Unified Greedy Approximability Beyond Submodular Maximization

Yann Disser(✉) and David Weckbecker

TU Darmstadt, Darmstadt, Germany
{disser,weckbecker}@mathematik.tu-darmstadt.de

Abstract. We consider classes of objective functions of cardinality-constrained maximization problems for which the greedy algorithm guarantees a constant approximation. We propose the new class of γ-α-augmentable functions and prove that it encompasses several important subclasses, such as functions of bounded submodularity ratio, α-augmentable functions, and weighted rank functions of an independence system of bounded rank quotient – as well as additional objective functions for which the greedy algorithm yields an approximation. For this general class of functions, we show a tight bound of $\frac{\alpha}{\gamma} \cdot \frac{e^\alpha}{e^\alpha - 1}$ on the approximation ratio of the greedy algorithm that tightly interpolates between bounds from the literature for functions of bounded submodularity ratio and for α-augmentable functions. In particular, as a by-product, we close a gap in [Math.Prog., 2020] by obtaining a tight lower bound for α-augmentable functions for all $\alpha \geq 1$. For weighted rank functions of independence systems, our tight bound becomes $\frac{\alpha}{\gamma}$, which recovers the known bound of $1/q$ for independence systems of rank quotient at least q.

Keywords: Greedy algorithm · Cardinality-constrained maximization · Approximation ratio · Independence system · Submodularity ratio · Augmentability

1 Introduction

We consider cardinality-constrained maximization problems of the form

$$\begin{aligned} \max\ & f(X) \\ \text{s.t. } & |X| \leq k \\ & X \subseteq U, \end{aligned}$$

with a *monotone* objective function $f\colon 2^U \to \mathbb{R}_{\geq 0}$ over a finite ground set U. Additional constraints of the form $X \in \mathcal{X}$ can be modeled by the monotone objective $f'(X) := \max\{f(Y) | Y \in 2^X \cap \mathcal{X}\}$. In this way, every combinatorial, cardinality-constrained maximization problem with monotone objective function

Supported by DFG grant DI 2041/2.

I. Ljubić et al. (Eds.): ISCO 2022, LNCS 13526, pp. 299–311, 2022.
https://doi.org/10.1007/978-3-031-18530-4_22

can be captured, and we adopt this framework throughout the paper.[1] For example, the maximum weighted matching problem on a graph $G = (V, E)$ with edge weights $w\colon E \to \mathbb{R}_{\geq 0}$ yields the objective function

$$f(X \subseteq E) = \max\{\sum_{e \in M} w(e) | M \subseteq X, M \text{ is a matching in } G\}.$$

We focus on the performance of the *greedy algorithm* that iteratively produces a solution $S^{\mathrm{G}}_{f,k} = \{x_1, \ldots, x_k\}$ with

$$x_i \in \arg\max_{x \in U \setminus \{x_1,\ldots,x_{i-1}\}} f(\{x_1, \ldots, x_{i-1}\} \cup \{x\}),$$

for all $i \in [k] := \{1, \ldots, k\}$, i.e., it adds elements such that the increase in objective value is maximized in each step. The greedy algorithm is inherently incremental and may be regarded as the most natural approach for incrementally building up infrastructures that support changing active solutions (in the sense of the definition $f'(X)$ above). While this algorithm is widely used in practical applications, greedy solutions can be arbitrarily far away from optimal (e.g., for the knapsack problem). A natural question in this context is, for which objective functions f the greedy algorithm gives a good solution. We are interested in characterizing these objective functions.

Note that we consider the *adaptive* greedy solution $S^{\mathrm{G}}_{f,k}$ as opposed to the *non-adaptive* greedy solution $\tilde{S}^{\mathrm{G}}_{f,k} := S^{\mathrm{G}}_{f,\min\{k,\bar{k}\}}$, where $\bar{k} \in [|U|]$ is the smallest cardinality such that $f(S^{\mathrm{G}}_{\bar{k}} \cup \{x\}) = f(S^{\mathrm{G}}_{\bar{k}})$ for all $x \in U \setminus S^{\mathrm{G}}_{\bar{k}}$. In other words, the non-adaptive greedy algorithm terminates as soon as it cannot improve the solution further. This non-adaptive variant of the greedy algorithm has often been considered in the early literature (e.g., [15,16,22,23]). Note, that for submodular functions, i.e., functions with $f(X \cup Y) + f(X \cap Y) \leq f(X) + f(Y)$ for all $X, Y \subseteq U$, there is no difference between these two variants, and for our purposes both variants are interchangeable in the following sense.

Formally, we measure the quality of the greedy algorithm on a set of objectives $\mathcal{F}$ by the approximation ratio $\sup_{f \in \mathcal{F}} \max_{k \in [|U_f|]} f(S^*_{f,k})/f(X_{f,k})$, where U_f is the ground set of the function $f \in \mathcal{F}$, $S^*_{f,k} \in \arg\max_{X \subseteq U : |X| \leq k} f(X)$ denotes an optimum solution of cardinality at most k, and $X_{f,k} \in \{S^{\mathrm{G}}_{f,k}, \tilde{S}^{\mathrm{G}}_{f,k}\}$ refers to the (non-)adaptive greedy solution of cardinality k. We claim that the approximation ratios of both variants of the greedy algorithm coincide. To see this, observe that the non-adaptive setting is more restrictive, and that every lower bound instance in the non-adaptive setting can be made adaptive by introducing additional elements that add a vanishingly small but positive objective value when added to every solution. This implies that all our bounds on the approximation ratio of the (adaptive) greedy algorithm immediately apply to both variants.

