The Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI-23)

Optimizing Multiple Simultaneous Objectives for Voting and Facility Location

Yue Han, Christopher Jerrett, Elliot Anshelevich

Rensselaer Polytechnic Institute, 110 Eighth Street, Troy, New York 12180
hany4@rpi.edu, jerrec@rpi.edu, eanshel@cs.rpi.edu

Abstract

We study the classic facility location setting, where we are given n clients and m possible facility locations in some arbitrary metric space, and want to choose a location to build a facility. The exact same setting also arises in spatial social choice, where voters are the clients and the goal is to choose a candidate or outcome, with the distance from a voter to an outcome representing the cost of this outcome for the voter (e.g., based on their ideological differences). Unlike most previous work, we do not focus on a single objective to optimize (e.g., the total distance from clients to the facility, or the maximum distance, etc.), but instead attempt to optimize several different objectives *simultaneously*. More specifically, we consider the *l-centrum* family of objectives, which includes the total distance, max distance, and many others. We present tight bounds on how well any pair of such objectives (e.g., max and sum) can be simultaneously approximated compared to their optimum outcomes. In particular, we show that for any such pair of objectives, it is always possible to choose an outcome which simultaneously approximates both objectives within a factor of $1 + \sqrt{2}$, and give a precise characterization of how this factor improves as the two objectives being optimized become more similar. For $q > 2$ different centrum objectives, we show that it is always possible to approximate all q of these objectives within a small constant, and that this constant approaches 3 as $q \to \infty$. Our results show that when optimizing only a few simultaneous objectives, it is always possible to form an outcome which is a significantly better than 3 approximation for all of these objectives.

1 Introduction

When working on optimization problems, it is often difficult to pick one single objective to optimize: in most real applications different parties care about many different objectives at the same time. For example, consider the classic setting when we are given n clients and m possible facility locations in some metric space, and want to choose a location to build a facility. Note that while the setting we consider is for facility location problems, the exact same setting also arises in spatial social choice (see e.g., Merrill III, Merrill, and Grofman (1999); Enelow and Hinich (1984); Anshelevich et al. (2021)), where voters are the clients and the goal is to choose a candidate or outcome located in some metric space, where the distance from a voter to an outcome represents the cost of this outcome for the voter (e.g., based on their ideological differences). When choosing where to build a facility (or which candidate to select) for the public good (e.g., where to build a new post office, supermarket, etc.), we may care about minimizing the average distance from users to the chosen facility (a utilitarian measure), or the maximum distance (an egalitarian measure), or many other measures of fairness or happiness. Focusing on just a single measure may not be useful for actual policy makers, who often want to satisfy multiple objectives simultaneously and in fact refuse to commit themselves to a single one, as many objectives have their own unique merits. In this paper we instead attempt to simultaneously minimize multiple objectives. For example, what if we care about *both* the average and the maximum distance to the chosen facility, and not just about some linear combination of the two? What if we want to choose a facility so that it is close to optimum in terms of the average distance from the users, and *at the same time* is also close to optimum in terms of the maximum distance? Is this even possible to do?

More specifically, we consider *l-centrum* problems (Slater 1978; Tamir 2001; Peeters 1998), where we are given a set of possible facilities $\mathcal{F}$ and a set of n clients $\mathcal{C}$ in an arbitrary metric space with distance function d. For each client $i \in \mathcal{C}$ there is a cost $d(i, j)$ if we choose to build facility $j \in \mathcal{F}$. Then the goal is to pick one facility from $\mathcal{F}$ such that it minimizes the sum of the l most expensive costs induced by the choice of facility location. Such problems generalize minimizing the total client cost ($l = n$), as well as the maximum client cost ($l = 1$). The latter may be considered a measure which is more fair to all the clients (since it makes sure that *all* clients have small cost, not just on average), but would have the drawback that a solution where all except a single client have low cost would be considered the same as a solution where they all have high cost, as long as the maximum cost stays the same. Because of this, some may argue that an objective where we consider only the costs of the worst 10 percent of the clients may be better. In this work, we sidestep questions about which objective is best entirely. Since each of the *l-centrum* objectives has its own advantages, our goal is to simultaneously approximate multiple such objectives. This idea of simultaneously approximating *l-centrum*

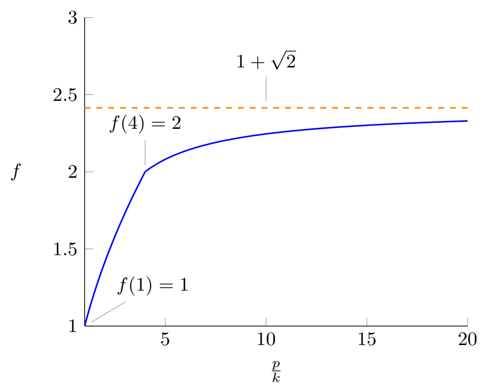

Figure 1: Plot of the tight upper bound of the simultaneous approximation ratio for c_k and c_p (denoted as function f) with respect to $\frac{p}{k}$. When k and p are similar this factor is small, and in fact when $p \leq 4k$ they can both be approximated within a factor of 2. As $\frac{p}{k}$ becomes larger, the worst-case approximation approaches $1+\sqrt{2}$.

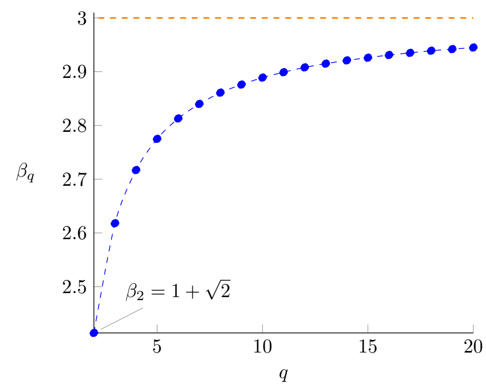

Figure 2: Plot of upper bound of the simultaneous approximation ratio β_q for q different centrum objectives.

problems as a method of creating "fair" outcomes was previously discussed in Kumar and Kleinberg (2006); Goel and Meyerson (2006); Goel, Hulett, and Krishnaswamy (2018), and was adapted from the idea of *approximate majorization* in (Bhargava, Goel, and Meyerson 2001).

