***$\times\{p_2, p_3\}$; ***9***$\times\{p_1, p_3\}, \{f_2, p_1\}$; ***8***$\times\{p_1, p_2\}$; ***7***$\times\{f_2, p_1, p_2, p_3\}$; ***6***$\times\{s_2\}$; ***4***$\times\{f_1, f_2\}, \{p_1, p_2, p_3, m_1, s_1\}$; ***3***$\times\{f_1\}, \{f_2, p_2, p_3\}$; ***2***$\times\{f_2, p_1, p_3\}, \{f_1, p_1\}, \{f_2, p_1, p_2\}, \{p_1, p_2, p_3, s_1\}$; ***1***$\times\{p_2, m_1\}, \{s_1\}, \{m_1\}, \{f_2, p_3\}, \{p_2, p_3, s_2\}, \{f_1, f_2, p_3\}, \{p_2, s_1\}, \{p_1, p_2, p_3, s_1, s_2\}, \{f_1, p_1, p_2, p_3\}$.

Most voters approve arguments in favor of PR. Some voters approve both arguments in favor of PR and FPTP, e.g., the ballot $\{f_1, p_2\}$. Note that this ballot is conflict-free. However, not all ballots are conflict-free, e.g., $\{f_2, p_1, p_2, p_3\}$ was approved by seven voters. Moreover, the ballot with the most voters, namely $\{p_1\}$, is not admissible.

Our main concern will be to find an outcome that is small, in that it contains few points of view, but also represents as many agents as possible. For this, we need to formally define what it means for an agent to be represented by a point of view or by an outcome.

Definition 5. *Let $\mathcal{S} = (F, N, \bar{A})$ be an ABSAF and let $\alpha \in [0 \ldots 1]$. A point of view $\pi \in \sigma(F)$ α-represents voter $i \in N$ iff*

$$rep_i(\pi) = \frac{|\pi \cap A(i)|}{|A(i)|} \geq \alpha.$$

An outcome $\Omega \subseteq \sigma(F)$ α-represents voter $i \in N$ iff $rep_i(\Omega) = \max_{\pi \in \Omega} rep_i(\pi) \geq \alpha$.

Example 3 (Example 2 continued). *Let $A(i) = \{f_2, p_1, p_2\}$, and consider the preferred extensions (points of view) $\pi_1 = \{f_1, f_2, m_1\}$, $\pi_2 = \{p_1, p_2, p_3, m_1\}$, and $\pi_3 = \{f_2, p_1, p_2, m_1\}$. Then $rep_i(\pi_1) = 1/3$, $rep_i(\pi_2) = 2/3$, and $rep_i(\pi_3) = 1$. Regarding outcomes of size 2 we have $rep_i(\{\pi_1, \pi_2\}) = 2/3$ and $rep_i(\{\pi_1, \pi_3\}) = rep_i(\{\pi_2, \pi_3\}) = 1$.*

In deliberative democracy, we usually assume that voters are rational agents [18] and thus, in principle, would arrive at a consistent and defendable point of view when considering all arguments for a sufficient amount of time, which we could call their *ideal* point of view. However, in practice, the vast majority of participants in an online discussion will not carefully consider all comments. Hence, we do not assume that ballots are complete or consistent. In light of this, we interpret $rep_i(\pi)$ as a measure of how consistent the observed voting behavior of i is with the assumption that π is her ideal point of view. Due to our rationality assumption, ideal points of view should at least be admissible. As we aim to represent voters with a small number of extensions, we can focus on preferred extensions, i.e., $\sigma = prf$, since this gives us subset-maximal coherent viewpoints. If an admissible viewpoint π contains all approved arguments of an agent, so does the preferred viewpoint $\pi' \supseteq \pi$. Finally, as two different preferred extensions must by definition be conflicting, each voter's ideal point of view can only coincide with one preferred extension. Thus, we define $rep_i(\Omega)$ as $\max_{\pi \in \Omega}(rep_i(\pi))$.

4 PERFECT REPRESENTATION

Ideally, we want to find an outcome that perfectly represents everyone, i.e., an outcome that 1-represents all voters. This is not always possible in practice, however. Indeed, Example 2 clearly demonstrates that real votes cannot be assumed to be conflict-free, in which case there can be no outcome that 1-represents everyone. But even if ballots are conflict-free, it may not be possible to fully represent all voters: consider the ABSAF $(F, N, \bar{A})$ from Figure 2. The labels above the arguments represent the voters approving of them, i.e., $A(1) = \{a, b\}$, $A(2) = \{a, c, d, e\}$. Note that $prf(F) = \{\{a, b\}, \{a, c\}\}$, but $rep_2(\{a, b\}) = 0.25$, and $rep_2(\{a, c\}) = 0.5$.

Additionally, deciding whether all voters can be perfectly represented is NP-complete. This follows from the fact that deciding credulous acceptance[2] is NP-complete for preferred semantics [28].

Definition 6. 1-Representability *is the following decision problem: given an ABSAF $\mathcal{S} = (F, N, \bar{A})$ and $k \in \{1, \ldots, |N|\}$, is there $\Omega \subseteq prf(F)$ with $|\Omega| \leq k$ such that $rep_i(\Omega) = 1$ for all $i \in N$?*

[1] https://github.com/compdemocracy/openData/tree/master/canadian-electoral-reform

[2] An argument x is credulously accepted (w.r.t. σ) in an AF F iff $\exists S \in \sigma(F)\colon x \in S$.

Proposition 1. 1-Representability *is* NP*-complete.* NP*-hardness holds even if there is only one voter, i.e., for* $n = 1$.

The above result also suggests that finding a straightforward characterization for representability is not possible in the general case. However, we can restrict approval-ballots and AFs to guarantee representation. Specifically, assuming conflict-free votes, we can give a lower-bound for α-representation by differentiating between arguments that defend themselves and arguments that do not.

Proposition 2. *Let* $\mathcal{S} = (F = (Arg, Att), N, \bar{A})$ *be an ABSAF. Let* $SD(i) = \{a \in A(i) \mid (b, a) \in Att \implies (a, b) \in Att\}$ *be the self-defending arguments approved by* $i \in N$*. If* $A(i) \in cf(F)$ *and* $\frac{|SD(i)|}{|A(i)|} \geq \alpha$ *for all* $i \in N$*, then there is* $\Omega \subseteq prf(F)$ *such that* $rep_i(\Omega) \geq \alpha$ *for all* $i \in N$*.*

If we restrict ourselves to symmetric AFs [20, 24], a natural subclass of AFs where $(a, b) \in Att$ implies $(b, a) \in Att$, every agent always approves self-defending arguments. Thus, by Proposition 2, if we are given conflict-free ballots and a symmetric AF F, there is an outcome $\Omega \subseteq prf(F)$ such that $rep_i(\Omega) = 1$ for all $i \in N$.

