

OOPS 2020

Mean field methods in high-dimensional statistics and nonconvex optimization

Lecturer: Andrea Montanari

Problem session leader: Michael Celentano

July 7, 2020

Problem Session 1

Problem 1: from Gordon's objective to the fixed point equations

Recall *Gordon's min-max problem* is

$$B^*(\mathbf{g}, \mathbf{h}) := \min_{\mathbf{u} \in \mathbb{R}^d} \max_{\mathbf{v} \in \mathbb{R}^n} \left\{ \frac{1}{n} \|\mathbf{v}\| \langle \mathbf{g}, \mathbf{u} \rangle + \frac{1}{n} \|\mathbf{u}\| \langle \mathbf{h}, \mathbf{v} \rangle - \frac{\sigma}{n} \langle \mathbf{w}, \mathbf{v} \rangle - \frac{1}{2n} \|\mathbf{v}\|^2 + \frac{\lambda}{\sqrt{n}} \|\boldsymbol{\theta}_0 + \mathbf{u}\|_1 \right\}. \quad (1)$$

In lecture, we claimed that by analyzing Gordon's objective we can show that the Lasso solution is described in terms of the solutions τ^*, β^* to the *fixed point equations*

$$\begin{aligned} \tau^2 &= \sigma^2 + \frac{1}{\delta} \mathbb{E} [(\eta(\Theta + \tau Z; \tau \lambda / \beta) - \Theta)^2], \\ \beta &= \tau \left(1 - \frac{1}{\delta} \mathbb{E} [\eta'(\Theta + \tau Z; \tau \lambda / \beta)] \right), \end{aligned}$$

where η is the solution of the 1-dimensional problem

$$\eta(y; \alpha) := \arg \min_{x \in \mathbb{R}} \left\{ \frac{1}{2} (y - x)^2 + \alpha |x| \right\} = (|x| - \alpha)_+ \text{sign}(x).$$

η is commonly known as *soft-thresholding*. In this problem, we will outline how to derive the fixed point equations from Gordon's min-max problem.

(a) Define $\tilde{\boldsymbol{\theta}}_0 = \sqrt{n} \boldsymbol{\theta}_0$ and $\tilde{\mathbf{u}} = \sqrt{n} \mathbf{u}$. Prove that $B^*(\mathbf{g}, \mathbf{h})$ has the same distribution as

$$\min_{\tilde{\mathbf{u}} \in \mathbb{R}^d} \max_{\beta \geq 0} \left\{ \left(\left\| \frac{\|\tilde{\mathbf{u}}\|}{\sqrt{n}} \frac{\mathbf{h}}{\sqrt{n}} - \frac{\sigma \mathbf{w}}{\sqrt{n}} \right\| - \frac{1}{n} \langle \mathbf{g}, \mathbf{u} \rangle \right) \beta - \frac{1}{2} \beta^2 + \frac{\lambda}{n} \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\}. \quad (2)$$

Hint: Let $\beta = \|\mathbf{v}\| / \sqrt{n}$

(b) Argue (heuristically) that we may approximate the optimization above by

$$\min_{\tilde{\mathbf{u}} \in \mathbb{R}^d} \max_{\beta \geq 0} \left\{ \left(\sqrt{\frac{\|\tilde{\mathbf{u}}\|^2}{n} + \sigma^2} - \frac{1}{n} \langle \mathbf{g}, \mathbf{u} \rangle \right) \beta - \frac{1}{2} \beta^2 + \frac{\lambda}{n} \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\}. \quad (3)$$

Remark: If we maximize over $\beta \geq 0$ explicitly, we get

$$\min_{\tilde{\mathbf{u}} \in \tilde{U}} \left\{ \frac{1}{2} \left(\sqrt{\frac{\|\tilde{\mathbf{u}}\|^2}{n} + \sigma^2} - \frac{\langle \mathbf{g}, \tilde{\mathbf{u}} \rangle}{n} \right)_+^2 + \frac{\lambda}{n} \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\}.$$

Note that the objective on the right-hand side is locally strongly convex around any point $\tilde{\mathbf{u}}$ at which the first term is positive. When $n < p$, the Lasso objective is nowhere locally strongly convex. This convenient feature of the new form of Gordon's problem is very useful for its analysis. We do not explore this further here.

