Bounding Inefficiency of Equilibria in Continuous Actions Games using Submodularity and Curvature

Pier Giuseppe Sessa
ETH Zürich

Maryam Kamgarpour
ETH Zürich

Andreas Krause
ETH Zürich

Abstract

Games with continuous strategy sets arise in several machine learning problems (e.g. adversarial learning). For such games, simple no-regret learning algorithms exist in several cases and ensure convergence to coarse correlated equilibria (CCE). The efficiency of such equilibria with respect to a *social function*, however, is not well understood. In this paper, we define the class of *valid utility games with continuous strategies* and provide efficiency bounds for their CCEs. Our bounds rely on the social function being a monotone DR-submodular function. We further refine our bounds based on the curvature of the social function. Furthermore, we extend our efficiency bounds to a class of non-submodular functions that satisfy approximate submodularity properties. Finally, we show that valid utility games with continuous strategies can be designed to maximize monotone DR-submodular functions subject to disjoint constraints with approximation guarantees. The approximation guarantees we derive are based on the efficiency of the equilibria of such games and can improve the existing ones in the literature. We illustrate and validate our results on a budget allocation game and a sensor coverage problem.

1 Introduction

Game theory is a powerful tool for modelling many real-world multi-agent decision making problems [7]. In machine learning, game theory has received substantial interest in the area of adversarial learning (e.g. generative adversarial networks [14]) where models are trained via games played by competing modules [2]. Apart from modelling interactions among agents, game theory is also used in the context of distributed optimization. In fact, specific games can be *designed* so that multiple entities can contribute to optimizing a common objective function [24, 22].

A game is described by a set of players aiming to maximize their individual payoffs which depend on each others' strategies. The efficiency of a joint strategy profile is measured with respect to a *social function*, which depends on the strategies of all the players. When the strategies for each player are uncountably infinite, the game is said to be *continuous*.

Continuous games describe a broad range of problems where integer or binary strategies may have limited expressiveness. In market sharing games [13], for instance, competing firms may invest continuous amounts in each market, or may produce an infinitely divisible product. Also, several integer problems can be generalized to continuous domains. For example, in budget allocation problems continuous amounts can be allocated to each media channel [4]. In machine learning, many games are naturally continuous [21].

1.1 Related work

Although continuous games are finding increasing applicability, from a theoretical viewpoint they are less understood than games with finitely many strategies. Recently, no-regret learning algorithms [7] have been proposed for continuous games under different set-ups [32, 30, 25]. Similarly to finite games [7], these no-regret dynamics converge to *coarse correlated equilibria* (CCEs) [30, 2], the weakest class of equilibria which includes pure Nash equilibria, mixed Nash equilibria and correlated equilibria. However, CCEs may be highly suboptimal for the social function. A central open question is to understand the (in)efficiency of such equilibria. Differently from the finite case, where bounds on such inefficiency are known for a large variety of games [28], in continuous games this question is not well understood.

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89.

To measure the inefficiency of CCEs arising from no-regret dynamics, [6] introduces the *price of total anarchy.* This notion generalizes the well-established price of anarchy (PoA) of [19] which instead measures the inefficiency of the worst pure Nash equilibria of the game. There are numerous reasons why players may not reach a pure Nash equilibrium [6, 29, 28]. In contrast, regret minimization can be done by each player via simple and efficient algorithms [6]. Recently, [28] generalizes the price of total anarchy defining the *robust* PoA which measures the inefficiency of any CCE (including the ones arising from regret minimization), and provides examples of games for which it can be bounded.

In the context of distributed optimization, where a game is designed to optimize a given objective [24], bounds on the robust price of anarchy find a similar importance. In this setting, a distributed scheme to optimize the social function is to let each player implement a no-regret learning algorithm based only on its payoff information. A bound on the robust PoA provides an approximation guarantee to such optimization scheme.

Bounds on the robust PoA provided by [28] mostly concern games with finitely many actions. A class of such games are the *valid utility games* introduced by [31]. In such games, the social function is a *submodular* set function and, using this property, [28] showed that the PoA bound derived in [31] indeed extends to all CCEs of the game. This class of games covers numerous applications including market sharing, facility location, and routing problems, and were used by [24] for distributed optimization. Strategies consist of selecting subsets of a ground set, and can be equivalently represented as binary decisions. Recently, authors in [23] extend the notion of valid utility games to integer domains. By leveraging properties of submodular functions over integer lattices, they show that the robust PoA bound of [28] extends to the integer case. The notion of submodularity has recently been extended to continuous domains, mainly in order to design efficient optimization algorithms [1, 4, 16]. To the best of author's knowledge, such notion has not been utilized for analyzing efficiency of equilibria of games over continuous domains.

1.2 Our contributions

We bound the robust price of anarchy for a subclass of continuous games, which we denote as *valid utility games with continuous strategies*. They are the continuous counterpart of the valid utility games introduced by [31] and [23] for binary and integer strategies, respectively. Our bounds rely on a particular game structure and on the social function being a monotone DR-submodular function [4, Definition 1]. Hence, we define the *curvature* of a monotone DR-submodular function on continuous domains, analyze its properties, and use it to refine our bounds. We also show that our bounds can be extended to non-submodular functions which have 'approximate' submodularity properties. This is in contrast with [31, 23] where only submodular social functions were considered. Finally, employing the machinery of [24], we show that valid utility games with continuous strategies can be designed to maximize non convex/non concave functions in a distributed fashion with approximation guarantees. Depending on the curvature of the function, the obtained guarantees can improve the ones available in the literature.

1.3 Notation

We denote by $\mathbf{e}_i, \mathbf{0}$, and $\mathbb{1}$, the i^{th} unit vector, null vector, and vector of all ones of appropriate dimensions, respectively. Given $n \in \mathbb{N}$, with $n \geq 1$, we define $[n] := \{1, \dots, n\}$. Given vectors $\mathbf{x}, \mathbf{y}$, we use $[\mathbf{x}]_i$ and x_i interchangeably to indicate the i^{th} coordinate of $\mathbf{x}$, and $(\mathbf{x}, \mathbf{y})$ to denote the vector obtained from their concatenation, i.e., $(\mathbf{x}, \mathbf{y}) := [\mathbf{x}^\top, \mathbf{y}^\top]^\top$. Moreover, for vectors of equal dimension, $\mathbf{x} \leq \mathbf{y}$ means $x_i \leq y_i$ for all i. Given $\mathbf{x} \in \mathbb{R}^n$ and $j \in \{0, \dots, n\}$, we define $[\mathbf{x}]_1^j := (x_1, \dots, x_j, 0, \dots, 0) \in \mathbb{R}^n$ with $[\mathbf{x}]_1^0 = \mathbf{0}$. A function $f : \boldsymbol{\mathcal{X}} \subseteq \mathbb{R}^n \to \mathbb{R}$ is *monotone* if, for all $\mathbf{x} \leq \mathbf{y} \in \boldsymbol{\mathcal{X}}$, $f(\mathbf{x}) \leq f(\mathbf{y})$. Moreover, f is *affine* if for all $\mathbf{x}, \mathbf{y} \in \boldsymbol{\mathcal{X}}$, $f(\mathbf{x} + \mathbf{y}) - f(\mathbf{x}) = f(\mathbf{y}) - f(\mathbf{0})$.

