On the Approximability of Dodgson and Young Elections

Ioannis Caragiannis* Jason A. Covey† Michal Feldman‡ Christopher M. Homan§
Christos Kaklamanis¶ Nikos Karanikolas‖ Ariel D. Procaccia** Jeffrey S. Rosenschein††

Abstract

The voting rules proposed by Dodgson and Young are both designed to find the alternative closest to being a Condorcet winner, according to two different notions of proximity; the score of a given alternative is known to be hard to compute under either rule.

In this paper, we put forward two algorithms for approximating the Dodgson score: an LP-based randomized rounding algorithm and a deterministic greedy algorithm, both of which yield an $\mathcal{O}(\log m)$ approximation ratio, where m is the number of alternatives; we observe that this result is asymptotically optimal, and further prove that our greedy algorithm is optimal up to a factor of 2, unless problems in $\mathcal{NP}$ have quasi-polynomial time algorithms. Although the greedy algorithm is computationally superior, we argue that the randomized rounding algorithm has an advantage from a social choice point of view.

Further, we demonstrate that computing any reasonable approximation of the ranking produced by Dodgson's rule is $\mathcal{NP}$-hard. This result provides a complexity-theoretic explanation of sharp discrepancies that have been observed in the Social Choice Theory literature when comparing Dodgson elections with simpler voting rules.

Finally, we show that the problem of calculating the Young score is $\mathcal{NP}$-hard to approximate by any factor. This leads to an inapproximability result for the Young ranking.

*Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece, email: caragian@ceid.upatras.gr

†Department of Computer Science, Rochester Institute of Technology, 102 Lomb Memorial Drive, Rochester, NY 14623-5603, email: jac8687@rit.edu

‡School of Business Administration and Center for the Study of Rationality, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: mfeldman@huji.ac.il

§Department of Computer Science, Rochester Institute of Technology, 102 Lomb Memorial Drive Rochester, NY 14623-5603, email: cmh@cs.rit.edu

¶Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece, email: kakl@ceid.upatras.gr

‖Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics University of Patras, 26500 Rio, Greece, email: nkaranik@ceid.upatras.gr

**Microsoft Israel R&D Center, Herzeliya 46725, Israel, email: arielpro@gmail.com

††School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel, email: jeff@cs.huji.ac.il

1 Introduction

The discipline of voting theory deals with the following setting: a group of n agents each ranks a set of m alternatives; one alternative is to be elected. The big question is: which alternative best reflects the social good? The French philosopher and mathematician Marie Jean Antoine Nicolas de Caritat, marquis de Condorcet, suggested the following intuitive criterion: the winner should be an alternative that beats every other alternative in a *pairwise election*, i.e., an alternative that a majority of the agents prefers over any other alternative. Sadly, it is fairly easy to see that the preferences of the majority may be cyclic, hence a *Condorcet winner* does not necessarily exist. This unfortunate phenomenon is known as the *Condorcet paradox* (see Black [5]).

In order to circumvent this result, several researchers have proposed to choose an alternative that is "as close as possible" to a Condorcet winner. Different notions of proximity can be considered, and yield different voting rules. One such notion was suggested in 1876 by Charles Dodgson, better known by his pen name Lewis Caroll, author of "Alice's Adventures in Wonderland". The *Dodgson score* [5] of an alternative, with respect to a given set of agents' preferences, is the minimum number of exchanges between adjacent alternatives in the agents' rankings one has to introduce in order to make the given alternative a Condorcet winner. A *Dodgson winner* is any alternative with a minimum Dodgson score.

Young [35] raised a second option: measuring the distance by agents. Specifically, the *Young score* of an alternative is the size of the largest subset of agents such that, if only these ballots are taken into account, the given alternative becomes a Condorcet winner. A *Young winner* is any alternative with the maximum Young score. Alternatively, one can perceive a Young winner as the alternative that becomes a Condorcet winner by removing the least number of agents.

Though these two voting rules sound appealing and straightforward, they are notoriously complicated to resolve. As early as 1989, Bartholdi, Tovey and Trick [2] showed that computing the Dodgson score is

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

$\mathcal{NP}$-complete, and that pinpointing a Dodgson winner is $\mathcal{NP}$-hard. This important paper was one of the first to introduce complexity-theoretic considerations to social choice theory. Hemaspaandra et al. [14] refined the abovementioned result by showing that the Dodgson winner problem is complete for Θ_2^p, the class of problems that can be solved by $\mathcal{O}(\log n)$ queries to an $\mathcal{NP}$ set. Subsequently, Rothe et al. [31] proved that the Young winner problem is also complete for Θ_2^p.

The abovementioned complexity results give rise to the agenda of *approximately* calculating an alternative's score, under the Dodgson and Young schemes. This is clearly an interesting computational problem, as an application area of algorithmic techniques.

However, from the point of view of social choice theory, it is not immediately apparent that an approximation of a voting rule is satisfactory, since an "incorrect" alternative—in our case, one that is not closest to a Condorcet winner—might be elected. Nevertheless, we argue that the use of such an approximation is strongly motivated. Indeed, at least in the case of the Dodgson and Young rules, the winner is an "approximation" in the first place, in instances where no Condorcet winner exists. Moreover, the approximation algorithm is equivalent to a new voting rule, which is guaranteed to elect an alternative that is not far from being a Condorcet winner. In other words, a perfectly sensible definition of a "socially good" winner, given the circumstances, is simply the alternative chosen by the approximation algorithm. Note that the approximation algorithm can be designed to satisfy the Condorcet criterion, i.e., always elect a Condorcet winner if one exists. This is always true for an approximation of the Dodgson score, as the Dodgson score of a Condorcet winner is zero. Moreover, approximation algorithms can be designed to satisfy other, less trivial, social choice desiderata, and hence may ultimately be considered socially sensible voting rules.