Note that our approach is very different from combining several objectives into a single one (e.g., by taking a weighted sum); we instead want to make sure that the chosen outcome is good with respect to each objective we are interested in simultaneously. More formally, for $1 \leq l \leq n$, we define the cost function for choosing facility $A \in \mathcal{F}$ to be $c_l(A)$, which is the sum of the top l distances from A to each client in $\mathcal{C}$. The *l-centrum* problem asks to minimize $c_l(A)$ with a fixed l value; denote the optimal facility location for this objective by O_l. Now suppose we have q such objectives that we want to optimize, such that $l \in \mathcal{K} = \{k_1, k_2, \cdots, k_q\}$. We then say a facility $A \in \mathcal{F}$ is a simultaneous α-approximation for all of the q objectives iff $c_l(A) \leq \alpha \cdot c_l(O_l)$ for all $l \in \mathcal{K}$.

1.1 Our Contributions

We first consider the setting where we attempt to optimize two objectives. These objectives could, for example, be minimizing the maximum distance $\max_i d(i,j)$ and the total distance $\sum_i d(i,j)$. Or more generally they can be two arbitrary centrum objectives with one objective being *k-centrum* and the other being *p-centrum* with $k \leq p$. We prove that for any such pair of objectives, it is always possible to choose an outcome $A \in \mathcal{F}$ which simultaneously approximates both objectives within a factor of $1+\sqrt{2}$. In fact, we provide a tight upper bound for how well any such pair of objectives can be approximated at the same time, as shown in Figure 1. Our results show that when two people disagree on which objective is the best to optimize, they can both be made relatively happy in our setting.

We then proceed to study optimizing more than two objectives at the same time. When optimizing q different centrum objectives, we prove that it is always possible to approximate all q of these objectives within a small constant; the plot of this upper bound is shown in Figure 2. When $q=2$ this factor coincides with our above bound for 2 objectives; as the number of objectives grows this value approaches 3. Thus our results show, that when optimizing only a few simultaneous objectives, it is always possible to form an outcome which is a significantly better than 3 approximation for all of these objectives.[1] Finally, in Section 5 we discuss the important special case when facilities can be placed at any client location, i.e., $\mathcal{F} = \mathcal{C}$, which for the social choice setting corresponds to all the voters also being possible candidates.

1.2 Related Work

Facility location, as well as spatial voting problems, are a huge area with far too much existing work to survey here; many variants have been studied with many different objectives (see for example Chan et al. (2021); Farahani, SteadieSeifi, and Asgari (2010) and the references therein). As discussed above, we want to simultaneously approximate multiple objectives. However, there are many other different approaches when considering multiple objectives, with a common one being converting multiple objectives into one objective and optimizing the new objective, as in Ehrgott and Gandibleux (2000); Farahani, SteadieSeifi, and Asgari (2010). As discussed in Section 5.1 in (Ehrgott and Gandibleux 2000), the most commonly used conversion is the weighted sum of the objectives (such as in, e.g., (Alamdari and Shmoys 2017; McGinnis and White 1978; Ohsawa 1999)), but of course there are many ways to combine several objectives. Since different people have different opinions and priorities, it is usually impossible to pick one combination that can make everyone satisfied. In addition, a

[1] If the goal is to approximate *all* the centrum objectives at the same time, then (Kumar and Kleinberg 2006) provided a 4-approximation for this, and the results in (Goel, Hulett, and Krishnaswamy 2018; Gkatzelis, Halpern, and Shah 2020) imply that this can be improved to 3.

good outcome for the new (combined) objective does not directly imply it is also good with respect to each of the original objectives, and in fact simultaneously approximating several objectives may be impossible even if the combined objective is approximable. On the other hand, a simultaneous α-approximation for all the objectives *does* imply an α-approximation of any convex combinations of the objectives, and thus forms a strictly stronger result.

In addition to forming a weighted sum of two objectives, Alamdari and Shmoys (2017) looked at a different notion of approximation for multiple objectives. They considered approximating the original objectives with respect to a specific feasible solution instead of their respective optimal solutions, and gave a polynomial time (4,8)-approximation algorithm for minimax and minisum for the special case when $\mathcal{F} = \mathcal{C}$. A simultaneous α-approximation for both minimax and minisum (which is what we form for $\alpha = 1+\sqrt{2}$) would also imply a (α, α)-approximation in their notion of approximation. Note that the result in Alamdari and Shmoys (2017) applies to selecting more than one facility, however, while in that setting we know it is impossible to form a bounded approximation ratio for both minimax and minisum simultaneously (Alamdari and Shmoys 2017; Goel and Meyerson 2006). Because of this, similarly to much of the work on this topic (see below), we focus on selecting a single facility while optimizing multiple simultaneous objectives. When considering multiple objectives, another approach is to consider efficient algorithms to find the Pareto optimal set[2] as discussed in Ehrgott and Gandibleux (2000), with specific examples such as Nickel et al. (2005); Roostapour, Kiarazm, and Davoodi (2015) for placing one facility. However, this approach generally does not consider how good those solutions are in comparison to each of the objectives, nor if there exists one that is good for all the objectives.