2 Problem formulation and examples

We consider a class of non-cooperative continuous games, where each player i chooses a vector $\mathbf{s}_i$ in its feasible strategy set $\mathcal{S}_i \subseteq \mathbb{R}^d_+$. We let N be the number of players, $\mathbf{s} = (\mathbf{s}_1, \dots, \mathbf{s}_N)$ be the vector of all the strategy profiles, i.e., the outcome of the game, and $\boldsymbol{\mathcal{S}} = \prod_{i=1}^N \mathcal{S}_i \subseteq \mathbb{R}^{Nd}_+$ be the joint strategy space. For simplicity, we assume each strategy $\mathbf{s}_i$ is d-dimensional, although different dimensions could exist for different players. Each player aims to maximize her payoff function $\pi_i : \boldsymbol{\mathcal{S}} \to \mathbb{R}$, which in general depends on the strategies of all the players. We let the social function be $\gamma : \mathbb{R}^{Nd}_+ \to \mathbb{R}_+$. For the rest of the paper we assume $\gamma(\mathbf{0}) = 0$. We denote such games with the tuple $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$. Given an outcome $\mathbf{s}$ we use the standard notation $(\mathbf{s}_i, \mathbf{s}_{-i})$ to denote the outcome where player i chooses strategy $\mathbf{s}_i$ and the other players select strategies $\mathbf{s}_{-i} = (\mathbf{s}_1, \dots, \mathbf{s}_{i-1}, \mathbf{s}_{i+1}, \dots \mathbf{s}_N)$.

A pure Nash equilibrium is an outcome $\mathbf{s} \in \boldsymbol{\mathcal{S}}$ such that

$$\pi_i(\mathbf{s}) \geq \pi_i(\mathbf{s}'_i, \mathbf{s}_{-i})\,,$$

for every player i and for every strategy $\mathbf{s}'_i \in \mathcal{S}_i$. A coarse correlated equilibrium (CCE) is a probability distribution σ over the outcomes $\boldsymbol{\mathcal{S}}$ that satisfies

$$\mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s})] \geq \mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s}_i', \mathbf{s}_{-i})] \,,$$

for every player i and for every strategy $\mathbf{s}_i' \in \mathcal{S}_i$. CCE's are the weakest class of equilibria and they include pure Nash, mixed Nash, and correlated equilibria [28].

Since each player selfishly maximizes her payoff, the outcome $\mathbf{s} \in \boldsymbol{\mathcal{S}}$ of the game is typically suboptimal for the social function γ. To measure such suboptimality, [28] introduced the *robust price of anarchy* (robust PoA) which measures the inefficiency of any CCE. Given $\mathcal{G}$, we let Δ be the set of all the CCEs of $\mathcal{G}$ and define the robust PoA as the quantity

$$PoA_{CCE} := \frac{\max_{\mathbf{s}\in\boldsymbol{\mathcal{S}}} \gamma(\mathbf{s})}{\min_{\sigma\in\Delta} \mathbb{E}_{\mathbf{s}\sim\sigma}[\gamma(\mathbf{s})]} \,.$$

It can be easily seen that $PoA_{CCE} \geq 1$. As discussed in the introduction, PoA_{CCE} has two important implications. In multi-agent systems, PoA_{CCE} bounds the efficiency of no-regret learning dynamics followed by the selfish agents. In fact, these dynamics converge to a CCE of the game [30, 2]. In the context of distributed optimization, no-regret learning algorithms can be implemented distributively to optimize a given function and PoA_{CCE} certifies the overall approximation guarantee. Bounds for PoA_{CCE}, however, were obtained mostly for games with finitely many actions [28].

In this paper, we are interested in upper bounding PoA_{CCE} for continuous games $\mathcal{G}$ defined above. To motivate our results, we present two relevant examples of such games. The first one is a budget allocation game, while in the second example a continuous game can be designed for distributed maximization in the spirit of [24]. We will come back to these examples in Section 3 and derive upper bounds for their respective PoA_{CCE}'s.

Example 1 (**Continuous budget allocation game**). A set of N advertisers enters a market consisting of a set of d media channels. By allocating (or investing) part of their budget in each advertising channel, the goal of each advertiser is to maximize the expected number of *activated* customers, i.e., customers who purchase her product. The market is described by a bipartite graph $G = (\mathcal{R} \cup \mathcal{T}, \mathcal{E})$, where the left vertices $\mathcal{R}$ denote channels and the right vertices $\mathcal{T}$ denote customers, with $d = |\mathcal{R}|$. For each advertiser i and edge $(r,t) \in \mathcal{E}$, $p_i(r,t) \in [0,1]$ is the probability that advertiser i activates customer t via channel r. Each advertiser chooses a strategy $\mathbf{s}_i \in \mathbb{R}^d_+$, which represents the amounts allocated (or invested) to each channel, subject to budget constraints $\mathcal{S}_i = \{\mathbf{s}_i \in \mathbb{R}^d : \mathbf{c}_i^\top \mathbf{s}_i \leq b_i, \mathbf{0} \leq \mathbf{s}_i \leq \bar{\mathbf{s}}_i\}$. This generalizes the set-up in [23], where strategies $\mathbf{s}_i$ are integer. Hence, we consider the continuous version of the game modeled by [23]. For every customer $t \in \mathcal{T}$ and advertiser $i \in [N]$, we define $\Gamma(t) = \{r \in \mathcal{R} : (r,t) \in \mathcal{E}\}$ and the quantity

$$P_i(\mathbf{s}_i, t) = 1 - \prod\nolimits_{r\in\Gamma(t)} (1 - p_i(r,t))^{[\mathbf{s}_i]_r} \,,$$

which is the probability that i activates t when the other advertisers are ignored. For each customer t, a permutation $\rho \in \mathcal{P}_N$ is drawn uniformly at random, where $\mathcal{P}_N$ is the set of all permutations of $[N]$. Then, according to ρ each advertiser sequentially attempts to activate customer t. Hence, for a given allocation $\mathbf{s} = (\mathbf{s}_1, \ldots, \mathbf{s}_N) \in \boldsymbol{\mathcal{S}} = \prod_{i=1}^N \mathcal{S}_i$, the payoff of each advertiser can be written in closed form as [23]:

$$\pi_i(\mathbf{s}) = \frac{1}{N!} \sum\nolimits_{t\in\mathcal{T}} \sum\nolimits_{\rho\in\mathcal{P}_N} P_i(\mathbf{s}_i, t) \prod_{j\prec_\rho i} (1 - P_j(\mathbf{s}_j, t)) \,,$$

where $j \prec_\rho i$ indicates that j precedes i in ρ. The term $\pi_i(\mathbf{s})$ represents the expected number of customers activated by advertiser i in allocation $\mathbf{s}$. The goal of the market analyst, which assumes the role of the game planner, is to maximize the expected number of customers activated. Hence, for any $\mathbf{s}$, the social function γ is

$$\gamma(\mathbf{s}) = \sum\nolimits_{i=1}^N \pi_i(\mathbf{s}) = \sum\nolimits_{t\in\mathcal{T}} \Big(1 - \prod\nolimits_{i=1}^N (1 - P_i(\mathbf{s}_i, t))\Big) \,.$$