From now on, we write $S^{\mathrm{G}}_k := S^{\mathrm{G}}_{f,k}$ and $S^*_k := S^*_{f,k}$, whenever f is clear from the context. In these terms, we are interested in characterizing the set of objectives for which the greedy algorithm has a bounded approximation ratio. Known

[1] Note that the objective function f may be computationally hard to evaluate. If we assume that the greedy algorithm has oracle access to f, it requires $O(|U|k)$ queries to the oracle.

examples include the objectives of maximum (weighted) (b-)matching, maximum (weighted) coverage, and many more [2,3,8,18,30], and we additionally introduce a multi-commodity flow problem (Sect. 2), where the greedy algorithm yields an approximation.

A well-known class of functions for which the greedy algorithm has a bounded approximation ratio of (exactly) $\frac{\mathrm{e}}{\mathrm{e}-1}$ are monotone, submodular functions [22]. This class includes the maximum coverage problem, but fails to capture many other greedily approximable settings. See Fig. 1 along with the following.

Das and Kempe [8] introduced the class of functions of bounded *submodularity ratio* as a generalization of submodular functions. Importantly, its definition depends on the greedy solutions for different cardinalities. We adapt and weaken the definition from [8] for consistency, by restricting ourselves to greedy solutions and by minimizing over all cardinalities.

Definition 1 ([8]). *The* weak submodularity ratio *of $f\colon 2^U \to \mathbb{R}_{\geq 0}$ is (using $\frac{0}{0} := 1$)*

$$\gamma(f) := \min_{X \in \{S_0^G, \ldots, S_{\bar{k}}^G\}, Y \subseteq U \setminus X} \frac{\sum_{y \in Y}(f(X \cup \{y\}) - f(X))}{f(X \cup Y) - f(X)} \in [0, 1].$$

Das and Kempe [8] showed an upper bound of $\frac{\mathrm{e}^\gamma}{\mathrm{e}^\gamma - 1}$ on the approximation ratio of the greedy algorithm for the set of all monotone functions with submodularity ratio at least $\gamma > 0$, and Bian et al. [3] extended this to a tight bound that is additionally parameterized by the curvature of the objective. Since submodular functions have submodularity ratio 1, this bound generalizes the submodular bound. Crucially, it is easy to verify that these results carry over to the set $\tilde{\mathcal{F}}_\gamma$ of all monotone functions with *weak* submodularity ratio at least $\gamma > 0$.[2]

Another generalization of submodularity was proposed by Bernstein et al. [2]. We extend the definition by a weakened variant in order to bring it more in line with Definition 1.

Definition 2 ([2]). *The function $f\colon 2^U \to \mathbb{R}_{\geq 0}$ is* (weakly) *α-augmentable for $\alpha \geq 1$, if, for every $X \subseteq U$ ($X \in \{S_0^G, \ldots, S_{\bar{k}}^G\}$) and $Y \subseteq U$ with $Y \not\subseteq X$, there exists an element $y \in Y \setminus X$ with*

$$f(X \cup \{y\}) - f(X) \geq \frac{f(X \cup Y) - \alpha f(X)}{|Y|}.$$

Bernstein et al. showed that the greedy algorithm has an approximation ratio of at most $\alpha \cdot \frac{\mathrm{e}^\alpha}{\mathrm{e}^\alpha - 1}$ on the set $\mathcal{F}_\alpha$ of monotone, α-augmentable functions, for $\alpha \geq 1$, and that this bound is tight for $\alpha \in \{1, 2\}$ and in the limit $\alpha \to \infty$. Since submodular functions are 1-augmentable, this bound again generalizes the submodular bound. The class of α-augmentable problems captures the objective of the maximum (weighted) α-dimensional matching problem, which is not submodular. In this paper, we introduce a natural α-commodity flow variant that

[2] Here and throughout we use the notation $\tilde{\mathcal{F}}$ as opposed to $\mathcal{F}$ to refer to a function class based on a *weak* definition.

is α-augmentable, and we prove a tight lower bound on the approximation ratio for all $\alpha \geq 1$.

Another well-known setting, besides submodularity, where the greedy algorithm has a bounded approximation ratio, are weighted rank functions of independence systems of bounded rank quotient [17]. An *independence system* is a tuple $(U, \mathcal{I} \subseteq 2^U)$, where $\mathcal{I}$ is closed under taking subsets and $\emptyset \in \mathcal{I}$. For a given weight function $w\colon U \to \mathbb{R}_{\geq 0}$, the *weighted rank function* of $(U, \mathcal{I})$ is given by $f(X) = \max\{\sum_{x \in Y} w(x) | Y \in \mathcal{I} \cap 2^X\}$. The *rank quotient* of an independence system $(U, \mathcal{I})$ is $q(U, \mathcal{I}) := \min_{X \subseteq U} \min_{B, B' \in \mathcal{B}(X)} |B|/|B'|$, where $\frac{0}{0} := 1$, and the set $\mathcal{B}(X)$ of all *bases* of some set $X \subseteq U$ is the set of inclusion-wise maximal subsets of $\mathcal{I} \cap 2^X$, i.e., $\mathcal{B}(X) := \{B \in \mathcal{I} \cap 2^X | \forall x \in X \setminus B\colon B \cup \{x\} \notin \mathcal{I}\}$. Jenkyns [15] and Korte and Hausmann [16] showed that the greedy algorithm has an approximation ratio of exactly $1/q$ on the set $\mathcal{F}_q$ of all weighted rank functions of independence systems with rank quotient at least $q > 0$.[3]

