Working Together: Committee Selection and the Supermodular Degree

Rani Izsak
Department of Computer Science and Applied Mathematics
Weizmann Institute of Science
Rehovot, Israel
ran.izsak@weizmann.ac.il

ABSTRACT

We introduce a voting rule for committee selection that captures positive correlation (synergy) between candidates. We argue that positive correlation can naturally happen in common scenarios that are related to committee selection. For example, in the movies selection problem, where prospective travelers are requested to choose the movies that will be available on their flight, it is reasonable to assume that they will tend to prefer voting for a movie in a series, only if they can watch also the former movies in that series. In elections to the parliament, it can be that two candidates are working extremely well together, so voters will benefit from being represented by both of them together.

In our model, the preferences of the candidates are represented by set functions, and we would like to maximize the total satisfaction of the voters. We show that although computing the best solution is $\mathcal{NP}$-hard, there exists an approximation algorithm with approximation guarantees that deteriorate gracefully with the amount of synergy between the candidates. This amount of synergy is measured by a natural extension of the supermodular degree [Feige and Izsak, ITCS 2013] that we introduce – the joint supermodular degree.

1. INTRODUCTION

Consider the following scenario (see, e.g., [9, 21]). An airline wishes to increase the satisfaction of the travelers by letting them choose the set of movies that will be available on their flight. It is decided to store on the airplane some fixed number k of movies. The airline surveys the preferences of the prospective passengers of the flight, and aims to make the best decision given their preferences. Two questions arise. First, how should the preferences of the prospective travelers be modeled? Second, given the preferences of the travelers, how should the set of movies be chosen? This problem of choosing some fixed number of candidates to the satisfaction of the voters is a fundamental problem. Generally speaking, in the k-COMMITEE SELECTION problem, we have a set V of n voters and a set C of m candidates, and we would like to select k candidates out of the m, such that the voters will be most satisfied. The answers to the two questions above vary in the literature. For example, by the Chamberlin-Courant rule we have a value for each of the candidates, by each of the voters, and the satisfaction of a voter is measured by the highest value she has for any elected candidate. The overall satisfaction is either the sum of the values of the voters or the value of the least satisfied voter (utilitarian [5] or egalitarian [2] variant, respectively). Other possibilities are to aggregate for every voter her value for every elected candidate or to give higher weight for candidates ranked higher by her (e.g. Borda rule). In a recent work, Skowron, Faliszewski and Lang [21] introduce an elegant model that captures the latter examples as well as others. They model the preferences of each voter by an intrinsic value for each of the candidates. Then, they calculate the value of a possible set of k candidates by a voter, by ordering her k intrinsic values for the k candidates, and multiplying them by some weight that corresponds to their rank in the order. This vector of weights is called "OWA operator" (Ordered Weighted Average). Skowron, Faliszewski and Lang [21] study their model for different restrictions on the OWA vector. Among their results, they show a $(1 - 1/e)$-approximation algorithm for the case of non-increasing weights OWA vectors, by showing it is captured by submodular set functions.[1]

However, none of the models above capture positive correlation (i.e. synergy) between specific candidates (see Section 2.2 for further discussion). Positive correlation can happen in various cases: from two candidates to the parliament that are working great together (see Woolley et al. [22] for a research about collective intelligence), to a series of movies that people tend to prefer watching the latter parts only after watching the former parts. In this paper we suggest a voting rule that captures positive correlation between specific candidates. Specifically, our answers to the two questions above are:

- The preferences of each of the candidates are modeled by a non-decreasing monotone set function from subsets of candidates to non-negative real numbers.
- A set of k candidates that maximizes the sum of values of the voters is elected.

We formally present our model in Section 4. In order to measure the amount of synergy between different candi-

Appears at: 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE 2017). Held as part of the Workshops at the 16th International Conference on Autonomous Agents and Multiagent Systems. May 8th-9th, 2017. Sao Paulo, Brazil.