Example 2 (**Sensor coverage with continous assignments**). Given a set of N autonomous sensors, we seek to monitor a finite set of d locations in order to maximize the probability of detecting an event. For each sensor, a continuous variable $x_i \in \mathbb{R}^d_+$ indicates the energy assigned (or time spent) to each location, subject to budget constraints $\mathcal{X}_i := \{\mathbf{x}_i \in \mathbb{R}^d : \mathbf{c}_i^\top \mathbf{x}_i \leq b_i, \mathbf{0} \leq \mathbf{x}_i \leq \bar{\mathbf{x}}_i\}$. This generalizes the well-known sensor coverage problem studied in [24] (and previous works), where $\mathbf{x}_i$'s are binary and indicate the locations sensor i is assigned to. The probability that sensor i detects an event in location r is $1 - (1 - p_i^r)^{[x_i]_r}$, with $0 \leq p_i^r \leq 1$, and it increases as more energy is assigned to the location. Hence, given a strategy $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_N)$, the joint probability of detecting an event in location r is

$$P(r, \mathbf{x}) = 1 - \prod\nolimits_{i\in[N]} (1 - p_i^r)^{[\mathbf{x}_i]_r} \,.$$

The goal of the planner is to maximize the probability of detecting an event

$$\gamma(\mathbf{x}) = \sum\nolimits_{r\in[d]} w_r \, P(r, \mathbf{x}) \,,$$

where w_r's represent the a priori probability that an event occurs in location r. As in [24], we can set up a continuous game $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$ where $\mathcal{S}_i = \mathcal{X}_i$ for each i, and π_i's are designed so that good monitoring solutions can be obtained when each player selfishly maximizes her payoff.

Proofs of the upcoming propositions and remarks are presented in Appendix A.

3 Main results

We derive PoA_{CCE} bounds for a sublass of continuous games $\mathcal{G}$ by extending the valid utility games considered in [31] and [23] to continuous strategy sets. Hence, in Definition 3 we will define the class of *valid utility games with continuous strategies*. As will be seen, the two problems described above and several other examples fall into this class. At the end of the section, we will show that valid utility games can be designed to maximize non-convex/non-concave objectives in a distributed fashion with approximation guarantees.

3.1 Robust PoA bounds

As in [31, 23], the PoA_{CCE} bounds obtained rely on the social function γ experiencing diminishing returns (DR). Differently from set functions, in continuous (and integer) domains, different notions of DR exist. Similarly to [23], our first main result relies on γ satisfying the strongest notion of DR, also known as `DR` property [4], which we define in Definition 1. Moreover, as in [31] our bound can be refined depending on the *curvature* of γ. While DR properties have been recently studied also in continuous domains, notions of curvature of a submodular function were only explored for set functions [10, 17] (see [3, Appendix C] for a comparison of the existing notions). Hence, in Definition 2 we define the curvature of a monotone DR-submodular function on continuous domains.

Definition 1 (`DR` property)**.** A function $f : \boldsymbol{\mathcal{X}} = \prod_{i=1}^n \mathcal{X}_i \to \mathbb{R}$ with $\mathcal{X}_i \subseteq \mathbb{R}$ is *DR-submodular* if for all $\mathbf{x} \leq \mathbf{y} \in \boldsymbol{\mathcal{X}}$, $\forall i \in [n], \forall k \in \mathbb{R}_+$ such that $(\mathbf{x} + k\mathbf{e}_i)$ and $(\mathbf{y} + k\mathbf{e}_i)$ are in $\boldsymbol{\mathcal{X}}$,

$$f(\mathbf{x} + k\mathbf{e}_i) - f(\mathbf{x}) \geq f(\mathbf{y} + k\mathbf{e}_i) - f(\mathbf{y})\,.$$

When restricted to binary sets $\boldsymbol{\mathcal{Z}} = \{0,1\}^n$, Definition 1 coincides with the standard notion of submodularity for set functions. An equivalent characterization of the `DR` property for a twice-differentiable function is that all the entries of its Hessian are non-positive [4]:

$$\forall \mathbf{x} \in \boldsymbol{\mathcal{X}}, \quad \frac{\partial^2 f(\mathbf{x})}{\partial x_i \partial x_j} \leq 0, \quad \forall i, j\,.$$

Definition 2 (curvature)**.** Given a monotone DR-submodular function $f : \boldsymbol{\mathcal{X}} \subseteq \mathbb{R}_+^n \to \mathbb{R}$, and a set $\boldsymbol{\mathcal{Z}} \subseteq \boldsymbol{\mathcal{X}}$ with $\mathbf{0} \in \boldsymbol{\mathcal{Z}}$, we define the curvature of f with respect to $\boldsymbol{\mathcal{Z}}$ by

$$\alpha(\boldsymbol{\mathcal{Z}}) = 1 - \inf_{\substack{\mathbf{x} \in \boldsymbol{\mathcal{Z}}, i \in [n]: \\ \mathbf{x} + k\mathbf{e}_i \in \boldsymbol{\mathcal{Z}}}} \lim_{k \to 0^+} \frac{f(\mathbf{x} + k\mathbf{e}_i) - f(\mathbf{x})}{f(k\mathbf{e}_i) - f(\mathbf{0})}\,.$$

Remark 1. For any monotone function $f : \mathbb{R}^n \to \mathbb{R}$ and $\forall \boldsymbol{\mathcal{Z}} \subseteq \mathbb{R}^n$ with $\mathbf{0} \in \boldsymbol{\mathcal{Z}}$, $\alpha(\boldsymbol{\mathcal{Z}}) \in [0,1]$.

When restricted to binary sets $\boldsymbol{\mathcal{Z}} = \{0,1\}^n$, Definition 2 coincides with the *total curvature* defined in [10]. Moreover, if f is montone DR-submodular and *differentiable*, its curvature with respect to a set $\boldsymbol{\mathcal{Z}}$ can be computed as:

$$\alpha(\boldsymbol{\mathcal{Z}}) = 1 - \inf_{\substack{\mathbf{x} \in \boldsymbol{\mathcal{Z}} \\ i \in [n]}} \frac{[\nabla f(\mathbf{x})]_i}{[\nabla f(\mathbf{0})]_i}\,.$$

Based on the previous definitions, we define the class of valid utility games with continuous strategies.

Definition 3. A game $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$ is a *valid utility game with continuous strategies* if:

i) The function γ is monotone DR-submodular.
ii) For each player i and for every outcome $\mathbf{s}$, $\pi_i(\mathbf{s}) \geq \gamma(\mathbf{s}) - \gamma(\mathbf{0}, \mathbf{s}_{-i})$.
iii) For every outcome $\mathbf{s}$, $\gamma(\mathbf{s}) \geq \sum_{i=1}^N \pi_i(\mathbf{s})$.