Our Results. Our goal is to unify and to generalize the above classes of functions on which the greedy algorithm has a bounded approximation ratio. To this end, we first observe that each one of the classes $\tilde{\mathcal{F}}_\gamma$, $\mathcal{F}_\alpha$, and $\mathcal{F}_q$ uniquely captures greedily approximable objectives (cf. Fig. 1). In particular, we construct a natural α-augmentable variant of multi-commodity flow that does not have bounded (weak) submodularity ratio (for $\alpha \in \mathbb{N} \setminus \{1\}$) and cannot be expressed as the maximization of a weighted rank function. Besides the α-dimensional matching problem, to our knowledge, the problem introduced in Sect. 2 is the only other natural α-augmentable problem to date.

Proposition 1. *For every $\gamma, q \in (0, 1)$ and $\alpha \geq 1$, it holds that*

$$\tilde{\mathcal{F}}_\gamma \not\subseteq (\mathcal{F}_\alpha \cup \mathcal{F}_q) \quad \text{and} \quad \mathcal{F}_\alpha \not\subseteq (\tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_q) \quad \text{and} \quad \mathcal{F}_q \not\subseteq (\tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_\alpha).$$

This motivates the following definition to consolidate all three classes.

Definition 3. *The function $f\colon 2^U \to \mathbb{R}_{\geq 0}$ is (weakly) γ-α-augmentable for $\gamma \in (0, 1]$ and $\alpha \geq \gamma$ if, for all sets $X \subseteq U$ ($X \in \{S_0^G, ..., S_k^G\}$) and all $Y \subseteq U$ with $Y \not\subseteq X$, there exists $y \in Y$ with*

$$f(X \cup \{y\}) - f(X) \geq \frac{\gamma f(X \cup Y) - \alpha f(X)}{|Y|}.$$

Note that we need to consider the weak variant of this definition if we hope to encompass the class $\tilde{\mathcal{F}}_\gamma$, which enforces its defining property only for "greedy sets", however, any upper bound on the approximation ratio immediately carries over to the same bound in the stronger definition. Also note that γ-α-augmentability only requires $\alpha \geq \gamma$, unlike α-augmentability where $\alpha \geq 1$. This is in line with the definitions of α-augmentability where $\gamma = 1$ and of the submodularity ratio where $\alpha = \gamma$. We let $\tilde{\mathcal{F}}_{\gamma,\alpha}$ denote the set of all weakly

[3] Note that we abuse notation, since, e.g., $\mathcal{F}_\alpha \neq \mathcal{F}_q$ for $\alpha = q = 1$. However, the set of functions we are referring to will always be clear by the naming of the indices.

γ-α-augmentable functions. The first part of our main result is that this set encompasses all functions in $\tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_\alpha \cup \mathcal{F}_q$ and captures additional functions (cf. Fig. 1). Formally, we show the following.

Theorem 1. *For every $\gamma, q \in (0,1]$, every $\gamma' \in (0,1)$, every $\alpha \geq 1$, and every $\alpha' \geq \gamma'$, it holds that*

$$\tilde{\mathcal{F}}_{\gamma,\max\{\alpha,1/q\}} \supseteq \tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_\alpha \cup \mathcal{F}_q \qquad \text{and} \qquad \tilde{\mathcal{F}}_{\gamma',\alpha'} \nsubseteq \tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_\alpha \cup \mathcal{F}_q.$$

Note that α' and γ' in Theorem 1 do not depend on α, γ and q. The second part of our main result is a tight bound on the approximation ratio of the greedy algorithm on $\tilde{\mathcal{F}}_{\gamma,\alpha}$ (cf. Theorem 6 and Proposition 4).

Theorem 2. *The approximation ratio of the greedy algorithm on the class $\tilde{\mathcal{F}}_{\gamma,\alpha}$ of monotone, weakly γ-α-augmentable functions, with $\gamma \in (0,1]$ and $\alpha \geq \gamma$, is exactly*

$$\frac{\alpha}{\gamma} \cdot \frac{\mathrm{e}^\alpha}{\mathrm{e}^\alpha - 1}.$$

Importantly, this bound recovers exactly the known bound for functions of bounded submodularity ratio, since $\tilde{\mathcal{F}}_\gamma \subseteq \tilde{\mathcal{F}}_{\gamma,\gamma}$, as well as the known bound for α-augmentable functions, since $\mathcal{F}_\alpha \subseteq \tilde{\mathcal{F}}_{1,\alpha}$. In that sense, our new bound interpolates tightly between these two bounds and generalizes them. In addition, our tight lower bound for $\tilde{\mathcal{F}}_{1,\alpha}$ is obtained with an α-augmentable function. This means that, in particular, we are able to close the gap left in [2], by showing a tight lower bound for α-augmentable objectives, for all $\alpha \geq 1$.

