

Utilitarians Without Utilities: Maximizing Social Welfare for Graph Problems Using Only Ordinal Preferences

Ben Abramowitz

Department of Computer Science
Rensselaer Polytechnic Institute
abramb@rpi.edu

Elliot Anshelevich

Department of Computer Science
Rensselaer Polytechnic Institute
eanshel@cs.rpi.edu

Abstract

We consider ordinal approximation algorithms for a broad class of utility maximization problems for multi-agent systems. In these problems, agents have utilities for connecting to each other, and the goal is to compute a maximum-utility solution subject to a set of constraints. We represent these as a class of graph optimization problems, including matching, spanning tree problems, TSP, maximum weight planar subgraph, and many others. We study these problems in the ordinal setting: latent numerical utilities exist, but we only have access to ordinal preference information, i.e., every agent specifies an ordering over the other agents by preference. We prove that for the large class of graph problems we identify, ordinal information is enough to compute solutions which are close to optimal, thus demonstrating there is no need to know the underlying numerical utilities. For example, for problems in this class with bounded degree b a simple ordinal greedy algorithm always produces a $(b + 1)$ -approximation; we also quantify how the quality of ordinal approximation depends on the sparsity of the resulting graphs. In particular, our results imply that ordinal information is enough to obtain a 2-approximation for Maximum Spanning Tree; a 4-approximation for Max Weight Planar Subgraph; a 2-approximation for Max-TSP; and a 2-approximation for various Matching problems.

1 Introduction

Human beings are terrible at expressing their feelings quantitatively. For example, when forming collaborations people may be able to order their peers from “best to collaborate with” to worst, but would have a difficult time assigning exact numeric values to the acuteness of these preferences. In other words, even when numerical (possibly latent) utilities exist, in many settings it is much more reasonable to assume that we only know *ordinal* preferences: every agent specifies the *order* of their preferences over the alternatives, instead of a numerical value for each alternative. Recently there has been a lot of work using such an *implicit utilitarian approach*, especially for matching and social choice (see Related Work), in situations where obtaining true numerical utilities may be difficult. Amazingly, as this line of work shows, it is often possible to design algorithms and mechanisms which perform well using only ordinal information.

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In fact, ordinal algorithms often perform almost as well as *omniscient* mechanisms which know the true underlying numerical utilities, instead of just the ordinal preferences induced by these utilities.

In this work we consider a relatively general network formation setting. All problems considered herein are modeled by an undirected complete graph $G = (\mathcal{N}, \mathcal{E})$, where the (symmetric) weight $w(x, y)$ of each edge $(x, y) \in \mathcal{E}$ represents the hidden utility of connecting agents $x, y \in \mathcal{N}$. The goal is to form a maximum-weight (i.e., maximum utility) graph which obeys some given constraints. For example, constraints may include bounds on the maximum degree, on component size, and many others. This framework includes such problems as matching, group formation, TSP, and many others as special cases; see Section 2 for example constraints and how they lead to different important settings.

In many settings we may not have access to the true edge weights $w(x, y)$. Instead, each agent $x \in \mathcal{N}$ reports a strict *preference ordering* over the other agents $\mathcal{N} - \{x\}$ with whom it can connect. We assume this ordering to be consistent with the latent weights, so that $w(x, y) > w(x, z)$ implies that x prefers y to z . While it is clearly impossible to form an optimal (i.e., maximum-utility) solution without direct knowledge of the edge weights, we show how to design good approximation algorithms for selecting a maximum weight subgraph of G , subject to a large set of constraints. As usual in this line of work, the measure of performance is simply the sum of the agent utilities. Our paper provides good approximations for a broad class of ordinal analogues to graph optimization problems representing utility maximization for multi-agent systems. Note that unlike *all* previous work mentioned here, we do not make additional assumptions about the structure of the edge weights: we *do not* assume either that the agent utilities are normalized (as in (Boutilier et al. 2015; Brandt et al. 2016; Caragiannis et al. 2017; Caragiannis and Procaccia 2011; Christodoulou et al. 2016)), nor that they form a metric space (as in (Anshelevich, Bhardwaj, and Postl 2015; Anshelevich and Postl 2017; Anshelevich and Sekar 2016a; 2016b; Caragiannis et al. 2016; Goel, Krishnaswamy, and Munagala 2017; Gross, Anshelevich, and Xia 2017; Skowron and Elkind 2017)). Thus, our results demonstrate how well one can perform using only ordinal information without additional assumptions.

ABC Systems More specifically, we define a class of constraints called *ABC Systems* which consists of three types of constraints. The first two are the familiar constraints which bound the maximum degree of each node and the maximum component size, or number of nodes in a connected component. The third constraint is a much more general requirement called attachment which only applies to nodes that are already in the same connected component. The maximization problem for an ABC System is to compute a maximum weight subgraph $S \subseteq \mathcal{E}$ of an undirected, complete graph $G = (\mathcal{N}, \mathcal{E})$ such that in S every node has degree at most b , every connected component has size (number of nodes) at most c , and $S \in \mathcal{A}$ for some arbitrary **attachment set** \mathcal{A} . A collection of subgraphs \mathcal{A} is an attachment set of G if the following properties hold for all subgraphs $F \subseteq \mathcal{E}$:

- 1) *Heredity*: If $F \in \mathcal{A}$ and $F' \subseteq F$ then $F' \in \mathcal{A}$.
- 2) *Attachment*: If $F \in \mathcal{A}$ and $F + e \notin \mathcal{A}$ for some $e = (u, v) \notin F$, then there is a (u, v) -path in F .¹

Note that all three of our constraints possess the heredity property which enables greedy heuristics, like the Ordinal Greedy algorithm we introduce in Section 3.1, to construct valid solutions. The intuition behind the attachment property is that if $F \in \mathcal{A}$ but $F + e \notin \mathcal{A}$, then e must have both endpoints in the same component. Therefore, the number of such edges whose addition would violate \mathcal{A} within any component of size x is bounded by $\frac{x \cdot \min(b, x-1)}{2} - (x-1)$, where x is bounded by c .