Intuitively, the conditions above ensure that the payoff for each player is at least her contribution to the social function and that optimizing γ is somehow bind to the goals of the players. Defining the set $\tilde{\boldsymbol{\mathcal{S}}} := \{\mathbf{x} \in \mathbb{R}_+^{Nd} \mid \mathbf{0} \leq \mathbf{x} \leq \mathbf{s}_{max}\}$, with $\mathbf{s}_{max}$ such that $\forall \mathbf{s}, \mathbf{s}' \in \boldsymbol{\mathcal{S}}$, $\mathbf{s}+\mathbf{s}' \leq \mathbf{s}_{max}$, we can establish the following main theorem.

Theorem 1. Let $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$ be a valid utility game with continuous strategies with social function $\gamma : \mathbb{R}_+^{Nd} \to \mathbb{R}_+$ having *curvature* $\alpha(\tilde{\boldsymbol{\mathcal{S}}}) \leq \alpha$. Then, $PoA_{CCE} \leq (1 + \alpha)$.

We will prove Theorem 1 in Section 4.2.

Remark 2. If $\mathcal{G}$ is a valid utility game with continuous strategies, then $PoA_{CCE} \leq 2$.

Remark 3. The notion of valid utility games above is an exact generalization of the one by [23] for integer strategy sets. Leveraging recent advances in 'approximate' submodular functions, in Section 5 we relax condition i) and derive PoA_{CCE} bounds for a strictly larger class of games.

Using Theorem 1, the following proposition upper bounds PoA_{CCE} of Example 1. Our bound depends on the activation probabilities $p_i(r,t)$'s and on the connectivity of the market G.

Proposition 1. The budget allocation game defined in Example 1 is a valid utility game with continuous strategies. Moreover, $PoA_{CCE} \leq 1 + \alpha < 2$ with $\alpha := 1-$

$$\min_{i \in [N], r \in [d]} \frac{\sum\limits_{t \in \mathcal{T}: r \in \Gamma(t)} \ln(1 - p_i(r,t)) \prod\limits_{j \in [N]} (1 - P_j(2\bar{\mathbf{s}}_j, t))}{\sum\limits_{t \in \mathcal{T}: r \in \Gamma(t)} \ln(1 - p_i(r,t))}\,.$$

In our more general continuous actions framework, the obtained bound strictly improves the bound of 2 by [23], since the curvature of the social function was not considered in [23]. We will visualize our bound in the experiments of Section 6.

Using Theorem 1, we now generalize Example 2 and show that valid utility games with continuous strategies can be designed to maximize monotone DR-submodular functions subject to decoupled constraints with approximation guarantees. The proposed optimization scheme will be used in Section 6 to solve an instance of the sensor coverage problem (Example 2).

3.2 Game-based monotone DR-submodular maximization

Consider the general problem of maximizing a monotone DR-submodular function $\gamma : \mathbb{R}^n \to \mathbb{R}_+$ subject to *decoupled* constraints $\boldsymbol{\mathcal{X}} = \prod_{i=1}^N \mathcal{X}_i \subseteq \mathbb{R}^n$. We can assume $\mathcal{X}_i \subseteq \mathbb{R}_+$ without loss of generality [4], since otherwise one could optimize γ over a shifted version of its constraints. Moreover, we assume $\gamma(\mathbf{0}) = 0$ for ease of exposition. Note that the class of monotone DR-submodular functions includes non concave functions. To find approximate solutions, we set up a game

$$\hat{\mathcal{G}} := (N, \{\hat{\mathcal{S}}_i\}_{i=1}^N, \{\hat{\pi}_i\}_{i=1}^N, \gamma)\,,$$

where for each player i, $\hat{\mathcal{S}}_i := \mathcal{X}_i$, and $\hat{\pi}_i(\mathbf{s}) := \gamma(\mathbf{s}) - \gamma(\mathbf{0}, \mathbf{s}_{-i})$ for every outcome $\mathbf{s} \in \boldsymbol{\mathcal{S}} = \boldsymbol{\mathcal{X}}$. By using DR-submodularity of γ, we can affirm the following.

Fact 1. $\hat{\mathcal{G}}$ is a valid utility game with continuous strategies.

Assume there exists $\mathbf{x}_{max} \in \mathbb{R}_+^n$ such that $\forall \mathbf{x}, \mathbf{x}' \in \boldsymbol{\mathcal{X}}$, $\mathbf{x} + \mathbf{x}' \leq \mathbf{x}_{max}$. Then, we denote with $\alpha(\tilde{\boldsymbol{\mathcal{X}}})$ the curvature of γ with respect to $\tilde{\boldsymbol{\mathcal{X}}} := \{\mathbf{x} \in \mathbb{R}_+^n \mid \mathbf{0} \leq \mathbf{x} \leq \mathbf{x}_{max}\}$ and let $\alpha \in [0, 1]$ be an upper bound for $\alpha(\tilde{\boldsymbol{\mathcal{X}}})$. If such $\mathbf{x}_{max}$ does not exist, we let $\alpha = 1$. Moreover, assume that for each player $i \in [N]$ there exists a no-regret algorithm [27, Sec. 3] to play $\hat{\mathcal{G}}$. That is, when $\hat{\mathcal{G}}$ is repeated over time, player i can ensure that $\frac{1}{T} \max_{s \in \hat{\mathcal{S}}_i} \sum_{t=1}^T \hat{\pi}_i(s, \mathbf{s}_{-i}^t) - \frac{1}{T} \sum_{t=1}^T \hat{\pi}_i(s_i^t, \mathbf{s}_{-i}^t) \to 0$ as $T \to \infty$, for any sequence $\{\mathbf{s}_{-i}^t\}_{t=1}^T$. We let D-noRegret be the distributed algorithm where such no-regret algorithms are simultaneously implemented for each player. We can establish the following corollary of Theorem 1 [1]

Corollary 1. Let $\mathbf{x}^\star = \arg\max_{\mathbf{x} \in \boldsymbol{\mathcal{X}}} \gamma(\mathbf{x})$. Then, D-noRegret converges to a distribution σ over $\boldsymbol{\mathcal{X}}$ such that $\mathbb{E}_{\mathbf{x} \sim \sigma}[\gamma(\mathbf{x})] \geq 1/(1 + \alpha)\gamma(\mathbf{x}^\star)$.

Note that the Frank-Wolfe variant of [4] can also be used to maximize γ with $(1 - e^{-1})$ approximations, under the additional assumption that $\boldsymbol{\mathcal{X}}$ is down-closed. For small α's, however, our guarantee can strictly improve the one by [4].

If $\boldsymbol{\mathcal{X}}$ is convex compact and γ is concave in each $\mathcal{X}_i$, then $\hat{\pi}_i$'s are concave in each x_i and the online gradient ascent algorithm by [32] ensures no-regret for each player [12]. Using Corollary 1, we show that the sensor coverage problem of Example 2 falls into this class and D-noRegret has approximation guarantees that depend on the sensing probabilities $P(r, \cdot)$'s.

Proposition 2. Consider the sensor coverage problem of Example 2 and assume we set-up the game $\hat{\mathcal{G}}$. Then, online gradient ascent [32] is a no-regret algorithm for each player. Moreover, D-noRegret has an expected approximation ratio of $1/(1 + \alpha)$, where $\alpha := \max_{r \in [d]} P(r, 2\bar{\mathbf{x}})$ and $\bar{\mathbf{x}} = (\bar{\mathbf{x}}_1, \dots, \bar{\mathbf{x}}_N)$.

