

CS395T: Continuous Algorithms

Homework II

Kevin Tian

Due date: February 20, 2025, start of class (3:30 PM).

Please list all collaborators on the first page of your solutions. Unless we have discussed and I have specified otherwise, homework is not accepted if it is not turned in by hand at the start of class, or turned in electronically on Canvas by then. Send me an email to discuss any exceptions.

1 Problem 1

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a differentiable, convex function, and for all $\lambda > 0$, let

$$\mathbf{x}_\lambda^* := \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) + \frac{\lambda}{2} \|\mathbf{x}\|_2^2.$$

- (i) Prove that for any $\lambda' > \lambda > 0$, we have $\|\mathbf{x}_{\lambda'}^*\|_2 \geq \|\mathbf{x}_\lambda^*\|_2$.
- (ii) Suppose $\|\nabla f(\mathbf{0}_d)\|_2 \leq L$ and let $R > 0$. Give a λ_R such that for any $\lambda \geq \lambda_R$, $\|\mathbf{x}_\lambda^*\|_2 \leq R$.

2 Problem 2

Let $\mathbf{M} \in \mathbb{S}_{>\mathbf{0}}^{d \times d}$, and let $\|\cdot\|$ be a norm on \mathbb{R}^d .

- (i) Prove that if $\mathbf{M}[\mathbf{v}, \mathbf{v}] \leq \|\mathbf{v}\|^2$ for all $\mathbf{v} \in \mathbb{R}^d$, then $\|\mathbf{M}\mathbf{v}\|_* \leq \|\mathbf{v}\|$ for all $\mathbf{v} \in \mathbb{R}^d$.¹
- (ii) Prove that $\mathbf{M}[\mathbf{v}, \mathbf{v}] \leq \|\mathbf{v}\|^2$ for all $\mathbf{v} \in \mathbb{R}^d$ iff $\mathbf{M}^{-1}[\mathbf{v}, \mathbf{v}] \geq \|\mathbf{v}\|_*^2$ for all $\mathbf{v} \in \mathbb{R}^d$.²

3 Problem 3

Let $\mathcal{X} \subseteq \mathbb{R}^d$, let $f : \mathcal{X} \rightarrow \mathbb{R}$ be convex, differentiable, and L -Lipschitz in $\|\cdot\|$, and let $\psi : \mathcal{X} \rightarrow \mathbb{R}$ be convex and admit subgradients everywhere in \mathcal{X} . Let $F(\mathbf{x}) = f(\mathbf{x}) + \psi(\mathbf{x})$.³ Finally, let $\varphi : \mathcal{X} \rightarrow \mathbb{R}$ be 1-strongly convex in $\|\cdot\|$ and of Legendre type. Consider iterating the update

$$\mathbf{x}_{t+1} \leftarrow \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} \{\langle \eta \nabla f(\mathbf{x}_t), \mathbf{x} \rangle + \eta \psi(\mathbf{x}) + D_\varphi(\mathbf{x} \|\mathbf{x}_t)\}, \text{ for } 0 \leq t < T, \eta > 0.$$

Give an initialization strategy for choosing \mathbf{x}_0 that depends only on ψ and \mathcal{X} (but not f), and prove that following your initialization strategy,

$$F(\bar{\mathbf{x}}) - F(\mathbf{x}^*) \leq \frac{\Theta}{\eta T} + \frac{\eta L^2}{2},$$

where $\bar{\mathbf{x}} = \frac{1}{T} \sum_{0 \leq t < T} \mathbf{x}_t$, $\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x})$, and $\Theta \geq \sup_{\mathbf{x}_0 \in \mathcal{X}} D_\varphi(\mathbf{x}^* \|\mathbf{x}_0)$.

¹This justifies that for general norms, the second claim in Lemma 14, Part II is true. Note that the last part of the proof of Lemma 6, Part II implicitly used this claim in the ℓ_2 case.

²Combined with $\nabla^2 f^*(\nabla f(\mathbf{x})) = (\nabla^2 f(\mathbf{x}))^{-1}$, shown in class when f, f^* are twice-differentiable, this gives a simple proof of smoothness-strong convexity duality for twice-differentiable convex functions in general norms.

³The purpose of this problem is to explore how to generalize the mirror descent framework in Theorem 2, Part III, to handle composite objectives just as gradient descent can, as shown by Section 5.2, Part II.

4 Problem 4

- (i) Prove that the following $2d \times 2d$ matrix is positive semidefinite, for any $a \in \mathbb{R}$ and $\mathbf{M} \in \mathbb{S}_{\succeq \mathbf{0}}^{d \times d}$:

$$\begin{pmatrix} a^2 \mathbf{M} & a \mathbf{M} \\ a \mathbf{M} & \mathbf{M} \end{pmatrix}.$$

- (ii) Prove that for any $\mathbf{a} \in \mathbb{R}^n$, and $\{\mathbf{M}_i\}_{i \in [n]} \subset \mathbb{S}_{\succeq \mathbf{0}}^{d \times d}$ satisfying $\sum_{i \in [n]} \mathbf{M}_i \preceq \mathbf{I}_d$, we have⁴

$$\left(\sum_{i \in [n]} \mathbf{a}_i \mathbf{M}_i \right)^2 \preceq \sum_{i \in [n]} \mathbf{a}_i^2 \mathbf{M}_i.$$

5 Problem 5

Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be differentiable, and let $\mathbf{x}^* := \operatorname{argmin}_{\mathbf{x} \in \mathbb{B}(R)} f(\mathbf{x})$ exist and be unique. Suppose that \mathcal{A} is a randomized algorithm which produces a point $\mathbf{x} \in \mathbb{B}(R)$ satisfying

$$\mathbb{E}[f(\mathbf{x}) - f(\mathbf{x}^*)] \leq \epsilon.$$

Let $\mathbf{g} : \mathbb{B}(R) \rightarrow \mathbb{R}^d$ be a randomized estimator satisfying

$$\mathbb{E} \mathbf{g}(\mathbf{x}) = \nabla f(\mathbf{x}), \quad \mathbb{E} \|\mathbf{g}(\mathbf{x})\|_2^2 \leq L^2, \quad \text{for all } \mathbf{x} \in \mathbb{B}(R).$$

- (i) Let $\alpha > 0$, $\delta \in (0, 1)$. Given $\mathbf{x}, \mathbf{x}' \in \mathbb{B}(R)$, design an algorithm which calls \mathbf{g} $O(\frac{L^2 R^2}{\alpha^2} \log \frac{1}{\delta})$ times, and estimates $f(\mathbf{x}) - f(\mathbf{x}')$ to additive error α with probability $\geq 1 - \delta$.⁵
- (ii) Give an algorithm which calls \mathcal{A} $O(\log \frac{1}{\delta})$ times and \mathbf{g} $O(\frac{L^2 R^2}{\epsilon^2} \operatorname{polylog}(\frac{1}{\delta}))$ times,⁶ and produces \mathbf{x} such that with probability $\geq 1 - \delta$,

$$f(\mathbf{x}) - f(\mathbf{x}^*) \leq 3\epsilon.$$

⁴This can be viewed as a matrix extension of the Cauchy-Schwarz inequality.

⁵It may be helpful to first write $f(\mathbf{x}) - f(\mathbf{x}')$ as an integral involving gradients.

⁶I am curious to see how low people can get the polylog :)