[1] A submodular set function is a function $f : 2^M \to \mathbb{R}^+$, such that for every $S' \subseteq S \subseteq M$, and every $j \in M$, $f(j \mid S') \geq f(j \mid S)$, where $f(j \mid S) = f(\{j\} \cup S) - f(S)$ is the marginal value of j with respect to S. That is, the marginal values are monotone non-increasing

dates, we extend the supermodular degree [10], by introducing the joint supermodular degree (Section 4.1).

We also study applications for the model. In Section 5, we justify the naturalness of the joint supermodular degree from an applicative view point. In Section 6, we demonstrate how preference elicitation can be practically done.

Finally, in Section 7, we study the computability of our voting rule. On the bright side, we show that although computing the optimum is, generally, $\mathcal{NP}$-hard, one can approximate the optimum with a guarantee that depends on the amount of synergy between different candidates, as measured by the joint supermodular degree. On the flip side, we show that the same results cannot be achieved for the supermodular degree.

2. PRELIMINARIES

The definitions below are taken from the works [10, 11]. Let C be a set of items (e.g. candidates in election, movies to watch on an airplane) and let $f : 2^C \rightarrow \mathbb{R}^+$ be a set function (e.g. of preferences of one of the voters). The following definition is standard.

Definition 1. *Let $c \in C$. The marginal set function $f_c : 2^{C\setminus\{c\}} \rightarrow \mathbb{R}^+$ is a function mapping each subset $S \subseteq C \setminus \{c\}$ to the marginal value of c given S:*

$$f_c(S) \stackrel{def}{=} f(S \cup \{c\}) - f(S) \ .$$

We denote the marginal value $f_c(S)$ by $f(c \mid S)$. For $S' = \{c_1, \dots, c_{|S'|}\} \subseteq C$ and $S \subseteq C \setminus S'$ we also use either of the notations $f(c_1, \dots, c_{|S'|} \mid S)$ or $f(S' \mid S)$ to indicate $f(S \cup S') - f(S)$.

The following definitions were introduced by Feige and Izsak [10].

Definition 2. *Let $c \in C$. The supermodular dependency set of c by f is the set of all items $c' \in C$ such that there exists $S \subseteq C \setminus \{c, c'\}$ such that $f(c \mid S \cup \{c'\}) > f(c \mid S)$. We denote the supermodular dependency set of c by $\mathcal{D}^+_f(c)$. We sometimes omit f, when it is clear from the context.*

Definition 3. *The supermodular degree of f is defined as $\mathcal{D}^+_f \stackrel{def}{=} \max_{c\in C} |\mathcal{D}^+_f(c)|$.*

2.1 Representation of set functions

Let $f : 2^C \rightarrow \mathbb{R}^+$ be a set function. Then, f associates values to $2^{|C|}$ possible subsets. If we want our algorithms to run in time polynomial in $|C|$, they, of course, cannot read an input that is exponential in C. Therefore, it is crucial to consider the representation of set functions. One common way to represent set functions is by queries. Another is by an explicit representation. In this section, we mention both.

Queries

The arguably simplest queries are the following.

Definition 4. *Value queries for f are defined as follows:*
Input: A subset $S \subseteq C$.
Output: $f(S)$.

That is, if we assume our algorithm has access to value queries for a given set function, we merely assume it can ask for the value of a subset by the function. Another type of queries that we use (see [10]) is the following.

Definition 5. *Supermodular queries for f are defined as follows:*
Input: An item (i.e. a candidate) $c \subseteq C$.
Output: $\mathcal{D}^+_f(c)$.

That is, given a candidate we can ask with whom she has a positive correlation as defined by the supermodular dependencies. In the context of movies, we can ask for a movie that is part of a series, what are the other movies in that series. See Section 5 for further discussion.