Related work. The agenda of approximating voting rules was recently pursued by Ailon et al. [1], Coppersmith et al. [9], and Kenyon-Mathieu and Schudy [17]. These works deal, directly or indirectly, with the Kemeny rank aggregation rule, which chooses a ranking of the alternatives instead of a single winning alternative. The Kemeny rule picks the ranking that has the maximum number of agreements with the agents' individual rankings regarding the correct order of pairs of alternatives. Ailon et al. improve the trivial 2-approximation algorithm to an involved, randomized algorithm that gives an 11/7-approximation; Kenyon-Mathieu and Schudy further improve the approximation, and obtain a PTAS.

Two recent works have directly put forward algorithms for the Dodgson winner problem [15, 23]. Both papers independently build upon the same basic idea: if the number of agents is significantly larger than the number of alternatives, and one looks at a uniform distribution over the preferences of the agents, with high probability one obtains an instance on which it is trivial to compute the Dodgson score of a given alternative. This directly gives rise to an algorithm that can usually compute the Dodgson score (under the assumption on the number of agents and alternatives). However, this is not an approximation algorithm in the usual sense, since the algorithm *a priori* gives up on certain instances, whereas an approximation algorithm is judged by its worst-case guarantees. In addition, this algorithm would be useless if the number of alternatives is not small compared with the number of agents.[1]

Betzler et al. [4] have investigated the parameterized computational complexity of the Dodgson and Young rules. The authors have devised a fixed parameter algorithm for exact computation of the Dodgson score, where the fixed parameter is the "edit distance," i.e., the number of exchanges. Specifically, if k is an upper bound on the Dodgson score of a given alternative, n is the number of agents, and m the number of alternatives, the algorithm runs in time $\mathcal{O}(2^k \cdot nk + nm)$. Notice that in general it may hold that $k = \Omega(nm)$. In contrast, computing the Young score is $W[2]$-complete; this implies that there is no algorithm that computes the Young score exactly, and whose running time is polynomial in nm and only exponential in k, where the parameter k is the number of remaining votes. These results complement ours nicely, as we shall also demonstrate that computing the Dodgson score is in a sense easier than computing the Young score, albeit in the context of approximation.

Putting computational complexity aside, several works by social choice theorists have considered comparing the ranking produced by Dodgson, i.e., the ordering of the alternatives by nondecreasing Dodgson score, with elections based on simpler voting rules. Such comparisons have always revealed sharp discrepancies. For example, the Dodgson winner can appear in any position in the Kemeny ranking [28] and in the ranking of any positional scoring rule [29] (e.g., Borda or Plurality), Dodgson rankings can be exactly the opposite of Borda [20] and Copeland rankings [18], while the winner

[1] This would normally not happen in political elections, but can certainly be the case in many other settings. For instance, consider a group of agents trying to reach an agreement on a joint plan, when multiple alternative plans are available. Specifically, think of a group of investors deciding which company to invest in.

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

of Kemeny of Slater elections can appear in any position of the Dodgson ranking [19].

More distantly related to our work is research that is concerned with exactly resolving hard-to-compute voting rules by heuristic methods. Typical examples include works regarding the Kemeny rule [8] and the Slater rule [7]. Another more remotely related field of research is concerned with finding approximate, efficient representations of voting rules, by eliciting as little information as possible; this line of research employs techniques from learning theory [25, 26].

Our results. In the context of approximating the Dodgson score, we devise an $\mathcal{O}(\log m)$ randomized approximation algorithm, where m is the number of alternatives. Our algorithm is based on solving the linear program proposed by Bartholdi et al. [2] and using randomized rounding. We then propose a second, deterministic and greedy, algorithm for the Dodgson score, with the same asymptotic approximation ratio. Although the latter algorithm is computationally superior in every way, we show that the former has the advantage of satisfying a flavor of monotonicity, which is a desirable property from a social choice point of view. We further observe that it follows from the work of McCabe-Dansted [22] that the Dodgson score cannot be approximated within sublogarithmic factors by polynomial-time algorithms unless $\mathcal{P} = \mathcal{NP}$. We prove a more explicit inapproximability result of $(1/2-\epsilon)\ln m$, under the assumption that problems in $\mathcal{NP}$ do not have algorithms running in quasi-polynomial time; this implies that the approximation ratio achieved by our greedy algorithm is optimal up to a factor of 2.

Some of the results mentioned above [28, 29, 18, 19, 20] establish that there are sharp discrepancies between the Dodgson ranking and the rankings produced by other rank aggregation rules. Some of these rules (e.g., Borda and Copeland) are polynomial-time computable, so the corresponding results can be viewed as negative results regarding the approximability of the Dodgson ranking by polynomial-time algorithms. We show that the problem of distinguishing between whether a given alternative is the unique Dodgson winner or in the last $O(\sqrt{m})$ positions in any Dodgson ranking is $\mathcal{NP}$-hard. This theorem provides a complexity-theoretic explanation for some of the observed discrepancies, but in fact is much wider in scope as it applies to any efficiently computable rank aggregation rule.

The problem of calculating the Young score seems at first glance simple compared with the Dodgson score (we discuss in Section 4 why this seems so). Therefore, we found the following result quite surprising: it is $\mathcal{NP}$-hard to approximate the Young score within any factor. Specifically, we show that it is $\mathcal{NP}$-hard to distinguish between the case where the Young score of a given alternative is 0, and the case where the score is greater than 0. As a corollary we obtain an inapproximability result for the Young ranking.

Structure of the paper. In Section 2, we introduce some notations and definitions. In Section 3, we present our upper and lower bounds for approximating Dodgson elections. In Section 4, we prove that the Young score and ranking are inapproximable.