Note that the obtained approximation ratio is strictly larger than $\frac{1}{2}$ and it increases when the number of sensors N or the detection probabilities decrease, a fact also noted in [24] for the binary setting. We compare the performance of D-noRegret and the Frank-Wolfe variant of [4] in Section 6.

A decentralized maximization scheme for submodular functions is also proposed in [26], albeit in a different setting. In [26], γ consists of a sum of *local* functions subject to a common down-closed convex contraint set, while we considered a generic objective γ subject to local constraints.

4 Analysis

In order to prove Theorem 1 and its extension to non-submodular functions (Section 5), we first review the main properties of submodularity in continuous domains and show a fundamental property of the curvature of a monotone DR-submodular function.

4.1 Submodularity and curvature on continuous domains

Submodularity in continuous domains has received recent attention for approximate maximization and minimization of non convex/non concave functions [4, 16, 1]. Submodular continuous functions are defined on subsets of $\mathbb{R}^n$ of the form $\boldsymbol{\mathcal{X}} = \prod_{i=1}^n \mathcal{X}_i$, where each $\mathcal{X}_i$ is a compact subset of $\mathbb{R}$. A function $f : \boldsymbol{\mathcal{X}} \to \mathbb{R}$ is *submodular* if for all $\mathbf{x} \in \boldsymbol{\mathcal{X}}$, $\forall i \neq j$ and $a_i, a_j > 0$ s.t. $x_i + a_i \in \mathcal{X}_i$, $x_j + a_j \in \mathcal{X}_j$, [1]

$$f(\mathbf{x} + a_i\mathbf{e}_i) - f(\mathbf{x}) \geq f(\mathbf{x} + a_i\mathbf{e}_i + a_j\mathbf{e}_j) - f(\mathbf{x} + a_j\mathbf{e}_j)\,.$$

The above property also includes submodularity of set functions, by restricting $\mathcal{X}_i$'s to $\{0, 1\}$, and over integer lattices, by restricting $\mathcal{X}_i$'s to $\mathbb{Z}_+$. We are interested, however, in submodular continuous functions, where $\mathcal{X}_i$'s are compact subsets of $\mathbb{R}$. As thoroughly studied for set functions, submodularity is related to diminishing return properties of f. However, differences exist when considering functions over continuous (or integer) domains. In particular, submodularity is equivalent to the following *weak* DR property [4].

Definition 4 (`weak DR` property)**.** A function $f : \boldsymbol{\mathcal{X}} \subseteq \mathbb{R}^n \to \mathbb{R}$ is *weakly DR-submodular* if, for all $\mathbf{x} \leq \mathbf{y} \in \boldsymbol{\mathcal{X}}$, $\forall i$ s.t. $x_i = y_i$, $\forall k \in \mathbb{R}_+$ s.t. $(\mathbf{x} + k\mathbf{e}_i)$ and $(\mathbf{y} + k\mathbf{e}_i)$ are in $\boldsymbol{\mathcal{X}}$,

$$f(\mathbf{x} + k\mathbf{e}_i) - f(\mathbf{x}) \geq f(\mathbf{y} + k\mathbf{e}_i) - f(\mathbf{y})\,.$$

[1] A similar version of the corollary can be obtained when no-α-regret [18, Definition 4] algorithms exist for each player, such as the ones by [9, 8] for online submodular maximization.

The `DR` property, which we defined in Definition 1 characterizes the full notion of diminishing returns and indentifies a subclass of submodular continuous functions. While `weak DR` and `DR` properties coincide for set functions, this is not the case for functions on integer or continuous lattices. As the next section reveals, the `weak DR` property of γ is indeed not sufficient to prove Theorem 1. However, it will be useful in Section 5 when we extend our results to non-submodular functions. In Appendix A.7 we provide equivalent characterizations of `weak DR` and `DR` properties and discuss submodularity for differentiable functions in Appendix A.8. The following proposition is key for the proof of Theorem 1.

Proposition 3. Consider a monotone DR-submodular function $f : \boldsymbol{\mathcal{X}} \subseteq \mathbb{R}^n_+ \to \mathbb{R}$, and a set $\boldsymbol{\mathcal{Z}} := \{\mathbf{x} \in \mathbb{R}^n : \mathbf{0} \leq \mathbf{x} \leq \mathbf{z}_{max}\} \subseteq \boldsymbol{\mathcal{X}}$. Then, for any $\mathbf{x}, \mathbf{y} \in \boldsymbol{\mathcal{Z}}$ such that $\mathbf{x} + \mathbf{y} \in \boldsymbol{\mathcal{Z}}$,

$$f(\mathbf{x} + \mathbf{y}) - f(\mathbf{x}) \geq (1 - \alpha(\boldsymbol{\mathcal{Z}}))[f(\mathbf{y}) - f(\mathbf{0})] \,,$$

where $\alpha(\boldsymbol{\mathcal{Z}})$ is the curvature of f with respect to $\boldsymbol{\mathcal{Z}}$.

4.2 Proof of Theorem 1

The proof uses submodularity of the social function similarly to [28, Example 2.6] and [23, Proposition 4]. However, it allows us to consider the curvature of γ. Differently from [23, Proposition 4], our proof does not rely on the structure of the strategy sets $\mathcal{S}_i$'s. The `weak DR` and `DR` properties are used separately in the proof, to show that the `weak DR` property of γ is not sufficient to obtain the results. This fact was similarly noted in [23] for the integer case.

To upper bound PoA_{CCE}, we first prove that for any pair of outcomes $\mathbf{s}, \mathbf{s}^\star \in \boldsymbol{\mathcal{S}}$,

$$\sum\nolimits_{i=1}^{N} \pi_i(\mathbf{s}_i^\star, \mathbf{s}_{-i}) \geq \gamma(\mathbf{s}^\star) - \alpha\, \gamma(\mathbf{s}) \,.$$

In the framework of [28], this means that $\mathcal{G}$ is a *$(1, \alpha)$-smooth* game. Then, few inequalities from [28] show that $PoA_{CCE} \leq (1 + \alpha)$.

The smoothness proof is obtained as follows. Consider any pair of outcomes $\mathbf{s}, \mathbf{s}^\star \in \boldsymbol{\mathcal{S}}$. For $i \in \{0, \ldots, N\}$ with a slight abuse of notation we define $[\mathbf{s}^\star]_1^i = (\mathbf{s}_1^\star, \ldots, \mathbf{s}_i^\star, \mathbf{0}, \ldots, \mathbf{0})$ with $[\mathbf{s}^\star]_1^0 = \mathbf{0}$, where $\mathbf{s}_j^\star$ is the strategy of player j in the outcome $\mathbf{s}^\star$. We have:

$$\begin{aligned}
&\sum\nolimits_{i=1}^{N} \pi_i(\mathbf{s}_i^\star, \mathbf{s}_{-i}) \geq \sum\nolimits_{i=1}^{N} \gamma(\mathbf{s}_i^\star, \mathbf{s}_{-i}) - \gamma(\mathbf{0}, \mathbf{s}_{-i}) \\
&\geq \sum\nolimits_{i=1}^{N} \gamma(\mathbf{s}_i^\star + \mathbf{s}_i, \mathbf{s}_{-i}) - \gamma(\mathbf{s}) \\
&\geq \sum\nolimits_{i=1}^{N} \gamma(\mathbf{s} + [\mathbf{s}^\star]_1^i) - \gamma(\mathbf{s} + [\mathbf{s}^\star]_1^{i-1}) \\
&= \gamma(\mathbf{s} + \mathbf{s}^\star) - \gamma(\mathbf{s}) \\
&\geq (1 - \alpha)\gamma(\mathbf{s}) + \gamma(\mathbf{s}^\star) - \gamma(\mathbf{s}) = \gamma(\mathbf{s}^\star) - \alpha\, \gamma(\mathbf{s}) \,.
\end{aligned}$$

