Name: ____________________ Student ID: ________________ Scores: ____________________

AIE6003 Homework 4
Instructor: Xiaopeng Li **Deadline**: May 6

Notice: If you are trying to invoke a result outside of this course, then you should put a reference to it. The questions with 1 star are easy, with 2 stars are medium, and with 3 stars are hard. Easy and medium level of questions are likely to occur in the exams, while hard questions are not likely to occur in the exams.

1*. (15 points) Denote $[m] := \{1, \ldots, m\}$ for any positive integer m. Let

$$B_*(\tau) := \{X \in \mathbb{R}^{n\times n} : \|X\|_* \le \tau\}, \qquad \tau > 0.$$

Throughout this problem, you may use without proof the result from our lecture slides.

(a) (4 points) Let $A = \sigma uv^\top$, where $\sigma > 0$ and $\|u\|_2 = \|v\|_2 = 1$. Consider the fully observed problem

$$\min_{X \in B_*(\tau)} \frac{1}{2}\|X - A\|_F^2.$$

Show that the unique optimizer is $X^\star = \min\{\sigma, \tau\}uv^\top$.

Solution: Since $A = \sigma uv^\top$ has singular values $(\sigma, 0, \ldots, 0)$, the Euclidean projection of A onto the nuclear-norm ball is obtained by projecting its singular-value vector onto the ℓ_1 ball of radius τ. Thus the projected singular values are $(\min\{\sigma, \tau\}, 0, \ldots, 0)$, with the same singular vectors. Therefore $X^\star = \min\{\sigma, \tau\}\, uv^\top$. The objective $X \mapsto \frac{1}{2}\|X - A\|_F^2$ is strictly convex, so this optimizer is unique.

(b) (4 points) Starting from $X_0 = 0$, consider the following update. Set $S_0 := \tau uv^\top$. Then choose the scalar $\eta^\star \in [0, 1]$ that gives the smallest objective value along the segment from 0 to S_0:

$$\eta^\star \in \operatorname*{argmin}_{0\le\eta\le1} \frac{1}{2}\|\eta S_0 - A\|_F^2, \qquad X_1 := \eta^\star S_0.$$

Show that $\eta^\star = \min\left\{\frac{\sigma}{\tau}, 1\right\}$ and hence $X_1 = X^\star$. Conclude that, for this special rank-one fully observed problem, this rank-one update solves the constrained problem in one iteration.

Solution: By the definition in the question, $S_0 = \tau uv^\top$. Since $\|uv^\top\|_F = 1$, the one-dimensional minimization problem is

$$\min_{0\le\eta\le1} \frac{1}{2}\|\eta S_0 - A\|_F^2 = \min_{0\le\eta\le1} \frac{1}{2}\|\eta\tau uv^\top - \sigma uv^\top\|_F^2 = \min_{0\le\eta\le1} \frac{1}{2}(\eta\tau - \sigma)^2.$$

The unconstrained minimizer is $\eta = \sigma/\tau$. Restricting to $[0, 1]$ gives $\eta^\star = \min\left\{\frac{\sigma}{\tau}, 1\right\}$. Hence

$$X_1 = \eta^\star S_0 = \eta^\star \tau uv^\top = \begin{cases} \sigma uv^\top, & \sigma \le \tau, \\ \tau uv^\top, & \sigma > \tau, \end{cases} = \min\{\sigma, \tau\}\, uv^\top = X^\star.$$

So this rank-one update solves the special fully observed rank-one problem in one step.

(c) (3 points) Now return to matrix completion. Define $M_{\text{coh}} = e_1 e_1^\top$ and $M_{\text{inc}} = \frac{1}{n}\mathbf{1}\mathbf{1}^\top$. Suppose $\Omega \subset [n] \times [n]$ is sampled uniformly without replacement with $|\Omega| = m$. Show that if $(1,1) \notin \Omega$, then $P_\Omega(M_{\text{coh}}) = 0$. Conclude that on this event, no algorithm can distinguish M_{coh} from the zero matrix using only the observations. Compute $\mathbb{P}((1,1) \notin \Omega)$.