So far, we attempted to represent *all* approved arguments for each voter. However, even if we assume votes to be conflict-free, some voters might approve of arguments that cannot occur together in a preferred extension, or even arguments that cannot be defended altogether. Voter 2 in Figure 2, for instance, approves two undefendable arguments (d and e). Thus, $rep_2(\{a, c\}) = 0.5$ even though $\{a, c\}$ contains *all* arguments in $A(2)$ that can *actually* occur in a viewpoint. Hence, one could argue that voter 2 is already represented as well as possible. We therefore introduce an alternative notion of representation that we call *core-representation.*

Definition 7. *Let* $\mathcal{S} = (F, N, \bar{A})$ *be an ABSAF. For* $i \in N$ *we let* $\mu(i) = \max_{\pi \in \sigma(F)} |\pi \cap A(i)|$*. If* $\mu(i) = 0$ *we let* $rep^c_i(\pi) = 1$*, otherwise*

$$rep^c_i(\pi) = \frac{|\pi \cap A(i)|}{\mu(i)}.$$

We say that $\pi \in \sigma(F)$ α*-core-represents* i *iff* $rep^c_i(\pi) \geq \alpha$*. Moreover,* $\Omega \subseteq \sigma(F)$ α*-core-represents* i *iff* $rep^c_i(\Omega) = \max_{\pi \in \Omega} rep^c_i(\pi) \geq \alpha$*.*

Using core-representation for the ABSAF from Figure 2, for voter 1 we get $rep^c_1(\{a, b\}) = 1$ and $rep^c_1(\{a, c\}) = 0.5$ while for voter 2 we get $rep^c_2(\{a, b\}) = 0.5$ and $rep^c_2(\{a, c\}) = 1$.

In contrast to regular representation, it is always possible to find an outcome that perfectly core-represents every voter.

Observation 3. *Given an ABSAF* $\mathcal{S} = (F, N, \bar{A})$*, the outcome* $\Omega = prf(F)$ *1-core-represents every voter* $i \in N$*. Thus, there is an outcome* $\Omega' \subseteq prf(F)$ *of size* $|\Omega'| = |N|$ *that 1-core-represents every* $i \in N$*.*

However, deciding whether there is a *small* outcome that perfectly core-represents every voter is harder than in the case of regular representation, for which this problem is NP-complete (cf. Proposition 1). We define 1-Core-Representability analogously to 1-Representability (cf. Definition 6), except that we ask for an outcome of size $|\Omega| \leq k$ that 1-core-represents every voter $i \in N$.

Theorem 4. *1-Core-Representability is* Θ_2P*-complete. Moreover,* Θ_2P*-hardness holds even if there are only two voters, i.e., for* $n = 2$*.*

Proof. Θ_2P-membership: let $\mathcal{S} = (F = (Arg, Att), N, \bar{A})$ be an ABSAF and $k \in \{1, \ldots, |N|\}$. Note that Θ_2P coincides with

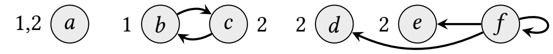

Figure 2: Voter 2 approves undefendable arguments.

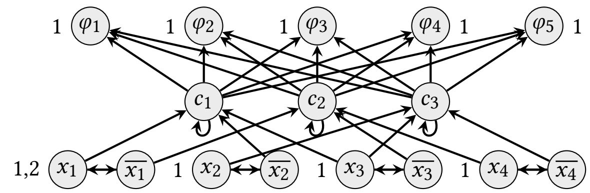

Figure 3: Construction used in the proof of Theorem 4.

$\mathsf{P}^{\mathsf{NP}}_{\|[2]}$ [12], the class of problems solvable in polynomial time with 2 rounds of parallel NP-oracle calls. In the first round, for every $i \in N$ and $m \in \{1, \ldots, |Arg|\}$, use an NP-oracle to decide if there is $\pi \in adm(F)$ such that $|\pi \cap A(i)| = m$. Then, for each $i \in N$, we compute $\mu(i) = \max_{\pi \in prf(F)} |\pi \cap A(i)| = \max_{\pi \in adm(F)} |\pi \cap A(i)|$. In the second round, use a single NP-oracle call to execute the following procedure: guess sets $\pi_1, \ldots, \pi_k \subseteq Arg$. Then, check that

- $\pi_j \in adm(F)$ for every $1 \leq j \leq k$,
- $|\pi_j \cap A(i)| = \mu(i)$ for every $1 \leq i \leq |N|$ and some $1 \leq j \leq k$.

Since $\pi_j \in adm(F)$ there is $\pi'_j \in prf(F)$ such that $\pi_j \subseteq \pi'_j$. For all $1 \leq i \leq |N|$ there is $1 \leq j \leq k$ such that $|\pi_j \cap A(i)| \leq |\pi'_j \cap A(i)| \leq \mu(i) = |\pi_j \cap A(i)|$, i.e., $|\pi'_j \cap A(i)| = \mu(i)$ and therefore $rep^c_i(\pi'_j) = 1$. Thus, for $\Omega = \{\pi'_1, \ldots, \pi'_k\}$ we have $rep^c_i(\Omega) = 1$ for every $i \in N$.

Θ_2P-hardness: via reduction from CardMaxSat [21], where we are given a propositional formula φ in CNF and a variable y, and ask if there is a cardinality-maximal model of φ containing y. Let X denote the set of variables in φ, and C the set of clauses. φ can be assumed to be satisfiable, since φ can be replaced by $\varphi' = \varphi \vee (\bigwedge_{x \in X} (\neg x))$, which can be transformed into CNF. Let $k = 1$ and construct ABSAF $\mathcal{S} = (F = (Arg, Att), N, \bar{A})$ as follows:

- $Arg = X \cup \{\overline{x} \mid x \in X\} \cup C \cup \varphi^*$, where $\varphi^* = \{\varphi_i \mid 1 \leq i \leq |X|+1\}$;
- $Att = \{(x, \overline{x}), (\overline{x}, x) \mid x \in X\} \cup \{(x, c) \mid x \in X, c \in C, x \in c\} \cup \{(\overline{x}, c) \mid x \in X, c \in C, \neg x \in c\} \cup \{(c, c), (c, \varphi_i) \mid c \in C, \varphi_i \in \varphi^*\}$;
- $N = \{1, 2\}$ with $A(1) = X \cup \varphi^*$ and $A(2) = \{y\}$.