The utility maximization objective for ABC Systems encompasses a wide variety of well-known problems central to algorithm design. The examples we address in this paper include Max Weight b -Matching, Maximum Weight Spanning Tree, Maximum Traveling Salesperson, and Max Weight Planar Subgraph. Our results also encapsulate many other interesting optimization problems for ABC Systems which we will not discuss directly, like finding the maximum weight subgraph with minimum girth k , maximum cycle length l , or which excludes a variety of graph minors (including all 2-edge-connected minors). As we show, all such problems are amenable to knowing only secondhand ordinal information, instead of the true numerical utilities.

1.1 Our Contributions

Most algorithmic techniques for maximizing utility for the full-information setting do not translate to the ordinal information setting. These typically rely on non-local information, like comparisons between weights of non-adjacent edges, or comparing the total weights of sets of edges. This is not possible using only ordinal information. Even the fundamental, and well-studied (Vince 2002; Edmonds 1971; Rado 1942), *Omniscient Greedy* algorithm, which adds edges in strictly non-increasing order of their weight, cannot be executed using only ordinal information. Instead, we focus on the natural *Ordinal Greedy* algorithm (defined in Section 3.1), which adds edges iteratively as long as the edge (x, y) being added is the most preferred edge for both x and y out of all the possible edges which could be added at that

¹We use the “+” and “-” notation when adding or removing a single edge to or from a set.

time. Ordinal Greedy has some very nice properties: in addition to being natural and providing high-utility solutions (as we prove in this paper), it also always creates *pairwise stable* solutions: no pair of agents would have incentive to destroy some of their links and form a new link connecting them. Note that the performance of Ordinal Greedy can be very different from Omniscient Greedy: see Example 1 in Section 3.1 for intuition of why this must be.

In this paper, we analyze the performance of Ordinal Greedy for many ABC Systems (see Table 1). We first prove that for general ABC Systems, Ordinal Greedy always produces a solution with weight at most factor $b+1$ away from optimum, and that this factor is tight. In other words, for general problems including all those in Table 1, as long as the number of connections for each node must be bounded by some small b , then using only ordinal information it is possible to compete with the best possible solution, and thus with any algorithm which knows the true numerical utilities. Such results tell us that when b is small, there is no need to find out the hidden edge weights/utilities; knowing the ordinal preferences is good enough.

Second, we show that by relaxing the component size constraint (i.e., setting c to be unbounded) we can achieve significant improvements. For convenience, we call ABC Systems with the component size constraint relaxed *AB Systems*. We prove that as long as any solution formed by such an AB System is guaranteed to be at most d -sparse, then Ordinal Greedy forms a solution within a factor of $d+1$ from optimum. Since the sparsity of a graph is at most half the average degree in any subgraph, we know that $d \leq \frac{b}{2}$. Therefore an AB System is at worst $(\frac{b}{2} + 1)$ -approximable, giving us a factor of 2 improvement over general ABC Systems.

This result is more powerful than it may first appear, as many important constraints yield sparse solutions. For example, since all tours and trees are 1-sparse, Ordinal Greedy provides a 2-approximation for both Maximum Traveling Salesperson and Maximum Weight Spanning Tree. And since all planar graphs are 3-sparse, we obtain a 4-approximation for Max Weight Planar Subgraph which uses only ordinal information.

Lastly, we consider Max Weight b -Matching, in which the only constraint is that each agent can be matched with at most b others. This is simply an ABC System with unbounded c and \mathcal{A} being all possible sets. We prove that our approximation factor drops to a constant 2, regardless of the value of b .

To prove the ordinal approximations above, we first demonstrate that for any ABC System (and any graphic system with the heredity property defined earlier) Ordinal Greedy achieves its worst approximation on an instance with weight function $w : \mathcal{E} \rightarrow \{0, 1\}$. We use this fact heavily to establish our approximation bounds, and believe it to be of independent interest. Note that similar results for Omniscient Greedy have relied critically on the fact that it selects edges in strictly non-increasing order by weight. Clearly this does not and cannot hold for the ordinal setting, as it is even possible for the minimum weight edge of the graph to be selected before the maximum weight edge. Because

Maximization Problem	Ordinal Greedy	Omniscient Greedy	Best Known
ABC System	$b + 1$	$b + 1$	-
AB System	$\max\{2, d + 1\}$	$\max\{2, d + 1\}$	-
Spanning Tree	2	1	1
Planar Subgraph	4	3	72/25 (Calinescu et al.)
Traveling Salesperson	2	2	9/7 (Paluch, Mucha, and Madry)
b -Matching	2	2	1

Table 1: Here we compare our results for Ordinal Greedy, known results for Omniscient Greedy, and the best known polynomial-time algorithm with full-information. All of our bounds are tight except for the one on Planar Subgraph.

of this, our proofs require completely new approaches and techniques.