Solution: The coherent matrix $M_{\text{coh}} = e_1 e_1^\top$ has only one nonzero entry, namely $(M_{\text{coh}})_{11} = 1$. Therefore, if $(1,1) \notin \Omega$, all observed entries of M_{coh} are zero, so $P_\Omega(M_{\text{coh}}) = 0$. On this event, the observations from M_{coh} are identical to the observations from the zero matrix. No algorithm using only $P_\Omega(M)$ can distinguish these two cases. Because Ω is a uniformly sampled subset of size m from n^2 entries,

$$\mathbb{P}\big((1,1) \in \Omega\big) = \frac{m}{n^2}, \quad \mathbb{P}\big((1,1) \notin \Omega\big) = 1 - \frac{m}{n^2}.$$

(d) (4 points) Let $n = 40$, $\tau = 1$, and

$$f(X) = \frac{1}{2}\|P_\Omega(X - Y)\|_F^2, \qquad Y = P_\Omega(M).$$

For each trial, sample the same index set Ω for both matrices, with $m \in \{2n, 4n, 6n, 8n\}$. Run Frank-Wolfe from $X_0 = 0$ for 200 iterations with $\eta_t = \frac{2}{t+2}$. If the gradient is zero, you may leave the iterate unchanged. Report the average relative error

$$\frac{\|X_{200} - M\|_F}{\|M\|_F}$$

over 20 trials, and plot the average error versus m for both M_{inc} and M_{coh}. Briefly explain why the coherent case is much harder even though both matrices have rank 1 and nuclear norm 1.

Solution: The important thing is that if the gradient is zero, the iterate should be left unchanged; otherwise a numerical SVD routine may return an arbitrary singular-vector pair for the zero matrix and create an artificial movement.

For the coherent matrix, in expectation the final error is close to

$$\mathbb{P}\big((1,1) \notin \Omega\big) = 1 - \frac{m}{n^2},$$

because if $(1,1)$ is observed then the first Frank-Wolfe step recovers $e_1 e_1^\top$, while if $(1,1)$ is not observed the data are identical to zero. For $n = 40$, this gives approximately $0.95, 0.90, 0.85, 0.80$ for $m = 2n, 4n, 6n, 8n$. The incoherent matrix spreads its mass over all entries, so every random sample carries some information about the singular vectors and scale. This is why the incoherent error should decrease much more as m increases.

2**. (15 points) Consider the LASSO problem

$$\min_{\beta \in \mathbb{R}^d} F(\beta) := \frac{1}{2}\|A\beta - b\|_2^2 + \lambda\|\beta\|_1,$$

where $A = [a_1, \ldots, a_d] \in \mathbb{R}^{n \times d}$, $b \in \mathbb{R}^n$, and $\lambda > 0$.

(a) (3 points) Show that $\beta = 0$ is optimal if and only if $\|A^\top b\|_\infty \leq \lambda$.

Solution: The Lasso optimality condition is

$$A^\top(A\beta - b) + \lambda s = 0, \qquad s \in \partial\|\beta\|_1.$$

At $\beta = 0$, this becomes

$$-A^\top b + \lambda s = 0, \qquad s \in [-1,1]^d.$$

Such an s exists if and only if

$$\left\| \frac{A^\top b}{\lambda} \right\|_\infty \le 1 \iff \|A^\top b\|_\infty \le \lambda.$$

(b) (4 points) Let $S \subset [d]$, and let $\hat{\theta} \in \mathbb{R}^{|S|}$ be an optimizer of the restricted Lasso problem

$$\min_{\theta \in \mathbb{R}^{|S|}} \frac{1}{2}\|A_S\theta - b\|_2^2 + \lambda\|\theta\|_1.$$

Pad $\hat{\theta}$ with zeros outside S to obtain $\tilde{\beta} \in \mathbb{R}^d$. Show that if $|a_i^\top(A\tilde{\beta} - b)| \le \lambda$ for all $i \notin S$, then $\tilde{\beta}$ is a global minimizer of the full Lasso problem.

Solution: Let $r := A\tilde{\beta} - b = A_S\hat{\theta} - b$. Since $\hat{\theta}$ solves the restricted problem, for every $i \in S$ we have $a_i^\top r \in -\lambda\partial|\tilde{\beta}_i|$. For $i \notin S$, we have $\tilde{\beta}_i = 0$, and the assumed certificate gives $|a_i^\top r| \le \lambda$. This is equivalent to

$$a_i^\top r \in -\lambda[-1,1] = -\lambda\partial|0| = -\lambda\partial|\tilde{\beta}_i|.$$

Therefore the coordinate-wise Lasso optimality condition holds for every $i = 1, \dots, d$. Hence $\tilde{\beta}$ is globally optimal for the full Lasso problem.