(c) Argue that the quantity in Eq. (3) is equal to

$$\max_{\beta \geq 0} \min_{\tau \geq 0} \left\{ \frac{\sigma^2 \beta}{2\tau} + \frac{\tau \beta}{2} - \frac{\beta^2}{2} + \frac{1}{n} \min_{\tilde{\mathbf{u}} \in \mathbb{R}^d} \left\{ \frac{\beta}{2\tau} \|\tilde{\mathbf{u}}\|^2 - \beta \langle \mathbf{g}, \tilde{\mathbf{u}} \rangle + \lambda \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\} \right\}. \quad (4)$$

Hint: Recall the identity $\sqrt{x} = \min_{\tau \geq 0} \left\{ \frac{x}{2\tau} + \frac{\tau}{2} \right\}$.

(d) Write

$$\hat{\mathbf{u}} := \arg \min_{\tilde{\mathbf{u}} \in \mathbb{R}^d} \left\{ \frac{\beta}{2\tau} \|\tilde{\mathbf{u}}\|^2 - \beta \langle \mathbf{g}, \tilde{\mathbf{u}} \rangle + \lambda \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\}$$

in terms of the soft-thresholding operator.

(e) Compute

(i) the derivative of $\min_{\tilde{\mathbf{u}} \in \mathbb{R}^d} \left\{ \frac{\beta}{2\tau} \|\tilde{\mathbf{u}}\|^2 - \beta \langle \mathbf{g}, \tilde{\mathbf{u}} \rangle + \lambda \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\}$ with respect to τ .

(ii) the derivative of $\min_{\tilde{\mathbf{u}} \in \mathbb{R}^d} \left\{ \frac{\beta}{2\tau} \|\tilde{\mathbf{u}}\|^2 - \beta \langle \mathbf{g}, \tilde{\mathbf{u}} \rangle + \lambda \|\tilde{\boldsymbol{\theta}}_0 + \tilde{\mathbf{u}}\|_1 \right\}$ with respect to β .

(f) Write $\frac{1}{n} \mathbb{E}[\|\hat{\mathbf{u}}\|^2]$ and $\frac{1}{n} \mathbb{E}[\langle \mathbf{g}, \hat{\mathbf{u}} \rangle]$ as an expectation over the random variables $(\Theta, Z) \sim \hat{\mu}_{\boldsymbol{\theta}_0} \otimes \mathbf{N}(0, 1)$. For the latter, rewrite it using Gaussian integration by parts.

(g) Take the derivative of the objective in Eq. (4) with respect to τ and with respect to β . Show that setting the expectations of these derivatives to 0 is equivalent to the fixed point equations.

Problem 2: Gordon's objective for max-margin classification

The use of Gordon's technique extends well beyond linear models. In logistic regression, we receive iid samples according to

$$y_i \sim \text{Rad}(f(\langle \mathbf{x}_i, \boldsymbol{\theta}_0 \rangle)), \quad \mathbf{x}_i \sim \mathbf{N}(0, \mathbf{I}_d),$$

$$f(x) = \frac{\exp(x)}{\exp(x) + \exp(-x)}.$$

For a certain $\delta^* > 0$ the following occurs: when $n/p \rightarrow \delta < \delta^*$, $n, p \rightarrow \infty$, with high probability there exists $\boldsymbol{\theta}$ such that $y_i \langle \mathbf{x}_i, \boldsymbol{\theta} \rangle > 0$ for all $i = 1, \dots, n$. In such a regime, the data is *linearly separable*. The max-margin classifier is defined as

$$\hat{\boldsymbol{\theta}} \in \arg \max_{\boldsymbol{\theta}} \left\{ \min_{i \leq n} y_i \langle \mathbf{x}_i, \boldsymbol{\theta} \rangle : \|\boldsymbol{\theta}\| \leq 1 \right\}, \quad (5)$$

and the value of the optimization problem, which we denote by $\kappa(\mathbf{y}, \mathbf{X})$, is called the *maximum margin*. To simplify notation, we assume in this problem that $\|\boldsymbol{\theta}_0\| = 1$. In this problem, we outline how to set up Gordon's problem for max-margin classification. The analysis of Gordon's objective is complicated, and we do not describe it here. See Montanari, Ruan, Sohn, Yan (2019+). "The generalization error of max-margin linear classifiers: High-dimensional asymptotics in the overparameterized regime." [arxiv:1911.01544](#).