1.2 Related Work

Historically, it has been common to approach problems in the ordinal setting with a normative view by designing mechanisms which satisfy axiomatic properties, like stability or truthfulness. These axiomatic properties are useful in many applications, but do not provide a quantitative measure of the quality of a solution. The notion of distortion and the implicit utilitarian framework were first introduced by Procaccia and Rosenschein (2006) in the context of voting to provide such a measure. Since then the distortion, or approximation factor, of various ordinal utility maximization mechanisms has been studied, particularly for matchings (Anshelevich and Sekar 2016a; 2016b; Anshelevich and Zhu 2017; Christodoulou et al. 2016; Caragiannis et al. 2016) and social choice (Anshelevich, Bhardwaj, and Postl 2015; Anshelevich and Postl 2017; Boutilier et al. 2015; Brandt et al. 2016; Caragiannis et al. 2017; Caragiannis and Procaccia 2011; Goel, Krishnaswamy, and Munagala 2017; Gross, Anshelevich, and Xia 2017; Skowron and Elkind 2017).

Our work is unlike that in social choice, since we consider network formation problems where agent preferences are expressed over one another. In the context of matchings, Anshelevich and Sekar (2016a; 2016b) develop various matching algorithms as a black-box to provide approximations for a variety of matching and clustering problems under the implicit utilitarian view. Additionally, Christodoulou et al. (2016) and Caragiannis et al. (2016) provide results for one-sided matchings and Anshelevich and Zhu (2017) consider bipartite matchings. However, *all previous work* on approximation for utility maximization mentioned above either assumes the underlying weights form a metric space (Anshelevich, Bhardwaj, and Postl 2015; Anshelevich and Postl 2017; Anshelevich and Sekar 2016a; 2016b; Anshelevich and Zhu 2017; Goel, Krishnaswamy, and Munagala 2017; Gross, Anshelevich, and Xia 2017; Skowron and Elkind 2017; Caragiannis et al. 2016) or are normalized (Boutilier et al. 2015; Brandt et al. 2016; Caragiannis et al. 2017; Caragiannis and Procaccia 2011; Christodoulou et al. 2016), with only two exceptions. The first exception is that Maximum Traveling Salesperson yields a 2-approximation without the metric assumption (Anshelevich and Sekar 2016b). We prove this result as part

of a much more general theorem using much more general techniques. The second is the result developed in the full-information setting by Preis (1999) and affirmed in the ordinal setting by Anshelevich and Sekar (2016a), that Max Weight Matching yields a 2-approximation without any assumptions on the weights. We generalize this result to all b -Matchings instead of only $b = 1$.

The work most similar to ours is (Anshelevich and Sekar 2016a) which bounds the distortion of ordinal mechanisms for several problems, including Maximum Traveling Salesperson, but relies heavily on the assumption that the weights obey the triangle inequality. For the Ordinal Greedy algorithm, Anshelevich and Sekar (2016a) show the metric assumption implies that any two edges which are both most preferred by their respective endpoints at some iteration must be within a factor of 2 of one another, even if they are not adjacent. By contrast, this non-local information is unavailable to us in our model. Our paper is unique in that we identify a large class of problems for which assumptions on the weights are unnecessary to achieve good approximations to optimum with only ordinal information.

The Omniscient Greedy algorithm has been studied extensively. In fact, it is known to be optimal on exactly the set of independence systems (any system with the heredity property) which are matroids (Edmonds 1971; Rado 1942), which includes Maximum Spanning Trees. Korre and Hausmann (1978) showed that Omniscient Greedy provides good approximations for many independence systems, including matching and symmetric TSP. Dyer, Foulds, and Frieze (1985) further demonstrated that Omniscient Greedy provides a tight 3-approximation for Max Weight Planar Subgraph. These results were later reformulated as k -extendibility by Mestre (2006), who applies this idea to a diverse set of problems, including b -Matching. Unfortunately, the proofs for all results just mentioned rely critically on Omniscient Greedy selecting edges in strictly non-increasing order, making them untenable in the ordinal setting. No ordinal algorithm can yield optimum solutions, even for matroids. However, as we show in Table 1, our results compete well with the best known polynomial-time algorithms for ABC Systems.

2 Model and Problem Statements

The input for all problems in this paper is a set \mathcal{N} of agents (nodes) of size n , and a strict preference ordering for each $x \in \mathcal{N}$ over the edges adjacent to x . The preference order-

ings reported by each agent are induced by a set of hidden symmetric weights $w(x, y) = w(y, x)$ for all $x, y \in \mathcal{N}$. The set of hidden weights corresponds to an undirected, complete graph $G = (\mathcal{N}, \mathcal{E})$ with non-negative weight function $w : \mathcal{E} \rightarrow \mathbb{R}^+$. The transitive relation of the individual preference orderings for all agents determines a partial ordering σ over all edges; note that some pairs of edges may end up being incomparable in σ (see Example 1). The preference ordering σ is said to be **consistent** with the hidden weights if $\forall x, y, z \in \mathcal{N}$, if x prefers y to z then it must be that $w(x, y) \geq w(x, z)$. If an edge e_1 is known to be at least as large as edge e_2 according to this partial ordering σ , then we will say that e_1 *dominates* e_2 in σ . The problems we consider are optimization problems where for an instance (w, σ) the objective is to compute the subgraph of G with maximum total edge weight, subject to a set of constraints, knowing only σ .

For weight function w , we let OPT_w be the optimal solution for the weights prescribed by w , and we let $w(OPT_w)$ be the total weight of the optimal solution evaluated by w . Likewise, we use S to denote our constructed solution and $w(S)$ its weight. Our approximation factor for a problem is therefore $\alpha = \max_{(w, \sigma)} \frac{w(OPT_w)}{w(S)}$, where S is any solution returned by our algorithm for (w, σ) .