(c) (4 points) Now consider the group Lasso problem

$$\min_{\beta \in \mathbb{R}^d} \frac{1}{2}\|A\beta - b\|_2^2 + \lambda \sum_{g=1}^m \|\beta_{B_g}\|_2,$$

where $B_1, \dots, B_m$ partition $[d]$. Let $\mathcal{G} \subset [m]$ be a set of active groups (nonzero variables), and let $\tilde{\beta}$ be an optimizer of the restricted problem on the coordinates in $\bigcup_{g \in \mathcal{G}} B_g$, padded with zeros on the inactive groups. Prove that $\tilde{\beta}$ is globally optimal for the full group Lasso whenever $\|A_{B_g}^\top(A\tilde{\beta} - b)\|_2 \le \lambda$ for $g \notin \mathcal{G}$.

Solution: Let again $r := A\tilde{\beta} - b$. For an active group $g \in \mathcal{G}$, restricted optimality gives

$$A_{B_g}^\top r + \lambda s_g = 0, \qquad s_g \in \partial\|\tilde{\beta}_{B_g}\|_2.$$

For an inactive group $g \notin \mathcal{G}$, we have $\tilde{\beta}_{B_g} = 0$. The subdifferential of the Euclidean norm at zero is

$$\partial\|0\|_2 = \{s : \|s\|_2 \le 1\}.$$

The assumed certificate $\|A_{B_g}^\top r\|_2 \le \lambda$ allows us to choose

$$s_g = -\frac{1}{\lambda} A_{B_g}^\top r, \qquad \|s_g\|_2 \le 1.$$

Thus for every group, active or inactive, there exists $s_g \in \partial\|\tilde{\beta}_{B_g}\|_2$ such that $A_{B_g}^\top r + \lambda s_g = 0$. Stacking these group conditions gives

$$0 \in A^\top(A\tilde{\beta} - b) + \lambda\partial\left(\sum_{g=1}^m \|\tilde{\beta}_{B_g}\|_2\right),$$

which is the global optimality condition for the convex group Lasso problem.

(d) (4 points) Suppose the certificate in part (b) fails. Namely, for some $i \notin S$, $|a_i^\top(A\tilde{\beta} - b)| > \lambda$. Let $r := A\tilde{\beta} - b$ and consider changing only coordinate i: $\beta(\delta) = \tilde{\beta} + \delta e_i$. Derive the exact minimizer of $F(\beta(\delta))$ over $\delta \in \mathbb{R}$, and show that the objective decreases strictly. Quantify the decrease in terms of $a_i^\top r$, λ, and $\|a_i\|_2$.

Solution: Let $c := a_i^\top r$ and $q := \|a_i\|_2^2$. Since $|c| > \lambda$, we must have $q > 0$. Along the one-coordinate path $\beta(\delta) = \tilde{\beta} + \delta e_i$,

$$F(\beta(\delta)) = F(\tilde{\beta}) + c\delta + \frac{q}{2}\delta^2 + \lambda|\delta|.$$

The exact minimizer is the scalar soft-thresholding update

$$\delta^+ = -\frac{1}{q}S_\lambda(c), \qquad S_\lambda(c) = \operatorname{sign}(c)(|c| - \lambda)_+.$$

Because $|c| > \lambda$, this update is nonzero. Substituting the minimizer gives

$$F(\tilde{\beta} + \delta^+ e_i) - F(\tilde{\beta}) = -\frac{(|c| - \lambda)^2}{2q} = -\frac{(|a_i^\top r| - \lambda)^2}{2\|a_i\|_2^2} < 0.$$

3***. (20 points) For $\mu \geq 0$, consider

$$f_\mu(x, y) = \frac{\mu}{2}\|x\|_2^2 + x^\top A y - \frac{\mu}{2}\|y\|_2^2, \qquad A \in \mathbb{R}^{m \times n}.$$

Let $L := \|A\|_{\mathrm{op}}$. Define

$$z = \begin{pmatrix} x \\ y \end{pmatrix}, \qquad K = \begin{pmatrix} 0 & A \\ -A^\top & 0 \end{pmatrix}, \qquad F(z) = \begin{pmatrix} \nabla_x f_\mu(x, y) \\ -\nabla_y f_\mu(x, y) \end{pmatrix},$$

(a) (5 points) Show that $F(z) = (\mu I + K)z$. For $\mu > 0$, prove that F is μ-strongly monotone and ℓ-Lipschitz with $\ell = \sqrt{\mu^2 + L^2}$. Conclude that the saddle point is unique and equals $z^\star = 0$.