Figure 3 shows the ABSAF constructed from clauses $c_1 = (x_1 \vee \neg x_2 \vee x_3)$, $c_2 = (\neg x_1 \vee \neg x_3 \vee x_4)$, $c_3 = (x_2 \vee x_3 \vee \neg x_4)$ with $y = x_1$.

(φ, y) is a yes-instance of CardMaxSat iff there is $\pi \in prf(F)$ such that $rep^c_i(\{\pi\}) = 1$ for all $i \in N$ (see full paper for details). □

If the number of arguments/extensions in our given ABSAF is small, however, we can deal with representability more efficiently.

Proposition 5. *1-Representability and 1-Core-Representability are FPT with respect to the number of arguments in the given ABSAF.*

Proof. Let $\mathcal{S} = (F = (Arg, Att), N, \bar{A})$ and $k \in \{1, \ldots, |N|\}$. Let $m = |prf(F)|$. Note that m is in $O^*(3^{|Arg|/3}) \subseteq O(2^{|Arg|})$ [25], and that $prf(F)$ can be enumerated in $O^*(3^{|Arg|/3})$ time [34]. Moreover, we assume $k \leq m$ since $\Omega \subseteq prf(F)$ for any outcome Ω.

For core-representation, determine $\mu(i)$ for all $i \in N$ by enumerating every $\pi \in prf(F)$ and computing $|\pi \cap A(i)|$. $\mu(i)$ is the size of the largest intersection. This runs in $O^*(3^{|Arg|/3})$ time.

Then, check if there is a k-tuple Ω of preferred extensions such that $rep_i(\Omega) = 1$ (resp. $rep^c_i(\Omega) = 1$) for every $i \in N$. There are $\binom{m}{k}$ such k-tuples, i.e., this can be done in $O^*(\binom{m}{k}) \subseteq O^*(m^k)$ time. ◻

The results in the remainder of the paper apply to both regular representation (Definition 5) and core-representation (Definition 7).

We have seen that we generally cannot 1-represent all voters, but we can 1-core-represent them. If the AF is small enough, we can even do so efficiently. On the other hand, when dealing with larger instances, we may require a huge number of preferred extensions to 1-(core-)represent all voters which does not help to understand the discussion. In this case, we might instead aim for an optimal representation with a fixed number of preferred extensions.

5 OPTIMIZING REPRESENTATION

Let us now consider the question how to optimally represent the voters in an ABSAF with a fixed number of preferred extensions. One idea, commonly used in social choice theory and referred to as the *Utilitarian rule*, is to pick an outcome Ω maximizing the average representation across all voters, i.e., $\sum_{i \in N} rep_i(\Omega)$.[3] The second standard approach, referred to as the *Egalitarian rule*, is to pick an outcome maximizing the representation of the least-represented voter, i.e., $\min_{i \in N} rep_i(\Omega)$.

A family of rules generalizing these ideas is based on *ordered weighted averaging* (OWA) vectors [52]. Given an outcome Ω, let $\vec{s}(\Omega) = (s_1, \dots, s_n)$ be the vector $(rep_1(\Omega), \dots, rep_n(\Omega))$ sorted in non-decreasing order (i.e., s_1 is the least-represented voter). For a non-increasing vector of non-negative weights $\vec{w} = (w_1, \dots, w_n)$, where $w_1 > 0$, the corresponding OWA rule is defined as:

$$\mathrm{OWA}_{\vec{w}}(\mathcal{S}) \in \operatorname*{arg\,max}_{\Omega \subseteq prf(F)\colon |\Omega| \leq k} \vec{w} \cdot \vec{s}(\Omega).$$

Here, $\cdot$ is the dot product. The Utilitarian and Egalitarian rules are given by $\vec{w} = (1, \dots, 1)$ and $\vec{w} = (1, 0, \dots, 0)$, respectively. We will also consider the *Harmonic rule* based on the vector $(1, 1/2, \dots, 1/n)$. This sequence of weights is often used to achieve proportionality in multiwinner voting [40]. OWA-rules in general have been studied, e.g., in the context of multiple referenda/issues [1, 6, 39].

An alternative approach is simply to maximize the number of voters that are 1-represented. We call this rule MaxCov, defined as:

$$\mathrm{MaxCov}(\mathcal{S}) \in \operatorname*{arg\,max}_{\Omega \subseteq prf(F)\colon |\Omega| \leq k} |\{i \in N\colon rep_i(\Omega) = 1\}|.$$

Example 4. *Consider an ABSAF with voters $N = \{1, 2, 3, 4\}$ and extensions $prf(F) = \{\pi_1, \pi_2, \pi_3\}$. Say that we want to pick an outcome of size $k = 1$ and that the representation scores are as follows:*

Ω	$rep_1(\Omega)$	$rep_2(\Omega)$	$rep_3(\Omega)$	$rep_4(\Omega)$
$\{\pi_1\}$	0	1	1	1
$\{\pi_2\}$	0.4	0.5	1	1
$\{\pi_3\}$	0.5	0.5	0.5	0.5

It can be verified that Utilitarian and MaxCov select $\{\pi_1\}$, Harmonic selects $\{\pi_2\}$, and Egalitarian selects $\{\pi_3\}$.

Example 5. *Selecting outcomes of size $k = 2$ for the Canada electoral reform data (cf. Example 2), Utilitarian and MaxCov pick $\Omega_1 = \{\{p_1, p_2, p_3, s_1, s_2\}, \{f_2, p_1, p_2, m_1\}\}$ while Harmonic and Egalitarian pick $\Omega_2 = \{\{p_1, p_2, p_3, s_1, s_2\}, \{f_1, f_2, m_1\}\}$. The minimum and average (core-)representation-scores across all voters are:*

	min	avg	min (core)	avg (core)
Ω_1	0	0.9378	0	0.9706
Ω_2	0.5	0.9360	0.5	0.9697

The average representation for Utilitarian is only marginally larger than for Egalitarian/Harmonic in this instance.

Note that all rules achieve perfect representation if possible.

Proposition 6. *The Egalitarian rule returns an outcome of size k in which every voter is α-represented iff such an outcome exists. Moreover, every OWA rule, as well as MaxCov, returns an outcome of size k in which every voter is 1-represented iff such an outcome exists.*

Moreover, all of the rules considered here are computationally intractable.[4] This follows from Proposition 1 and 6.