2 Preliminaries

Let $N = \{1, \ldots, n\}$ be a set of agents, and let A be the set of *alternatives*. We denote $|A| = m$, and denote the alternatives themselves by letters, such as $a \in A$. Indices referring to agents appear in superscript. Each agent $i \in N$ holds a binary relation R^i over A that satisfies irreflexivity, asymmetry, transitivity and totality. Informally, R^i is a ranking of the alternatives. Let $L = L(A)$ be the set of all rankings over A; we have that each $R^i \in L$. We denote $R^N = \langle R^1, \ldots, R^n \rangle \in L^N$, and refer to this vector as a *preference profile*. We may also use Q^i to denote the preferences of agent i, in cases where we want to distinguish between two different rankings R^i and Q^i. For sets of alternatives $B_1, B_2 \subseteq A$, we write $B_1 R^i B_2$ if for all $a \in B_1$ and $b \in B_2$, aR^ib.

Let $a, b \in A$. Denote $\{i \in N : aR^ib\}$ as $P(a, b)$. We say that *a beats b in a pairwise election* if $|P(a, b)| > n/2$, that is, a is preferred to b by the majority of agents. A *Condorcet winner* is an alternative that beats every other alternative in a pairwise election.

The *Dodgson score* of a given alternative a^*, with respect to a given preference profile R^N, is the least number of exchanges between adjacent alternatives in R^N needed to make a^* a Condorcet winner. For instance, let $N = \{1, 2, 3\}$, $A = \{a, b, c\}$, and let R^N be given by:

R^1	R^2	R^3
a	b	a
b	a	c
c	c	b

In this example, the Dodgson score of a is 0 (a is a Condorcet winner), the score of b is 1, and the score of c is 3. Bartholdi et al. [2] have shown that computing the Dodgson score is an $\mathcal{NP}$-complete problem.

The *Young score* of a^* with respect to R^N is the size of the largest subset of agents for whom a^* is a Condorcet winner. This is the definition given by Young himself [35], and used in subsequent works [31]. If for every nonempty subset of agents a^* is not a Condorcet winner, its Young score is 0. In the above example, the

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

Young score of a is 3, the score of b is 1, and the score of c is 0.

Notice that, equivalently, a Young winner is an alternative such that one has to remove the least number of agents in order to make it a Condorcet winner. However, these two definitions are not equivalent in the context of approximation; we employ the former (original, prevalent) definition, but touch on the latter as well.

As the Young winner problem is known to be intractable [31], the Young score problem must also be hard; otherwise, we would be able to calculate the scores of all the alternatives efficiently, and identify the alternatives with maximum score.

3 Approximability of Dodgson

We begin by presenting our approximation algorithms for the Dodgson score. Let us first introduce some common notations.

Let $a^* \in A$ be a distinguished alternative, whose Dodgson score we wish to compute. Define the *deficit* of a^* with respect to $a \in A$, simply denoted $\text{def}(a)$ when the identity of a^* is clear, as the number of additional agents that must rank a^* above a in order for a^* to beat a in a pairwise election. For instance, if 4 agents prefer a to a^* and only one agent prefers a^* to a, then $\text{def}(a) = 2$. If a^* beats a in a pairwise election (namely a^* is preferred by the majority of agents) then $\text{def}(a) = 0$. We say that alternatives $a \in A$ with $\text{def}(a) > 0$ are *alive*. Alternatives that are not alive, i.e., $\text{def}(a) = 0$, are *dead*.

3.1 A Randomized Rounding algorithm.

Bartholdi et al. [2] provide an integer linear programming (ILP) formulation for the Dodgson score. The number of constraints and variables in their program depends solely on the number of alternatives. Therefore, if the number of alternatives is constant, the program is solvable in polynomial time using the algorithm of Lenstra [21]. However, if the number of alternatives is not constant, the LP is of gargantuan size.[2]

Fortunately, it is easy to modify the abovementioned ILP to obtain a program of polynomial size. As before, let $a^* \in A$ be the alternative whose score we wish to compute. Let the variables of the program be $x^i_j \in \{0,1\}$ for all $i \in N$ and $j \in \{0,\ldots,m-1\}$; $x^i_j = 1$ if and only if a^* is moved upward, or *pushed*, by j positions in the ranking of agent i. Define constants $e^i_{ja} \in \{0,1\}$, for all $i \in N$, $j \in \{0,\ldots,m-1\}$, and $a \in A \setminus \{a^*\}$, which depend on the given preference profile; $e^i_{ja} = 1$ iff pushing a^* by j positions in the ranking of agent i makes a^* gain an *additional* vote against a (note that $e^i_{ja} = 0$ for all j if $a^* R^i a$). The ILP that computes the Dodgson score of a^* is given by:

$$
\begin{aligned}
&\text{minimize} && \sum_{i,j} j \cdot x^i_j \\
(3.1)\quad &\text{s.t.} && \forall i \in N,\ \sum_j x^i_j = 1 \\
& && \forall a \in A \setminus \{a^*\},\ \sum_{i,j} x^i_j e^i_{ja} \geq \text{def}(a) \\
& && \forall i \in N, \forall j \in \{0,\ldots,m-1\},\ x^i_j \in \{0,1\}
\end{aligned}
$$

This ILP can be relaxed by requiring merely that $0 \leq x^i_j \leq 1$ for all i and j. The resulting linear program (LP) can be solved efficiently.

We are now ready to present our randomized rounding algorithm.

Randomized Rounding Algorithm
Input: An alternative a^* whose Dodgson score we wish to estimate, and a preference profile $R^N \in L^N$.
Output: An approximation of the Dodgson score of a^*.
The algorithm:

1. Solve the relaxed LP given by (3.1) to obtain a solution $\vec{x}$.
2. For $k = 1,\ldots,\alpha \cdot \log m$ (where $\alpha > 0$ is a constant to be chosen later)
   - For all $i \in N$, randomly and independently (from other agents and other iterations) choose a value X^i_k, such that $X^i_k = j$ with probability x^i_j.
3. For all $i \in N$, set $X^i_{max} = \max_k X^i_k$.
4. Let $\mathcal{X}'$ be the solution which moves a^* upward in the ranking of i by X^i_{max} positions; return $\text{cost}(\mathcal{X}') = \sum_{i \in N} X^i_{max}$.

