From Linear to Linearizable Optimization: A Novel Framework with Applications to Stationary and Non-stationary DR-submodular Optimization

Mohammad Pedramfar Vaneet Aggarwal
McGill University/Mila Purdue University

NEURAL INFORMATION PROCESSING SYSTEMS

Introduction

The problem of optimizing DR-submodular functions over a convex set has attracted considerable interest in machine learning and theoretical computer science. Example applications include experimental design, resource allocation, influence maximization, mean-field inference in probabilistic models, and MAP inference in determinantal point processes (DPPs), among others.

For $\gamma \in (0,1]$, a differentiable function $f : [0,1]^d \to \mathbb{R}_{\geq 0}$ is called *γ-weakly continuous DR-submodular* if for all $\mathbf{x}, \mathbf{y} \in [0,1]^d$ with $\mathbf{x} \geq \mathbf{y}$, we have $\gamma \nabla f(\mathbf{x}) \leq \nabla f(\mathbf{y})$.

Online optimization

The online optimization game could be modeled as a game between an agent and an adversary. At each time-step $1 \leq t \leq T$, the agent plays an action $\mathbf{x}_t$, then the adversary selects a function f_t and a query oracle for this function. Finally the agent then queries the query oracle. The feedback is called *bandit/semi-bandit* if the query oracle returns (an estimate of) the value/gradient of f_t at the point it is being queried and the agent only queries at the point of action, i.e., $\mathbf{x}_t$. The adversary is called *oblivious* if it selects the sequence of functions before the first action by the agent. We use $\mathrm{Adv}_i^{\mathrm{o}}(\mathbf{F}, B)$ to denote an oblivious adversary over function class $\mathbf{F}$ with i-th order query oracles that return values that are bounded by B and we replace the superscript with f to denote fully adaptive adversaries.

Following [1], in order to handle different notions of regret with the same approach, for an agent $\mathcal{A}$, adversary Adv, compact set $\mathcal{U} \subseteq \mathcal{K}^T$, approximation coefficient $0 < \alpha \leq 1$ and $1 \leq a \leq b \leq T$, we define *regret* as

$$\mathcal{R}^{\mathcal{A}}_{\alpha,\mathrm{Adv}}(\mathcal{U})[a,b] := \sup_{\mathcal{B} \in \mathrm{Adv}} \mathbb{E}\left[\alpha \max_{\mathbf{u}=(\mathbf{u}_1,\cdots,\mathbf{u}_T)\in\mathcal{U}} \sum_{t=a}^{b} f_t(\mathbf{u}_t) - \sum_{t=a}^{b} f_t(\mathbf{x}_t)\right],$$

where the expectation is over the randomness of the algorithm and the query oracle.

Static adversarial regret or simply *regret* corresponds to $a = 1$, $b = T$ and $\mathcal{U} = \mathcal{K}_\star^T := \{(\mathbf{x}, \cdots, \mathbf{x}) \mid \mathbf{x} \in \mathcal{K}\}$. When $a = 1$, $b = T$ and $\mathcal{U}$ contains only a single element then it is referred to as the *dynamic regret*. *Adaptive regret*, is defined as $\max_{1\leq a\leq b\leq T} \mathcal{R}^{\mathcal{A}}_{\alpha,\mathrm{Adv}}(\mathcal{K}_\star^T)[a,b]$. We drop a, b and $\mathcal{U}$ when the statement is independent of their value or their value is clear from the context.

Linearizable functions

Let $\mathcal{K} \subseteq \mathbb{R}^d$ be a convex set, $\mathbf{F}$ be a function class over $\mathcal{K}$. We say the function class $\mathbf{F}$ is *upper-quadratizable* if there are maps $\mathfrak{g} : \mathbf{F} \times \mathcal{K} \to \mathbb{R}^d$ and $h : \mathcal{K} \to \mathcal{K}$ and constants $\mu \geq 0$, $0 < \alpha \leq 1$ and $\beta > 0$ such that

$$\alpha f(\mathbf{y}) - f(h(\mathbf{x})) \leq \beta \left(\langle \mathfrak{g}(f, \mathbf{x}), \mathbf{y} - \mathbf{x} \rangle - \frac{\mu}{2}\|\mathbf{y} - \mathbf{x}\|^2 \right).$$

As a special case, when $\mu = 0$, we say $\mathbf{F}$ is *upper-linearizable*.