Now that we have discussed different approaches for multiple objective optimization, we want to consider a set of objectives that is appropriate for our setting. Two of the most commonly studied objectives for facility location problems are minimax (minimize maximum distance) and minisum (minimize sum of distances), see the survey Chan et al. (2021). In this work we study a more general version of these two objectives named the *l-centrum* objectives. *l-centrum* problems were first introduced by Slater (1978) and were later studied in other literature such as Tamir (2001); Peeters (1998). This set of problems is also a subset of the *ordered k-median* problems (Sornat 2019; Chakrabarty and Swamy 2018, 2019; Byrka, Sornat, and Spoerhase 2018; Kalcsics et al. 2002). In fact, *ordered k-median* objectives can be represented as convex combinations of the *l-centrum* objectives as discussed in Sornat (2019); Chakrabarty and Swamy (2018, 2019). This means that if we were to combine any *l-centrum* objectives by convex combination into a new objective, Chakrabarty and Swamy (2019) gives a $(5 + \epsilon)$ approximation for this new objective.

Note that the above work on *l-centrum* problems only considered the approximation ratio for a single objective at a time, but our goal is to approximate multiple objectives simultaneously. For this goal, Kumar and Kleinberg (2006) showed the existence of a simultaneous 4-approximation for all *l-centrum* objectives, and the results in Goel, Hulett, and Krishnaswamy (2018); Gkatzelis, Halpern, and Shah (2020) imply a similar 3-approximation. We provide a much simpler mechanism for obtaining this 3-approximation than the one in Gkatzelis, Halpern, and Shah (2020).[3] However, our main focus is on improving the upper bound with respect to the *number* of *l-centrum* objectives to be simultaneously approximated. We provide better approximations when the number of objectives to be approximated is small (instead of being *all* the *l-centrum* objectives), with a much more detailed tight analysis for two objectives.

Similar questions about facility location and voting have also been studied in mechanism design. For instance, there has been a significant amount of work in mechanism design considering the approximation ratio for strategy-proof mechanisms for placing a single facility (Walsh 2021; Alon et al. 2010; Feldman, Fiat, and Golomb 2016; Tang et al. 2020). Note that much of the previous work in this area studied only the 1D real line metric (e.g., Walsh (2021); Aziz et al. (2021)), while we look at general arbitrary metric spaces. For simultaneously approximating two objectives, Walsh (2021) showed that it is always possible to obtain a (3,3)-approximation for minimax and minisum for clients and facilities on a line. In addition, they also showed that no deterministic and strategy-proof mechanism can do better than 3-approximation for either of the two objectives.

Finally, our work also applies to spatial voting instead of facility location, where voters and candidates are located in a metric space, and the goal is to choose a candidate which minimizes some objectives over the voters (Merrill III, Merrill, and Grofman 1999; Enelow and Hinich 1984). Perhaps the most relevant work to ours in this space is the work on *distortion*, where instead of knowing the voters' exact locations, each voter only provides their ordinal preferences for the candidates (see the survey Anshelevich et al. (2021)). As part of this work, Anshelevich and Postl (2017); Feldman, Fiat, and Golomb (2016) showed that Random Dictatorship has an approximation ratio of 3 in a general metric space for minisum. More generally, the results in Gkatzelis, Halpern, and Shah (2020) imply a simultaneous 3-approximation for all l-centrum objectives, by choosing a candidate using a somewhat complex, but deterministic, mechanism. One of our goals is to improve this upper bound of 3 for simultaneously approximating multiple objectives. Because of this, just as in most work on distortion, our main focus is not on strategyproofness. Instead we study how well multiple objectives can be approximated simultaneously, even if we are given all the necessary information.

[2]Such set contains Pareto optimal solutions, or efficient solutions such that in any of these solutions, none of the objectives can be improved without simultaneously worsening any other objectives (Ehrgott and Gandibleux 2000; Farahani, SteadieSeifi, and Asgari 2010).

[3]Although mostly concerned with different questions, the 3-approximation mechanism we use can also be obtained as an easy consequence of the results in Goel, Hulett, and Krishnaswamy (2018).

2 Preliminaries and Notation

Consider the facility location problem where we are given the set of client locations $\mathcal{C}$ of size n, and the set of possible facility locations $\mathcal{F}$ in a metric space d. We want to select a location $A \in \mathcal{F}$ to place a facility such that the placement would simultaneously minimize some set of objectives. The kind of objectives that we are particularly interested in is the summation of the top k distances from the clients to the chosen facility location. More formally, suppose we order the n clients $a^1, a^2, a^3, \ldots, a^n$ so that

$$d(a^1, A) \geq d(a^2, A) \geq \ldots \geq d(a^n, A).$$

Then, define the k-centrum objective $c_k(A)$ to be the cost for choosing facility location A when considering the top k client-facility distances, i.e.,

$$c_k(A) = \sum_{i=1}^{k} d(a^i, A).$$

Our goal is to minimize $c_k(A)$ for multiple k simultaneously. Denote the optimal facility location for the k-centrum objective by O_k. It will be useful to denote by a_k^i the client that is i'th farthest from O_k, i.e.,

$$d(a_k^1, O_k) \geq d(a_k^2, O_k) \geq \ldots \geq d(a_k^n, O_k).$$

We are given a set of objectives to minimize, represented by a a set of distinct positive integers $\mathcal{K} = \{k_1, k_2, \cdots, k_q\}$, with each of its elements less than or equal to n. This means that we want to simultaneously minimize all objectives c_{k_i} for $k_i \in \mathcal{K}$. We slightly abuse notation and refer to $\mathcal{K}$ as the set of objectives, and say that an objective c_k is in $\mathcal{K}$ when $k \in \mathcal{K}$. However, in order to simultaneously minimize the objectives in $\mathcal{K}$, we would have to make some trade-offs such that the chosen facility location may not be the optimal location for some of the objectives.[4] We thus define the approximation ratio for choosing facility location A with respect to objective c_k as

$$\alpha_k(A) = \frac{c_k(A)}{c_k(O_k)} \geq 1$$

Therefore, by choosing facility location A, we would obtain a $(\alpha_{k_1}(A), \alpha_{k_2}(A), \cdots, \alpha_{k_q}(A))$ approximation for minimizing the set of objectives $\mathcal{K}$. As discussed in the Introduction, our goal is to establish that we can always choose some A so that all these values are small simultaneously.

