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

Pier Giuseppe Sessa
ETH Zürich
sessap@control.ee.ethz.ch

Maryam Kamgarpour
ETH Zürich
maryamk@ethz.ch

Andreas Krause
ETH Zürich
krausea@ethz.ch

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 satisfying recently introduced notions of submodularity over continuous domains. 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 and problem formulation

Game theory is a powerful tool for modelling many real-world multi-agent decision making problems [6]. In machine learning, game theory has received substantial interest in the area of adversarial learning (e.g. generative adversarial networks [11]) 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, games can be *designed* so that multiple entities can contribute to optimizing a common objective function [18, 16]. 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 (e.g. market sharing [10], or budget allocations [4]). In machine learning, many games are naturally continuous [15].

In this work we consider continuous N-players 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 $\mathbf{s} = (\mathbf{s}_1, \dots, \mathbf{s}_N)$ be the vector of all the strategy profiles and $\boldsymbol{\mathcal{S}} = \prod_{i=1}^N \mathcal{S}_i \subseteq \mathbb{R}^{Nd}_+$ be the joint strategy space. Each player aims to maximize her payoff function $\pi_i : \boldsymbol{\mathcal{S}} \to \mathbb{R}$, and we let the social function be $\gamma : \mathbb{R}^{Nd}_+ \to \mathbb{R}_+$ with $\gamma(\mathbf{0}) = 0$, where $\mathbf{0}$ is the null vector. 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 $\mathbf{s}_{-i} = (\mathbf{s}_1, \dots, \mathbf{s}_{i-1}, \mathbf{s}_{i+1}, \dots \mathbf{s}_N)$.

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

*Under review by the 22^{nd} International Conference on Artificial Intelligence and Statistics (AISTATS 2019)

Abstract for: Smooth Games Optimization and Machine Learning Workshop (NIPS 2018), Montréal, Canada.

[6] have been proposed for continuous games under different set-ups [23, 21, 19]. Similarly to finite games [6], these no-regret dynamics converge to *coarse correlated equilibria*, the weakest class of equilibria [21, 2]. A coarse correlated equilibrium (CCE) is a probability distribution σ over the outcomes $\boldsymbol{\mathcal{S}}$ that satisfies

$$\mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s})] \geq \mathbb{E}_{\mathbf{s}\sim\sigma}[\pi_i(\mathbf{s}_i', \mathbf{s}_{-i})] \quad \forall i \in \{1, \dots, N\}, \forall \mathbf{s}_i' \in \mathcal{S}_i\,.$$

However, CCEs may be highly suboptimal for the social function. A central open question is to understand the (in)efficiency of such equilibria.

To measure the inefficiency of CCEs arising from no-regret dynamics, [5] introduces the *price of total anarchy*. Recently, [20] generalizes this notion defining the *robust* price of anarchy (robust PoA) which measures the inefficiency of any CCE (including the ones arising from regret minimization). Given $\mathcal{G}$, we let Δ be the set of all CCEs of $\mathcal{G}$ and define the robust PoA:

$$PoA_{CCE} := \frac{\max_{\mathbf{s}\in\boldsymbol{\mathcal{S}}} \gamma(\mathbf{s})}{\min_{\sigma\in\Delta} \mathbb{E}_{\mathbf{s}\sim\sigma}[\gamma(\mathbf{s})]} \geq 1\,.$$

A bound on PoA_{CCE} hence, has two important implications. First, in multi-agent systems, PoA_{CCE} bounds the efficiency of no-regret learning dynamics followed by the selfish agents. Second, 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 on the robust PoA provided by [20] mostly concern games with finitely many actions. A class of such games are the *valid utility games* introduced by [22]. Strategies consist of selecting subsets of a ground set, and can be equivalently represented as binary decisions. Recently, authors in [17] extend the notion of valid utility games to integer domains. The PoA bounds obtained in [22] and [17] rely on the social function being a *submodular* set function and a submodular function over integer lattices, respectively. Recently, the notion of submodularity has been extended to continuous domains [1, 4, 12]. However, to the best of author's knowledge, it has not been utilized for analyzing efficiency of equilibria of games with continuous strategies.

2 Main results

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 [22] and [17] 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]. Moreover, as in [22] 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 [7, 13]. Hence, in Definition 2 we define the curvature of a monotone DR-submodular function on continuous domains. We also show that our bounds can be extended to non-submodular functions which have 'approximate' submodularity properties. This is in contrast to [22, 20, 17] where only submodular social functions were considered. Finally, motivated by the machinery of [18] in finite actions games, 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.

Notation. We denote by $\mathbf{e}_i$ the i^{th} unit vector of appropriate dimension. 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 $\mathbf{x}$. Moreover, for vectors of equal dimension, $\mathbf{x} \leq \mathbf{y}$ means $x_i \leq y_i$ for all i. A function $f : \boldsymbol{\mathcal{X}} \subseteq \mathbb{R}^n \to \mathbb{R}$ is *monotone* if, for all $\mathbf{x} \leq \mathbf{y} \in \boldsymbol{\mathcal{X}}$, $f(\mathbf{x}) \leq f(\mathbf{y})$.

2.1 Robust PoA bounds

Before defining the class of valid utility games with continuous strategies, we define DR-submodular functions and their curvature on continuous domains.