(a) Show that

$$\kappa(\mathbf{y}, \mathbf{X}) \geq \kappa \text{ if and only if } \min_{\|\boldsymbol{\theta}\| \leq 1} \|(\kappa \mathbf{1} - (\mathbf{y} \odot \mathbf{X}\boldsymbol{\theta}))_+\|_2 = 0.$$

Argue that

$$\begin{aligned} \min_{\|\boldsymbol{\theta}\| \leq 1} \frac{1}{\sqrt{d}} \|(\kappa \mathbf{1} - (\mathbf{y} \odot \mathbf{X}\boldsymbol{\theta}))_+\|_2 &= \min_{\|\boldsymbol{\theta}\| \leq 1} \max_{\|\boldsymbol{\lambda}\| \leq 1, \lambda \geq 0} \frac{1}{\sqrt{d}} \boldsymbol{\lambda}^\top (\kappa \mathbf{1} - \mathbf{y} \odot \mathbf{X}\boldsymbol{\theta}) \\ &= \min_{\|\boldsymbol{\theta}\| \leq 1} \max_{\|\boldsymbol{\lambda}\| \leq 1, \lambda \odot \mathbf{y} \geq 0} \frac{1}{\sqrt{d}} \boldsymbol{\lambda}^\top (\kappa \mathbf{y} - \mathbf{X}\boldsymbol{\theta}). \end{aligned}$$

(b) Why can't we use Gordon's inequality to compare the preceding min-max problem to

$$\min_{\|\boldsymbol{\theta}\| \leq 1} \max_{\|\boldsymbol{\lambda}\| \leq 1, \mathbf{y} \odot \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{d}} (\kappa \boldsymbol{\lambda}^\top \mathbf{y} + \|\boldsymbol{\lambda}\| \mathbf{g}^\top \boldsymbol{\theta} + \|\boldsymbol{\theta}\| \mathbf{h}^\top \boldsymbol{\lambda})?$$

(c) Let $\tilde{\mathbf{x}} = \mathbf{X}\boldsymbol{\theta}_0$. Show that the min-max problem is equivalent to

$$\min_{\|\boldsymbol{\theta}\| \leq 1} \max_{\|\boldsymbol{\lambda}\| \leq 1, \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{d}} \boldsymbol{\lambda}^\top (\kappa \mathbf{1} - (\mathbf{y} \odot \tilde{\mathbf{x}}) \langle \boldsymbol{\theta}_0, \boldsymbol{\theta} \rangle - \mathbf{y} \odot \mathbf{X} \Pi_{\boldsymbol{\theta}_0^\perp} \boldsymbol{\theta}).$$

Here $\Pi_{\boldsymbol{\theta}_0^\perp}$ is the projection operator onto the orthogonal complement of the space spanned by $\boldsymbol{\theta}_0$. Argue that

$$\begin{aligned} \mathbb{P} \left(\min_{\|\boldsymbol{\theta}\| \leq 1} \max_{\|\boldsymbol{\lambda}\| \leq 1, \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{d}} \boldsymbol{\lambda}^\top (\kappa \mathbf{1} - (\mathbf{y} \odot \tilde{\mathbf{x}}) \langle \boldsymbol{\theta}_0, \boldsymbol{\theta} \rangle - \mathbf{y} \odot \mathbf{X} \Pi_{\boldsymbol{\theta}_0^\perp} \boldsymbol{\theta}) \leq t \right) \\ \leq 2\mathbb{P} \left(\min_{\|\boldsymbol{\theta}\| \leq 1} \max_{\|\boldsymbol{\lambda}\| \leq 1, \boldsymbol{\lambda} \geq 0} \frac{1}{\sqrt{d}} \left(\boldsymbol{\lambda}^\top (\kappa \mathbf{1} - (\mathbf{y} \odot \tilde{\mathbf{x}}) \langle \boldsymbol{\theta}_0, \boldsymbol{\theta} \rangle) + \|\boldsymbol{\lambda}\| \mathbf{g}^\top \Pi_{\boldsymbol{\theta}_0^\perp} \boldsymbol{\theta} + \|\Pi_{\boldsymbol{\theta}_0^\perp} \boldsymbol{\theta}\| \mathbf{h}^\top \boldsymbol{\lambda} \right) \leq t \right), \end{aligned}$$

and likewise for the comparison of the probabilities that the min-max values exceed t , where $\mathbf{g} \sim \mathbf{N}(0, \mathbf{I}_d)$ and $\mathbf{h} \sim \mathbf{N}(0, \mathbf{I}_n)$ independent of everything else.

(d) What is the limit in Wasserstein-2 distance of $\frac{1}{n} \sum_{i=1}^n \delta_{(y_i, \tilde{x}_i, h_i)}$?