An explicit representation

Another way to represent set functions is by an explicit representation. For example, any set function can be represented in a unique way by a hypergraph with weighted edges (see [1, 6, 8]). In this representation, a vertex is introduced for each of the items in the ground set of f. The weights in the sub-hypergraph induced by a set of vertices sum up exactly to the value of the subset with the respective items, by f. To see how weights can be allocated, consider the following iterative process. To hyperedges of size 1, we allocate weights that are the values of the respective singleton subsets. Note that this allocation of weights to hyperedges is unique. Then, for hyperedges of size 2, we allocate weights that are the difference between the value of the respective subset and the sum of the weights of their two singleton subsets. Note that this allocation is unique, as well. Also note that after iteration ℓ, the values by the hypergraph representation are correct for subsets of size up to ℓ. We proceed iteratively till we arrive to the unique edge of size $|C|$, and then we have a representation of the set function for any size of subset.

A succinct representation.

We say that a representation of a set function is succinct if its size is polynomially bounded by the size of the ground set of the function. Note that in the hypergraph representation, we can list only the edges of value different from 0. So, sometimes this representation can be succinct. In particular, for additive set functions we clearly allocate non-zero values only for the hyperedges of size 1.

2.2 Related work

We list here some of the voting rules from the literature, mostly based on Masthoff [16], and also on the works [7, 9, 15, 16, 17, 21].

- Plurality: When electing a single candidate, plurality means selecting the candidate who is ranked first among the candidates, for the highest number of voters. When "ranked first" can mean that by the voting rule, preferences are ranks of candidates, or alternatively, that there are values for the candidates by the different voters that are used in order to get the candidates' ranks. In order to use this rule for choosing k candidates, one can just repeat it k times, while removing the winner at each iteration.

- Utilitarian: Each voter has a value for each of the candidates, and these values are summed up. The k candidates with the largest sums win.

- Borda [4]: This voting rule assumes the preferences of the candidates are modeled as a list of ranks, and it converts this list to values, with higher values for higher ranks: $m-1, m-2, \ldots, 0$ for ranks $1, \ldots, m$, respectively (m is the number of candidates). These values are summed up and highest scores win, similarly to the utilitarian rule above.

- Copeland: The score of a candidate is the number of pairwise elections she wins (by plurality) minus the number of pairwise elections she loses (ties do not count). Values are again, summed up, and higher scores win.

- Maximin: The score of a candidate c with respect to a candidate c' is the number of voters that prefer c over c' (we denote it by $\text{score}_c(c')$. The score of a candidate c is the minimum score of c with respect to a candidate (i.e. $\text{argmin}_{c'} \text{score}_c(c')$). For example, if for a candidate c, there exists a candidate that is preferred by all of the voters, then c will get a value of 0. If for a candidate c, *all* the voters prefer it over *all* the candidates, then (and only then) she will get the maximal score of n (i.e. the number of voters).

- Approval voting: Each voter either approves or disapproves every candidate. The k candidates with largest number of approvals win.

Positional scoring.

Positional scoring is a bunch of voting rules, where the preferences of the voters are just an ordering of the candidates and the rule is defined by a vector of size m of values corresponding to positions by the voters. The total value of a candidate is the sum of these values of the voters. Note that plurality is a positional scoring rule with the vector $(1, 0, \ldots, 0)$ and Borda is a positional scoring rule with the vector $(m-1, m-2, \ldots, 0)$. There is also a rule called "Veto" where the vector is $(1, \ldots, 1, 0)$, so a voter actually chooses one candidate she prefers *not* to include in the selected committee.

Weighted aggregation of preferences of a voter.