Corollary 1. *The approximation ratio of the greedy algorithm on the class $\mathcal{F}_\alpha$ of monotone, α-augmentable functions is exactly $\alpha \cdot \frac{\mathrm{e}^\alpha}{\mathrm{e}^\alpha - 1}$ for all $\alpha \geq 1$.*

Finally, we are also able to show a tight bound of α/γ for γ-α-augmentable, weighted rank functions on independence systems. Since $\mathcal{F}_q \subseteq \tilde{\mathcal{F}}_{1,1/q}$ (by Theorem 1), our bound recovers exactly the known bound of $1/q$ for the approximation ratio of the greedy algorithm when the rank quotient is bounded from below by $q > 0$. This means that the class of monotone, weakly γ-α-augmentable functions truly unifies and generalizes the three classes $\tilde{\mathcal{F}}_\gamma$, $\mathcal{F}_\alpha$, and $\mathcal{F}_q$ of greedily approximable functions (cf. Fig. 1). Note that, in particular, the lower bound is tight already for α-augmentable functions, which implies a tight bound of α for the approximation ratio of the greedy algorithm on α-augmentable weighted rank functions.

Theorem 3. *Let $\mathcal{F}_{\mathrm{IS}} := \bigcup_{q \in (0,1]} \mathcal{F}_q$ be the set of weighted rank functions on some independence system. The approximation ratio of the greedy algorithm on the class $\tilde{\mathcal{F}}_{\gamma,\alpha} \cap \mathcal{F}_{\mathrm{IS}}$, with $\gamma \in (0,1]$ and $\alpha \geq \gamma$, is exactly $\frac{\alpha}{\gamma}$.*

The proofs of all results can be found in the full version [9].

Related Work. We can view our cardinality-constrained maximization framework as a special case of maximization over an independence system. In particular, the cardinality-constraint can be expressed as a uniform matroid contraint [17].

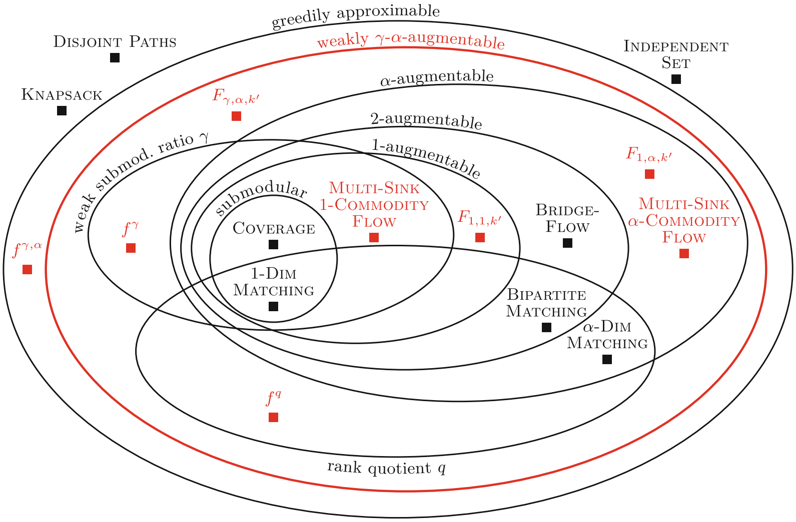

Fig. 1. Relation of the different problem classes. Newly introduced classes and problems are marked in red. The parameter k' is chosen sufficiently large, depending on γ and α. (Color figure online)

From that perspective, the most basic, non-trivial setting is the maximization of a linear (i.e., modular) objective over an independence system. Regarding the approximation ratio of the greedy algorithm, this classic setting is equivalent to the maximization of a weighted rank function, as considered in Theorem 3. This is easy to see by considering the non-adaptive variant of the greedy algorithm, and by observing that the greedy solution is guaranteed to remain feasible while the algorithm makes progress.

In that sense, the perfomance of the greedy algorithm for weighted rank function maximization has extensively been studied in the past. Rado [25] showed that the greedy algorithm is optimal for all weight functions if the underlying independence system is a matroid, and Edmonds [10] established the reverse implication. Jenkyns [15] extended this result by showing an upper bound of $1/q$ for the approximation ratio of the greedy algorithm on independence systems with rank quotient q, and Korte and Hausmann [16] gave a tight lower bound. Years later, Mestre [21] independently proved this tight bound for the subclass of k-extendible independence systems. Bouchet [4] gave a different generalization of the result by Rado and Edmonds by showing that the greedy algorithm remains optimal on symmetrical matroids.

Another prominent setting is the maximization of a submodular function over an independence system. Again, this includes cardinality-constrained maximization of a submodular objective, which is equivalent to submodular maximization over a uniform matroid. Nemhauser, Wolsey, and Fisher [23] showed that the greedy algorithm has a tight approximation ratio of $\frac{e}{e-1}$ for maximizing a monotone, submodular function under a cardinality-constraint. Krause et al. [19] observed that the approximation ratio is unbounded when maximizing the minimum of two monotone, submodular functions. Non-monotone submodular maximization over a cardinality-constraint (and knapsack constraints) was considered by Lee et al. [20]. Feldman et al. [14] analyzed a variant of the continuous greedy algorithm [28] and showed an upper bound on its approximation ratio of $(1/e - o(1))^{-1}$. This bound for the non-monotone case with cardinality-constraint was later improved by Buchbinder et al. [5] and Ene and Nguyen [11] by further adapting the (continuous) greedy algorithm. For maximizing a submodular function subject to k-extentible system and k-systems constraints, Feldman et al. [12,13] considered three variants of the greedy algorithm, a repeated greedy, a sample greedy and a simultaneous greedy. They were able to show approximation ratios of $k + O(1)$ for k-extendible system constraints and $k + O(\sqrt{k})$ for k-system constraints.