Algorithm 1: Online Maximization By Quadratization - $\mathtt{OMBQ}(\mathcal{A}, \mathcal{G}, h)$

Input: horizon T, semi-bandit algorithm $\mathcal{A}$, query algorithm $\mathcal{G}$ for $\mathfrak{g}$, the map $h : \mathcal{K} \to \mathcal{K}$
for $t = 1, 2, \ldots, T$ **do**
  Play $h(\mathbf{x}_t)$ where $\mathbf{x}_t$ is the action chosen by $\mathcal{A}$
  The adversary selects f_t and a first order query oracle for f_t
  Run $\mathcal{G}$ with access to $\mathbf{x}_t$ and the query oracle for f_t to calculate $\mathbf{o}_t$
  Return $\mathbf{o}_t$ as the output of the query oracle to $\mathcal{A}$
end

Theorem 1. Let $\mathcal{A}$ be an algorithm for online optimization with semi-bandit feedback. Also let $\mathbf{F}$ be a function class over $\mathcal{K}$ that is quadratizable with $\mu \geq 0$ and maps $\mathfrak{g}$ and h, and let $\mathcal{A}' = \mathtt{OMBQ}(\mathcal{A}, \mathcal{G}, h)$. If $\mathcal{G}$ is a query algorithm for $\mathfrak{g}$ that returns unbiased estimates and its output is bounded by B_1, then we have

$$\mathcal{R}^{\mathcal{A}'}_{\alpha,\mathrm{Adv}_1^{\mathrm{o}}(\mathbf{F},B_1)} \leq \beta \mathcal{R}^{\mathcal{A}}_{1,\mathrm{Adv}_1^{\mathrm{f}}(\mathbf{Q}_\mu[B_1])},$$

where $\mathbf{Q}_\mu[B_1] := \{q | q := \mathbf{y} \mapsto \langle \mathbf{o}, \mathbf{y} - \mathbf{x}\rangle - \frac{\mu}{2}\|\mathbf{y} - \mathbf{x}\|^2, \mathbf{o} \in \mathbb{R}^d, \|\mathbf{o}\| \leq B_1\}$.

Monotone functions over general convex sets

Lemma 1. [[2]] Let $f : [0,1]^d \to \mathbb{R}$ be a non-negative γ-weakly monotone DR-submodular function. Then, for all $\mathbf{x}, \mathbf{y} \in [0,1]^d$, we have

$$\frac{\gamma^2}{1+\gamma^2} f(\mathbf{y}) - f(\mathbf{x}) \leq \frac{\gamma}{1+\gamma^2} \langle \nabla f(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

Theorem 2. Let $\mathbf{F}$ be an M_1-Lipschitz function class over a convex set $\mathcal{K} \subseteq [0,1]^d$ where every $f \in \mathbf{F}$ may be extended to a function described in the above lemma. Then, for any $B_1 \geq M_1$, we have

$$\mathcal{R}^{\mathcal{A}}_{\frac{\gamma^2}{1+\gamma^2},\mathrm{Adv}_1^{\mathrm{o}}(\mathbf{F},B_1)} \leq \frac{\gamma}{1+\gamma^2} \mathcal{R}^{\mathcal{A}}_{1,\mathrm{Adv}_1^{\mathrm{f}}(\mathbf{Q}_0[B_1])}.$$

Monotone functions over convex sets containing origin

Lemma 2. [[3]] Let $f : [0,1]^d \to \mathbb{R}$ be a non-negative γ-weakly monotone DR-submodular differentiable function and let $F : [0,1]^d \to \mathbb{R}$ be the function defined by $F(\mathbf{x}) := \int_0^1 \frac{\gamma e^{\gamma(z-1)}}{(1-e^{-\gamma})z}(f(z * \mathbf{x}) - f(\mathbf{0}))dz$. Then F is differentiable and, if the random variable $\mathcal{Z} \in [0,1]$ is defined by the law

$$\forall z \in [0,1], \quad \mathbb{P}(\mathcal{Z} \leq z) = \int_0^z \frac{\gamma e^{\gamma(u-1)}}{1-e^{-\gamma}} du, \qquad (1)$$

then we have $\mathbb{E}\left[\nabla f(\mathcal{Z} * \mathbf{x})\right] = \nabla F(\mathbf{x})$. Moreover, we have

$$(1 - e^{-\gamma}) f(\mathbf{y}) - f(\mathbf{x}) \leq \frac{1 - e^{-\gamma}}{\gamma} \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