Skowron, Faliszewski and Lang [21] introduced the following family of voting rules for choosing k out of m candidates. The preferences of the voters are intrinsic values for the different candidates, and additionally, there is a vector of size k that is called OWA (ordered weighted average). When calculating the value for a set of k candidates by the preferences of a single voter, we do the following. We order the k candidates by their values according to the voter, in an increasing order of values, and then we sum up the values multiplied by the OWA vector (inner product). That is, every value is multiplied by a weight appearing in the OWA vector that corresponds to the rank of the candidate by the voter. To calculate the overall value of a subset of k candidates, we sum up the values of this set of candidates by the voters (utilitarian model). Skowron, Faliszewski and Lang [21] show that when the OWA vector is non-increasing (that is higher ranked candidates by a voter are multiplied by higher (or equal) weights), then the preferences of the voters can be represented by a submodular set function, and therefore a $(1-1/e)$-approximation guarantee can be achieved in polynomial time, by using the classical algorithm of Fisher, Nemhauser and Wolsey [14]. When the OWA vector is not non-increasing, some positive correlation between the candidates can happen, but not between specific candidates. For example, in the min OWA vector $(0, \ldots, 0, 1)$, only the worst candidate in the selected committee counts. This means, roughly speaking, that all the candidates should be adequate by a voter in order to have an adequate score by her. In terms of set functions, it means as follows. The marginal value of a candidate is 0 with respect to any committee that contains a worse (or equal) candidate. The marginal value of a candidate with respect to a committee that contains only better candidates is the difference between the intrinsic values of the new candidate and of the worst candidate in the committee. For example, adding a candidate with an intrinsic value of 1 to a committee, when the worse candidate in it has an intrinsic value of 10 means a marginal value of (-9). On the other hand, if there is also a candidate with an intrinsic value of 2 in the committee, then the marginal value of the new one will be (-1). That is, the marginal value of the new candidate increased because of the inclusion of the candidate with a value of 2. However, it is clear that this does not model synergy between these two candidates. Moreover, positive correlation between *specific* candidates cannot be modeled using OWA vectors, as described above, since they cannot relate to specific candidates differently. This means that in scenarios like the movies example described earlier, a positive correlation within a series of movies cannot be modeled. In this sense, our model adds new possibilities with respect to the model of Skowron, Faliszewski and Lang [21].

Another relevant model was studied by Fishburn and Pekec [13]. Fishburn and Pekec [13] studied an approval voting model, where each of the voters can approve a few candidates, and a committee is approved by a voter if it contains a sufficient number of candidates that are approved by the voter.

3. OUR CONTRIBUTION

This paper introduces a new model for voting rules, based on set functions, together with the required conceptual framework. Our model can be used to model both synergy between candidates (i.e. compliments) and substitutes (e.g., two candidates that each of them is worth 1 and both of them together are worth 1, as well). Since general set functions might be highly complex, we introduce the joint supermodular degree, which we see as a natural extension of the supermodular degree [10]. We demonstrate applications for our model in Section 5. In particular, we suggest practical preference elicitation that is tailored for the joint supermodular degree in Section 6.

Finally, in Section 7, we show how the joint supermodular degree enables one to easily use existing algorithms for function maximization that are tailored for the supermodular degree to achieve approximations for our voting rule. Since there exist such algorithms both for offline and online settings, one can use either and immediately get approximation guarantees for our voting rule in the corresponding setting. Moreover, future algorithms for the supermodular degree can also be easily used by our framework, to get computational results for committee selection. Conceptually speaking, the result of the approximation algorithms can also be seen as the voting rule itself (see Skowron, Faliszewski and Lang [21]). We complement our algorithmic result with a proof of computational hardness.

To the best of our knowledge, our results represent the first voting rules that capture synergy between specific candidates.

4. THE MODEL

We formally define our model. Let $V = \{v_1, \ldots, v_n\}$ be a set of n voters, let C be a set of m candidates and let k be an integer. Let $f_1, \ldots, f_n : 2^C \to \mathbb{R}^+$ be preference (set) functions, associated with the voters $v_1, \ldots, v_n$, respectively. We assume that the preferences functions are normalized (i.e., $\forall_i f_i(\varnothing) = 0$) and non-decreasing monotone (i.e., $\forall_{i, S' \subseteq S \subseteq M} f_i(S') \leq f_i(S)$). Our aim is to choose a set $C_{max} \subseteq C$ of size k that maximizes the satisfaction of the voters by their personal preferences:

$$C_{max} = \underset{S \subseteq C \mid |S| = k}{\operatorname{argmax}} \sum_{i=1}^{n} f_i(S) \ .$$

We refer to this problem as (the) k-COMMITEE SELECTION problem and to the selected subset as the selected committee. Note that this problem can be seen as a voting rule. Alternatively, an approximation algorithm to this problem can be seen as the voting rule (see also Skowron, Faliszewski and Lang [21]).