The first inequality follows from condition ii) of valid utility games as per Definition 3. The second inequality from γ being DR-submodular (and using Proposition 5 in Appendix A.7). The third inequality from γ being weakly DR-submodular (and using Proposition 4 in Appendix A.7). The last inequality follows since, by Proposition 3,

$$\gamma(\mathbf{s} + \mathbf{s}^\star) - \gamma(\mathbf{s}^\star) \geq (1 - \alpha(\tilde{\boldsymbol{\mathcal{S}}}))[\gamma(\mathbf{s}) - \underbrace{\gamma(\mathbf{0})}_{=0}] \,,$$

and $\alpha(\tilde{\boldsymbol{\mathcal{S}}}) \leq \alpha$.

For completeness we report the steps of [28] to prove that $PoA_{CCE} \leq (1 + \alpha)$. Let $\mathbf{s}^\star = \arg\max_{\mathbf{s} \in \boldsymbol{\mathcal{S}}} \gamma(\mathbf{s})$. Then, for any CCE σ of $\mathcal{G}$ we have

$$\begin{aligned}
\mathbb{E}_{\mathbf{s}\sim\sigma}[\gamma(\mathbf{s})] &\geq \sum_{i=1}^{N} \mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s})] \geq \sum_{i=1}^{N} \mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s}_i^\star, \mathbf{s}_{-i})] \\
&\geq \gamma(\mathbf{s}^\star) - \alpha\, \mathbb{E}_{\mathbf{s}\sim\sigma}[\gamma(\mathbf{s})] \,,
\end{aligned}$$

where the first inequality is due to condition iii) of valid utility games as per Definition 3, the second inequality holds from σ being a CCE, and the last one since $\mathcal{G}$ is $(1, \alpha)$-smooth. Moreover, linearity of expectation was used throughout. From the inequalities above it holds that for any CCE σ of $\mathcal{G}$, $\gamma(\mathbf{s}^\star)/\mathbb{E}_{\mathbf{s}\sim\sigma}[\gamma(\mathbf{s})] \leq 1 + \alpha$. Hence $PoA_{CCE} \leq 1 + \alpha$.

Remark 4. Although Theorem 1 requires DR-submodularity of γ over $\mathbb{R}_+^{Nd}$ (for simplicity), only DR-submodularity over $\tilde{\boldsymbol{\mathcal{S}}}$ was used. In case γ is DR-submodular only over $\boldsymbol{\mathcal{S}}$, one could consider $\tilde{\gamma} : \mathbb{R}_+^{Nd} \to \mathbb{R}_+$ defined as $\tilde{\gamma}(\mathbf{s}) = \gamma(\min(\mathbf{s}, \mathbf{s}_{max}))$ which is DR-submodular over $\mathbb{R}_+^{Nd}$. This can be proved using DR-submodularity and monotonicity of γ over $\tilde{\boldsymbol{\mathcal{S}}}$. The same smoothness proof is obtained with $\tilde{\gamma}$ in place of γ since the two functions are equal over $\boldsymbol{\mathcal{S}}$. However, the curvature of $\tilde{\gamma}$ with respect to $\tilde{\boldsymbol{\mathcal{S}}}$ is 1 and therefore a bound of 2 for PoA_{CCE} is obtained.

5 Extension to the non-submodular case

In many applications [3], functions are close to being submodular, where this closedness has been measured in term of *submodularity ratio* [11] (for set functions) and *weak-submodularity* [16] (on continuous domains). Accordingly, in this section we relax condition i) of valid utility games (Definition 3) and provide bounds for PoA_{CCE} when the social function γ is not necessarily DR-submodular. This case was never considered for the valid utility games of [31, 23]. We relax the `weak DR` property of γ with the following definition.

Definition 5. Given a game $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$ with γ monotone, we define *generalized submodularity ratio* of γ as the largest scalar η such that for any pair of outcomes $\mathbf{s}, \mathbf{s}' \in \boldsymbol{\mathcal{S}}$,

$$\sum\nolimits_{i=1}^{N} \gamma(\mathbf{s}_i + \mathbf{s}_i', \mathbf{s}_{-i}) - \gamma(\mathbf{s}) \geq \eta\big[\gamma(\mathbf{s} + \mathbf{s}') - \gamma(\mathbf{s})\big] \,.$$

It is straightforward to show that $\eta \in [0,1]$. Moreover, as stated in Appendix B (Proposition 8), if γ is weakly DR-submodular then γ has generalized submodularity ratio $\eta = 1$. When strategies $\mathbf{s}_i$ are scalar (i.e., $d = 1$), Definition 5 generalizes the submodularity ratio by [11] to continuous domains[2].

In addition, we relax the `DR` property of γ as follows.

Definition 6. Given a game $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$, we say that γ is *playerwise DR-submodular* if for every player i and vector of strategies $\mathbf{s}_{-i}$, $\gamma(\cdot, \mathbf{s}_{-i})$ is DR-submodular.

Analogously to Definition 1, if γ is twice-differentiable, it is playwerwise DR-submodular iff for every $i \in [N]$

$$\forall \mathbf{s} \in \boldsymbol{\mathcal{S}}, \quad \frac{\partial^2 \gamma(\mathbf{s})}{\partial [\mathbf{s}_i]_l \partial [\mathbf{s}_i]_m} \leq 0, \quad \forall\, l, m \in [d]\,.$$

While Definition 5 concerns the interactions among different players, Definition 6 requires that γ is DR-submodular with respect to each individual player. When the social function γ is DR-submodular, then it is also playerwise DR-submodular. Moreover, since the `DR` property is stronger than `weak DR`, γ has generalized submodularity ratio $\eta = 1$. If γ is not DR-submodular, however, the notions of Definition 5 and Definition 6 are not related. We visualize their differences in the following example.