Maximization of a monotone, submodular function over a matroid was considered by Vondrák [28] and by Calinescu et al. [6], who showed that the continuous greedy algorithm has an approximation ratio of $\frac{\mathrm{e}}{\mathrm{e}-1}$ in this setting. Nemhauser, Wolsey, and Fisher [23], showed an upper bound of $p + 1$ for the regular greedy algorithm when maximizing over the intersection of p matroids. A generalization of this upper bound to the setting o maximizing subject to a p-system constraint was later proven by Calinescu et al. [6]. Conforti and Cornuejols [7] gave an upper bound of $p + c$ depending on the curvature c of the monotone submodular function – this interpolates between the submodular bound of [23] ($c = 1$) and the linear bound of [16] ($c = 0$). Vondrák [29] showed that the continuous greedy algorithm has an approximation ratio of at most $c\frac{\mathrm{e}^c}{\mathrm{e}^c-1}$ over an arbitrary matroid, and Sviridenko, Vondrák, and Ward [27] showed an improved upper bound of $\frac{\mathrm{e}}{\mathrm{e}-c}$ for the approximation ratio of a modified continuous greedy algorithm over a uniform matroid (i.e., a cardinality-constraint).

Other variants of the problem setting include the maximization of a monotone, submodular function over a knapsack constraint [26], and robust submodular maximization [1,24].

2 Weak Submodularity Ratio, α-Augmentability, and Independence Systems

In this section, we give an idea how to prove Proposition 1, i.e., how to separate the function classes $\tilde{\mathcal{F}}_\gamma$, $\mathcal{F}_\alpha$, and $\mathcal{F}_q$.

We start by introducing a natural α-commodity flow problem that models, e.g., production processes where output is limited by availability of all components. The objective of this problem is (exactly) α-augmentable, but, for $\alpha \in \mathbb{N} \setminus \{1\}$, does not have a bounded (weak) submodularity ratio and cannot be expressed as a weighted rank function over an independence system. This problem also gives a tight lower bound for the approximation ratio of the greedy algorithm on α-augmentable functions, for $\alpha \in \mathbb{N}$. We will extend this lower bound to all $\alpha \geq 1$ in Sect. 3.1, and thus close a gap left by [2].

Definition 4. *Let $G = (V, E)$ be a directed graph with source $s \in V$, sinks $T \subseteq V$, and arc capacities $\mu \colon E \to \mathbb{R}_{\geq 0}$. We define an s-T-flow to be a function $\vartheta \colon E \to \mathbb{R}_{\geq 0}$ that satisfies*

$$\begin{aligned} \vartheta(e) &\leq \mu(e) && \forall e \in E && \textit{(capacity constraint)},\\ \mathrm{ex}_\vartheta(v) &= 0 && \forall v \in V \setminus (\{s\} \cup T) && \textit{(flow conservation)},\\ \mathrm{ex}_\vartheta(t) &\geq 0 && \forall t \in T && \textit{(T are sinks)}, \end{aligned}$$

where (using $\delta^+(v) := (\{v\} \times V) \cap E$, $\delta^-(v) := (V \times \{v\}) \cap E$) the excess of a vertex $v \in V$ is defined as $\mathrm{ex}_\vartheta(v) := \sum_{e \in \delta^-(v)} \vartheta(e) - \sum_{e \in \delta^+(v)} \vartheta(e)$.

We extend this notion to multi-commodity flows, where each commodity has an independent capacity function.

Definition 5. *Let $\alpha \in \mathbb{N}$ and $G = (V, E)$ be a graph, let $s \in V$ and $T \subseteq V$, and let $\vec{\mu} = (\mu_i \colon E \to \mathbb{R}_{\geq 0})_{i \in [\alpha]}$ be capacity functions.* A *multicommodity-flow in G w.r.t. $\vec{\mu}$ is a tuple $\vec{\vartheta} = (\vartheta_1, ..., \vartheta_\alpha)$, where ϑ_i is an $s - T - flow$ in G with respect to capacities μ_i. The minimum-excess of the sink vertex $t \in T$ in $\vec{\vartheta}$ is*

$$\mathrm{minex}_{\vec{\vartheta}}(t) := \min_{i \in [\alpha]} \mathrm{ex}_{\vartheta_i}(t).$$

For convenience, we let $\mu(u, v) := \mu((u, v))$, $\vartheta(u, v) := \vartheta((u, v))$, and we let $\mathrm{ex}_\vartheta(V') := \sum_{v \in V'} \mathrm{ex}_\vartheta(v)$ for $V' \subseteq V$, and $\mathrm{minex}_{\vec{\vartheta}}(T') := \sum_{t \in T'} \mathrm{minex}_{\vec{\vartheta}}(t)$ for $T' \subseteq T$ in the following.

An instance of the problem Multi-Sink α-Commodity Flow, for $\alpha \in \mathbb{N}$, is given by a tuple $(G, s, T, \vec{\mu})$, where $G = (V, E)$ is a directed graph, $s \in V$ is a source vertex, $T \subseteq V$ contains sink vertices, and $\vec{\mu} = (\mu_i \colon E \to \mathbb{R}_{\geq 0})_{i \in [\alpha]}$ are capacity functions. The problem is to find a subset of sinks $X \subseteq T$ with $|X| = k$ that maximizes the objective function

$$f(X) = \max_{\vec{\vartheta} \in \mathcal{M}_{G,\mu}} \mathrm{minex}_{\vec{\vartheta}}(X),$$

where $\mathcal{M}_{G,\mu}$ denotes the set of all multicommodity-flows in G w.r.t. capacities $\vec{\mu}$.