We remark that if a^* is a Condorcet winner from the outset, clearly the algorithm will calculate a score of 0 (with probability 1). Therefore, if we defined a new (randomized) voting rule, which elects the alternative with minimal score according to the algorithm, this voting rule would satisfy the Condorcet criterion.

THEOREM 3.1. *For any input a^* and R^N with m alternatives, the randomized rounding algorithm returns a $4\alpha \cdot \log m$-approximation of the Dodgson score of a^* with probability at least 1/2.*

[2] Note that there is also an efficient solution if the number of agents n is constant; indeed, brute force search requires checking $\mathcal{O}(m^n)$ possibilities.

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

The proof of the theorem is quite similar to the analysis of the randomized rounding algorithm for Set Cover [34, pp. 120-122], with one prominent additional argument, namely the application of Lemma 3.1.

Proof. Fix some iteration k of the algorithm's for loop. Let $X^i = X^i_k$, $i \in N$, be independent discrete random variables such that $X^i = j$ with probability x^i_j. Consider the sequence of exchanges induced by the variables X^i, i.e., each agent $i \in N$ moves a^* upward by j places with probability x^i_j. As a result of the constraint $\forall i \in N$, $\sum_j x^i_j = 1$, these are legal random variables. Moreover, let $\mathcal{X}$ be the chosen sequence of exchanges, and denote the optimal fractional solution of the LP by $\mathrm{OPT}_f = \sum_{i,j} j \cdot x^i_j$; it holds that

$$\mathbb{E}[\mathrm{cost}(\mathcal{X})] = \mathbb{E}\left[\sum_{i \in N} X^i\right] = \mathrm{OPT}_f \quad . \tag{3.2}$$

Now, fix some alternative $a \neq a^*$. We wish to bound the probability that a^* does not beat a after the exchanges given by $\mathcal{X}$ are made in R^N.

Let Y^i, $i \in N$, be independent Bernoulli trials, such that $Y^i = 1$ iff aR^ia^*, and a^* is moved above a in the preferences of agent i. In other words, $Y^i = 1$ if agent i becomes an additional agent that ranks a^* above a as a result of the exchanges. We want to provide an upper bound on $\Pr[\sum_{i \in N} Y^i < \mathrm{def}(a)]$. Denote

$$p^i = \sum_{j:\ e^i_{ja}=1} x^i_j \quad .$$

Notice that $Y^i = 1$ with probability p^i, so $\mathbb{E}[\sum_i Y^i] = \sum_i p^i$. Moreover, by the constraint $\forall a \in A \setminus \{a^*\}$, $\sum_{i,j} x^i_j e^i_{ja} \geq \mathrm{def}(a)$, we have that $\sum_i p^i \geq \mathrm{def}(a)$. We now employ a deceivingly intuitive but nontrivial result:

Lemma 3.1. (Jogdeo and Samuels [16]) *Let $Y^1, \ldots, Y^n$ be independent heterogeneous Bernoulli trials. Suppose that $\mathbb{E}[\sum_i Y^i]$ is an integer. Then*

$$\Pr\left[\sum_i Y^i < \mathbb{E}\left[\sum_i Y^i\right]\right] < 1/2 \quad .$$

Since $\mathrm{def}(a)$ is an integer, and $\mathbb{E}[\sum_i Y^i] = \sum_i p^i \geq \mathrm{def}(a)$, it follows from the lemma that:

$$\Pr[a \text{ not beaten in } \mathcal{X}] = \Pr\left[\sum_i Y^i < \mathrm{def}(a)\right] < 1/2 \quad .$$

At this point, we choose the value of the constant α to be such that $2^{\alpha \log m} \geq 4m$. Note that if $m \geq 4$, we can choose $\alpha \leq 2$. As in the algorithm, set $X^i_{max} = \max_k X^i_k$. Denote by $\mathcal{X}'$ the induced sequence of exchanges. It holds that a is not beaten in a pairwise election under $\mathcal{X}'$ only if a is not beaten under the exchanges obtained in each one of the $\alpha \cdot \log m$ individual iterations. Therefore,

$$\Pr[a \text{ not beaten in } \mathcal{X}'] < \left(\frac{1}{2}\right)^{\alpha \cdot \log m} \leq \frac{1}{4m} \quad .$$

By the union bound we get:[3]

$$\Pr[a^* \text{ is not a Condorcet winner in } \mathcal{X}'] \leq m \cdot \frac{1}{4m} = 1/4 \tag{3.3}$$

$X^i_1, \cdots, X^i_{\alpha \log m}$ are i.i.d. random variables; it holds that

$$X^i_{max} = \max_k X^i_k \leq \sum_k X^i_k \quad ,$$

and thus

$$\mathbb{E}\left[X^i_{max}\right] \leq \mathbb{E}\left[\sum_k X^i_k\right] = \alpha \cdot \log m \cdot \mathbb{E}[X^i_1] \quad . \tag{3.4}$$

Therefore, by the linearity of expectation,

$$\begin{aligned}
\mathbb{E}[\mathrm{cost}(\mathcal{X}')] &= \mathbb{E}\left[\sum_i X^i_{max}\right] \\
&\leq \alpha \cdot \log m \cdot \mathbb{E}\left[\sum_i X^i_1\right] \\
&= \alpha \cdot \log m \cdot \mathbb{E}[\mathrm{cost}(\mathcal{X})] \\
&= \alpha \cdot \log m \cdot \mathrm{OPT}_f \\
&\leq \alpha \cdot \log m \cdot \mathrm{OPT} \quad ,
\end{aligned}$$

where OPT is the Dodgson score of a^*, i.e., the optimal integral solution to the ILP (3.1).