Theorem 3. Let $\mathcal{K} \subseteq [0,1]^d$ be a convex set containing the origin and let $\mathbf{F}$ be an M_1-Lipschitz function class over $\mathcal{K}$ where every $f \in \mathbf{F}$ may be extended to a function described in the above lemma. Then, for any $B_1 \geq M_1$, we have

$$\mathcal{R}^{\mathcal{A}'}_{1-e^{-\gamma},\mathrm{Adv}_1^{\mathrm{o}}(\mathbf{F},B_1)} \leq \frac{1-e^{-\gamma}}{\gamma} \mathcal{R}^{\mathcal{A}}_{1,\mathrm{Adv}_1^{\mathrm{f}}(\mathbf{Q}_0[B_1])}$$

where $\mathcal{A}' = \mathtt{OMBQ}(\mathcal{A}, \mathtt{BQMO}, \mathrm{Id})$.

Algorithm 2: $\mathtt{BQMO}$

Input: First order query oracle, point $\mathbf{x}$
Sample $z \in [0,1]$ according to Equation (1)
Return the output of the query oracle at $z * \mathbf{x}$

Non-monotone functions over general convex sets

Lemma 3. [[4]] Let $f : [0,1]^d \to \mathbb{R}$ be a non-negative continuous DR-submodular differentiable function and let $\underline{\mathbf{x}} \in \mathcal{K}$. Define $F : [0,1]^d \to \mathbb{R}$ as the function $F(\mathbf{x}) := \int_0^1 \frac{2}{3z(1-\frac{z}{2})^3} \left(f\left(\frac{z}{2} * (\mathbf{x} - \underline{\mathbf{x}}) + \underline{\mathbf{x}}\right) - f(\underline{\mathbf{x}}) \right) dz$. Then F is differentiable and, if the random variable $\mathcal{Z} \in [0,1]$ is defined by the law

$$\forall z \in [0,1], \quad \mathbb{P}(\mathcal{Z} \leq z) = \int_0^z \frac{1}{3(1-\frac{u}{2})^3} du, \qquad (2)$$

then we have $\mathbb{E}\left[\nabla f\left(\frac{\mathcal{Z}}{2} * (\mathbf{x} - \underline{\mathbf{x}}) + \underline{\mathbf{x}}\right)\right] = \nabla F(\mathbf{x})$. Moreover, we have

$$\frac{1 - \|\underline{\mathbf{x}}\|_\infty}{4} f(\mathbf{y}) - f\left(\frac{\mathbf{x} + \underline{\mathbf{x}}}{2}\right) \leq \frac{3}{8} \langle \nabla F(\mathbf{x}), \mathbf{y} - \mathbf{x} \rangle.$$

Theorem 4. Let $\mathcal{K} \subseteq [0,1]^d$ be a convex set, $\underline{\mathbf{u}} \in \mathcal{K}$, $h := \|\underline{\mathbf{u}}\|_\infty$ and let $\mathbf{F}$ be an M_1-Lipschitz function class over $\mathcal{K}$ where every $f \in \mathbf{F}$ may be extended to a function described in the above lemma. Then, for any $B_1 \geq M_1$, we have

$$\mathcal{R}^{\mathcal{A}'}_{\frac{1-h}{4},\mathrm{Adv}_1^{\mathrm{o}}(\mathbf{F},B_1)} \leq \frac{3}{8} \mathcal{R}^{\mathcal{A}}_{1,\mathrm{Adv}_1^{\mathrm{f}}(\mathbf{Q}_0[B_1])},$$

where $\mathcal{A}' = \mathtt{OMBQ}(\mathcal{A}, \mathtt{BQN}, \mathbf{x} \mapsto \frac{\mathbf{x}+\underline{\mathbf{x}}}{2})$.