Proposition 7. *Computing an outcome that is optimal with respect to a given OWA-rule or MaxCov is* NP*-hard.*

While intractable, $\mathrm{OWA}_{\vec{w}}(\mathcal{S})$ and MaxCov rules are FPT: analogously to Proposition 5, we can first enumerate all preferred extensions in $O^*(3^{|Arg|/3})$ time and then enumerate all outcomes of size k in $O^*(\binom{m}{k}) \subseteq O^*(m^k)$ time, where $m = |prf(F)|$, allowing us to simply choose the best outcome w.r.t. the given rule.

In practice, however, a runtime of $O^*(\binom{m}{k})$ can be impractical, as we will see in Section 7. Thus, we introduce greedy variants for all of the above rules as follows. Consider first the greedy variant $\mathrm{GreedOWA}_{\vec{w}}$ of a rule $\mathrm{OWA}_{\vec{w}}$. Assuming that we have already selected ℓ viewpoints $\pi_1, \dots, \pi_\ell$, we pick the $(\ell + 1)$-th as follows:

$$\pi_{\ell+1} \in \operatorname*{arg\,max}_{\pi \in prf(F) \setminus \{\pi_1, \dots, \pi_\ell\}} \vec{w} \cdot \vec{s}(\{\pi_1, \dots, \pi_\ell, \pi\}).$$

We stop as soon as k points of view have been selected. The greedy variant of MaxCov, called GreedCov, is defined analogously. Observe that GreedCov approximates MaxCov with a factor of $1 - 1/e$, which follows directly from the approximation guarantee of the greedy algorithm for the Maximum Coverage Problem [37].

The greedy rules are also intractable, just like their non-greedy variants (cf. Proposition 7). This follows from Proposition 1.

Proposition 8. *Computing an outcome that is optimal with respect to a given greedy OWA-rule or GreedCov is* NP*-hard.*

However, we can improve upon the FPT-algorithm for the non-greedy rules since we do not need to enumerate all outcomes. Rather, it suffices to enumerate all points of view whenever we pick a new point of view in the greedy procedure. Thus, we can first enumerate all preferred extensions in $O^*(3^{|Arg|/3})$ time and then execute the greedy procedure in $O^*(mk)$, where $m = |prf(F)|$.

If we assume that we are given the preferred extensions, e.g., by precomputation via powerful argumentation solvers [29, 43], we have a runtime of $O^*(mk)$ for the greedy rules. This is polynomial in m since an outcome can contain at most m viewpoints, i.e., $k \leq m$. Assuming P $\neq$ NP, such an FPT-algorithm does not exist for the

[3] We assume that an arbitrary tiebreaking mechanism is used for all rules.

[4] Note that the problem of computing an optimal outcome w.r.t a given rule is an optimization problem, not a decision problem. We show that these optimization problems are NP-hard, i.e., that we can not solve them in polynomial time unless P = NP.

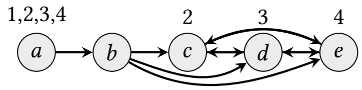

Figure 4: Counterexample SJR.

main non-greedy OWA-rules. For Utilitarian and Harmonic this can be shown by a reduction from the Chamberlin–Courant rule used in multiwinner voting [16, 40], and for Egalitarian via a reduction from Hitting Set [38]. See the full paper for details.

Proposition 9. *Assuming* $\mathsf{P} \neq \mathsf{NP}$*, there is no algorithm that, given an ABSAF* $\mathcal{S} = (F, N, \bar{A})$ *and given* $prf(F)$*, returns an outcome for* $\mathcal{S}$ *that is optimal with respect to the Utilitarian, Egalitarian, or Harmonic rule in polynomial time with respect to* $m = |prf(F)|$*.*

We have now introduced several voting rules. Next, we will compare them axiomatically and in numerical simulations.

6 JUSTIFIED REPRESENTATION

One of the key axioms of multiwinner voting is justified representation [3], which requires that each sufficiently large group of voters sharing the same opinion is represented in the outcome. This is also a natural desideratum in our setting. To formalize it, we first need to define what it means for a group to have the same opinion.

Definition 8. *Let* $\mathcal{S} = (F, N, \bar{A})$ *be an ABSAF with* $n = |N|$ *voters. We call a group of voters* $N' \subseteq N$ 1-representable *iff there is* $\pi \in prf(F)$ *such that for all* $i \in N'$ *we have* $rep_i(\pi) = 1$*.*

If such a group contains more than n/k voters, where k is the number of preferred extensions we want to select, then this group, arguably, deserves to be fully represented.

Definition 9. *An outcome* $\Omega \subseteq prf(F)$ *of size* $|\Omega| = k$ *satisfies* Strong Justified Representation *(SJR) iff for every* 1*-representable group* $N' \subseteq N$ *of size* $|N'| \geq n/k$ *there is a point of view* $\pi \in \Omega$ *such that* $rep_i(\pi) = 1$ *for every* $i \in N'$*.*

Observe that Definitions 8 and 9 could be written using the notion of core-representation (see Definition 7) instead. All results that we present in this section hold in either case.

As can be seen quite easily, if the approval sets of all voters are preferred extensions, then SJR can be satisfied.

Proposition 10. *Let* $\mathcal{S} = (F, N, \bar{A})$ *be an ABSAF. If* $A(i) \in prf(F)$ *for all* $i \in N$*, then for any* $k \in \{1, \dots, |prf(F)|\}$ *we can find an outcome* Ω *of size* $|\Omega| = k$ *that satisfies SJR.*

Unfortunately, once we drop this unreasonably strong assumption, the result does not hold. Consider the ABSAF from Figure 4. Voter groups $\{1, 2\}, \{1, 3\}, \{1, 4\}$ are 1-representable, but no outcome of size $k = 2$ satisfies SJR (see full paper for details).

Proposition 11. *It cannot be guaranteed that, given an ABSAF* $\mathcal{S} = (F, N, \bar{A})$ *and* $k \in \{1, \dots, |N| - 1\}$*, there is an outcome* Ω *of size* $|\Omega| = k$ *that satisfies SJR (even if votes are assumed to be admissible).*

This motivates us to introduce the following weakening of SJR, where we only require that at least one voter of every sufficiently large 1-representable group must be 1-represented. Analogous restrictions have also been proposed in the context of multiwinner voting [40] and Participatory Budgeting [47].

Definition 10. *An outcome* $\Omega \subseteq prf(F)$ *with* $|\Omega| = k$ *satisfies* Justified Representation *(JR) iff for every* 1*-representable group* $N' \subseteq N$ *with* $|N'| \geq n/k$ *there is* $\pi \in \Omega$ *such that* $rep_i(\pi) = 1$ *for some* $i \in N'$*.*

In order for an outcome Ω of size k to satisfy JR, the set of voters in N' that are not 1-represented by Ω must be smaller than n/k. Indeed, if n/k or more voters of a 1-representable group N' were not 1-represented by Ω, then $N'' = \{i \in N' \mid \forall \pi \in \Omega\ rep_i(\pi) < 1\}$ would be a 1-representable group of size n/k that violates JR.