3 Simultaneously Approximating Two Objectives

We will first consider the case where there are only two objectives. Let $|\mathcal{C}| = n, 1 \leq k < p \leq n$, we then wish to simultaneously minimize $c_k(A)$ and $c_p(A)$. Our goal is to find some $A \in \mathcal{F}$ such that both $\alpha_k(A)$ and $\alpha_p(A)$ are small. In fact, with this goal in mind, we can obtain the following result:

Theorem 3.1. *For $1 \leq k < p \leq n$, given the optimal facility location O_k that minimizes c_k and the optimal facility location O_p that minimizes c_p, we have that $\alpha_k(O_p) \leq \frac{1}{\alpha_p(O_k)} + 2$.*

Proof sketch. Suppose $1 \leq k < p \leq n$, let A be a facility location such that $A \in \mathcal{F}$. In the full version (Han, Jerrett, and Anshelevich 2022) we show two simple lemmas showing that $\sum_{i=1}^{p} d(a_k^i, A) \leq c_p(A)$ and $c_p(A) \leq \frac{p}{k} \cdot c_k(A)$. Note that these lemmas imply that $\sum_{i=1}^{p} d(a_k^i, O_k) \leq c_p(O_k)$ and $c_p(O_k) \leq \frac{p}{k} \cdot c_k(O_k)$. Then, note that all the clients and possible facility locations are located in metric space d, so we have

$$\begin{aligned} c_k(O_p) &= \sum_{i=1}^{k} d(a_p^i, O_p) \\ &\leq \sum_{i=1}^{k} \left[d(a_p^i, O_k) + d(O_k, O_p)\right] \\ &= \sum_{i=1}^{k} d(a_p^i, O_k) + k \cdot d(O_k, O_p). \end{aligned}$$

Then, since we have $\sum_{i=1}^{p} d(a_k^i, O_k) \leq c_p(O_k)$ and triangle inequality induced by metric space d, we have

$$\begin{aligned} c_k(O_p) &\leq c_k(O_k) + k \cdot d(O_k, O_p) \\ &\leq c_k(O_k) + \frac{k}{p}\left[\sum_{i=1}^{p} d(a_k^i, O_k) + \sum_{i=1}^{p} d(a_k^i, O_p)\right] \\ &\leq c_k(O_k) + \frac{k}{p}\left[c_p(O_k) + c_p(O_p)\right] \\ &= c_k(O_k) + \frac{k}{p}\left[c_p(O_k) + c_p(O_k) \cdot \frac{c_p(O_p)}{c_p(O_k)}\right] \\ &= c_k(O_k) + \frac{k}{p}\left[c_p(O_k) + \frac{c_p(O_k)}{\alpha_p(O_k)}\right] \\ &= c_k(O_k) + \frac{k}{p}\left[\left(1 + \frac{1}{\alpha_p(O_k)}\right) c_p(O_k)\right]. \end{aligned}$$

Recall that $c_p(O_k) \leq \frac{p}{k} \cdot c_k(O_k)$, then we can see that

$$\begin{aligned} c_k(O_p) &\leq c_k(O_k) + \frac{k}{p}\left[\left(1 + \frac{1}{\alpha_p(O_k)}\right) \cdot \frac{p}{k} c_k(O_k)\right] \\ &= c_k(O_k) + \left(1 + \frac{1}{\alpha_p(O_k)}\right) c_k(O_k) \\ &= \left(2 + \frac{1}{\alpha_p(O_k)}\right) c_k(O_k). \end{aligned}$$

Now, divide both side by $c_k(O_k)$, then

$$\begin{aligned} \frac{c_k(O_p)}{c_k(O_k)} &\leq \left(2 + \frac{1}{\alpha_p(O_k)}\right) \frac{c_k(O_k)}{c_k(O_k)} \\ \alpha_k(O_p) &\leq \frac{1}{\alpha_p(O_k)} + 2, \end{aligned}$$

[4] Here note that we assume $c_1(A) > 0$; this means that the cost function is always positive. Otherwise we should just choose the facility location A such that $c_1(A) = 0$. This indicates that all clients are at most 0 distance away from A, which means that all clients are at the same location as A given they are located in a metric space. In other words, A would be the optimal facility location for all objectives in $\mathcal{K}$.

as desired. □

The above theorem indicates that by picking either O_k or O_p, the values of $\alpha_k(O_p)$ and $\alpha_p(O_k)$ cannot be simultaneously large. In other words, either setting A to be O_k would ensure that both $\alpha_k(A)$ and $\alpha_p(A)$ are small or setting A to be O_p would.

Note that Theorem 3.1 immediately implies the following corollaries:

Corollary 3.1.1. *For $1 \leq k < p \leq n$,*

1. *By choosing the optimal facility location O_p that minimizes c_p, we obtain a $(3, 1)$ approximation for simultaneously minimizing c_k and c_p.*
2. *There always exists a facility location $A \in \mathcal{F}$ such that choosing A would give a $1 + \sqrt{2}$ approximation both for minimizing c_k and minimizing c_p. In fact, we would either get a $(1, 1 + \sqrt{2})$ approximation by choosing O_k or a $(1 + \sqrt{2}, 1)$ approximation by choosing O_p. In other words, at least one of $\alpha_k(O_p)$ or $\alpha_p(O_k)$ is always less than or equal to $1 + \sqrt{2}$.*

The above results show that it is always possible to approximate any pair of our objectives to within a factor of $1 + \sqrt{2}$ of optimum. However, it is natural to think that there exists some relationship between this approximation factor, and how similar the objectives are. Naturally, as the difference between k and p becomes smaller, we would expect that both $\alpha_p(O_k)$ and $\alpha_k(O_p)$ would also become smaller.