Algorithm 3: $\mathtt{BQN}$

Input: First order query oracle, point $\mathbf{x}$
Sample $z \in [0,1]$ according to Equation 2
Return the output of the query oracle at $\frac{z}{2} * (\mathbf{x} - \underline{\mathbf{x}}) + \underline{\mathbf{x}}$

Other meta-algorithms

We extend the applicability of meta-algorithms FOTZO, STB and FOTZO-2P in [1] to all α-regret (as opposed to 1-regret). Given an algorithm designed for stochastic first order feedback, FOTZO converts it to an algorithm that require stochastic zeroth-order feedback and FOTZO-2P converts it to an algorithm that require deterministic zeroth-order feedback. If the algorithm is semi-bandit, STB converts it to a bandit algorithm. We also introduce a new meta-algorithm, namely SFTT converts algorithms that are designed for full-information feedback into algorithms that only require trivial query (e.g. semi-bandit/bandit).

Applications

The figure below captures the applications that are mentioned in the tables. To obtain a result from the graph, let $\mathcal{A}$ be one of SO-OGA ([5]) or IA ([6]) and select a directed path that has the following properties: (i) The path starts at one of the three nodes on the left. (ii) The path must be at least of length 1 and the edges must be the same color. (iii) If $\mathcal{A}$ is IA, the path should not contain SFTT or OTB.

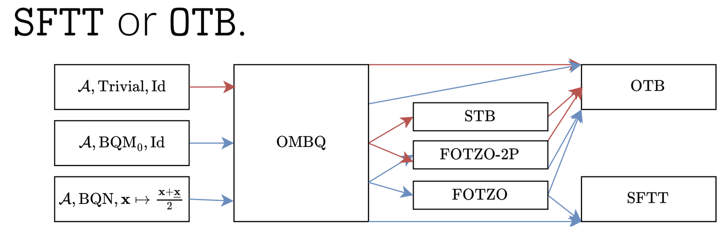

Online results

F	Set	Feedback			Reference	Appx.	# of queries	$\log_T$(α-regret)
Monotone	$0 \in \mathcal{K}$	∇F	Full Information	stoch.	[3] ‡ (*)	$1-e^{-\gamma}$	1	$1/2$
					[7]	$1-e^{-1}$	$T^\theta(\theta \in [0,1/2])$	$2/3-\theta/3$
					Corollary 7-c	$1-e^{-\gamma}$	1	$1/2$
			Semi-bandit	stoch.	[7]	$1-e^{-1}$	-	$3/4$
					Corollary 7-c	$1-e^{-\gamma}$	-	$2/3$
		F	Full Information	det.	Corollary 7-c	$1-e^{-1}$	2	$1/2$
				stoch.	[7]	$1-e^{-1}$	$T^\theta(\theta \in [0,1/4])$	$4/5-\theta/5$
					Corollary 7-c	$1-e^{-\gamma}$	1	$3/4$
			Bandit	det.	[8] ‡‡	$1-e^{-1}$	-	$3/4$
					[4] ‡(*)	$1-e^{-\gamma}$	-	$4/5$
				stoch.	[7]	$1-e^{-1}$	-	$5/6$
					Corollary 7-c	$1-e^{-\gamma}$	-	$4/5$
	general	∇F	Full Information	stoch.	[7]	$1/2$	$T^\theta(\theta \in [0,1/2])$	$2/3-\theta/3$
			Semi-bandit	stoch.	[9]‡(*)	$\gamma^2/(1+\gamma^2)$	-	$1/2$
					[7]	$1/2$	-	$3/4$
					Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	-	$1/2$
		F	Full Information	det.	Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	2	$1/2$
				stoch.	[7]	$1/2$	$T^\theta(\theta \in [0,1/4])$	$4/5-\theta/5$
			Bandit	stoch.	[7]	$1/2$	-	$5/6$
					Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	-	$3/4$
Non-Monotone	general	∇F	Full Information	stoch.	[7]	$(1-h)/4$	$T^\theta(\theta \in [0,1/2])$	$2/3-\theta/3$
					[4] ‡(*)	$(1-h)/4$	1	$1/2$
					Corollary 7-d	$(1-h)/4$	1	$1/2$
			Semi-bandit	stoch.	[7]	$(1-h)/4$	-	$3/4$
					Corollary 7-d	$(1-h)/4$	-	$2/3$
		F	Full Information	det.	Corollary 7-d	$(1-h)/4$	2	$1/2$
				stoch.	[7]	$(1-h)/4$	$T^\theta(\theta \in [0,1/4])$	$4/5-\theta/5$
					Corollary 7-d	$(1-h)/4$	1	$3/4$
			Bandit	det.	[4] ‡(*)	$(1-h)/4$	-	$4/5$
				stoch.	[7]	$(1-h)/4$	-	$5/6$
					Corollary 7-d	$(1-h)/4$	-	$4/5$