Fortunately, in contrast to SJR, JR can always be satisfied. Indeed, in contrast to multiwinner voting it can be satisfied by just maximizing the number of 1-represented voters.

Theorem 12. *MaxCov and GreedCov satisfy JR.*

Proof. Let Ω with $|\Omega| = k$ be the outcome of MaxCov. Assume there is a 1-representable group N' of size at least n/k represented by π' such that no voter in N' is 1-represented by Ω. Thus, for the set of 1-represented voters $N^1_\Omega := \{i \in N \mid \exists \pi \in \Omega\ rep_i(\pi) = 1\}$ we have $|N^1_\Omega| \leq n - |N'| \leq n - n/k$. Let $\pi_1, \pi_2, \dots, \pi_k$ be an arbitrary enumeration of the elements of Ω. We partition N^1_Ω into the sets

$$N^1_{\pi_j} := \{i \in N^1_\Omega \mid rep_i(\pi_j) = 1 \land \forall \ell < j\ rep_i(\pi_\ell) < 1\},$$

i.e., $N^1_{\pi_j}$ is the set of voters for which π_j is the minimal extension 1-representing them. Note that $\sum_{j=1}^k |N^1_{\pi_j}| = |N^1_\Omega| \leq n - n/k < n$. Thus, there is $\ell \leq k$ such that $|N^1_{\pi_\ell}| < n/k \leq |N'|$. Observe that

$$|N^1_\Omega| = \sum_{j=1}^k |N^1_{\pi_j}| < \left(\sum_{j=1}^k |N^1_{\pi_j}|\right) - |N^1_{\pi_\ell}| + |N'| \leq |N^1_{\Omega \cup \{\pi'\} \setminus \{\pi_\ell\}}|$$

where the last inequality holds because no voter in N' is represented by Ω. This contradicts the assumption that Ω was chosen by MaxCov. The case for GreedCov is contained in the full paper. □

In contrast, rules that do not explicitly maximize the number of 1-represented agents do not satisfy JR (see full paper for the proof).

Proposition 13. *No OWA or greedy OWA rule is guaranteed to return an outcome that satisfies JR, irrespective of the tie-breaking used, and even if votes are admissible.*

Theorem 12 and Proposition 13 together show that MaxCov (resp. GreedCov) does not belong to the class of OWA rules (resp. greedy OWA rules) and that MaxCov outperforms all OWA rules with respect to the key axiom of justified representation.

7 EXPERIMENTS

In this section, we present our experiments. We test our rules on a concrete scenario, and find that, on our data, the Utilitarian and Harmonic rule achieve good representation, and that their greedy versions scale to practically relevant inputs sizes.

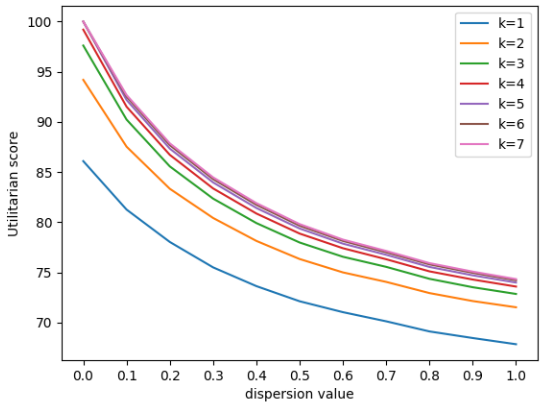

(a) Varying k experiment.

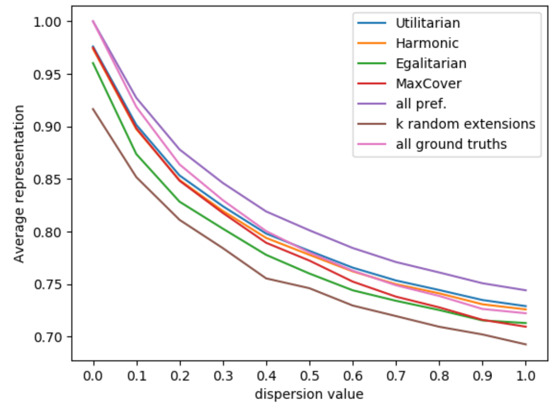

(b) Average representation.

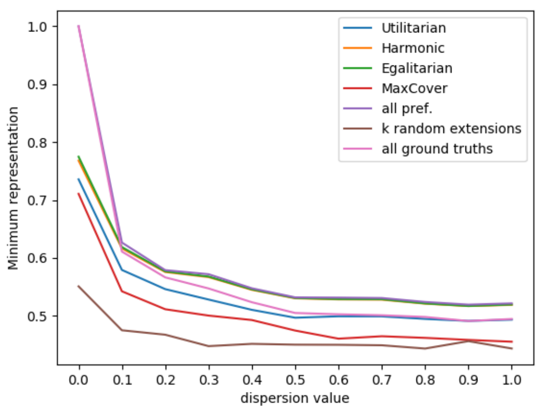

(c) Minimal representation.

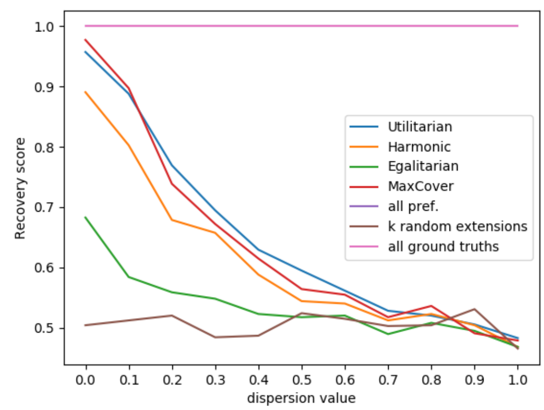

(d) Recovery score.

Figure 5: Results of the experiments.

Setup. We implemented our framework in Python. To compute the preferred extensions, we use `ASPARTIX` [29]. We ran our experiments on a device equipped with two Intel Xeon E5-2650 v4 (12 cores @ 2.20GHz) CPUs and 256GB of RAM. Since there is no dataset of ABSAFs, we generate our data; we will now sketch our generation method. See the full paper for details.