Example 3. Consider a game with $N = 2$, $d = 2$, and γ twice-differentiable. Let η be the generalized submodularity ratio of γ. Assume the Hessian of γ satisfies one of the three cases below, where with '+' or '−' we indicate the sign of its elements:

$$\underset{1.}{\begin{pmatrix} - & - & - & - \\ - & - & - & - \\ - & - & - & - \\ - & - & - & - \end{pmatrix}} \quad \underset{2.}{\begin{pmatrix} + & - & - & - \\ - & - & - & - \\ - & - & - & - \\ - & - & - & + \end{pmatrix}} \quad \underset{3.}{\begin{pmatrix} - & - & - & + \\ - & - & - & - \\ - & - & - & - \\ + & - & - & - \end{pmatrix}}$$

From the previous definitions, the function γ is playerwise DR-submodular iff all the entries highlighted in red are non-positive, while η depends on all the off-diagonal entries. In case 1., all the entries are negative, hence γ is DR-submodular. Thus, it is playerwise DR-submodular and has generalized submodularity ratio $\eta = 1$. In case 2., all off-diagonal entries are negative, hence γ is weakly DR-submodular (see Appendix A.8) and thus $\eta = 1$. However, γ is not playerwise DR-submodular since some highlighted entries are positive. In case 3., γ is playerwise DR-submodular and its generalized submodularity ratio depends on its parameters.

Note that only case 1. of the previous example satisfies the conditions of Theorem 1. However, the following Theorem 2 is applicable also to a subset of functions which fall in case 3. The proof can be found in Appendix B.

Theorem 2. Let $\mathcal{G} = (N, \{\mathcal{S}_i\}_{i=1}^N, \{\pi_i\}_{i=1}^N, \gamma)$ be a game where γ is monotone, playerwise DR-submodular and has generalized submodularity ratio $\eta > 0$. Then, if conditions ii) and iii) of Definition 3 are satisfied, $PoA_{CCE} \leq (1+\eta)/\eta$.

In light of the previous comments, when γ is DR-submodular Theorem 2 yields a bound of 2 which is always higher than $(1+\alpha)$ from Theorem 1. This is because the notion of curvature in Definition 2 cannot be used in the more general setting of Theorem 2 since γ may not be DR-submodular.

In Appendix B we show that examples of functions with generalized submodularity ratio $1 > \eta > 0$ are products of monotone weakly DR-submodular functions and monotone affine functions. As a consequence, the following generalization of Example 2 falls into the set-up of Theorem 2.

Sensor coverage problem with non-submodular objective. Consider the sensor coverage problem defined in Example 2, where the weights w_r's are monotone affine functions $w_r : \mathbb{R}_+^{Nd} \to \mathbb{R}_+$ rather than constants. For instance, the probability that an event occurs in location r can increase with the average amount of energy allocated to that location. That is, $\gamma(\mathbf{x}) = \sum_{r\in[d]} w_r(\mathbf{x}) P(r, \mathbf{x})$ with $w_r(\mathbf{x}) = \mathbf{a}_r \frac{\sum_{i=1}^N [\mathbf{x}_i]_r}{N} + b_r$. To maximize γ one could set up a game $\mathcal{G}$ where condition ii) of Definition 3 is satisfied with equality, as shown in Section 3.2. In Appendix B.4 we show that γ has generalized submodularity ratio $1 > \eta > 0$, it is playerwise DR-submodular, and that $\gamma(\mathbf{x}) \geq \frac{1}{2} \sum_{i=1}^N \pi_i(\mathbf{x})$ for every $\mathbf{x}$, which is a weaker version of condition iii). Nevertheless, using Theorem 2 and the last proof steps of Section 4.2 we prove that $PoA_{CCE} \leq (1 + 0.5\eta)/0.5\eta$. We also show that γ is concave in each $\mathcal{X}_i$. Therefore a distributed implementation of online gradient ascent maximizes γ up to $0.5\eta/(1 + 0.5\eta)$ approximations.

We remark that our definitions of curvature, submodularity ratio, and Theorems 1-2 can also be applied to games and optimizations over integer domains, i.e., when $\mathcal{S}_i \subseteq \mathbb{Z}_+^d$ and γ is defined on integer lattices.

6 Experimental results

In this section we analyze the examples defined in Section 2 using the developed framework.

6.1 Continuous budget allocation game

We consider $N = 10$ advertisers in a market with $d = 100$ channels and $|\mathcal{T}| = 10'000$ customers. For the budget constraints we select $b_i = 1$, $\bar{\mathbf{s}}_i = \mathbb{1}$ and each entry of $\mathbf{c}_i$ is sampled uniformly at random from $[0, 1]$. For each $i \in [N], r \in \mathcal{R}, t \in \mathcal{T}$, $p_i(r,t)$ is drawn uni-

[2] In Appendix B we define an exact generalization of the submodularity ratio by [11] to continuous domains. We relate it to Definition 5 and compare it to the ratio by [16].

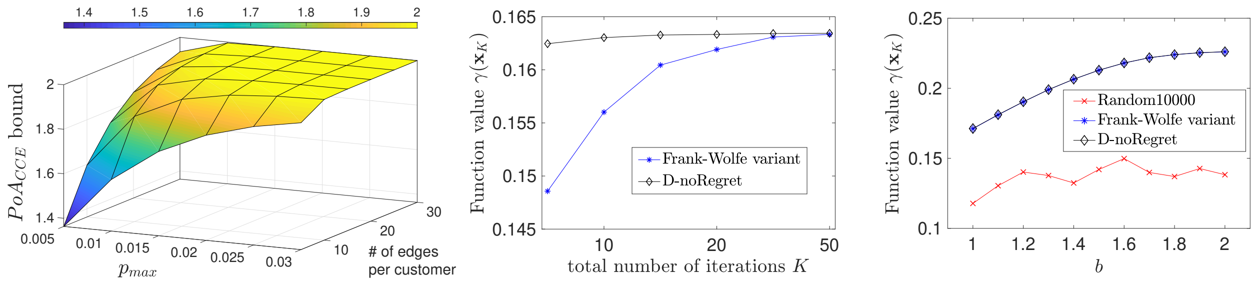

(a) Budget allocation game

(b) Sensor coverage problem

Figure 1: a) Bounds for PoA_{CCE}, as a function of p_{max} and the number of edges connected to each customer. The bounds strictly improve the bound of 2 provided by [23] which does not depend on any of such parameters. b) Performance of Frank-Wolfe variant and D-noRegret for $K = 3000$ iterations. Left: $\gamma(\mathbf{x}_K)$ as function of K. Right: $\gamma(\mathbf{x}_K)$ as function of the budgets b, with $K = 3000$. D-noRegret shows faster convergence, but for $K = 3000$ the two algorithms perform equally.

formly at random in $[0.8, 1]p_{max}$. In Figure 1a we visualize the bound for PoA_{CCE} obtained in Proposition 1 for different values of p_{max} and the number of random edges connected to each customer. The chosen ranges ensure that a sufficient fraction of customers will be activated. For instance, for $p_{max} = 0.01$ and drawing 20 random edges for each customer, we obtained [3] an expected number of 2270 activated customers. This is in line with [23], where problem parameters were chosen such that $\frac{1}{5}$ of the customers are activated. As visible, the bound decreases when the activation probabilities decrease or when less edges are connected to each customer and can strictly improve the bound of 2 provided by [23].