Theorem 4. *For every $\alpha \in \mathbb{N}$, the objective of* Multi-Sink α-Commodity Flow *is monotone and* α-augmentable.

For $\alpha = 2$, Multi-Sink α-Commodity Flow problem is equivalent to the BridgeFlow problem considered in [2]. We generalize the tight lower bound construction for BridgeFlow to arbitrary $\alpha \in \mathbb{N}$ (see full version [9] for details).

With this, we obtain a lower bound for the approximation ratio of the greedy algorithm on $\mathcal{F}_\alpha$ for $\alpha \in \mathbb{N}$ that tightly matches the upper bound of [2], i.e., we obtain Corollary 1 for $\alpha \in \mathbb{N}$. In particular, it follows that the objective of Multi-Sink α-Commodity Flow is not β-augmentable for any $\beta < \alpha$. We will generalize the lower bound to all $\alpha \geq 1$ in Sect. 3.1.

Theorem 5. *For $\alpha \in \mathbb{N}$, the greedy algorithm has an approximation ratio of at least $\alpha \frac{e^\alpha}{e^\alpha - 1}$ for* Multi-Sink α-Commodity Flow.

2.1 Separating Function Classes

We are now ready to show Proposition 1 for $\alpha \in \mathbb{N} \setminus \{1\}$. The case $\alpha \geq 1$ will be addressed in Sect. 3.1.

In order to seperate $\mathcal{F}_\alpha$ for $\alpha \in \mathbb{N} \setminus \{1\}$, we have shown that the objective of Multi-Sink α-Commodity Flow does not have a (weak) submodularity ratio bounded away from zero, and cannot be represented as the weighted rank function of some independence system. Separating $\tilde{\mathcal{F}}_\gamma$ and $\mathcal{F}_q$ for every $\gamma, q \in (0,1)$ is done by defining a simple according function.

3 γ-α-Augmentability

In this section, we argue that the class $\tilde{\mathcal{F}}_{\gamma,\alpha}$ of weakly γ-α-augmentable functions unifies and generalizes the classes $\tilde{\mathcal{F}}_\gamma$, $\mathcal{F}_\alpha$, and $\mathcal{F}_q$. We start by proving the first half of Theorem 1. The second half will be shown in Sect. 3.1, together with lower bounds for the approximation ratio of the greedy algorithm.

Since (weak) γ-α-augmentability implies (weak) γ'-α'-augmentability for all $\gamma \geq \gamma'$ and $\alpha \leq \alpha'$, the following proposition implies the first part of Theorem 1.

Proposition 2. *For every $\gamma, q \in (0,1]$, and every $\alpha \geq 1$, it holds that*

$$\tilde{\mathcal{F}}_{1,\alpha} \supseteq \mathcal{F}_\alpha \qquad \text{and} \qquad \tilde{\mathcal{F}}_{\gamma,\gamma} \supseteq \tilde{\mathcal{F}}_\gamma \qquad \text{and} \qquad \tilde{\mathcal{F}}_{\gamma,\gamma/q} \supseteq \mathcal{F}_q.$$

Having shown that $\tilde{\mathcal{F}}_{\gamma,\alpha}$ subsumes the other three classes of functions, we now prove the upper bound of Theorem 2 for this class. Observe that the upper bound trivially carries over to the class of monotone, γ-α-augmentable (not weakly) functions.

Theorem 6. *The approximation ratio of the greedy algorithm on the class $\tilde{\mathcal{F}}_{\gamma,\alpha}$ of monotone, weakly γ-α-augmentable functions, with $\gamma \in (0,1]$ and $\alpha \geq \gamma$, is at most*

$$\frac{\alpha}{\gamma} \cdot \frac{\mathrm{e}^\alpha}{\mathrm{e}^\alpha - 1}.$$

3.1 A Critical Function

To obtain the tight lower bound of Theorem 2 for weakly γ-α-augmentable problems and to separate this class from $\tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_\alpha \cup \mathcal{F}_q$, we introduce a function that is inspired by a construction in [3] for the submodularity ratio.

We fix $\gamma \in (0,1]$ and $\alpha \geq \gamma$. Let $k \in \mathbb{N}$ with $k > \alpha$, and let $A = \{a_1, ..., a_k\}$ and $B = \{b_1, ..., b_k\}$ be disjoint sets. We set $U = A \cup B$, define $\xi_i := \frac{1}{k}(\frac{k-\alpha}{k})^{i-1}$ and let $h(x) := \frac{\gamma^{-1}-1}{k-1}x^2 + \frac{k-\gamma^{-1}}{k-1}x$. For our purpose, the important facts about h are $h(0) = 0$, $h(1) = 1$, $h(k) = \frac{k}{\gamma}$ and that h is convex. With this in mind, we define the function $F_{\gamma,\alpha,k} \colon 2^U \to \mathbb{R}_{\geq 0}$ by

$$F_{\gamma,\alpha,k}(X) = \max_{X' \subseteq X} \Big\{ \frac{h(|\{b_1\} \cap X'| \cdot |B \cap X'|)}{k} \Big(1 - \alpha \sum_{\substack{i \in [k]: \\ a_i \in A \cap X'}} \xi_i\Big) + \sum_{\substack{i \in [k]: \\ a_i \in A \cap X'}} \xi_i \Big\}.$$

We show that our modification of the function introduced in [3] retains the same structure in regard to greedy solutions.