4.1 The joint supermodular degree

We introduce the following natural extensions of the definitions of Feige and Izsak [10] to a collection of set functions.

DEFINITION 6. *Let $f_1, \ldots, f_t$ be set functions for some $t \in \mathbb{N}$ and let $c \in C$. The joint supermodular dependency set of c by $f_1, \ldots, f_t$ is $\bigcup_{i=1}^{t} \mathcal{D}^+_{f_i}(c)$.*

DEFINITION 7. *The joint supermodular degree of $f_1, \ldots, f_t$ is the maximum cardinality among the cardinalities of joint dependency sets of items of C by $f_1, \ldots, f_t$.*

The main property of the joint supermodular degree that we use is that the sum function of functions with joint supermodular degree of at most d has supermodular degree of at most d.

We think this definition is natural for voting rules, since it means that positive correlation between the candidates can be modeled, when it is inherent to the candidates themselves, and not to the perspective of the voters about them.

For example, if a candidate is working well together with 2 other candidates, then each of the voters has the possibility to give these 3 candidates or any subset of them a score that is higher than the sum of their individual scores. However, if a candidate does not work well with some other candidate, then none of the voters has the possibility to give them together a score that is higher than the sum of their individual scores. That is, the set of other candidates that the candidate has synergy with depends on her. The decision of whether to take this into account depends on each of the voters. So, the supermodular dependency set of a candidate c, by any of the preference functions of the voters, will contain only other candidates that have synergy (i.e. are working well together) with c.

We discuss applications of our model with respect to the joint supermodular degree in Section 5. In particular, we suggest preference elicitation in Section 6.

5. APPLICATIONS

We discuss in this section applications of our model, together with the joint supermodular degree. Specifically, we demonstrate its merits for two real world examples (see [9]).

- Parliamentary elections: In voting to the parliament, it is possible that candidates complement each other, and work better together. It was actually shown by Woolley et al. [22] that there is a measure for the collective intelligence of a group of people that is different from the intelligence quantities of different people in the group. So, it seems reasonable to allow the voters to give extra value for choosing *together* a pair of candidates that are known to work well together on, e.g., suggesting complex laws in the parliament. Note that the fact that two candidates are working well together is related to the candidates and not to the voters, and indeed, the joint supermodular degree of the voters will reflect the synergies between the candidates.
- Movie selection: Consider the problem of choosing k movies to be available on an airplane (passengers can watch on their flight movies from the selected set). It seems reasonable that people would prefer to watch latter parts of a series only after the former. Moreover, it might be unreasonable to consider a series of movies as one movie, if, e.g., physical storage is a limitation. Then, it is plausible to give the prospective passengers the possibility to give higher values for movies in the series, given that all the former are selected, as well. Additionally, movie selection can admit submodular behaviour (i.e. substitutes). For example, since the time of the flight is bounded, the number of movies one can watch out of the k selected movies is bounded, as well. This means that, if for example, $k = 100$ and the time of the flight allows one passenger to watch up to 5 movies, then any movie out of the k that is not among the 5 best for that passenger is redundant for her. So her value will not increase given that we add to the selected set other great movies. On the other hand, we do want to allow k to be large enough to allow different passengers to enjoy different movies. The latter behaviour is submodular. Synergy between selected movies is supermodular. Our model enables one to express such preferences. Furthermore, submodularity does not hurt the approximation guarantees, since it does not increase the joint supermodular degree of the preference functions (see Section 7).