Offline results

F	Set	Feedback		Reference	Appx.	Complexity
Monotone	$0 \in \mathcal{K}$	∇F	stoch.	[10]	$1-e^{-\gamma}$	$O(1/\epsilon^3)$
				[11]	$1-e^{-\gamma}$	$O(1/\epsilon^2)$
				[3] ‡	$1-e^{-\gamma}$	$O(1/\epsilon^2)$
				Corollary 7-c	$1-e^{-\gamma}$	$O(1/\epsilon^2)$
		F	det.	[12]	$1-e^{-\gamma}$	$O(1/\epsilon^3)$
				Corollary 7-c	$1-e^{-\gamma}$	$O(1/\epsilon^2)$
			stoch.	[12]	$1-e^{-\gamma}$	$O(1/\epsilon^5)$
				Corollary 7-c	$1-e^{-\gamma}$	$O(1/\epsilon^4)$
	general	∇F	stoch.	[2]‡	$\gamma^2/(1+\gamma^2)$	$O(1/\epsilon^2)$
				[12]	$\gamma^2/(1+\gamma^2)$	$\tilde{O}(1/\epsilon^3)$
				Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	$O(1/\epsilon^2)$
		F	det.	[13]	$\gamma^2/(1+\gamma^2)$	$\tilde{O}(1/\epsilon^3)$
				Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	$O(1/\epsilon^2)$
			stoch.	[13]	$\gamma^2/(1+\gamma^2)$	$\tilde{O}(1/\epsilon^5)$
				Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	$O(1/\epsilon^4)$
Non-Monotone	general	∇F	stoch.	[12]	$\frac{\gamma(1-\gamma h)}{\gamma'-1}\left(\frac{1}{2}-\frac{1}{2\gamma'}\right)$	$O(1/\epsilon^3)$
				[4] ‡	$(1-h)/4$	$O(1/\epsilon^2)$
				Corollary 7-d	$(1-h)/4$	$O(1/\epsilon^2)$
		F	det.	[12]	$\frac{\gamma(1-\gamma h)}{\gamma'-1}\left(\frac{1}{2}-\frac{1}{2\gamma'}\right)$	$O(1/\epsilon^3)$
				Corollary 7-d	$(1-h)/4$	$O(1/\epsilon^2)$
			stoch.	[12]	$\frac{\gamma(1-\gamma h)}{\gamma'-1}\left(\frac{1}{2}-\frac{1}{2\gamma'}\right)$	$O(1/\epsilon^5)$
				Corollary 7-d	$(1-h)/4$	$O(1/\epsilon^4)$