In fact, we can form tighter bounds than what we have shown in Theorem 3.1. We begin by looking at the case where p is at least twice as large as k. The result follows.

Theorem 3.2. *For $1 \leq k < p \leq n$, $\frac{k}{p} \leq \frac{1}{2}$, given the optimal facility location O_k that minimizes c_k and the optimal facility location O_p that minimizes c_p, we have that $\alpha_k(O_p) \leq \frac{1}{\alpha_p(O_k)} + 2 - 2 \cdot \frac{k}{p}$.*

See the full version (Han, Jerrett, and Anshelevich 2022) of this paper for the full proof.

Note that this result is in a form similar to what we have shown in Theorem 3.1 but with an offset of $-2 \cdot \frac{k}{p}$. What this means is that if we know the value of $\alpha_p(O_k)$, then the value of $\alpha_k(O_p)$ would be further restricted by this offset comparing to the result in Theorem 3.1. In addition, we can also see that the bigger $\frac{k}{p}$ becomes, the smaller the right-hand side value of the inequality becomes. In other words, assume that $\frac{k}{p} \leq \frac{1}{2}$, as the difference between k and p becomes smaller, the upper bound of $\alpha_k(O_p)$ would also becomes smaller given a fixed value of $\alpha_p(O_k)$.

However, we still have not found the relationship among $\frac{k}{p}, \alpha_k(O_p)$ and $\alpha_p(O_k)$ when $\frac{1}{2} < \frac{k}{p} \leq 1$. In order to show the underlying relationship between these values, we will utilize a different (much simpler) method from what we have been using for Theorem 3.1 and Theorem 3.2 (see the full version Han, Jerrett, and Anshelevich (2022)), which yields the following results.

Theorem 3.3. *For $1 \leq k < p \leq n$, given the optimal facility location O_k that minimizes c_k and the optimal facility location O_p that minimizes c_p, we have that $\alpha_k(O_p) \leq \frac{p}{k} \cdot \frac{1}{\alpha_p(O_k)}$ and $\alpha_p(O_k) \leq \frac{p}{k} \cdot \frac{1}{\alpha_k(O_p)}$.*

Interestingly, since the proof for Theorem 3.3 (see the full version of this paper (Han, Jerrett, and Anshelevich 2022) for the full proof) does not use any properties of metric spaces, it is true even under non-metric spaces. In addition, note that as the difference between p and k becomes smaller, the value of $\frac{p}{k}$ becomes smaller. This means that the upper bound for $\alpha_k(O_p)$ would also become smaller if the value of $\alpha_p(O_k)$ is given. And vice versa, the upper bound for $\alpha_p(O_k)$ would become smaller if the value of $\alpha_k(O_p)$ is given. Now that we have obtained bounds over all $\frac{k}{p} \in (0, 1]$, which is equivalent to $\frac{p}{k} \in [1, \infty)$ from Theorem 3.2 and Theorem 3.3, we can conclude the following results:

Theorem 3.4. *For $1 \leq k < p \leq n$, let $x = \frac{p}{k}$. Define $f : [1, \infty) \to \mathbb{R}$ as*

$$f(x) = \begin{cases} \sqrt{x} & 1 \leq x \leq 4 \\ 1 - x^{-1} + \sqrt{x^{-2} - 2x^{-1} + 2} & x > 4 \end{cases}$$

For some fixed x, let $\beta = f(x)$. Then there exists a facility location A in $\mathcal{F}$ such that choosing A would give a β approximation simultaneously for both minimizing c_k and minimizing c_p. In fact, we would either get a $(1, \beta)$ approximation by choosing O_k or a $(\beta, 1)$ approximation by choosing O_p. In other words, at least one of $\alpha_k(O_p)$ and $\alpha_p(O_k)$ is less than or equal to β. Moreover, this result is tight: for each x we give an instance such that all locations in $\mathcal{F}$ are no better than a β approximation for at least one of the objectives.

The above theorem means that $f(\frac{p}{k})$ is a tight upper bound for the approximation ratio we can obtain for two simultaneous objectives in a general metric space.

As a result, as shown in Figure 1, the function $f(x), x = \frac{p}{k}$ we obtained from Theorem 3.4 is continuous and monotone increasing over $\frac{p}{k} \geq 1$. In addition, note that when the difference between k and p is sufficiently large such that the value of $\frac{p}{k}$ approaches $+\infty$, we have $\beta \approx f(\infty) = 1 + \sqrt{2}$, which matches the second result in Corollary 3.1.1. Moreover, note that as the value of $\frac{p}{k}$ approaches 1, the value of β also approaches $f(1) = 1$. This shows that the upper bound of the smaller value of $\alpha_k(O_p)$ and $\alpha_p(O_k)$ would approach 1 as the difference between k and p becomes smaller as we expected. In other words, there must always exist an outcome such that the approximation ratio for both c_k and c_p is between 1 and $1 + \sqrt{2}$, this bound is tight, and in fact choosing either O_k or O_p is enough to achieve it.