By Markov's inequality we have that

$$\Pr[\mathrm{cost}(\mathcal{X}') > \mathrm{OPT} \cdot 4\alpha \cdot \log m] \leq 1/4 \quad . \tag{3.5}$$

We now apply the union bound once again on (3.3) and (3.5), and obtain that with probability at least 1/2, a^* is a Condorcet winner under $\mathcal{X}'$ and, at the same time, $\mathrm{cost}(\mathcal{X}') \leq \mathrm{OPT} \cdot 4 \cdot \alpha \cdot \log m$. This completes the proof of Theorem 3.1. □

Note that it is possible to verify in polynomial time whether the output of the algorithm is, at the same time, a valid solution (i.e., a^* is a Condorcet winner) and a $4\alpha \cdot \log m$-approximation (by comparing with OPT_f). Therefore, it is possible to repeat the algorithm from scratch to improve the probability of success.

[3] Strictly speaking, we can use $m - 1$ instead of m.

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

3.2 A Deterministic Combinatorial Algorithm. In this section, we present a deterministic, combinatorial, greedy algorithm for approximating the Dodgson score of a given alternative. Consider, once again, a special alternative a^*, and recall that a live alternative is one with a positive deficit. In each step, the algorithm selects the most cost-effective push of alternative a^* in the preference of some agent. The *cost-effectiveness* of pushing a^* in the preference of an agent $i \in N$ is the ratio between the total number of positions a^* is moved upwards in the preference of i *compared with the original profile* R^N, and the number of *currently* live alternatives that a^* overtakes as a result of this push. For example, if for some agent the algorithm raises a^* by one position where the alternative over which a^* is raised is dead, and later by a second position that causes a^* to overtake a live alternative, then the cost-effectiveness of the push is two and not one, since a^* ends up being two positions higher than its original position and only overtakes one live alternative.

After selecting the most cost-effective push, the algorithm decreases $\text{def}(a)$ by one for each live alternative a that a^* overtakes. Alternatives $a \in A$ with $\text{def}(a) = 0$ become dead. The algorithm terminates when no live alternatives remain. The input and output of the algorithm are as before.

Greedy Algorithm:

1. Let A' be the set of live alternatives, namely those alternatives $a \in A$ with $\text{def}(a) > 0$.
2. While $A' \neq \emptyset$:
   - Perform the most cost-effective push, namely push a^* in the preferences of agent $i \in N$ in a way that minimizes the ratio between the *total* number of positions moved upwards in the preferences of i and the number of currently live alternatives overtaken by a^*.
   - Recalculate A'.
3. Return the number of exchanges performed.

By the definition of the algorithm, it is clear that it produces a profile where a^* is a Condorcet winner. It is important to notice that, as is the case with the randomized rounding algorithm, if a^* is initially a Condorcet winner then the algorithm calculates a Dodgson score of zero, so as a voting rule the algorithm satisfies the Condorcet criterion.

THEOREM 3.2. *For any input a^* and R^N with m alternatives, the greedy algorithm returns an H_{m-1}-approximation of the Dodgson score of a^*, where H_k is the k-th harmonic number.*

We may view the problem of approximating the Dodgson score as the following covering problem with different covering requirements and constraints. The ground set is the set of live alternatives. For each live alternative $a \in A \setminus \{a^*\}$, its deficit $\text{def}(a)$ is in fact its covering requirement, i.e., the number of different sets it has to belong to in the final cover. For each agent $i \in N$ that ranks a^* in place r^i, we have a subcollection $\mathcal{S}^i$ consisting of the sets S^i_k for $k = 1, \ldots, r^i - 1$, where the set S^i_k contains the (initially) live alternatives that appear in positions $r^i - k$ to $r^i - 1$ in the preference of agent i. The set S^i_k has cost k. Now, the covering problem to be solved is the following. We wish to select at most one set from each of the different subcollections so that each alternative $a \in A \setminus \{a^*\}$ appears in at least $\text{def}(a)$ sets and the total cost of the selected sets is minimized. The optimal cost is the Dodgson score of a^* and, hence, the cost of any approximate cover that satisfies the covering requirements and the constraints is an upper bound on the Dodgson score.

In terms of this covering problem, the greedy algorithm mentioned above can be thought of as working as follows. In each step, it selects the most cost-effective set where the cost-effectiveness of a set is defined as the ratio between the cost of the set and the number of live alternatives it covers that have not been previously covered by sets belonging to the same subcollection. For these live alternatives, the algorithm decreases their covering requirements at the end of the step. The algorithm terminates when all alternatives have died (i.e., their covering requirement has become zero). The output of the algorithm consists of the maximum-cost sets that were picked from each subcollection.

The proof of Theorem 3.2, given in the full version of the paper [6], uses the dual fitting technique. We remark that the foregoing covering problem that we use (and its analysis) is closely related to the Constrained Set Multicover problem considered in Rajagopalan and Vazirani [27] (see also [34, pp. 112–116]), with the additional constraint that at most one set has to be selected from each subcollection.

3.3 Interlude: On the Desirability of Approximation Algorithms as Voting Rules. In Section 1 we stated that an approximation algorithm for the Dodgson score should be considered as a new voting rule. This implies that our approximation algorithms should be compared according to two conceptually different, but not orthogonal, dimensions: their algorithmic properties and their social choice properties. Our greedy algorithm is clearly superior to the randomized rounding algorithm in terms of algorithmic properties: the former is combinatorial whereas the latter is LP-

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

based; the former is deterministic whereas the latter is randomized. In the sequel we suggest, however, that the latter has some desirable properties from a social choice point of view. It is important to note at this point that randomized voting rules are considered legitimate in the social choice literature (see, e.g., [13, 10]), hence our randomized rounding algorithm may be considered a valid voting rule.

In most algorithmic mechanism design settings [24], such as combinatorial auctions or scheduling, one usually seeks approximation algorithms that are truthful, i.e., the agents cannot benefit by lying. However, the well-known Gibbard-Satterthwaite Theorem [12, 32] precludes voting rules that are both truthful and reasonable, in a sense. Therefore, other desiderata are looked for in voting rules.