Recall the definition of ABC Systems. Given constraints \mathcal{A}, b, c we can say without loss of generality that $b \leq c - 1$, because if any node has c or more neighbors in a component, this component would have to be of size greater than c . When $b = c - 1$ this effectively removes the node degree constraint. Similarly, when $c = n$, the component size is effectively unbounded. Therefore, $1 \leq b < c \leq n$. Likewise, when $\mathcal{A} =$ all subgraphs of G , this effectively annuls the attachment set constraint. Some specific problems which we consider in this paper are as follows.

Max ABC: $c \leq n, b < c, \mathcal{A} =$ any attachment set of G

Max AB: $c = n, b < c, \mathcal{A} =$ any attachment set of G

Maximum Spanning Tree: $c = n, b = c - 1, \mathcal{A} =$ all acyclic subgraphs of G

Maximum Traveling Salesperson: $c = n, b = 2, \mathcal{A} =$ all subgraphs of G without non-Hamiltonian cycles

Max Weight Planar Subgraph: $c = n, b = c - 1, \mathcal{A} =$ all planar subgraphs of G

Max Weight b -Matching: $c = n, b < c, \mathcal{A} =$ all subgraphs of G

Max Weight Matching: $c = n, b = 1, \mathcal{A} =$ all subgraphs of G

3 Algorithmic Framework

In this section we define the Ordinal Greedy algorithm and reveal some of its salient properties. Rather than limit ourselves only to ABC Systems, in this section we consider general graphic *independence systems*. An independence system for our setting is a pair $(\mathcal{E}, \mathcal{L})$ where \mathcal{E} corresponds to the set of edges in some graph and \mathcal{L} is a collection of subsets of \mathcal{E} such that if $F \in \mathcal{L}$ and $F' \subseteq F$ then $F' \in \mathcal{L}$.

The sets in \mathcal{L} are called independent. It is easy to see that all ABC Systems are independence systems because their three constraints possess this heredity property. Let \mathcal{B} denote the set of all subgraphs in which all nodes have degree at most b and let \mathcal{C} denote the set of all subgraphs in which all connected components have size at most c . Our Max ABC problem can be restated as: Given a graph $G = (\mathcal{N}, \mathcal{E})$, attachment set \mathcal{A} , degree limit b and component size limit c , compute the maximum weight subgraph in $\mathcal{L} = \mathcal{A} \cap \mathcal{B} \cap \mathcal{C}$.

3.1 The Ordinal Greedy Algorithm

In an ordinal setting, algorithms only have access to a partial ordering, or set of preference orderings, which provide strictly local information about the preferences of each agent. This precludes the use of algorithms which require comparisons between the weights of non-adjacent edges. In fact, it is not difficult to see that no ordinal algorithm can be guaranteed to compute the optimal solution for even simple settings, e.g., forming a matching (Anshelevich and Sekar 2016a). However, the Ordinal Greedy algorithm defined below performs well in this setting because it relies on strictly local information. Ordinal Greedy starts from the empty set and builds up a sequence of intermediate solutions by adding locally optimal edges at each iteration which do not violate a set of constraints, i.e., preserve independence. To understand how this heuristic is applied to the ordinal setting, we must formalize what it means for an edge to be locally optimal.

Definition 1. Undominated Edge

Given a set E of edges, $(u, v) \in E$ is undominated if for all (u, x) and (v, y) in E , $w(u, v) \geq w(u, x)$ and $w(u, v) \geq w(v, y)$.

At this point it is important to make several observations. First, every edge set E has at least one undominated edge, because its maximum weight edge must be undominated. However, there may be undominated edges which are not globally maximum. Second, for any edge set E it is straightforward to find at least one undominated edge using only the partial ordering σ (see (Anshelevich and Sekar 2016a) for details). Undominated edges are either of the form (u, v) where u and v are each other's most preferred neighbor, or form cycles in which each subsequent node is the first choice of the previous one, and thus all edges in the cycle have the same weight.

What follows is a general purpose Ordinal Greedy algorithm, which starts from the empty set and iteratively selects undominated edges from the set of remaining edges which do not violate the constraints in question. The algorithm uses the partial ordering σ to determine which edges are undominated at each iteration. The algorithm concludes when there are no edges left which can be added to the subgraph without violating the constraints, so the final solution S is maximal in this sense.

We refer to the iteration at which an edge $e = (u, v)$ is removed from E as the **critical iteration** of e . When the inputs to our algorithm $(\mathcal{E}, \mathcal{L})$ characterize an ABC System, there are exactly four cases which may occur at the critical iteration of edge e :

Algorithm 1: Ordinal Greedy

Input: Edge set \mathcal{E} , partial ordering σ , collection of valid subgraphs \mathcal{L}
Initialize $S = \emptyset, E = \mathcal{E}$;
while $E \neq \emptyset$ **do**
 Pick an undominated edge $e = (u, v) \in E$ and
 add it to the intermediate solution: $S \leftarrow S + e$;
 Remove e from E ;
 Remove all edges f from E such that $S + f \notin \mathcal{L}$;
end
Output: Return S

- 1) e is added to the ordinal greedy solution S
- 2) e is removed from E because $S + e \notin \mathcal{A}$
- 3) e is removed from E because $S + e \notin \mathcal{B}$ (where \mathcal{B} = sets of edges with any degree $\leq b$)
- 4) e is removed from E because $S + e \notin \mathcal{C}$ (where \mathcal{C} = sets of edges with any component size $\leq c$)

For cases 2-4 we say e was **eliminated due to \mathcal{A}, \mathcal{B} , or \mathcal{C}** . If an edge $e = (u, v)$ was eliminated due to \mathcal{A} , the attachment property implies there must be a (u, v) -path in the intermediate solution at its critical iteration. In other words, u and v are already in the same connected component in S at this iteration. If $e = (u, v)$ was eliminated due to \mathcal{B} , either u or v must already have degree exactly b at this iteration. If $e = (u, v)$ was eliminated due to \mathcal{C} , u and v must already be in disjoint connected components whose cumulative size is greater than c . Note that in these three cases, an edge can only be eliminated if at least one adjacent edge of equal or greater weight has already been added to the intermediate solution, and all adjacent edges already added to the intermediate solution must be of equal or greater weight (since only undominated edges are added to our solution).