4 Simultaneously Approximating Multiple Objectives

Now that we have found a tight upper bound for the approximation ratio for two objectives, we want to see what would happen if we have more objectives. Assume we have a set of $q \geq 2$ distinct integers in $[1, n]$, $|\mathcal{C}| = n$, $\mathcal{K} = \{k_1, k_2, \cdots, k_q\}$, arranged in increasing order such that $k_1 < k_2 < \cdots < k_q$. Then, the set of objectives that we would like to simultaneously optimize is $\mathcal{K}$ [5]. First, immediately from Corollary 3.1.1, we can get the following result:

Theorem 4.1. *Consider the optimal facility location O_n that minimizes c_n. By picking O_n, we obtain a 3 approximation for all other objectives c_k for $k \leq n$.*

While the above theorem gives us a very simple way of obtaining a 3-approximation for all objectives (and in fact can also be obtained as a simple consequence of the results in Goel, Hulett, and Krishnaswamy (2018)), we are interested in a more fine-grained analysis of when better approximations are possible. What if we are only interested in approximating a few objectives simultaneously, instead of all of them (in other words, what if $|\mathcal{K}|$ is small)? Or what if the set $\mathcal{K}$ has some nice properties? Towards answering these questions, we first make the following observation from Theorem 3.3:

Corollary 4.1.1. *For $1 \leq k < p \leq n$, we have $\alpha_k(O_p) \leq \frac{p}{k}$ and $\alpha_p(O_k) \leq \frac{p}{k}$.*

Corollary 4.1.1 indicates that when the difference between any two objectives c_k and c_p is sufficiently small, both $\alpha_p(O_k)$ and $\alpha_k(O_p)$ would also be small. One direct observation we can see from this is when $\mathcal{K} = \{k, k+1, \cdots, 4k-1, 4k\}$ for some $1 \leq k \leq n$, by picking O_{2k} we can get a 2 approximation for every other objective in $\mathcal{K}$. The reason for this is because $2k$ is of a factor of 2 larger than the smallest element in $\mathcal{K}$ and of a factor of 2 smaller than largest element in $\mathcal{K}$. Then by Corollary 4.1.1 we must have for any $k \in \mathcal{K} \setminus \{2k\}$, $\alpha_k(O_{2k}) \leq 2$. Note that this is a better result than what we have shown in Theorem 4.1 but is only true for special cases when none of the objectives in $\mathcal{K}$ are very different. Now, we want to see if we can obtain a better result for multiple objectives in general. To do this, we will first construct a graph representation for this problem.

We construct a complete directed graph $G = (V, E)$ as follows. First, for each $k \in \mathcal{K}$, we will make a node representing O_k, which is the optimal facility location for objective c_k. For simplicity, we will denote this node by O_k. Then, for every pair $i, j \in \mathcal{K}$, $i \neq j$, we will make two edges (O_i, O_j) and (O_j, O_i) with weight $\alpha_j(O_i)$ and $\alpha_i(O_j)$ respectively. As an example, Figure 3 is an illustration of G for three objectives c_i, c_j and c_k. Note that for every node O_k, $k \in \mathcal{K}$, by choosing O_k, we would get a $(\alpha_{k_1}(O_k), \alpha_{k_2}(O_k), \cdots, \alpha_{k_q}(O_k))$ approximation for minimizing the set of objectives $\mathcal{K}$ but these values are exactly the weights of all the edges going out of node O_k. Therefore, instead of looking at individual approximation ratios and their relationship, we will utilize this graph representation.

Our goal is to find some value $\beta_q < 3$ such that for $|\mathcal{K}| = q$, there must exist some $k \in \mathcal{K}$ such that choosing O_k would be at worst a β_q approximation for every other objective in $\mathcal{K}$. Note that the objectives in $\mathcal{K}$ can be arbitrarily far apart, for example they could include the max c_1 and the sum c_n objectives. In order to find such β_q, we will first consider some β_j that satisfies some useful properties as follows.

[5] As defined in Section 2, the set of objectives considered is $\{c_{k_1}, c_{k_2}, \cdots, c_{k_q}\}$, but is denoted by the set of integers $\mathcal{K}$.

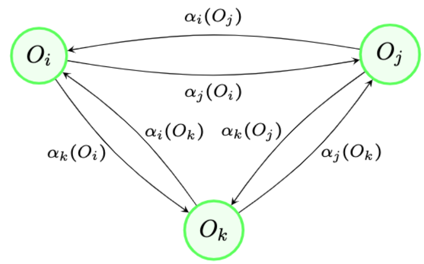

Figure 3: An example graph representation of three objectives, c_i, c_j and c_k. The green circles are nodes, arrows are directed edges with their respective weight labeled.

Lemma 4.2. *For every $j \geq 2, j \in \mathbb{N}$, there exists an unique β_j such that:*

1. $(\beta_j - 2)^{j-1}\beta_j = 1$

2. $1 + \sqrt{2} \leq \beta_j < 3$

Moreover, it holds that $\beta_{j+1} \geq \beta_j$ for $j \geq 2$.

See the full version (Han, Jerrett, and Anshelevich 2022) of this paper for the full proof. Then, we can obtain the following results.

Theorem 4.3. *Let $|\mathcal{K}| = q$, and consider β_q as defined in Lemma 4.2. Then there must exist some $k \in \mathcal{K}$ such that choosing O_k would be at worst a β_q approximation for every other objective in $\mathcal{K}$.*

Proof sketch. We begin by showing that for the graph representation G defined above, there does not exist a cycle of size j such that all of the edges in the cycle have weight strictly larger than β_j as defined in Lemma 4.2. To do this, as shown in the full version (Han, Jerrett, and Anshelevich 2022) of this paper, we have

$$c_k(O_k) \geq \gamma_p^k \cdot \frac{k}{p} \cdot c_p(O_p)$$

$$\gamma_p^k = \begin{cases} \alpha_p(O_k) & k < p \\ \alpha_p(O_k) - 2 & k > p \end{cases}$$

We then proceed by assuming there exists a cycle of size j in G such that all of its edges have weight larger than or equal to some β. More formally, for simplicity, denote the edges involved in this cycle by $(O_{k_1}, O_{k_2}), (O_{k_2}, O_{k_3}), \cdots, (O_{k_{j-1}}, O_{k_j}), (O_{k_j}, O_{k_1})$, where $\{k_1, k_2, \cdots, k_j\} \subseteq \mathcal{K}$. So we have $\alpha_{k_2}(O_{k_1}) \geq \beta, \alpha_{k_3}(O_{k_2}) \geq \beta, \cdots, \alpha_{k_j}(O_{k_{j-1}}) \geq \beta, \alpha_{k_1}(O_{k_j}) \geq \beta$. Now, combine these inequalities with the relationship between $c_k(O_k)$ and $c_p(O_p)$ discussed above, we can conclude that $(\beta - 2)^{j-1}\beta \leq 1$, which implies that $\beta \leq \beta_j$. This means that one of the $\alpha's$ has to be smaller than or equal to β_j. Therefore, we can conclude that there does not exist a cycle of size j such that all of the edges in the cycle have weight strictly larger than β_j.