6.2 Sensor coverage with continous assignments

To maximize the probability γ of detecting an event, we compare the performance of D-noRegret with the Frank-Wolfe variant by [4] and a hit-and-run sampler [20] Random10000 which samples 10'000 random feasible points. We choose $N = 5$ sensors and $d = 30$ locations. For each budget constraint $\mathcal{X}_i$, entries of $\mathbf{c}_i$ are chosen uniformly at random in $\frac{1}{d}[1, 3]$, $b_i = 1$, and $\bar{\mathbf{x}}_i = \mathbb{1}$. For each $i \in [N]$ and $r \in [d]$, we select $p_i^r = p = 0.05$ and w_r uniformly at random in $[0, 1]$ such that $\sum_{r=1}^{d} w_r = 1$. Under this choice, $\alpha \approx 0.4$, and D-noRegret has approximation guarantee ≈ 0.71 which is greater than $(1 - e^{-1}) \approx 0.63$. We initialize both Frank-Wolfe and D-noRegret at $\mathbf{x}_0 = 0$ and run them for $K \in \{10, 20, 50, 100, 500, 3000\}$ iterations. Since the constraints are decoupled, also the Frank-Wolfe variant can be implemented distributively. Step sizes for both algorithms are chosen costant and proportional to $1/K$ and $1/\sqrt{K}$ as per [4] and [12, Lemma 3.2], respectively. In Figure 1b we compare the values of $\gamma(\mathbf{x}_K)$ as a function of the number of iterations K (left plot). Moreover, for $K = 3000$, we compare the performance of the algorithms when we enlarge the constraints $\mathcal{X}_i$ by choosing $b_i = b$ for each i, with $b \in \{1, 1.1, 1.2, \ldots, 2\}$ (right plot). As visible, D-noRegret shows faster convergence than Frank-Wolfe variant. However, for $K = 3000$ the two algorithms return the same values. Average computation times per iteration are 0.019 s and 0.009 s for Frank-Wolfe and D-noRegret, respectively, on a 16 Gb machine at 3.1 GHz using Matlab.

7 Conclusions and future work

We bounded the robust price of anarchy for a subclass of continuous games, denoted as valid utility games with continuous strategies. Our bound relies on a particular structure of the game and on the social function being monotone DR-submodular. We introduced the notion of curvature of a monotone DR-submodular function and refined the bound using this notion. In addition, we extended the obtained bounds to a class of non-submodular functions. We showed that valid utility games can be designed to maximize monotone DR-submodular functions subject to disjoint constraints. For a subclass of such functions, our approximation guarantees improve the ones in the literature. We demonstrated our results numerically via a continuous budget allocation game and a sensor coverage problem. In light of the obtained approximation guarantees, we believe that the introduced notion of curvature of a monotone DR-submodular function can be used to tighten existing guarantees for constrained maximization. Currently, we are studying the tightness of the obtained bounds and their applicability to several continuous games such as auctions.

[3] Since γ is monotone DR-submodular and $\boldsymbol{\mathcal{S}}$ is downclosed, we used the Frank-Wolfe variant by [4] to maximize γ up to $(1 - e^{-1})$ approximations.

Acknowledgements

This work was gratefully supported by Swiss National Science Foundation, under the grant SNSF 200021_172781, and by the European Union's Horizon 2020 ERC grant 815943.

References

[1] Francis Bach. Submodular functions: from discrete to continuous domains. *Mathematical Programming*, Feb 2018.

[2] David Balduzzi, Sebastien Racaniere, James Martens, Jakob Foerster, Karl Tuyls, and Thore Graepel. The mechanics of n-player differentiable games. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 354–363, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.

[3] Andrew An Bian, Joachim M. Buhmann, Andreas Krause, and Sebastian Tschiatschek. Guarantees for greedy maximization of non-submodular functions with applications. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 498–507, 2017.

[4] Andrew An Bian, Baharan Mirzasoleiman, Joachim M. Buhmann, and Andreas Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, pages 111–120, 2017.

[5] Jeff Bilmes. Submodular functions, their optimization and applications, 2011.

[6] Avrim Blum, MohammadTaghi Hajiaghayi, Katrina Ligett, and Aaron Roth. Regret minimization and the price of total anarchy. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 373–382, New York, NY, USA, 2008. ACM.

[7] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.

[8] Lin Chen, Christopher Harshaw, Hamed Hassani, and Amin Karbasi. Projection-free online optimization with stochastic gradient: From convexity to submodularity. In *ICML*, volume 80 of *JMLR Workshop and Conference Proceedings*, pages 813–822. JMLR.org, 2018.

[9] Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization. In *AISTATS*, volume 84 of *Proceedings of Machine Learning Research*, pages 1896–1905. PMLR, 2018.

[10] Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: Tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete Applied Mathematics*, 7(3):251 – 274, 1984.

[11] Abhimanyu Das and David Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 1057–1064, 2011.

[12] Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: Gradient descent without a gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 385–394, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.

[13] M. X. Goemans, Li Li, V. S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 24(5):1020–1033, May 2006.

[14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.

[15] Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, pages 360–367, New York, NY, USA, 2008. ACM.

[16] S. Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *NIPS*, pages 5843–5853, 2017.

[17] Rishabh Iyer and Jeff Bilmes. Submodular optimization with submodular cover and submodular knapsack constraints. In *Proceedings of the 26th*

International Conference on Neural Information Processing Systems - Volume 2, NIPS'13, pages 2436–2444, USA, 2013. Curran Associates Inc.

[18] Sham M. Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. In *Proceedings of the Thirty-ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 546–555, New York, NY, USA, 2007. ACM.

[19] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*, STACS'99, pages 404–413, Berlin, Heidelberg, 1999. Springer-Verlag.

[20] D.P. Kroese, T. Taimre, and Z.I. Botev. *Handbook of Monte Carlo Methods*. Wiley Series in Probability and Statistics. Wiley, 2013.

[21] Kyowoon Lee, Sol-A Kim, Jaesik Choi, and Seong-Whan Lee. Deep reinforcement learning in continuous action spaces: a case study in the game of simulated curling. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2937–2946, Stockholmsmässan, Stockholm Sweden, 10–15 Jul 2018. PMLR.

[22] Na Li and Jason R. Marden. Designing games for distributed optimization. *J. Sel. Topics Signal Processing*, 7(2):230–242, 2013.

[23] Takanori Maehara, Akihiro Yabe, and Ken-Ichi Kawarabayashi. Budget allocation problem with multiple advertisers: A game theoretic view. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, pages 428–437. JMLR.org, 2015.

[24] Jason R. Marden and Adam Wierman. Distributed welfare games. *Operations Research*, 61(1):155–168, 2013.

[25] Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, Mar 2018.

[26] Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. Decentralized submodular maximization: Bridging discrete and continuous settings. In *ICML*, 2018.

[27] Barnabé Monnot and Georgios Piliouras. Limits and limitations of no-regret learning in games. *Knowledge Eng. Review*, 32:e21, 2017.

[28] Tim Roughgarden. Intrinsic robustness of the price of anarchy. *J. ACM*, 62(5):32:1–32:42, November 2015.

[29] Noah D. Stein, Pablo A. Parrilo, and Asuman Ozdaglar. Correlated equilibria in continuous games: Characterization and computation. *Games and Economic Behavior*, 71(2):436 – 455, 2011.

[30] Gilles Stoltz and Gábor Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, 59(1):187 – 208, 2007.

[31] Adrian Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, FOCS '02, pages 416–, Washington, DC, USA, 2002. IEEE Computer Society.

[32] Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on International Conference on Machine Learning*, ICML'03, pages 928–935. AAAI Press, 2003.