There are limiting values of \mathcal{A} , b , and c for which elimination due to these constraints cannot occur. When $c = n$, no edge can be eliminated due to \mathcal{C} because all nodes can be in the same connected component. Additionally, when \mathcal{A} is the set of all subgraphs of G , no edge can be eliminated due to \mathcal{A} . If adding an edge would violate more than one constraint, we say that it was eliminated in order of priority $\mathcal{C} \rightarrow \mathcal{B} \rightarrow \mathcal{A}$. For example, when $b = c - 1$, no edge can be eliminated due to \mathcal{B} , because for a node to reach degree b the size of its component must be exactly c and we say that any incident edge would be eliminated due to \mathcal{C} . As the following sections show, the approximation factor of Ordinal Greedy for an ABC System depends on which cases of elimination can occur.

Notice that the performance of the Ordinal Greedy algorithm can deviate significantly from the Omniscient Greedy algorithm in the full-information setting (which we call ‘‘Omniscient Greedy’’ because it knows the underlying edge weights and can choose the edge with maximum weight at each iteration). Consider the following example.

Example 1. Suppose the graph $G = (\mathcal{N}, \mathcal{E})$ is constructed as follows. Let $\mathcal{N} = \{u_1, \dots, u_k, v_1, \dots, v_k\}$. Let $w(u_i, v_i) = 1 + \epsilon$ for $i \leq k$ for some infinitesimal ϵ . Let $w(u_i, u_j) = 1$ for all $i \neq j$. Let $w(v_i, v_j) = \epsilon$ for all $i \neq j$. Let all other

edges have weight 0. Consider the ABC System corresponding to finding a Maximum-Weight Spanning Tree. It is clear that Omniscient Greedy will find the optimum solution with weight $w(OPT_w) = k(1 + \epsilon) + (k - 1)$.

Now consider Ordinal Greedy. Suppose Ordinal Greedy begins by selecting (u_i, v_i) for $i \leq k$, which are all undominated at the beginning of the algorithm. Once these edges have been selected, edges of the form (u_i, u_j) and (v_i, v_j) become undominated for $i \neq j$. Now, if an edge (u_i, u_j) or (v_i, v_j) is selected, the other must be eliminated at that iteration, since taking it would form a cycle. Notice that since we only have access to ordinal information, there is *no possible way* for Ordinal Greedy to tell which of these edges is better: they are both edges which are most preferred by their endpoints, even though one secretly has weight 1 and the other only ϵ . In other words, these edges are incomparable in the partial preference order σ . Suppose Ordinal Greedy proceeds by selecting (v_i, v_{i+1}) for $i < k$. Then the Ordinal Greedy solution formed has weight $w(S) = k(1 + \epsilon) + (k - 1)\epsilon$. This example shows that (in the limit) it is not possible for Ordinal Greedy to always result in solutions better than a factor of 2 away from optimum, even though Omniscient Greedy can easily compute the true optimum solution. As we show in this paper, however, despite its knowledge handicap, Ordinal Greedy can often produce surprisingly good results.

3.2 Properties of Ordinal Greedy

For any independence system in the full-information setting, the Omniscient Greedy algorithm has been shown to achieve its worst approximation on an instance with a **binary weight function** $\bar{w} : \mathcal{E} \rightarrow \{0, 1\}$ (Korte and Hausmann 1978). However, previous proofs have relied crucially on the fact that Omniscient Greedy selects edges in strictly non-increasing order by weight, which is not possible with only ordinal information. We offer a new proof to show that even in the ordinal setting, Ordinal Greedy always achieves its worst approximation factor on an instance with a binary weight function $\bar{w} : \mathcal{E} \rightarrow \{0, 1\}$ for any graphic independence system. This theorem will allow us to prove approximation bounds for ABC and AB Systems later in this paper.

Theorem 1. *For any graphic independence system $(\mathcal{E}, \mathcal{L})$, for any instance (w, σ) with weight function $w : \mathcal{E} \rightarrow \mathbb{R}^+$ and partial ordering σ consistent with w , there exists an instance (\bar{w}, σ) with weight function $\bar{w} : \mathcal{E} \rightarrow \{0, 1\}$ such that σ is consistent with \bar{w} and the worst-case ratio of the optimal solution to an Ordinal Greedy solution is at least as large as for (w, σ) .*