Then, recall that in G, for some node O_k, $k \in \mathcal{K}$, every edge that goes out of O_k into some O_p has weight representing the approximation ratio using O_k with respect to objective c_p, denoted by $\alpha_p(O_k)$. Since the graph G is complete, we will show the above theorem by showing that there must exist some O_k, $k \in \mathcal{K}$, $|\mathcal{K}| = q$ such that all of the edges leaving O_k have weight less than or equal to β_q. We will show this by contradiction. Suppose otherwise, for every $k \in \mathcal{K}$, we must have at least one of the edges going out of O_k having weight larger than β_q. Denote this set of edges by E', note that $|E'| = q$. We will then consider these edges. Here note that since each node would not have an edge going into itself and every edge in E' has a distinct starting node, a subset of E' must be able to form a cycle of size $j \le q$, denote the cycle by C_j. As discussed before we know that G cannot have a cycle of size j such that all of it edge weights are larger than β_j. However, by Lemma 4.2, since we have $j \le q$, we must also have $\beta_q \ge \beta_j$ so all the edges in C_j have weight larger than β_j, which is a contradiction. Therefore, we can conclude that there must exist some $O_k, k \in \mathcal{K}$, such that all of the edges leaving O_k have weight less than or equal to β_q. This means that choosing O_k would be at worst a β_q approximation for every other objective in $\mathcal{K}$. □

Theorem 4.3 indicates that β_q defined in Lemma 4.2 is the upper bound of the approximation ratio for simultaneously approximating all the objectives in $\mathcal{K}$, which is exactly what we are looking for.

Note that unlike Theorem 4.1, Theorem 4.3 shows that an outcome can always be chosen with approximation ratio better than 3 for every other objective (see Figure 2). This is true since we have shown that $\beta_q < 3$ in Lemma 4.2. In fact, note that $\beta_2 = 1 + \sqrt{2}$, which matches with the tight bound we obtained in Theorem 3.4 for approximating two simultaneous objectives in the case where the difference between the two objectives is allowed to be arbitrary. As the number of objectives we care about grows, so does the approximation factor, and it can become strictly larger than $1 + \sqrt{2}$, as we show in the following proposition:

Proposition 4.1. *There exists an instance with 3 objectives such that all possible facility locations in $\mathcal{F}$ result in at least β_3 approximation for at least one of the three objectives.*

5 Choosing a Facility Location from Client Locations

So far in this paper we considered the general case when only some set of locations $\mathcal{F}$ allow the building of a facility. It is also interesting to consider the easier case, as was done in much of existing work (Chakrabarty and Swamy 2019; Alamdari and Shmoys 2017; Chakrabarty and Swamy 2018), when facilities are allowed to be built at any client location, i.e., when $\mathcal{C} = \mathcal{F}$. A lot of results become simpler with this assumption; for example it is easy to see that choosing any client location is immediately a 2-approximation for the c_1 (max distance) objective. Thus, choosing O_n immediately gives a $(2, 1)$ approximation for the c_1 (max) and c_n (sum) objectives, while for general $\mathcal{F}$ we have shown that nothing better than $1 + \sqrt{2}$ simultaneous approximation is possible. In fact, as an easy extension to the results shown in Gkatzelis, Halpern, and Shah (2020) (using the fact that when $\mathcal{C} = \mathcal{F}$, *decisiveness* as defined in that paper equals zero, see Proposition 6 in the arXiv version of their paper), we can always get a 2 simultaneous approximation for arbitrarily many *l-centrum* objectives. This is instead of the 3-approximation bound which we know holds for the general case when $\mathcal{C} \neq \mathcal{F}$. In this section we further extend our analysis from Section 3 to the case when $\mathcal{C} = \mathcal{F}$.

For simultaneously approximating two objectives, combined with the result implied in Gkatzelis, Halpern, and Shah (2020), when $\mathcal{C} = \mathcal{F}$, the simultaneous approximation bound is the same as in Figure 1 for $\frac{p}{k} \le 4$, but for $\frac{p}{k} > 4$ it simply equals 2. It is not difficult to show that this bound is tight, using similar examples as our bounds from Section 3 for general facility locations $\mathcal{F}$. Note that unlike all the previous results in this paper, the outcome that is chosen may *not* one of the optimum outcomes for the objectives in $\mathcal{K}$; to obtain the best simultaneous approximation it is often necessary to choose an outcome which is sub-optimal for all individual objectives.

6 Conclusion

We have shown that, when selecting a facility according to multiple competing interests, it is always possible to form an outcome approximating several competing objectives, at least as long as these objectives are one of the l-centrum objectives. For instance, both minimizing the maximum cost and minimizing the total cost can be simultaneously approximated within a ratio of $1+\sqrt{2}$. We can in fact extend the obtained upper bound of the approximation ratio to a broader range of problems. As discussed in the Introduction, if we can get an α simultaneous approximation ratio for a set of *l-centrum* objectives $\mathcal{K}$, then we can get an α-approximation for the corresponding *ordered 1-median* problems such that their objectives can be represented as convex combinations of the objectives in $\mathcal{K}$. This implies that there always exists a 3 approximation for all the *ordered 1-median* problems, and in fact all convex combinations of any q l-centrum objectives can be simultaneously approximated within our ratio β_q, with for example $\beta_2 = 1 + \sqrt{2}$.