To generate AFs, we use `AFBenchGen2` [15] with the Barabási-Albert sampling algorithm [4], which produces scale-free graphs. To generate votes, we use a mixed Mallows model [42, 49]. Here, there are some central votes (*ground truths*), of which the ballots are noisy signals. Given a ground truth S, the probability of sampling a vote S' is proportional to $\Phi^{d(S,S')}$, where d is the distance measure $d(S, S') = |S'| - |S' \cap S|$ and $\Phi \in [0, 1]$ is the *dispersion parameter*.

We generate 50 frameworks with 10 preferred extensions each. For each AF F, we sample 5 ground truths from $prf(F)$, and for each ground truth, we sample 20 ballots centered around it. We run the experiments for $\Phi \in \{0, 0.1, \ldots, 0.9, 1\}$, and for each AF and value of Φ, we generate 5 ABSAFs, for a total of $50 \cdot 11 \cdot 5 = 2750$ ABSAFs.

We ran all experiments for both the base definition of representation (cf. Definition 5) and core-representation (cf. Definition 7). Here, we show results for the former; the others are comparable.[5]

[5] All the plots omitted in this section can be found in the full paper.

In our first experiment, given a rule, we compute the score of its chosen outcome over all generated AFs and for $k \in \{1, \ldots, 7\}$. For each value of k, we average the results over each value of Φ. Since there are 5 ground truths, our expectation (for low dispersion) is that, for $k = 5$, our method should be able to represent the voters well. By considering $k < 5$ we investigate how robust our approach is (i.e., whether the score dips) when selecting outcomes smaller than the number of ground truths. Conversely, by looking at $k > 5$, we check whether we gain any performance by outcomes larger than the ground truth. The results of this experiment are shown in Figure 5a. We ran this for Utilitarian, Harmonic and Egalitarian. Here, we show results for Utilitarian (the others are comparable).

Next, we run the Utilitarian, Egalitarian, Harmonic and MaxCov rules on our data with $k = 3$. We look at the following metrics, all of which we average over the different dispersion parameters:

- The average and minimal representation score (over all voters) given by the selected outcome;
- The "recovery score" of an outcome, computed as $r(\Omega)/k$, where $r(\Omega)$ is the number of ground truths included in Ω. Intuitively, this captures how well a rule can recover the ground truths around which the voters' preferences center;
- The approximation ratio of the greedy variant of each rule.

Table 1: Performance experiment.

	non-greedy		greedy	
# pref. ext.	*successes*	*time*	*successes*	*time*
$\{8, \ldots, 12\}$	30/30	0.03	30/30	0.01
$\{13, \ldots, 17\}$	30/30	0.73	30/30	0.02
$\{18, \ldots, 22\}$	30/30	5.63	30/30	0.03
$\{23, \ldots, 27\}$	1/30	17.24	30/30	0.06
$\{28, \ldots, 32\}$	0/30	-	30/30	0.12

We report the results in Figures 5b, 5c and 5d. For the greedy approximation experiment, we do not show the full plot here, but only the minimum of the (average) approximation ratio (across all dispersion parameters): for Utilitarian, Harmonic, Egalitarian and MaxCov we have a minimum value of 99.26%, 96.61%, 89.41% and 98.75%, respectively. Observe that we ran all the above also for $k \in \{2, 5\}$; the results are comparable to what we show.

Finally, we compare the runtime of the greedy and non-greedy rules. We generate AFs with increasing values for $m = |prf(F)|$, starting with 30 AFs with $8 \leq m \leq 12$, then 30 AFs with $13 \leq m \leq 17$, etc. For each AF we sample approximately 100 voters, centered around $\lfloor m/4 \rfloor$ ground truths.[6] We run the Utilitarian and greedy-Utilitarian rules on each input, looking for outcomes of size $k = \lfloor m/4 \rfloor$, with a timeout of 45 seconds. We increase m until no instance out of 30 is solved within the timeout. For the instances that terminate before timeout, we report the average runtime. The results are shown in Table 1. We also ran the greedy algorithm on larger AFs, to assess its limits. On a framework with 256 extensions, we could find an outcome of size $k = 64$ in 22 seconds. With 512 extensions and $k = 128$, we exceed the 45-seconds timeout.

Discussion. Let us now analyze the results. Looking at Figure 5a we see the following: (1) As Φ grows, performance drops at the same rate for all k; (2) Our method seems quite robust: if we select outcomes slightly smaller than the true number of ground truths (i.e., $k \in \{3, 4\}$), we do not observe substantial performance drops; (3) Conversely, if we select outcomes larger than the number of ground truths ($k > 5$), we do not gain significant advantages.

Next, looking at Figures 5b, 5c and 5d, we note that all rules outperform the k-random-extensions baseline. This is promising, although this baseline achieves quite high performance in some metrics (Figure 5b). Next, although the "all preferred extensions" baseline is (as expected) always best performing, the performance of "all ground truths" decreases, being even outperformed by some rules. This is reasonable: as Φ increases, votes become less similar to the ground truths. Moreover, we can see from Figures 5b and 5c that Harmonic seems to be a good compromise between Utilitarian and Egalitarian. In particular, Harmonic performs almost as well as Egalitarian in the minimum representation metric. Furthermore, we can see that MaxCov performs poorly for minimum representation, and its performance decreases starkly as Φ increases for average representation. Finally, in Figure 5d, Utilitarian and MaxCov seem the best performing rules; however, Harmonic performs comparably.

[6]The exact number of voters depends on the value of $\lfloor m/4 \rfloor$, and is never less than 96. Regardless, the number of voters has a limited impact on the runtime.

In general, despite the axiomatic advantages of MaxCov (Theorem 12), in our setup, Utilitarian seems to perform better w.r.t. most metrics.[7] Depending on the application scenario, we can recommend the Utilitarian rule, with the Harmonic rule being a sensible alternative, compromising between efficiency and fairness.

Finally, we notice that the greedy versions of our rules offer a good approximation ratio, always well above 90% (with the exception of Egalitarian). Moreover, from Table 1 we can see, as hinted by our theoretical findings (Proposition 9), that the greedy algorithm drastically improves over the performance of the non-greedy one, and seems to scale quite well. These findings combined suggest that the greedy algorithms are a scalable and well-performing method to apply our framework in real-world scenarios.

8 CONCLUSION

Summary. We presented a new framework for representing different viewpoints in online discussions by combining approval voting and abstract argumentation. In this framework, citizens can both propose arguments and vote on them. We then use argumentation theory to find the maximal, consistent and defendable points of view. For smaller instances, we can then efficiently find a set that represents the *defendable cores* of all voter's ballots. For larger instances, we propose different methods for picking a small set of most representative points of view and compared them axiomatically and in simulations. Axiomatically, MaxCov showed the best behavior and can be recommended when justified representation is desired. If justified representation is not necessary, Utilitarian and Harmonic can be recommended according to our experiments, with the former being more efficient and the latter being fairer.