Proof Sketch. Please see the full version of this paper for detailed proof; here we only provide a short proof sketch. Suppose on instance (w, σ) the ratio between the optimal solution OPT_w and solution S constructed by Ordinal Greedy is $\frac{w(OPT_w)}{w(S)} = \delta$. Our goal is to construct a binary weight function \bar{w} such that $\frac{\bar{w}(OPT_{\bar{w}})}{\bar{w}(S)} \geq \delta$. When δ is infinite, constructing \bar{w} is straightforward, so we only consider finite values of δ . First we create a weight function \hat{w} by raising the weights of all edges not in S as much as possible without

altering the weights of the edges of S , such that σ remains consistent with \hat{w} . Since Ordinal Greedy selects S and none of its edge weights have changed, and the edge weights of OPT_w cannot have decreased, then $\frac{\hat{w}(OPT_w)}{\hat{w}(S)} \geq \delta$. From \hat{w} we carefully create \bar{w} by proving that there must exist a subset of edges to which we can assign weight 1 and let all other edges have weight 0, such that σ is consistent with \bar{w} and $\frac{\bar{w}(OPT_{\bar{w}})}{\bar{w}(S)} \geq \frac{\hat{w}(OPT_w)}{\hat{w}(S)} \geq \delta$. \square

Another nice property is that every solution constructed by Ordinal Greedy is pairwise stable. Pairwise stability means that no pair of agents x, y has incentive to collude to add edge (x, y) by each giving up some of their edges in the Ordinal Greedy solution S . Either this exchange would decrease the total utility of one of the agents, or adding (x, y) is infeasible even after sacrificing the other edges. Here we assume that the utility of a node x in solution S is simply the total weight of edges in S incident on x .

Theorem 2. *Any solution S constructed by Ordinal Greedy on an independence system is pairwise stable.*

4 Ordinal Approximation for ABC Systems

In this section we bound the worst-case performance of Ordinal Greedy compared to the optimal solution for any ABC System. We use α_{ABC} to denote the approximation factor, or the ratio of the optimal solution to the worst possible Ordinal Greedy solution for any ABC System.

Unlike Example 1 in Section 3.1 for the maximum spanning tree problem, Ordinal Greedy does not provide a constant approximation factor for all ABC Systems. However, it does always provide a finite approximation which depends on the degree limit b . To simplify notation, since the optimal solution here is only evaluated using the same weight function used to generate it, we refer to the total weight of the optimal solution $w(OPT_w)$ as $w(OPT)$. Here we show that $\alpha_{ABC} \leq b + 1$ for any ABC System and provide a family of examples where $\frac{w(OPT)}{w(S)} = b + 1$ to show that $b + 1$ is a tight bound on the approximation factor. In later sections, we explore classes of ABC Systems in which Ordinal Greedy achieves a better worst-case approximation.

Note that this result is quite general. As we discussed, ABC Systems include many varied constraints, some quite difficult to approximate. Our result in this section states that, even for extremely complex \mathcal{A} and constraints on component size c , as long as the maximum allowed degree of any node is small, then it is possible to form a good approximation to the true optimum solution while only knowing ordinal information instead of the true edge weights.

Theorem 3. *For any ABC System, the Ordinal Greedy algorithm always produces a solution within a factor of $(b+1)$ of the optimal solution, and this bound is tight.*

Proof Sketch. See the full version of this paper for a detailed proof; here we only provide a short proof sketch. We proceed via a charging argument. We wish to charge the weight of the edges of OPT to the edges of S such that all edges of OPT are fully charged somewhere, and no edge of S receives a charge greater than $b + 1$ times its

weight. However, unlike Omniscient Greedy in the full-information setting, we cannot assume that any eliminated edge of OPT has weight smaller than all edges of the Ordinal Greedy solution S which were added before its critical iteration. This prohibits us from using the methods in previous work. Thankfully, due to Theorem 1 we know that if Ordinal Greedy produces a solution within a factor of $(b + 1)$ of optimal for all instances with binary weight functions, then this holds for all instances. We therefore assume that all weights are $\{0, 1\}$, and can now charge any edge of OPT to any weight 1 edge of S , but must ensure that no weight is charged to any edges of S with weight 0.

To ensure no edge of S is charged more than $b + 1$ times its weight, we look at the connected components of S with only weight 1 edges, and charge all edges of OPT to these components. Specifically, we design the following charging scheme. Let (u, v) be an edge of OPT where $w(u, v) = 1$. Let P_u and P_v be the connected components containing u and v in the subgraph of S containing only weight 1 edges. We charge the weight of $w(u, v)$ between P_u and P_v based on what occurs at the critical iteration of (u, v) . If $(u, v) \in S$ then $P_u = P_v$, so charge its full weight to this component. If (u, v) was eliminated due to \mathcal{A} , charge its full weight to either P_u or P_v arbitrarily. Note that while the attachment property of \mathcal{A} ensures that u and v are in the same connected component in S at this iteration, this does *not* imply that $P_u = P_v$ because all (u, v) -paths in S may contain a weight 0 edge. If (u, v) was eliminated due to \mathcal{B} , one of its endpoints must have a degree of exactly b in S at this iteration, so charge its full weight to the component containing this endpoint. If (u, v) was eliminated due to \mathcal{C} , we split the charge between P_u and P_v based on the size of the connected components in S containing u and v at this iteration. Let q_u and q_v be the sizes of the connected components in S containing u and v at the critical iteration of (u, v) . Charge $\frac{q_u - 1}{q_u + q_v - 2}$ to P_u and $\frac{q_v - 1}{q_u + q_v - 2}$ to P_v . In all four cases we have ensured the full weight of the edge of OPT has been charged between the components containing its endpoints.

The rest of the proof involves arguing that each such component P_u of size p is charged a total of at most $(b + 1)(p - 1)$ using the above charging scheme. Since such a component must contain at least $p - 1$ edges with weight 1, this completes the proof of the upper bound because it shows that the total number of edges in OPT with weight 1 is at most $b + 1$ times the number of such edges in S .