Definition 1 (DR property)**.** A function $f : \boldsymbol{\mathcal{X}} \subseteq \mathbb{R}^n \to \mathbb{R}$ 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 still in $\boldsymbol{\mathcal{X}}$,

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

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 $0 \in \boldsymbol{\mathcal{Z}}$, we define the curvature of f with respect to $\boldsymbol{\mathcal{Z}}$ by

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

Remark 1. For any monotone function $f : \mathbb{R}^n \to \mathbb{R}$ and $\forall \boldsymbol{\mathcal{Z}} \subseteq \mathbb{R}^n$ with $0 \in \boldsymbol{\mathcal{Z}}$, $\alpha(\boldsymbol{\mathcal{Z}}) \in [0, 1]$. When restricted to binary sets $\boldsymbol{\mathcal{Z}} = \{0, 1\}^n$, Definition 2 coincides with the *total curvature* defined in [7]. 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 $\tilde{\boldsymbol{\mathcal{S}}} := \{\mathbf{x} \in \mathbb{R}_+^{Nd} \mid \mathbf{0} \leq \mathbf{x} \leq \mathbf{s}_{max}\}$ such that $\forall \mathbf{s}, \mathbf{s}' \in \boldsymbol{\mathcal{S}}$, $\mathbf{s} + \mathbf{s}' \leq \mathbf{s}_{max}$, we can establish Theorem 1.

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, $PoA_{CCE} \leq (1 + \alpha)$.

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

The notion of valid utility games above is an exact generalization of the one by [17] for integer strategy sets. However, the curvature of γ was not used to refine the bound from $PoA_{CCE} \leq 2$. Moreover, leveraging recent advances in 'approximate' submodular functions, in the next section we relax condition i) and derive PoA_{CCE} bounds for a strictly larger class of games.

2.2 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* [8] (for set functions) and *weak-submodularity* [12] (on continuous domains). Motivated by this, we relax the DR property required in condition i) with the following two definitions.

Definition 4. 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 $\mathbf{s}, \mathbf{s}' \in \boldsymbol{\mathcal{S}}$,

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

Clearly, $\eta \in [0, 1]$. Moreover, it is not hard to show that if γ is DR-submodular then γ has generalized submodularity ratio $\eta = 1$ [2].

Definition 5. 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.

While Definition 4 concerns the interactions among different players, Definition 5 requires γ DR-submodular with respect to each individual player. When the social function γ is DR-submodular, then it is also playerwise DR-submodular and has generalized submodularity ratio $\eta = 1$. If γ is not DR-submodular, the notions of Definition 4 and Definition 5 are two independent properties of γ and none implies the other. We can affirm the following theorem.

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 $\mathcal{G}$ is a valid utility game 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.

2.3 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 $\mathcal{X}_i \subseteq \mathbb{R}_+^n$ and $\gamma(\mathbf{0}) = 0$ without loss of generality [4]. Note that the class of monotone DR-submodular functions includes non concave functions. To find approximate solutions, we set up a game

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

[2] In the full version, we show that when $d = 1$ Definition 4 generalizes the submodularity ratio by [8] to continuous domains. Moreover, it is sufficient that γ has the `weak DR` property [4, Definition 2] to have $\eta = 1$.

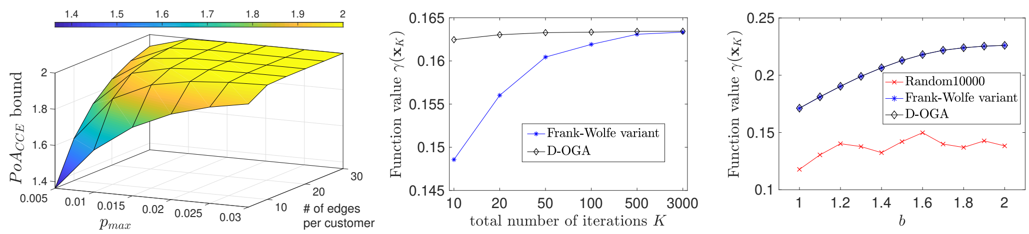

(a) Budget allocation game

(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 [17] which does not depend on any of such parameters. b) Performance of FRANK-WOLFE variant and D-NOREGRET for K 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.

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 exist $\mathbf{x}_{max} \in \mathbb{R}^n_+$ 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 exists, we let $\alpha = 1$. Moreover, assume that for each player $i \in [N]$ there exists a no-regret algorithm to play $\hat{\mathcal{G}}$. For instance, if $\boldsymbol{\mathcal{X}}$ is convex 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 [23] ensures no-regret for each player [9]. 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 [3].

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 by [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 strictly improves the one by [4].

Note: To prove our main results, in the full version of the paper we prove two DR properties equivalent to the ones from the literature and a fundamental property of the introduced notion of curvature.

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.

3 Examples and experiments

Continuous budget allocation game. We show that the continuous version of the integer budget allocation game defined by [17] (see [17] for details) is a valid utility game with continuous strategies. As visible in Figure 1a, our PoA_{CCE} bound depends on the *maximum activation probability* p_{max} and the number of edges connected to each customer and strictly improves the bound of 2 by [17].

Sensor coverage with continuous assignments. We generalize the sensor coverage problem by [18] (see [18] for details) assigning for each sensor a continuous amount of energy to each location. The problem of maximizing the probability of detecting an event falls into the set-up of Section 2.3, where online gradient ascent is a no-regret algorithm for each player. As visible in Figure 1b, D-NOREGRET converges faster than FRANK-WOLFE variant. However, for $K = 3000$ iterations the two algorithms perform equally. Finally, we show that a more general non-submodular version of such problem falls under the hypothesis of Theorem 2 and provide approximation guarantees.

[3] A similar version of the corollary can be obtained when no-ϵ-regret [14] algorithms exist for each player.

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] 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.

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

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] 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.

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

[13] 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.

[14] 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.

[15] 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.

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

[17] 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.

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

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

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

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

[22] 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.

[23] 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.