Formalizing Discussions. In this paper, we assume discussions to be already formalized as AFs. In practice, when using real-world data, this formalization is a crucial step. In Example 2, we manually created the AF representing the canadian election reform discussion. However, while manual annotation is feasible for small discussions, a fair moderation is important, as to avoid any subjective biases. The formalization of larger debates, instead, may require argument mining techniques that can convert (often previously annotated) natural language text into AFs. While such methods are being worked on [13, 35, 45], there are still challenges to be overcome [7].

Future Work. Our framework could be extended by also allowing disapprovals, which are commonly seen in practice. Moreover, our general approach is independent of the choice of abstract argumentation for identifying consistent sets of comments. Thus, we plan to investigate the effect of using other mechanisms instead.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their feedback. This research has been supported by the Austrian Science Fund (FWF) (through projects P32830 and J4581), the Vienna Science and Technology Fund (WWTF) (through project ICT19-065), and by the European Union's Horizon 2020 research and innovation programme (under grant agreement 101034440).

[7]Observe that we also tested for an additional metric, namely, the percentage of voters that are 1-represented by the selected outcome. Clearly, MaxCov was the best performing rule, but Utilitarian was essentially equally good. Therefore, even from the angle of perfect representation, in practice, Utilitarian seems comparable to MaxCov.