We have been careful to emphasize that both the randomized rounding algorithm and the greedy algorithm satisfy the Condorcet property. Let us now consider the *monotonicity* property, one of the major desiderata on the basis of which voting rules are compared. Many different notions of monotonicity can be found in the literature; for our purposes, a (score-based) voting rule is *weakly monotonic* if and only if pushing an alternative in the preferences of the agents cannot worsen the score of the alternative, that is, increase it when a lower score is desirable (as in Dodgson), or decrease it when a higher score is desirable. All prominent score-based voting rules (positional scoring rules, Copeland, Maximin) are weakly monotonic; it is straightforward to see that the Dodgson and Young rules are weakly monotonic as well.

We first claim that our randomized rounding algorithm, or, more accurately, a slight variant thereof, is weakly monotonic. Indeed, consider the variant of the algorithm where $\mathcal{X}'$ is the solution that moves a^* upward in the ranking of i by $\sum_k X_k^i$ positions rather than $\max_k X_k^i$; the cost of this solution is

$$\text{cost}(\mathcal{X}') = \sum_k \sum_{i \in N} X_k^i \quad .$$

It is easy to verify (see (3.4)) that the exact same worst-case approximation bound holds for this variant as well (although in practice its approximation ratio would usually be significantly worse).

Now, consider a situation where a^* is moved upwards in the preferences of the agents. It is obvious that this decreases the value of OPT_f. In addition, for every k, we have $\mathbb{E}\left[\sum_i X_k^i\right] = \text{OPT}_f$. Therefore, by the linearity of expectation, the expected cost of the solution produced by the algorithm $\mathbb{E}\left[\sum_k \sum_{i \in N} X_k^i\right]$ decreases as well.

In contrast, let us now consider the greedy algorithm. We design a preference profile and a push of a^* that demonstrate that the algorithm is not weakly monotonic. Agents 1 through 6 vote according to the profile R^N given in Figure 1(a). The positions marked by "." are placeholders for the rest of the alternatives, in some arbitrary order. Let $A' = \{a_1, \ldots, a_4\}$, $A'' = \{b_1, \ldots, b_{17}\}$. Notice that $\text{def}(a) = 1$ for all $a \in A'$, $\text{def}(b) = 0$ for all $b \in A''$. The optimal sequence of exchanges moves a^* all the way to the top of the preferences of agent 2, with a cost of seven. The greedy algorithm, given this preference profile, indeed chooses this sequence.

On the other hand, consider the profile $(R^1, R^2, Q^3, Q^4, Q^5, Q^6)$ given in Figure 1(b) (where the position of a^* was improved by two positions in the preferences of agents 3 through 6). First notice that the deficits have not changed compared to the profile R^N. The greedy algorithm would in fact push a^* to the top of the preferences of agents 6, 5, 4, and 3 (in this order), with a total cost of ten. Note that the optimal solution still has a cost of seven.

The following stronger notion of monotonicity is often considered in the literature: pushing a winning alternative in the preferences of the agents cannot harm it, that is, cannot make it lose the election. We say that a voting rule that satisfies this property is *strongly monotonic*.[4] Interestingly, Dodgson itself is not strongly monotonic [33], a fact that is considered by many to be a serious flaw. However, this does not preclude the existence of an approximation algorithm for the Dodgson score that is strongly monotonic as a voting rule. An intriguing open question is the existence of such algorithms with a good approximation ratio.

Additionally, there are other prominent social choice properties that are often considered, e.g., *homogeneity*: a voting rule is said to be homogeneous if duplicating the electorate does not change the outcome of the election. We leave the comparison of our two algorithms on the basis of additional social choice desiderata, as well as more general questions regarding the design of socially desirable approximation algorithms, for future work.

3.4 Lower Bounds. McCabe-Dansted [22] gives a polynomial-time reduction from the Minimum Dominating Set problem to the Dodgson score problem with the following property: given a graph G with k vertices, the reduction creates a preference profile with $n = \Theta(k)$ agents and $m = \Theta(k^4)$ alternatives, such that the size of the minimum dominating set of G is $\lfloor k^{-2} \text{sc}_D(a^*) \rfloor$, where $\text{sc}_D(a^*)$ is the Dodgson score of a distinguished

[4] This is often simply referred to as *monotonic* in the literature

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

R^1	R^2	R^3	R^4	R^5	R^6
a_4	a_4				
a_3	a_3	a_4			
a_2	a_2	b_4	a_3		
a_1	a_1	b_5	b_9	a_2	
.	b_1	b_6	b_{10}	b_{13}	a_1
.	b_2	b_7	b_{11}	b_{14}	b_{16}
.	b_3	b_8	b_{12}	b_{15}	b_{17}
.	a^*	a^*	a^*	a^*	a^*
.	.	.	.	.	.
a^*	.	.	.	.	.

(a) Original Profile.

R^1	R^2	Q^3	Q^4	Q^5	Q^6
a_4	a_4				
a_3	a_3	a_4			
a_2	a_2	b_4	a_3		
a_1	a_1	b_5	b_9	a_2	
.	b_1	b_6	b_{10}	b_{13}	a_1
.	b_2	a^*	a^*	a^*	a^*
.	b_3	b_7	b_{11}	b_{14}	b_{16}
.	a^*	b_8	b_{12}	b_{15}	b_{17}
.	.	.	.	.	.
a^*	.	.	.	.	.

(b) Improvement of a^*.

Figure 1: The greedy algorithm is not weakly monotonic: an example.

alternative $a^* \in A$. Since the Minimum Dominating Set problem is known to be $\mathcal{NP}$-hard to approximate to within logarithmic factors [30], this implies that the Dodgson score problem is also hard to approximate to a factor of $\Omega(\log m)$. Due to the relation of Minimum Dominating Set to Minimum Set Cover, using an inapproximability result due to Feige [11], the explicit inapproximability bound can become $\left(\frac{1}{4} - \epsilon\right) \ln m$ under the assumption that problems in $\mathcal{NP}$ do not have quasi-polynomial-time algorithms.[5] This means that our algorithms are asymptotically optimal.