However, there still exist questions left unanswered. For example, we do not know if the upper bound of the simultaneous approximation ratio for more than 3 objectives is tight, or if a better approximation is possible. More generally, it would be interesting to see if other types of objectives can be simultaneously approximated for these facility location and voting settings.

Acknowledgements

This work was partially supported by NSF award CCF-2006286.

References

Alamdari, S.; and Shmoys, D. 2017. A bicriteria approximation algorithm for the k-center and k-median problems. In *International Workshop on Approximation and Online Algorithms*, 66–75. Springer.

Alon, N.; Feldman, M.; Procaccia, A. D.; and Tennenholtz, M. 2010. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 35(3): 513–526.

Anshelevich, E.; Filos-Ratsikas, A.; Shah, N.; and Voudouris, A. A. 2021. Distortion in Social Choice Problems: The First 15 Years and Beyond. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence Survey Track.*

Anshelevich, E.; and Postl, J. 2017. Randomized social choice functions under metric preferences. *Journal of Artificial Intelligence Research*, 58: 797–827.

Aziz, H.; Lam, A.; Lee, B. E.; and Walsh, T. 2021. Strategyproof and proportionally fair facility location. *arXiv preprint arXiv:2111.01566*.

Bhargava, R.; Goel, A.; and Meyerson, A. 2001. Using approximate majorization to characterize protocol fairness. *ACM SIGMETRICS Performance Evaluation Review*, 29(1): 330–331.

Byrka, J.; Sornat, K.; and Spoerhase, J. 2018. Constant-factor approximation for ordered k-median. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, 620–631.

Chakrabarty, D.; and Swamy, C. 2018. Interpolating between k-Median and k-Center: Approximation Algorithms for Ordered k-Median. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

Chakrabarty, D.; and Swamy, C. 2019. Approximation algorithms for minimum norm and ordered optimization problems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 126–137.

Chan, H.; Filos-Ratsikas, A.; Li, B.; Li, M.; and Wang, C. 2021. Mechanism Design for Facility Location Problems: A Survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, 4356–4365. International Joint Conferences on Artificial Intelligence Organization.

Ehrgott, M.; and Gandibleux, X. 2000. A survey and annotated bibliography of multiobjective combinatorial optimization. *OR-spektrum*, 22(4): 425–460.

Enelow, J. M.; and Hinich, M. J. 1984. *The spatial theory of voting: An introduction*. CUP Archive.

Farahani, R. Z.; SteadieSeifi, M.; and Asgari, N. 2010. Multiple criteria facility location problems: A survey. *Applied mathematical modelling*, 34(7): 1689–1709.

Feldman, M.; Fiat, A.; and Golomb, I. 2016. On voting and facility location. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, 269–286.

Gkatzelis, V.; Halpern, D.; and Shah, N. 2020. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 1427–1438.

Goel, A.; Hulett, R.; and Krishnaswamy, A. K. 2018. Relating Metric Distortion and Fairness of Social Choice Rules. In *NetEcon'18: Proceedings of the 13th Workshop on Economics of Networks, Systems and Computation*.

Goel, A.; and Meyerson, A. 2006. Simultaneous optimization via approximate majorization for concave profits or convex costs. *Algorithmica*, 44(4): 301–323.

Han, Y.; Jerrett, C.; and Anshelevich, E. 2022. Optimizing Multiple Simultaneous Objectives for Voting and Facility Location. *arXiv preprint arXiv:2212.03467*.

Kalcsics, J.; Nickel, S.; Puerto, J.; and Tamir, A. 2002. Algorithmic results for ordered median problems. *Operations Research Letters*, 30(3): 149–158.

Kumar, A.; and Kleinberg, J. 2006. Fairness measures for resource allocation. *SIAM Journal on Computing*, 36(3): 657–680.

McGinnis, L. F.; and White, J. A. 1978. A single facility rectilinear location problem with multiple criteria. *Transportation Science*, 12(3): 217–231.

Merrill III, S.; Merrill, S.; and Grofman, B. 1999. *A unified theory of voting: Directional and proximity spatial models*. Cambridge University Press.

Nickel, S.; Puerto, J.; Rodríguez-Chía, A.; and Weissler, A. 2005. Multicriteria planar ordered median problems. *Journal of Optimization Theory and Applications*, 126(3): 657–683.

Ohsawa, Y. 1999. A geometrical solution for quadratic bicriteria location models. *European Journal of Operational Research*, 114(2): 380–388.

Peeters, P. H. 1998. Some new algorithms for location problems on networks. *European Journal of Operational Research*, 104(2): 299–309.

Roostapour, V.; Kiarazm, I.; and Davoodi, M. 2015. Deterministic algorithm for 1-median 1-center two-objective optimization problem. In *International Conference on Topics in Theoretical Computer Science*, 164–178. Springer.

Slater, P. J. 1978. Centers to centroids in graphs. *Journal of graph theory*, 2(3): 209–222.

Sornat, K. 2019. *Approximation Algorithms for Multiwinner Elections and Clustering Problems*. Ph.D. thesis, University of Wrocaw.

Tamir, A. 2001. The k-centrum multi-facility location problem. *Discrete Applied Mathematics*, 109(3): 293–307.

Tang, Z.; Wang, C.; Zhang, M.; and Zhao, Y. 2020. Mechanism design for facility location games with candidate locations. In *International Conference on Combinatorial Optimization and Applications*, 440–452. Springer.

Walsh, T. 2021. Strategy proof mechanisms for facility location at limited locations. In *Pacific Rim International Conference on Artificial Intelligence*, 113–124. Springer.