We then provide a family of examples to show this bound is tight. Omniscient Greedy has the same worst-case solution as Ordinal Greedy on this family of examples, so our ordinal approximation competes well despite its knowledge handicap. Note that in the example yielding the lower bound of $\alpha_{ABC} \geq b + 1$, all edges are eliminated due to \mathcal{C} . The following section demonstrates that when the component size constraint is relaxed by allowing $c = n$, the approximation factor improves significantly. \square

5 AB Systems and Important Special Cases

In this section, we bound the performance of Ordinal Greedy on ABC Systems where $c = n$, effectively removing the

component size constraint. We then discuss some common examples of maximization problems on AB Systems, including Max Spanning Tree, Max TSP, and Max Planar Subgraph. To improve our bound from $b + 1$ we invoke the notion of sparsity.

Definition 2. Sparsity

A graph S is d -sparse if for all subgraphs $F \subseteq S$ containing $V(F)$ nodes and $E(F)$ edges $\frac{E(F)}{V(F)} \leq d$ and for any $\hat{d} < d$ there exists a subgraph $\hat{F} \subseteq S$ such that $\frac{E(\hat{F})}{V(\hat{F})} > \hat{d}$.

Suppose our attachment set \mathcal{A} and degree limit b imply that any feasible solution must be d -sparse. Note that this sparsity is implied by our constraints, and is not a separate constraint. Our main result in this section is that, for any graph collection which is guaranteed to be d -sparse, ordinal information is enough to produce good approximations. Specifically, we prove a bound of $d + 1$ for such settings. Since the sparsity corresponds to an upper bound on average degree of the nodes, it is always true that $d \leq \frac{b}{2}$, and so when $c = n$, this immediately reduces the approximation factor from $b + 1$ to $\frac{b}{2} + 1$. Even for large b , however, there are many natural classes of graphs that are always sparse, including planar graphs, scale-free graphs, graphs of small arboricity or treewidth, and many others. As we discuss in the next section, this result allows us to provide extremely strong guarantees for many important problems.

Theorem 4. *For any ABC System where the components can be of any size and the constraints imply that any feasible solution must be d -sparse, the Ordinal Greedy algorithm always produces a solution within a factor of $\max\{2, (d + 1)\}$ of the optimal solution, and this bound is tight.*

Proof Sketch. See the full version of this paper for detailed proof. As with our proof of Theorem 3 for general ABC Systems, we only need to consider instances with weights $\{0, 1\}$ due to Theorem 1. However, the charging schemes and proofs for ABC Systems and AB Systems differ significantly. To lower the approximation factor from $b + 1$ to $\max\{2, d + 1\}$, we have to be more selective about where we charge the edges of OPT . For simplicity, we first assign the edges of OPT to their endpoints, before considering the total charge to all the nodes in any component. Since OPT is d -sparse, the edges of OPT can be assigned to their endpoints such that each node is assigned at most d edges. We then take such an assignment and for each edge of OPT eliminated due to \mathcal{B} we change its assignment, if necessary, to the node which caused its elimination. Let P be a component of the subgraph of S containing only weight 1 edges, and suppose $p = |P|$. Then, similarly to the proof of Theorem 3, we must show that this component will be charged at most $\max\{2, d + 1\}w(P)$, but unlike before, components may be charged more than $(d + 1)(p - 1)$ if $w(P) > p - 1$.

Now we consider two cases based on whether any node in a component was charged an edge of OPT eliminated due to \mathcal{B} . If there such a node in a component, then it must be possible to distribute the charge on the nodes over the edges of the component directly so that no edge is charged more than $\max\{2, d + 1\}$ times its weight. If there is no such node, then we show that at least one node in the component

must be charged at most $p - 1$ and the rest are charged at most $d(p - 1)$, cumulatively providing a charge at most $(d + 1)(p - 1)$ which can then be distributed over the edges in the component. Once again, we provide a family of examples to show that this bound is tight. \square

5.1 Important Cases of AB Systems

Theorem 4 establishes that for AB Systems in which solutions are always sparse, ordinal algorithms don't perform much worse than ones which know the true underlying edge weights. While our result in the previous section is quite general, it is worth noting how it applies to many important problems which happen to be special cases of AB Systems. Since all tours and cycles are 1-sparse, and all planar graphs are at most 3-sparse, we immediately arrive at the following corollaries:

Corollary 1. *Ordinal Greedy always computes a 2-approximation for Maximum Weight Spanning Tree, and this bound is tight.*

Corollary 2. *Ordinal Greedy always computes a 2-approximation for Maximum Traveling Salesperson, and this bound is tight.*

Corollary 3. *Ordinal Greedy always computes a 4-approximation for Max Weight Planar Subgraph.*

See Example 1 and the full version of this paper for tight lower bound examples for Maximum Spanning Tree and Maximum Traveling Salesperson, respectively.

More generally, the same arguments can be applied to any problem where the goal is to find maximum-weight subgraphs with some excluded minor, finding maximum-weight graphs with small treewidth or arboricity, as well as a variety of other graph problems.

6 b -Matching

For any ABC System where $c = n$ and $\mathcal{A} =$ all subgraphs of G , the only constraint is that each node must have degree at most b . This is equivalent to the well-known problem of Max Weight b -Matching. In this case, the approximation provided by Ordinal Greedy improves greatly over general AB Systems. In fact, it provides a strict 2-approximation regardless of the value of b .

Theorem 5. *For any ABC System on graph G , where $c = n$ and $\mathcal{A} =$ all subgraphs of G , Ordinal Greedy always constructs a solution within a factor of 2 of the optimal solution. This bound is tight.*

Please see the full version of this paper for the proof. Note that $b = 1$ is the problem Max Weight Matching, and our result generalizes the results from (Anshelevich and Sekar 2016a) and (Preis 1999).