Proposition 3. *For $i \in [k]$, the greedy algorithm picks the element a_i in iteration i, and, for $i \in [2k] \setminus [k]$, the greedy algorithm picks the element b_{i-k} in iteration i.*

With this, we can show that $F_{\gamma,\alpha,k}$ is weakly γ-α-augmentable.

Lemma 1. *For every $\gamma \in (0, 1]$, every $\alpha \geq \gamma$, and every $k \in \mathbb{N}$ with $k > \alpha$, it holds that $F_{\gamma,\alpha,k} \in \tilde{\mathcal{F}}_{\gamma,\alpha}$.*

It is straightforward to bound the approximation ratio of the greedy algorithm for $F_{\gamma,\alpha,k}$.

Proposition 4. *The approximation ratio of the greedy algorithm for maximizing the function $F_{\gamma,\alpha,k}$, with $\gamma \in (0, 1]$, $\alpha \geq \gamma$ and $k \in \mathbb{N}$ with $k > \alpha$, is at least*

$$\frac{\alpha}{\gamma} \frac{1}{1 - (1 - \frac{\alpha}{k})^k}.$$

Now, the tight lower bound of Theorem 2 follows in the limit $k \to \infty$.

It turns out that, for $\gamma = 1$, the function $F_{\gamma,\alpha,k}$ is α-augmentable. Together with Proposition 4, this extends the lower bound of Theorem 5 to all $\alpha \geq 1$ and thus proves Corollary 1.

For k large enough, it can even be shown that $F_{1,\alpha,k} \notin (\tilde{\mathcal{F}}_\gamma \cup \mathcal{F}_q)$ for fixed $\gamma, q \in (0, 1]$, which separates $\mathcal{F}_\alpha$ for $\alpha \geq 1$, and thus closes the gap left in Sect. 2.1.

3.2 γ-α-Augmentability on Independence Systems

To tightly capture the class $\mathcal{F}_q$ of weighted rank functions on independence systems, we show a stronger bound for the approximation ratio of the greedy algorithm on monotone, (weakly) γ-α-augmentable functions. In particular, it was already shown in [2] that the objective function of α-Dimensional Matching is (exactly) α-augmentable, while the greedy algorithm yields an approximation ratio of α, which beats the upper bound of $\alpha \cdot \frac{e^\alpha}{e^\alpha - 1}$ for this case. We show that this can be explained by the fact that α-Dimensional Matching can be represented via a weighted rank function over an indepencence system. We first show the upper bound of Theorem 3.

Proposition 5. *Let $\mathcal{F}_{\text{IS}} := \bigcup_{q \in (0,1]} \mathcal{F}_q$ be the set of weighed rank functions on some independence system. The approximation ratio of the greedy algorithm on the class $\tilde{\mathcal{F}}_{\gamma,\alpha} \cap \mathcal{F}_{\text{IS}}$ is at most $\frac{\alpha}{\gamma}$, for every $\gamma \in (0, 1]$ and $\alpha \geq \gamma$.*

The lower bound of Theorem 3 follows directly from the well-known tight bound of $1/q$ for $\mathcal{F}_q$ [15] and the fact that every weighted rank function over an independence system with rank quotient q is γ-$\frac{\gamma}{q}$-augmentable, by Proposition 2.

4 Outlook

The vision guiding our work is to precisely characterize the set of cardinality-constrained maximization problems for which the greedy algorithm yields an approximation, and to tightly bound the corresponding approximation ratio.

In this paper, we have made progress towards this goal by unifying and generalizing important classes of greedily approximable maximization problems, and by providing tight bounds on the approximation ratio for the resulting generalized class of problems. While this brings us closer to a full characterization, there are still settings that are not captured by (weak) γ-α-augmentability.

Proposition 6. *For $\gamma \in (0, 1]$ and $\alpha \geq \gamma$, there exists a monotone function $f^{\gamma,\alpha}$ that is not weakly γ-α-augmentalbe, and for which the greedy algorithm computes an optimum solution.*

Proof (Sketch). Let U be any ground set of size $|U| > \frac{1}{\gamma}$ and consider the objective function $f^{\gamma,\alpha} \colon 2^U \to \mathbb{R}_{\geq 0}$ with

$$f^{\gamma,\alpha}(X) = |X|^2,$$

and show that it is not weakly γ-α-augmentable. Yet, picking elements in any order is obviously optimal. □

We leave it as an open problem to find a natural generalization of weak γ-α-augmentability that captures a larger set of greedily approximable objectives. The challenge is to find a meaningful generalization in terms of a natural definition that does not directly depend on the behavior of the greedy algorithm, but rather enforces some structural property of the objective function. In that sense, the dependency of weak γ-α-augmentability on the greedy solutions $S_0^{\mathrm{G}}, \ldots, S_{\bar{k}}^{\mathrm{G}}$ is a significant flaw. Note that we needed to introduce this dependency in order to encompass settings with bounded (weak) submodularity ratios, since the definition of the latter depends on the greedy solutions as well. Importantly, our upper bound on the approximation ratio of the greedy algorithm carries over to the stronger notion of γ-α-augmentability that requires the defining property to hold for *all* sets X, and not just the greedy solutions. Our tight lower bound does not immediately translate to this, more restrictive, definition, and it remains an open problem to construct a tight lower bound in this setting as well.