6. PREFERENCE ELICITATION

Consider the movies selection example. When a prospective passenger is asked to express her preferences about possible movies, it seems unreasonable to require her to specify her values for all the exponentially many possibilities. We briefly demonstrate a simple user interface to elicit users' preferences in that case, while enabling them to benefit from the possibility of expressing positive correlations.

The user interface will be as follows. Each of the prospective passengers will be able to give a value for each of the possible movies (these are the values of the singleton subsets). In addition, the prospective passengers will be able to add for each of the movies other values – the marginal values of a movie, with respect to a subset of its joint supermodular dependency set (i.e., other movies in the same series). In order to select such a subset of the movies, a list of the movies in the joint supermodular dependency set will be presented, and a passenger will be able to select the relevant movies (e.g. by checking them by a 'V'). In order to enforce the preference functions of the prospective passengers to be well defined (i.e. a single value for each of the subsets), we will let the prospective passengers check by a 'V' only the movies that were former to a movie in a series.

Note that the supermodular dependency is symmetric (see [10] for a proof). So, in a series of movies, also the former movies are dependent on the latter movies. As an example, one can think of two movies, where each of them is worth 1, but the second one is worth 10 with respect to the first. Then, both movies together are worth 11, and the marginal contribution of each of them with respect to the other is 10, instead of 1 (as it is with respect to the empty set).

Generally speaking, this example interface can be extended in any way that enforces the preference functions to be well defined (e.g. by ordering the items and letting the prospective passengers to check a dependency by 'V' only if it is before the current item in that ordering).

To see the power of combining supermodular dependencies with submodular behaviour, note that we can also ask each passenger how many movies she would like to watch in her flight (with a maximum that depends on the duration of the flight), and then calculate as her preference, the best subset of that number of movies, from any input subset of movies.

Note that it is easy to emulate both value and supermodular queries using such a representation, and then to use the algorithms of Feldman and Izsak [11], as described in Section 7.

7. COMPUTATIONAL RESULTS

The following theorem shows that there exists an approximation algorithm with approximation guarantee that is linear in the amount of synergy between the candidates, as measured by the joint supermodular degree of the preference functions of the voters. For submodular set functions, the result described by the theorem coincides with the optimal result for submodular set functions of Fisher, Nemhauser and Wolsey [14] that is used by Skowron, Faliszewski and Lang [21].

Theorem 1. *When the joint supermodular degree of the preferences functions of the voters is d, the k-committee selection problem admits an approximation algorithm with guarantee $(1 - e^{-1/(d+1)}) \geq 1/(d+2)$. The algorithm gets access to the preference functions by value queries and supermodular queries, and its running time is $\texttt{Poly}(n, m, 2^d)$.*

Note that the above result captures the example of movies selection from the introduction (see Section 5 for further discussion). Note also that the proof of the above result applies to the case of committee selection subject to a *general* matroid constraint (cardinality constraint is a special case of a matroid constraint), but with an approximation guarantee of $1/(d + 2)$, by using the respective algorithm of Feldman and Izsak [11].

Moreover, one can use the algorithms of Feldman and Izsak [12] in order to get an online (secretary like) version of Theorem 1, when the candidates arrive one by one in an online fashion, and we need to decide on the spot, irrevocably, whether to elect a candidate or not, based on the preferences of the voters (for exact details of the model, see [12]). As an example, consider hiring a team to a project, where each of the candidates meets with a few interviewers. Then, an optimal team of candidates should be hired, according to the preferences of the interviewers.

By using the algorithm of Feldman and Izsak [12] for a cardinality constraint, one gets an approximation guarantee polynomial in the joint supermodular degree. Any approximation guarantee that depends only on the joint supermodular degree gives a constant approximation guarantee, if the candidates admit synergy only with a constant number of other candidates (e.g. if there is a positive correlation only within series of movies, and all the series suggested are of length up to 3). See also Oren and Lucier [18] for a different secretary like model.