Online non-stationary results

F	Set	Feedback			Reference	Appx.	regret type	α-regret
Monotone	$0 \in \mathcal{K}$	∇F	Full Information	stoch.	Corollary 8-c	$1-e^{-\gamma}$	dynamic	$T^{1/2}(1+P_T)^{1/2}$
					Corollary 7-c	$1-e^{-\gamma}$	adaptive	$T^{1/2}$
			Semi-bandit	stoch.	Corollary 7-c	$1-e^{-\gamma}$	adaptive	$T^{2/3}$
		F	Full Information	det.	Corollary 8-c	$1-e^{-\gamma}$	dynamic	$T^{1/2}(1+P_T)^{1/2}$
					Corollary 7-c	$1-e^{-\gamma}$	adaptive	$T^{1/2}$
				stoch.	Corollary 8-c	$1-e^{-\gamma}$	dynamic	$T^{3/4}(1+P_T)^{1/2}$
					Corollary 7-c	$1-e^{-\gamma}$	adaptive	$T^{3/4}$
			Bandit	stoch.	Corollary 7-c	$1-e^{-\gamma}$	adaptive	$T^{4/5}$
	general	∇F	Semi-bandit	stoch.	Corollary 8-b	$\gamma^2/(1+c\gamma^2)$	dynamic	$T^{1/2}(1+P_T)^{1/2}$
					Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	adaptive	$T^{1/2}$
		F	Full Information	det.	Corollary 8-c	$\gamma^2/(1+c\gamma^2)$	dynamic	$T^{1/2}(1+P_T)^{1/2}$
					Corollary 7-c	$\gamma^2/(1+c\gamma^2)$	adaptive	$T^{1/2}$
			Bandit	stoch.	Corollary 8-b	$\gamma^2/(1+c\gamma^2)$	dynamic	$T^{3/4}(1+P_T)^{1/2}$
					Corollary 7-b	$\gamma^2/(1+c\gamma^2)$	adaptive	$T^{3/4}$
Non-Monotone	general	∇F	Full Information	stoch.	Corollary 8-d	$(1-h)/4$	dynamic	$T^{1/2}(1+P_T)^{1/2}$
					Corollary 7-d	$(1-h)/4$	adaptive	$T^{1/2}$
			Semi-bandit	stoch.	Corollary 7-d	$(1-h)/4$	adaptive	$T^{2/3}$
		F	Full Information	det.	Corollary 8-d	$(1-h)/4$	dynamic	$T^{1/2}(1+P_T)^{1/2}$
					Corollary 7-d	$(1-h)/4$	adaptive	$T^{1/2}$
				stoch.	Corollary 8-d	$(1-h)/4$	dynamic	$T^{3/4}(1+P_T)^{1/2}$
					Corollary 7-d	$(1-h)/4$	adaptive	$T^{3/4}$
			Bandit	stoch.	Corollary 7-d	$(1-h)/4$	adaptive	$T^{4/5}$

References

[1] Mohammad Pedramfar and Vaneet Aggarwal. A unified framework for analyzing meta-algorithms in online convex optimization. *arXiv preprint arXiv:2402.08621*, 2024.

[2] Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, 2017.

[3] Qixin Zhang, Zengde Deng, Zaiyi Chen, Haoyuan Hu, and Yu Yang. Stochastic continuous submodular maximization: Boosting via non-oblivious function. In *Proceedings of the 39th International Conference on Machine Learning*, 2022.

[4] Qixin Zhang, Zongqi Wan, Zengde Deng, Zaiyi Chen, Xiaoming Sun, Jialin Zhang, and Yu Yang. Boosting gradient ascent for continuous DR-submodular maximization. *arXiv preprint arXiv:2401.08330*, 2024.

[5] Dan Garber and Ben Kretzu. New projection-free algorithms for online convex optimization with adaptive regret guarantees. In *Proceedings of Thirty Fifth Conference on Learning Theory*, pages 2326–2359. PMLR, 2022.

[6] Lijun Zhang, Shiyin Lu, and Zhi-Hua Zhou. Adaptive online learning in dynamic environments. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

[7] Mohammad Pedramfar, Yididiya Y. Nadew, Christopher John Quinn, and Vaneet Aggarwal. Unified projection-free algorithms for adversarial DR-submodular optimization. In *The Twelfth International Conference on Learning Representations*, 2024.

[8] Zongqi Wan, Jialin Zhang, Wei Chen, Xiaoming Sun, and Zhijie Zhang. Bandit multi-linear dr-submodular maximization and its applications on adversarial submodular bandits. In *International Conference on Machine Learning*, 2023.

[9] Lin Chen, Hamed Hassani, and Amin Karbasi. Online continuous submodular maximization. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, April 2018.

[10] Aryan Mokhtari, Hamed Hassani, and Amin Karbasi. Stochastic conditional gradient methods: From convex minimization to submodular maximization. *The Journal of Machine Learning Research*, 21(1):4232–4280, 2020.

[11] Hamed Hassani, Amin Karbasi, Aryan Mokhtari, and Zebang Shen. Stochastic conditional gradient++: (non)convex minimization and continuous submodular maximization. *SIAM Journal on Optimization*, 30(4):3315–3344, 2020.

[12] Mohammad Pedramfar, Christopher Quinn, and Vaneet Aggarwal. A unified approach for maximizing continuous γ-weakly DR-submodular functions. *optimization-online preprint optimization-online:25915*, 2024.

[13] Mohammad Pedramfar, Christopher John Quinn, and Vaneet Aggarwal. A unified approach for maximizing continuous DR-submodular functions. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.