References

1. Anari, N., Haghtalab, N., Naor, S., Pokutta, S., Singh, M., Torrico, A.: Structured robust submodular maximization: offline and online algorithms. In: Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS), pp. 3128–3137 (2019)
2. Bernstein, A., Disser, Y., Groß, M., Himburg, S.: General bounds for incremental maximization. Math. Program. **191**, 953–979 (2020). https://doi.org/10.1007/s10107-020-01576-0

3. Bian, A., Buhmann, J., Krause, A., Tschiatschek, S.: Guarantees for greedy maximization of non-submodular functions with applications. In: Proceedings of the 34th International Conference on Machine Learning (ICML), pp. 498–507 (2017)
4. Bouchet, A.: Greedy algorithm and symmetric matroids. Math. Program. **38**(2), 147–159 (1987). https://doi.org/10.1007/BF02604639
5. Buchbinder, N., Feldman, M., Naor, J., Schwartz, R.: Submodular maximization with cardinality constraints. In: Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1433–1452 (2014). https://doi.org/10.1137/1.9781611973402.106
6. Calinescu, G., Chekuri, C., Pál, M., Vondrák, J.: Maximizing a monotone submodular function subject to a matroid constraint. SIAM J. Comput. **40**(6), 1740–1766 (2011). https://doi.org/10.1137/080733991
7. Conforti, M., Cornuéjols, G.: Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the Rado-Edmonds theorem. Discret. Appl. Math. **7**, 251–274 (1984)
8. Das, A., Kempe, D.: Approximate submodularity and its applications: subset selection, sparse approximation and dictionary selection. J. Mach. Learn. Res. **19**, 1–34 (2018)
9. Disser, Y., Weckbecker, D.: Unified greedy approximability beyond submodular maximization. arXiv:2011.00962 (2020)
10. Edmonds, J.: Matroids and the greedy algorithm. Math. Program. **1**, 127–136 (1971). https://doi.org/10.1007/BF01584082
11. Ene, A., Nguyen, H.L.: Constrained submodular maximization: beyond 1/e. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), pp. 248–257 (2016). https://doi.org/10.1109/FOCS.2016.34
12. Feldman, M., Harshaw, C., Karbasi, A.: Greed is good: near-optimal submodular maximization via greedy optimization. In: Proceedings of the 30th Conference on Learning Theory (COLT), pp. 758–784 (2017)
13. Feldman, M., Harshaw, C., Karbasi, A.: Simultaneous greedys: a swiss army knife for constrained submodular maximization. arXiv:2009.13998 (2020)
14. Feldman, M., Naor, J., Schwartz, R.: A unified continuous greedy algorithm for submodular maximization. In: 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS), pp. 570–579 (2011). https://doi.org/10.1109/FOCS.2011.46
15. Jenkyns, T.: The efficacy of the "greedy" algorithm. In: Proceedings of the 7th Southeastern Conference on Combinatorics, Graph Theory and Computing, pp. 341–350 (1976)
16. Korte, B., Hausmann, D.: An analysis of the greedy heuristic for independence systems. Ann. Discrete Math. **2**, 65–74 (1978)
17. Korte, B., Vygen, J.: Combinatorial Optimization: Theory and Algorithms. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-24488-9
18. Krause, A., Golovin, D.: Submodular function maximization. Tractability Pract. Approach. Hard Prob. **3**, 71–104 (2011). https://doi.org/10.1017/CBO9781139177801.004
19. Krause, A., McMahan, H., Guestrin, C., Gupta, A.: Robust submodular observation selection. J. Mach. Learn. Res. **9**, 2761–2801 (2008)
20. Lee, J., Mirrokni, V., Nagarajan, V., Sviridenko, M.: Non-monotone submodular maximization under matroid and knapsack constraints. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), pp. 323–332 (2009). https://doi.org/10.1145/1536414.1536459

21. Mestre, J.: Greedy in approximation algorithms. In: Azar, Y., Erlebach, T. (eds.) ESA 2006. LNCS, vol. 4168, pp. 528–539. Springer, Heidelberg (2006). https://doi.org/10.1007/11841036_48
22. Nemhauser, G., Wolsey, L., Fisher, M.: An analysis of approximations for maximizing submodular set functions - I. Math. Program. **14**, 265–294 (1978). https://doi.org/10.1007/BF01588971
23. Nemhauser, G., Wolsey, L., Fisher, M.: An analysis of approximations for maximizing submodular set functions - II. Math. Program. Stud. **8**, 73–87 (1978)
24. Orlin, J., Schulz, A., Udwani, R.: Robust monotone submodular function maximization. Math. Program. **172**, 505–537 (2015). https://doi.org/10.1007/s10107-018-1320-2
25. Rado, R.: A theorem on independence relations. Q. J. Math. **13**, 83–89 (1942). https://doi.org/10.1093/qmath/os-13.1.83
26. Sviridenko, M.: A note on maximizing a submodular set function subject to a knapsack constraint. Oper. Res. Lett. **32**(1), 41–43 (2004). https://doi.org/10.1016/S0167-6377(03)00062-2
27. Sviridenko, M., Vondrák, J., Ward, J.: Optimal approximation for submodular and supermodular optimization with bounded curvature. In: Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1134–1148 (2015)
28. Vondrák, J.: Optimal approximation for the submodular welfare problem in the value oracle model. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC), pp. 67–74 (2008). https://doi.org/10.1145/1374376.1374389
29. Vondrák, J.: Submodularity and curvature: the optimal algorithm. RIMS Kôkyûroku Bessatsu **B23**, 253–266 (2010)
30. Williamson, D., Shmoys, D.: The Design of Approximation Algorithms. Cambridge University Press, Cambridge (2011)