Solution: We have

$$\nabla_x f_\mu(x, y) = \mu x + Ay, \qquad \nabla_y f_\mu(x, y) = A^\top x - \mu y.$$

Therefore

$$F(z) = \begin{pmatrix} \mu x + Ay \\ -A^\top x + \mu y \end{pmatrix} = \left(\mu I + \begin{pmatrix} 0 & A \\ -A^\top & 0 \end{pmatrix}\right) z = (\mu I + K)z.$$

Let z_1, z_2 be arbitrary and set $h = z_1 - z_2$. Since $K^\top = -K$, $\langle Kh, h\rangle = 0$. Thus

$$\langle F(z_1) - F(z_2), z_1 - z_2\rangle = \langle(\mu I + K)h, h\rangle = \mu\|h\|_2^2.$$

So F is μ-strongly monotone.

For Lipschitzness,

$$\|F(z_1) - F(z_2)\|_2 = \|(\mu I + K)h\|_2.$$

Moreover,

$$(\mu I + K)^\top(\mu I + K) = \mu^2 I + K^\top K,$$

because $K^\top + K = 0$. Hence

$$\|\mu I + K\|_{\mathrm{op}}^2 = \mu^2 + \|K\|_{\mathrm{op}}^2.$$

Since $\|K\|_{\text{op}} = \|A\|_{\text{op}} = L$, we get

$$\|F(z_1) - F(z_2)\|_2 \le \sqrt{\mu^2 + L^2}\,\|z_1 - z_2\|_2.$$

Finally, when $\mu > 0$, if $F(z) = 0$, then taking inner product with z gives

$$0 = \langle F(z), z\rangle = \mu\|z\|_2^2.$$

Therefore $z = 0$. Thus the saddle point is unique and equals $z^\star = 0$.

(b) (6 points) Simultaneous GDA with step size $\eta > 0$ is

$$z_{t+1} = z_t - \eta F(z_t) = \big((1-\eta\mu)I - \eta K\big)z_t.$$

Use $K^\top = -K$ to prove the exact identity

$$\|z_{t+1}\|_2^2 = (1-\eta\mu)^2\|z_t\|_2^2 + \eta^2\|Kz_t\|_2^2.$$

Deduce that GDA is a global contraction for every initialization if and only if

$$0 < \eta < \frac{2\mu}{\mu^2 + L^2}.$$

Find the best fixed step size and the corresponding worst-case contraction factor.

Solution: The update is $z_{t+1} = \big((1-\eta\mu)I - \eta K\big)z_t$. Therefore

$$\begin{aligned}\|z_{t+1}\|_2^2 &= \|(1-\eta\mu)z_t - \eta K z_t\|_2^2 \\ &= (1-\eta\mu)^2\|z_t\|_2^2 + \eta^2\|Kz_t\|_2^2 - 2\eta(1-\eta\mu)\langle z_t, Kz_t\rangle.\end{aligned}$$

Because $K^\top = -K$, $\langle z_t, Kz_t\rangle = 0$. Hence

$$\|z_{t+1}\|_2^2 = (1-\eta\mu)^2\|z_t\|_2^2 + \eta^2\|Kz_t\|_2^2.$$

Using $\|Kz_t\|_2 \le L\|z_t\|_2$, the worst-case squared contraction factor is

$$q(\eta) = (1-\eta\mu)^2 + \eta^2 L^2 = 1 - 2\eta\mu + \eta^2(\mu^2 + L^2).$$

This is strictly less than 1 exactly when

$$-2\eta\mu + \eta^2(\mu^2 + L^2) < 0 \iff 0 < \eta < \frac{2\mu}{\mu^2 + L^2},$$

given $\eta > 0$. The condition is also necessary for global contraction because there are vectors achieving $\|Kz\|_2 = L\|z\|_2$.
The best fixed step size minimizes $q(\eta)$:

$$q'(\eta) = -2\mu + 2\eta(\mu^2 + L^2) = 0 \implies \eta^\star = \frac{\mu}{\mu^2 + L^2}.$$

The corresponding worst-case squared contraction factor is

$$q(\eta^\star) = 1 - \frac{\mu^2}{\mu^2 + L^2} = \frac{L^2}{\mu^2 + L^2}.$$

Equivalently, the worst-case contraction factor in norm is $L/\sqrt{\mu^2 + L^2}$.