In the following, we present an alternative and more natural reduction directly from Minimum Set Cover that allows us to obtain a better explicit inapproximability bound. This bound implies that our greedy algorithm is optimal up to a factor of 2. The proof of the following theorem is given in the full version of the paper [6].

THEOREM 3.3. *There exists $\beta > 0$ such that it is $\mathcal{NP}$-hard to approximate the Dodgson score of a given alternative in an election with m alternatives to within a factor of $\beta \ln m$. Furthermore, for any $\epsilon > 0$, there is no polynomial-time $\left(\frac{1}{2} - \epsilon\right) \ln m$-approximation for the Dodgson score of a given alternative unless problems in $\mathcal{NP}$ of input size k have algorithms running in time $k^{\mathcal{O}(\log \log k)}$.*

A related question is the approximability of the Dodgson ranking, that is, the ranking of alternatives given by ordering them by nondecreasing Dodgson score. To the best of our knowledge, no rank aggregation function, which maps preferences profiles to rankings of the alternatives, is known to provably produce rankings that are close to the Dodgson ranking [28, 29, 18, 19, 20] (see the survey of related work in Section 1).

Our next result establishes that efficient approximation algorithms are unlikely to exist unless $\mathcal{P} = \mathcal{NP}$, by proving that the problem of distinguishing between whether a given alternative is the unique Dodgson winner or in the last $O(\sqrt{m})$ positions is $\mathcal{NP}$-hard.

THEOREM 3.4. *Given a preference profile with m alternatives and an alternative a^*, it is $\mathcal{NP}$-hard to decide whether a^* is a Dodgson winner or has rank at least $m - 6\sqrt{m}$ in any Dodgson ranking.*

Our proof, given in the full version of the paper [6], uses a reduction from Minimum Vertex Cover in 3-regular graphs and exploits a very weak statement concerning its inapproximability (marginally stronger than its $\mathcal{NP}$-hardness) that follows from the work of Berman and Karpinski [3]. The approach is similar to the proof of Theorem 3.3, albeit considerably more involved. This result provides a complexity-theoretic explanation for the sharp discrepancies observed in the Social Choice Theory literature when comparing Dodgson elections with simpler, efficiently computable, voting rules.

4 Approximability of Young

Recall that the Young score of a given alternative $a^* \in A$ is the size of the largest subset of agents for which a^* is a Condorcet winner.

It is straightforward to obtain a simple ILP for the Young score problem. As before, let $a^* \in A$ be the alternative whose Young score we wish to compute. Let the variables of the program be $x^i \in \{0, 1\}$ for all $i \in N$; $x^i = 1$ iff agent i is included in the subset of agents for a^*. Define constants $e_a^i \in \{-1, 1\}$ for all $i \in N$ and $a \in A \setminus \{a^*\}$, which depend on the given preference

[5] Both inapproximability bounds have not been explicitly observed by McCabe-Dansted.

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

profile; $e_a^i = 1$ iff agent i ranks a^* higher than a. The ILP that computes the Young score of a^* is given by:

$$\begin{aligned} &\text{maximize} && \sum_{i \in N} x^i \\ &\text{subject to} && \forall a \in A \setminus \{a^*\},\ \sum_{i \in N} x^i e_a^i \geq 1 \\ & && \forall i \in N,\ x^i \in \{0,1\} \end{aligned} \tag{4.6}$$

The ILP (4.6) for the Young score is seemingly simpler than the one for the Dodgson score, given as (3.1). This might seem to indicate that the problem can be easily approximated by similar techniques. Therefore, the following result is quite surprising.

THEOREM 4.1. *It is $\mathcal{NP}$-hard to approximate the Young score by any factor.*

This result becomes more self-evident when we notice that the Young score has the rare property of being nonmonotonic as an optimization problem, in the following sense: given a subset of agents that make a^* a Condorcet winner, it is not necessarily the case that a smaller subset of the agents would satisfy the same property. This stands in contrast to many approximable optimization problems, in which a solution which is worse than a valid solution is also a valid solution. Consider the Set Cover problem, for instance: if one adds more subsets to a valid cover, one obtains a valid cover. The same goes for the Dodgson score problem: if a sequence of exchanges makes a^* a Condorcet winner, introducing more exchanges (that push a^* upwards) on top of the existing ones would not undo this fact.

In order to prove the inapproximability of the Young score, we define the following problem.

NonEmptySubset
Instance: An alternative a^*, and a preference profile $R^N \in L^N$.
Question: Is there a nonempty subset of agents $C \subseteq N$, $C \neq \emptyset$, for which a^* is a Condorcet winner?

To prove Theorem 4.1, it is sufficient to prove that NonEmptySubset is $\mathcal{NP}$-hard. Indeed, this implies that it is $\mathcal{NP}$-hard to distinguish whether the Young score of a given alternative is zero or greater than zero, which directly entails that the score cannot be approximated.

LEMMA 4.1. NonEmptySubset *is $\mathcal{NP}$-complete.*

The proof of the lemma appears in the full version of the paper [6]. A short discussion is in order. Theorem 4.1 states that the Young score cannot be efficiently approximated to any factor. The proof shows that, in fact, it is impossible to efficiently distinguish between a zero and a nonzero score. However, the proof actually shows more: it constructs a family of instances, where it is hard to distinguish between a score of zero and almost $2m/3$. Now, if one looks at an alternative formulation of the Young score problem where all the scores are scaled by an additive constant, it is no longer true that it is hard to approximate the score to *any* factor; however, the proof still shows that it is hard to approximate the Young score, even under this alternative formulation, to a factor of $\Omega(m)$.

As noted in Section 2, one can imagine another alternative formulation of the Young score. Indeed, one might ask: given a preference profile, what is the *smallest* number of agents that must be *removed* in order to make a^* a Condorcet winner? This minimization problem, where the score is the number of agents that are removed, is referred to as the *Dual Young score* by Betzler et al. [4]. Of course, a Young winner according to the primal formulation is always a winner according to the dual formulation, and vice versa. Notice that it is easy to obtain an ϵn-approximation under the dual formulation for any constant $\epsilon > 0$ by enumerating all subsets of agents of size at least $n - 1/\epsilon$ and checking whether a^* is the Condorcet winner in the preferences of these agents. However, we conjecture that the dual Young score is hard to approximate significantly better; we leave this issue for future work.