REFERENCES

- [1] Georgios Amanatidis, Nathanaël Barrot, Jérôme Lang, Evangelos Markakis, and Bernard Ries. 2015. Multiple Referenda and Multiwinner Elections Using Hamming Distances: Complexity and Manipulability. In *Proceedings of the 14th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2015)*. IFAAMAS.
- [2] Edmond Awad, Richard Booth, Fernando Tohmé, and Iyad Rahwan. 2017. Judgement aggregation in multi-agent argumentation. *J. Log. Comput.* 27, 1 (2017), 227–259.
- [3] Haris Aziz, Markus Brill, Vincent Conitzer, Edith Elkind, Rupert Freeman, and Toby Walsh. 2017. Justified Representation in Approval-Based Committee Voting. *Social Choice and Welfare* 48, 2 (2017), 461–485.
- [4] Albert-László Barabási and Réka Albert. 1999. Emergence of Scaling in Random Networks. *Science* 286, 5439 (1999), 509–512.
- [5] Pietro Baroni, Martin Caminada, and Massimiliano Giacomin. 2018. Abstract Argumentation Frameworks and Their Semantics. In *Handbook of Formal Argumentation*. College Publications, Chapter 4, 159–236.
- [6] Nathanaël Barrot, Jérôme Lang, and Makoto Yokoo. 2017. Manipulation of Hamming-based approval voting for multiple referenda and committee elections. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2017)*. IFAAMAS.
- [7] Ringo Baumann, Gregor Wiedemann, Maximilian Heinrich, Ahmad Dawar Hakimi, and Gerhard Heyer. 2020. The Road Map to FAME: A Framework for Mining and Formal Evaluation of Arguments. *Datenbank-Spektrum* 20, 2 (2020), 107–113.
- [8] Dorothea Baumeister, Daniel Neugebauer, and Jörg Rothe. 2021. Collective Acceptability in Abstract Argumentation. *FLAP* 8, 6 (2021), 1503–1542.
- [9] Michael Bernreiter, Jan Maly, Oliviero Nardi, and Stefan Woltran. 2024. Combining Voting and Abstract Argumentation to Understand Online Discussions. arXiv:2402.05895 [cs.GT]
- [10] Michael Bernreiter, Jan Maly, Oliviero Nardi, and Stefan Woltran. 2024. *Supplementary Material for "Combining Voting and Abstract Argumentation to Understand Online Discussions"*. https://doi.org/10.5281/zenodo.10571624
- [11] Stefano Bistarelli and Francesco Santini. 2021. Weighted Argumentation. *FLAP* 8, 6 (2021), 1589–1622.
- [12] Samuel R. Buss and Louise Hay. 1991. On Truth-Table Reducibility to SAT. *Inf. Comput.* 91, 1 (1991), 86–102.
- [13] Elena Cabrio and Serena Villata. 2012. Generating Abstract Arguments: A Natural Language Approach. In *Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, September 10-12, 2012 (Frontiers in Artificial Intelligence and Applications, Vol. 245)*. IOS Press, 454–461.
- [14] Martin Caminada and Gabriella Pigozzi. 2011. On Judgment Aggregation in Abstract Argumentation. *Autonomous Agents and Multi-Agent Systems* 22 (2011), 64–102.
- [15] Federico Cerutti, Massimiliano Giacomin, and Mauro Vallati. 2016. Generating Structured Argumentation Frameworks: AFBenchGen2. In *Computational Models of Argument - Proceedings of COMMA 2016, Potsdam, Germany, 12-16 September, 2016 (Frontiers in Artificial Intelligence and Applications, Vol. 287)*, Pietro Baroni, Thomas F. Gordon, Tatjana Scheffler, and Manfred Stede (Eds.). IOS Press, 467–468.
- [16] John R Chamberlin and Paul N Courant. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77, 3 (1983), 718–733.
- [17] Weiwei Chen and Ulle Endriss. 2019. Preservation of Semantic Properties in Collective Argumentation: The Case of Aggregating Abstract Argumentation Frameworks. *Artificial Intelligence* 269 (2019), 27–48.
- [18] Joshua Cohen. 2007. *Deliberative Democracy*. Palgrave Macmillan UK, London, 219–236. https://doi.org/10.1057/9780230591080_10
- [19] Sylvie Coste-Marquis, Caroline Devred, Sébastien Konieczny, Marie-Christine Lagasquie-Schiex, and Pierre Marquis. 2007. On the Merging of Dung's Argumentation Systems. *Artificial Intelligence* 171, 10 (2007), 730–753. Argumentation in Artificial Intelligence.
- [20] Sylvie Coste-Marquis, Caroline Devred, and Pierre Marquis. 2005. Symmetric Argumentation Frameworks. In *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 8th European Conference, ECSQARU 2005, Barcelona, Spain, July 6-8, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3571)*, Lluís Godo (Ed.). Springer, 317–328.
- [21] Nadia Creignou, Reinhard Pichler, and Stefan Woltran. 2018. Do Hard SAT-Related Reasoning Tasks Become Easier in the Krom Fragment? *Log. Methods Comput. Sci.* 14, 4 (2018).
- [22] Jérôme Delobelle, Adrian Haret, Sébastien Konieczny, Jean-Guy Mailly, Julien Rossit, and Stefan Woltran. 2016. Merging of Abstract Argumentation Frameworks. In *Proceedings of the 15th International Conference on Principles of Knowledge Representation and Reasoning (KR-2016)*. AAAI Press.
- [23] Phan Minh Dung. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artif. Intell.* 77, 2 (1995), 321–358.
- [24] Paul E. Dunne. 2007. Computational properties of argument systems satisfying graph-theoretic constraints. *Artif. Intell.* 171, 10-15 (2007), 701–729.
- [25] Paul E. Dunne, Wolfgang Dvorák, Thomas Linsbichler, and Stefan Woltran. 2015. Characteristics of multiple viewpoints in abstract argumentation. *Artif. Intell.* 228 (2015), 153–178.
- [26] Paul E. Dunne, Anthony Hunter, Peter McBurney, Simon Parsons, and Michael Wooldridge. 2011. Weighted Argument Systems: Basic Definitions, Algorithms, and Complexity Results. *Artificial Intelligence* 175, 2 (2011), 457–486.
- [27] Paul E. Dunne, Pierre Marquis, and Michael Wooldridge. 2012. Argument Aggregation: Basic Axioms and Complexity Results. In *Proceedings of the 4th International Conference on Computational Models of Argument (COMMA-2012)*. IOS Press, 129–140.
- [28] Wolfgang Dvořák and Paul E. Dunne. 2018. Computational Problems in Formal Argumentation and their Complexity. In *Handbook of Formal Argumentation*, Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre (Eds.). College Publications, Chapter 14, 631–687. also appears in IfCoLog Journal of Logics and their Applications 4(8):2557–2622.
- [29] Uwe Egly, Sarah A. Gaggl, and Stefan Woltran. 2008. Aspartix: Implementing Argumentation Frameworks Using Answer-Set Programming. In *Proceedings of the 24th International Conference on Logic Programming (ICLP-2008)*. Springer, 734–738.
- [30] Ulle Endriss. 2016. Judgment Aggregation. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 17.
- [31] Sara Fish, Paul Gölz, David C. Parkes, Ariel D. Procaccia, Gili Rusak, Itai Shapira, and Manuel Wüthrich. 2023. Generative Social Choice. arXiv:2309.01291 [cs.GT]
- [32] Roslyn Fuller. 2022. *SDI Digital Democracy Report*. Solonian Democracy Institute.
- [33] Jordi Ganzer, Natalia Criado, Maite Lopez-Sanchez, Simon Parsons, and Juan A. Rodriguez-Aguilar. 2023. A Model to Support Collective Reasoning: Formalization, Analysis and Computational Assessment. *Journal of Artificial Intelligence Research* 77 (2023), 1021–1086.
- [34] Serge Gaspers and Ray Li. 2019. Enumeration of Preferred Extensions in Almost Oriented Digraphs. In *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany (LIPIcs, Vol. 138)*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 74:1–74:15.
- [35] Pierpaolo Goffredo, Elena Cabrio, Serena Villata, Shohreh Haddadan, and Jhonatan Torres Sanchez. 2023. DISPUTool 2.0: A Modular Architecture for Multi-Layer Argumentative Analysis of Political Debates. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*. AAAI Press, 16431–16433.
- [36] Daniel Halpern, Gregory Kehne, Ariel D. Procaccia, Jamie Tucker-Foltz, and Manuel Wüthrich. 2023. Representation with Incomplete Votes. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*. AAAI Press.
- [37] Dorit S. Hochbaum. 1997. Approximation Algorithms for NP-Hard Problems. *SIGACT News* 28, 2 (1997), 40–52.
- [38] Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA (The IBM Research Symposia Series)*. Plenum Press, New York, 85–103.
- [39] Martin Lackner, Jan Maly, and Oliviero Nardi. 2023. Free-Riding in Multi-Issue Decisions. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2023)*. IFAAMAS.
- [40] Martin Lackner and Piotr Skowron. 2023. *Multi-Winner Voting with Approval Preferences*. Springer International Publishing. https://doi.org/10.1007/978-3-031-09016-5
- [41] João Leite and João G. Martins. 2011. Social Abstract Argumentation. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*, Toby Walsh (Ed.). IJCAI/AAAI, 2287–2292.
- [42] Colin L. Mallows. 1957. Non-Null Ranking Models. *Biometrica* 44, 1 (1957), 114–130.
- [43] Andreas Niskanen and Matti Järvisalo. 2020. μ-toksia: An Efficient Abstract Argumentation Reasoner. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12-18, 2020*. 800–804.
- [44] OECD. 2020. *Innovative citizen participation and new democratic institutions: Catching the deliberative wave*. OECD Publishing.
- [45] Raquel Mochales Palau and Marie-Francine Moens. 2009. Argumentation mining: the detection, classification and structure of arguments in text. In *The 12th International Conference on Artificial Intelligence and Law, Proceedings of the Conference, June 8-12, 2009, Barcelona, Spain*. ACM, 98–107.
- [46] Dominik Peters and Piotr Skowron. 2020. Proportionality and the Limits of Welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC-2020)*. 793–794.
- [47] Simon Rey and Jan Maly. 2023. The (Computational) Social Choice Take on Indivisible Participatory Budgeting. arXiv:2303.00621 [cs.GT]

[48] Christopher T. Small, Ivan Vendrov, Esin Durmus, Hadjar Homaei, Elizabeth Barry, Julien Cornebise, Ted Suzman, Deep Ganguli, and Colin Megill. 2023. Opportunities and Risks of LLMs for Scalable Deliberation with Polis. arXiv:2306.11932 [cs.SI]

[49] Stanisław Szufa, Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. 2020. Drawing a Map of Elections in the Space of Statistical Cultures. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-2020)*. IFAAMAS.

[50] Fernando A. Tohmé, Gustavo A. Bodanza, and Guillermo R. Simari. 2008. Aggregation of Attack Relations: A Social-Choice Theoretical Analysis of Defeasibility Criteria. In *Proceedings of the 5th International Symposium on Foundations of Information and Knowledge Systems (FoIKS-2008)*. Springer, 8–23.

[51] Klaus W. Wagner. 1990. Bounded Query Classes. *SIAM J. Comput.* 19, 5 (1990), 833–846.

[52] Ronald R. Yager. 1988. On Ordered Weighted Averaging Aggregation Operators in Multicriteria Decisionmaking. *IEEE Transactions on systems, Man, and Cybernetics* 18, 1 (1988), 183–190.