(c) (5 points) Now specialize to the scalar bilinear case $\mu = 0$, $m = n = 1$, $A = [\sigma]$, $\sigma > 0$. Let $w_t := x_t + iy_t$. Show that simultaneous GDA satisfies

$$w_{t+1} = (1 + i\eta\sigma)w_t,$$

so the last iterate expands for every $\eta > 0$ unless $w_t = 0$. For extragradient (EG),

$$z_{t+1/2} = z_t - \eta F(z_t), \qquad z_{t+1} = z_t - \eta F(z_{t+1/2}),$$

show that

$$w_{t+1} = \left(1 + i\eta\sigma - \eta^2\sigma^2\right)w_t.$$

Determine exactly for which $\eta > 0$ EG converges linearly, and find the step size in this interval that gives the smallest contraction factor.

Solution: In the scalar bilinear case, $F(x, y) = (\sigma y, -\sigma x)$. Simultaneous GDA gives

$$x_{t+1} = x_t - \eta\sigma y_t, \qquad y_{t+1} = y_t + \eta\sigma x_t.$$

Thus, with $w_t = x_t + iy_t$,

$$\begin{aligned} w_{t+1} &= x_t - \eta\sigma y_t + i(y_t + \eta\sigma x_t) \\ &= (1 + i\eta\sigma)(x_t + iy_t) = (1 + i\eta\sigma)w_t. \end{aligned}$$

Therefore $|w_{t+1}| = \sqrt{1 + \eta^2\sigma^2}\,|w_t|$, which expands for every $\eta > 0$ unless $w_t = 0$.
For extragradient, the half-step is

$$w_{t+1/2} = (1 + i\eta\sigma)w_t.$$

The second step uses the vector field at the half-step but starts again from w_t, so

$$w_{t+1} = w_t + i\eta\sigma w_{t+1/2} = w_t + i\eta\sigma(1 + i\eta\sigma)w_t.$$

Hence

$$w_{t+1} = (1 + i\eta\sigma - \eta^2\sigma^2)w_t.$$

The squared contraction factor is

$$|1 + i\eta\sigma - \eta^2\sigma^2|^2 = (1 - \eta^2\sigma^2)^2 + \eta^2\sigma^2.$$

Let $a := \eta\sigma > 0$. Then this equals

$$(1 - a^2)^2 + a^2 = 1 - a^2 + a^4.$$

EG converges linearly exactly when

$$1 - a^2 + a^4 < 1 \iff 0 < a < 1.$$

Therefore $0 < \eta < 1/\sigma$. To minimize the squared contraction factor $1 - a^2 + a^4$, differentiate:

$$\frac{d}{da}(1 - a^2 + a^4) = -2a + 4a^3 = 2a(2a^2 - 1).$$

The minimizer in $(0, 1)$ is $a = 1/\sqrt{2}$, so $\eta^\star = \frac{1}{\sqrt{2}\sigma}$. The best squared contraction factor is

$$1 - \frac{1}{2} + \frac{1}{4} = \frac{3}{4},$$

and the best contraction factor in norm is $\sqrt{3}/2$.

(d) (4 points) For the same scalar bilinear problem, consider the implicit proximal-point step

$$z_{t+1} = z_t - \eta F(z_{t+1}).$$

Writing again $w_t = x_t + iy_t$, solve this two-dimensional linear system explicitly and show that

$$w_{t+1} = \frac{1 + i\eta\sigma}{1 + \eta^2\sigma^2}\, w_t.$$

What is the contraction factor?

Solution: For the proximal-point step, $z_{t+1} = z_t - \eta F(z_{t+1})$. In complex notation, subtracting ηF corresponds to adding $i\eta\sigma$ times the point at which the vector field is evaluated. Thus

$$w_{t+1} = w_t + i\eta\sigma w_{t+1} \iff (1 - i\eta\sigma)w_{t+1} = w_t.$$

Therefore

$$w_{t+1} = \frac{1}{1 - i\eta\sigma} w_t = \frac{1 + i\eta\sigma}{1 + \eta^2\sigma^2} w_t.$$

The contraction factor is

$$\left| \frac{1 + i\eta\sigma}{1 + \eta^2\sigma^2} \right| = \frac{\sqrt{1 + \eta^2\sigma^2}}{1 + \eta^2\sigma^2} = \frac{1}{\sqrt{1 + \eta^2\sigma^2}} < 1$$

for every $\eta > 0$.