Finally, the strong inapproximability result for the Young score intuitively implies that the Young ranking cannot be approximated. The following corollary, whose proof (given in the full version of the paper [6]) is a straightforward variation on the proof of Lemma 4.1, shows that this is indeed the case. It can be viewed as an analog of Theorem 3.4 for Young.

COROLLARY 4.1. *For any constant $\epsilon > 0$, given a preference profile with m alternatives and an alternative a^*, it is $\mathcal{NP}$-hard to decide whether a^* has rank $\mathcal{O}(m^\epsilon)$ or is ranked in place m (that is, ranked last) in any Young ranking.*

5 Acknowledgments

The authors wish to thank Felix Brandt, Felix Fischer, and Nati Linial for helpful discussions. The work was done while Procaccia was at the Hebrew University of Jerusalem, and was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. The work of Procaccia and Rosenschein was partially supported by Israel Science Foundation grant #898/05. The work of Caragiannis, Kaklamanis, and Karanikolas was partially supported by the

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.

European Union under IST FET Integrated Project AEOLUS "Algorithmic Principles for Building Efficient Overlay Computers" and COST Action IC0602 "Algorithmic Decision Theory", and by a "Caratheodory" basic research grant from the University of Patras.

References

[1] N. Ailon, M. Charikar, and A. Newman. Aggregating inconsistent information: ranking and clustering. In *STOC*, pages 684–693, 2005.

[2] J. Bartholdi, C. A. Tovey, and M. A. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6:157–165, 1989.

[3] P. Berman and M. Karpinski. On some tighter inapproximability results. In *ICALP*, volume 1644 of *LNCS*, pages 200–209. Springer, 1999.

[4] N. Betzler, J. Guo, and R. Niedermeier. Parameterized computational complexity of Dodgson and Young elections. In *SWAT*, 2008.

[5] D. Black. *Theory of Committees and Elections*. Cambridge University Press, 1958.

[6] I. Caragiannis, J. A. Covey, M. Feldman, C. M. Homan, C. Kaklamanis, N. Karanikolas, A. D. Procaccia, and J. S. Rosenschein. On the approximability of Dodgson and Young elections. Technical report 2008-80, Leibniz Center for Research in Computer Science, 2008. Available from: http://leibniz.cs.huji.ac.il/tr/1137.full.pdf .

[7] V. Conitzer. Computing Slater rankings using similarities among candidates. In *AAAI*, pages 613–619, 2006.

[8] V. Conitzer, A. Davenport, and J. Kalagnanam. Improved bounds for computing Kemeny rankings. In *AAAI*, pages 620–626, 2006.

[9] D. Coppersmith, L. Fleischer, and A. Rudra. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *SODA*, pages 776–782, 2006.

[10] C. d'Aspremont and B. Peleg. Ordinal Bayesian incentive compatible representations of committees. *Social Choice and Welfare*, 5:261–279, 1988.

[11] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4):643–652, 1998.

[12] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.

[13] A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.

[14] E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.

[15] C. Homan and L. A. Hemaspaandra. Guarantees for the success frequency of an algorithm for finding Dodgson election winners. In *MFCS*, pages 528–539, 2006.

[16] K. Jogdeo and S. Samuels. Monotone convergence of binomial probabilities and a generalization of Ramanujan's equation. *Annals of Mathematical Statistics*, 39:1191–1195, 1968.

[17] C. Kenyon-Mathieu and W. Schudy. How to rank with few errors. In *STOC*, pages 95–103, 2007.

[18] C. Klamler. A comparison of the Dodgson method and the Copeland rule. *Economics Bulletin*, 4(8):1–7, 2003.

[19] C. Klamler. The Dodgson ranking and its relation to Kemeny's method and Slater's rule. *Social Choice and Welfare*, 23(1):91–102, 2004.

[20] C. Klamler. The Dodgson ranking and the Borda count: a binary comparison. *Mathematical Social Sciences*, 48(1):103–108, 2004.

[21] H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8:538–548, 1983.

[22] J. C. McCabe-Dansted. Approximability and computational feasibility of Dodgson's rule. Master's thesis, University of Auckland, 2006.

[23] J. C. McCabe-Dansted, G. Pritchard, and A. M. Slinko. Approximability of Dodgson's rule. In *COMSOC*, pages 331–344, 2006.

[24] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 9. Cambridge University Press, 2007.

[25] A. D. Procaccia, A. Zohar, Y. Peleg, and J. S. Rosenschein. Learning voting trees. In *AAAI*, pages 110–115, 2007.

[26] A. D. Procaccia, A. Zohar, and J. S. Rosenschein. Automated design of scoring rules by learning from examples. In *AAMAS*, pages 951–958, 2008.

[27] S. Rajagopalan and V. V. Vazirani. Primal-dual RNC approximation algorithms for set cover and covering integer programs. *SIAM Journal on Computing*, 28:526–541, 1999.

[28] T. C. Ratliff. A comparison of Dodgson's method and Kemeny's rule. *Social Choice and Welfare*, 18(1):79–89, 2001.

[29] T. C. Ratliff. A comparison of Dodgson's method and the Borda count. *Economic Theory*, 20(2):357–372, 2002.

[30] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and sub-constant error-probability PCP characterization of NP. In *STOC*, pages 475–484, 1997.

[31] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.

[32] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.

[33] N. Tideman. *Collective Decisions and Voting: The Potential for Public Choice*. Ashgate, 2006.

[34] V. V. Vazirani. *Approximation Algorithms*. Springer, 2001.

[35] H. P. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16:335–353, 1977.

Copyright © by SIAM.
Unauthorized reproduction of this article is prohibited.