7 Conclusion and Further Directions

In this paper we identify a large class of problems we call ABC Systems for which ordinal preference information is sufficient for algorithms to provide good approximations to optimal, even without access to cardinal utilities. Previous work has shown that if agent preferences form a metric

space, approximations for TSP and matching can improve in expectation (Anshelevich and Sekar 2016a). It remains to be seen how Ordinal Greedy performs on ABC Systems in expectation and how much the approximation factors for general ABC or AB Systems improve when this metric assumption holds. Also along the lines of previous work, it would be interesting to investigate whether truthful ordinal algorithms for ABC and AB Systems can compete with our non-truthful algorithm, much as (Anshelevich and Sekar 2016b) did for the problems first approached in (Anshelevich and Sekar 2016a). Lastly, we have seen that all solutions produced by Ordinal Greedy are pairwise stable, but it is unknown for our problems whether all pairwise stable solutions produce a good approximation to optimum (although it is easily seen to be true for MST and TSP).

Acknowledgments

This work was partially supported by NSF award CCF-1527497.

References

- Anshelevich, E., and Abramowitz, B. 2017. Utilitarians without utilities: Maximizing social welfare for graph problems using only ordinal preferences. (*Full Version*). Available at <http://www.cs.rpi.edu/~eanshel/pubs.html>.
- Anshelevich, E., and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research* 58:797–827.
- Anshelevich, E., and Sekar, S. 2016a. Blind, greedy, and random: algorithms for matching and clustering using only ordinal information. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, 383–389.
- Anshelevich, E., and Sekar, S. 2016b. Truthful mechanisms for matching and clustering in an ordinal world. In *International Conference on Web and Internet Economics (WINE)*.
- Anshelevich, E., and Zhu, W. 2017. Tradeoffs between information and ordinal approximation for bipartite matching. In *International Symposium on Algorithmic Game Theory (SAGT)*.
- Anshelevich, E.; Bhardwaj, O.; and Postl, J. 2015. Approximating optimal social choice under metric preferences. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 777–783.
- Boutilier, C.; Caragiannis, I.; Haber, S.; Lu, T.; Procaccia, A. D.; and Sheffet, O. 2015. Optimal social choice functions: A utilitarian view. *Artificial Intelligence* 227:190–213.
- Brandt, F.; Conitzer, V.; Endriss, U.; Procaccia, A. D.; and Lang, J. 2016. *Handbook of computational social choice*. Cambridge University Press.
- Calinescu, G.; Fernandes, C. G.; Karloff, H.; and Zelikovsky, A. 2003. A new approximation algorithm for finding heavy planar subgraphs. *Algorithmica* 36(2):179–205.
- Caragiannis, I., and Procaccia, A. D. 2011. Voting almost maximizes social welfare despite limited communication. *Artificial Intelligence* 175(9-10):1655–1671.
- Caragiannis, I.; Filos-Ratsikas, A.; Frederiksen, S. K. S.; Hansen, K. A.; and Tan, Z. 2016. Truthful facility assignment with resource augmentation: An exact analysis of serial dictatorship. In *International Conference on Web and Internet Economics (WINE)*.
- Caragiannis, I.; Nath, S.; Procaccia, A. D.; and Shah, N. 2017. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research* 58:123–152.
- Christodoulou, G.; Filos-Ratsikas, A.; Frederiksen, S. K. S.; Goldberg, P. W.; Zhang, J.; and Zhang, J. 2016. Social welfare in one-sided matching mechanisms. In *International Conference on Autonomous Agents and Multiagent Systems*, 30–50. Springer.
- Dyer, M.; Foulds, L.; and Frieze, A. 1985. Analysis of heuristics for finding a maximum weight planar subgraph. *European Journal of Operational Research* 20(1):102–114.
- Edmonds, J. 1971. Matroids and the greedy algorithm. *Mathematical programming* 1(1):127–136.
- Goel, A.; Krishnaswamy, A. K.; and Munagala, K. 2017. Metric distortion of social choice rules: Lower bounds and fairness properties. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 287–304. ACM.
- Gross, S.; Anshelevich, E.; and Xia, L. 2017. Vote until two of you agree: Mechanisms with small distortion and sample complexity. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 544–550.
- Korte, B., and Hausmann, D. 1978. An analysis of the greedy heuristic for independence systems. *Annals of Discrete Mathematics* 2:65–74.
- Mestre, J. 2006. Greedy in approximation algorithms. In *European Symposium on Algorithms*, volume 4168, 528–539. Springer.
- Paluch, K. E.; Mucha, M.; and Madry, A. 2009. A $7/9$ -approximation algorithm for the maximum traveling salesman problem. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques* 298–311.
- Preis, R. 1999. Linear time $1/2$ -approximation algorithm for maximum weighted matching in general graphs. In *STACS*, volume 99, 259–269. Springer.
- Procaccia, A. D., and Rosenschein, J. S. 2006. The distortion of cardinal preferences in voting. In *In The Tenth International Workshop on Cooperative Information Agents (CIA): Lecture Notes in Computer Science (LNCS)*, volume 4149, 317–331. Springer.
- Rado, R. 1942. A theorem on independence relations. *The Quarterly Journal of Mathematics* (1):83–89.
- Skowron, P. K., and Elkind, E. 2017. Social choice under metric preferences: Scoring rules and stv. In *Thirty-First AAAI Conference on Artificial Intelligence*, 706–712.
- Vince, A. 2002. A framework for the greedy algorithm. *Discrete Applied Mathematics* 121(1):247–260.