Additionally, we show a hardness result for the case of non-bounded joint supermodular degree, even when the supermodular degree of all the set functions is bounded by 1. For this, we use a reduction from the k-dense subgraph problem (see e.g. Bhaskara el al. [3]).

Definition 8. *The k-dense subgraph problem is the following. We are given as input a graph $G = (V, E)$ and an integer $k \in \mathbb{N}$, and our aim is to select k vertices such that the number of edges in their induced subgraph is maximized.*

This problem is NP-hard and it is highly believed it is hard to approximate it within any constant guarantee. Actually, no efficient algorithm is currently known that approximates it within a guarantee better than n^c, for some constant c (see e.g. [3, 19, 20]).

Theorem 2. *The k-commitee selection problem is at least as hard as the k-dense subgraph problem, even if the supermodular degree of the set functions is 1, and even if an explicit representation of the preference functions is given. This means, in particular, that it is $\mathcal{NP}$-hard[2] and $\mathcal{SSE}$-hard (see [19] and also [20]).*

Proof of Theorem 1. Let V be the set of n voters, let C be the set of m candidates, let k be the requested number of elected candidates and let $f_1, \ldots, f_n : 2^C \to \mathbb{R}^+$ be the preference functions of the voters. We prove that since the joint supermodular degree of $f_1, \ldots, f_n$ is upper bounded by d, then the supermodular degree of their summation function $f_\Sigma(S) \stackrel{\text{def}}{=} \sum_{i=1}^n f_i(S)$ is upper bounded by d, as well. Note that this would not be necessarily true if only the supermodular degree of $f_1, \ldots, f_n$ was bounded by d (or even by 1). Actually, Theorem 2 serves as a counter example to the latter for $d = 1$.

To prove the bound on the supermodular degree of the summation function f_Σ, we show that every supermodular dependency by f_Σ induces the same supermodular dependency by one of the f_is in the sum. Let $c, c' \in C$ and $S \subseteq C$ be such that $f_\Sigma(c \mid S \cup \{c'\}) > f_\Sigma(c \mid S)$. Then, by the definition of f_Σ, $\sum_{i=1}^n f_i(c \mid S \cup \{c'\}) > \sum_{i=1}^n f_i(c \mid S)$. So, $\exists_{1 \leq i \leq n}$ s.t. $f_i(c \mid S \cup \{c'\}) > f_i(c \mid S)$, as claimed.

Now, we can just use the algorithm of [11] for monotone function maximization subject to uniform matroid constraint (i.e. cardinality constraint) on the function f_Σ with a constraint k. Note that the latter algorithm gives an optimal approximation guarantee for submodular set functions, and generally its guarantee deteriorates linearly with the supermodular degree. Moreover, its running time is as required by the Theorem. This concludes the proof of Theorem 1. □

[2] NP-hardness is actually true also for submodular set functions, i.e. supermodular degree of 0.

PROOF OF THEOREM 2. The proof is somewhat similar to the proof of $\mathcal{SSE}$-hardness for maximizing set function subject to cardinality constraint, given by [11]. Given an algorithm for solving the k-commitee selection problem within approximation guarantee α, we show how to solve any input instance of the k-dense subgraph problem within approximation guarantee α. Let $G = (S, E)$ be an instance of the k-dense graph problem. Then, our set of candidates C will be S (the set of vertices of G). We also introduce a voter v_e for every edge $e = \{v_{e1}, v_{e2}\} \in E$ and let $V = \bigcup_{e \in E}\{v_e\}$. For every voter v_e, her preference set function is:

$$f_e = \begin{cases} 1 & \text{if } v_{e_1} \text{ and } v_{e_2} \text{ are both elected.} \\ 0 & \text{otherwise} \end{cases}$$

That is, in this instance of the k-committee selection problem, our aim is to find a subset of k candidates (where the set of candidates corresponds exactly to the set S of vertices of G), such that the number of pairs of candidates, that correspond to the preference functions of the voters, is maximized (where these pairs of candidates are exactly the edges E of G). This is exactly the k-dense subgraph problem. That is, given a solution to this instance of k-committee selection problem, we just output the subset of vertices of S that corresponds to the candidates in C that were selected, as a solution to the input instance of the k-dense subgraph problem. This gives us a feasible solution with the same value, and thus with the same approximation guarantee α. This concludes the proof of Theorem 2. □

8. CONCLUSIONS

We suggest a new voting rule for committee selection that enables the voters to express positive correlation between the candidates. We also introduce the joint supermodular degree that enables us to use existing computational results for the supermodular degree, and get efficient approximation algorithms for our voting rule. We see our work as a proof of concept, and hope that it will lead to further study of committee selection with positive correlation between the candidates.

Acknowledgments

Work supported in part by the Israel Science Foundation (grant No. 1388/16). I would like to thank Uri Feige for many useful discussions and for his contributions to this paper. I would like to thank Nimrod Talmon for useful discussions and for directing me to the paper of Skowron, Faliszewski and Lang [21]. I would also like to thank Moshe Babaioff, Shahar Dobzinski and Moshe Tennenholtz for useful discussions.

REFERENCES

[1] I. Abraham, M. Babaioff, S. Dughmi, and T. Roughgarden. Combinatorial auctions with restricted complements. In *EC*, pages 3–16, New York, NY, USA, 2012. ACM.

[2] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Artificial Intelligence Research*, 47:475–519, 2013.

[3] A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan. Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k-subgraph. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 201–210, 2010.

[4] J. Borda. Memmoire sur les elections au scrutine. *Histoire de l'Academie Royale des Sciences*, 1781.

[5] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the borda rule. *American Political Science Review*, 77:718–733, 1983.

[6] Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Multiagent resource allocation in k-additive domains: preference representation and complexity. *Annals of Operations Research*, 163:49–62, 2008.

[7] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *J. ACM*, 54(3), 2007.

[8] V. Conitzer, T. Sandholm, and P. Santi. Combinatorial auctions with k-wise dependent valuations. In *AAAI*, pages 248–254, 2005.

[9] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. In *AAMAS*, pages 53–60, 2014.

[10] U. Feige and R. Izsak. Welfare maximization and the supermodular degree. In *ITCS*, pages 247–256, 2013.

[11] M. Feldman and R. Izsak. Constrained monotone function maximization and the supermodular degree. In *APPROX-RANDOM*, pages 160–175, 2014.

[12] M. Feldman and R. Izsak. Building a good team: Secretary problems and the supermodular degree. In *SODA*, 2017.

[13] P. Fishburn and A. Pekec. Approval voting for committees: Threshold approaches. 2017. working paper.

[14] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey. An analysis of approximations for maximizing submodular set functions – II. In *Polyhedral Combinatorics*, volume 8 of *Mathematical Programming Study*, pages 73–87. North-Holland Publishing Company, 1978.

[15] J. Goldsmith, J. Lang, N. Mattei, and P. Perny. Voting with rank dependent scoring rules. In *AAAI*, pages 698–704, 2014.

[16] J. Masthoff. Group modeling: Selecting a sequence of television items to suit a group of viewers. *User Model. User-Adapt. Interact.*, 14(1):37–85, 2004.

[17] J. Masthoff. Group recommender systems: Combining individual models. In *Recommender Systems Handbook*, pages 677–702. 2011.

[18] J. Oren and B. Lucier. Online (budgeted) social choice. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, July 27 -31, 2014, Québec City, Québec, Canada.*, pages 1456–1462, 2014.

[19] P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. In *STOC*, pages 755–764, 2010.

[20] P. Raghavendra, D. Steurer, and M. Tulsiani. Reductions between expansion problems. In *IEEE Conference on Computational Complexity*, pages 64–73, 2012.

[21] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional

multirepresentation to group recommendation. In *AAAI*, pages 2131–2137, 2015.

[22] A. W. Woolley, C. F. Chabris, A. Pentland, N. Hashmi, and T. W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330:265–294, 2010.