SEEM5020 2026 Spring: Homework 8

This assignment is due at 1:30 pm, March 27, 2026. Please bring your solution and submit it at the beginning of the next class. Total weight of Homework 8: 5%.

- **Q1. [15 Points] Existence of an internal η-Net on the Unit Sphere.**

  Let $\mathcal{U}$ be an r-dimensional linear subspace of $\mathbb{R}^n$, and let $\mathcal{S} = \{\mathbf{u} \in \mathcal{U} : \|\mathbf{u}\|_2 = 1\}$ be the unit sphere of $\mathcal{U}$. An internal η-*net* of $\mathcal{S}$ is a finite set $\mathcal{N} \subseteq \mathcal{S}$ such that for every $\mathbf{u} \in \mathcal{S}$, there exists $\mathbf{v} \in \mathcal{N}$ with $\|\mathbf{u} - \mathbf{v}\|_2 \leq \eta$.

  For any $0 < \eta < 1$, prove that:

  (a) An internal η-net of $\mathcal{S}$ exists.

  (b) Any such η-net constructed by the greedy algorithm satisfies $|\mathcal{N}| \leq \left(1 + \frac{2}{\eta}\right)^r$.

  Hint: Start by constructing $\mathcal{N}$ greedily: pick an arbitrary point, then keep adding points that are far from all existing points, until no such point exists. For the size bound, consider packing balls of radius $\eta/2$ centered at points in $\mathcal{N}$ and compare volumes. Note: $\mathrm{Vol}_r(R)$ denotes the volume of an r-dimensional ball of radius R, and $\mathrm{Vol}_r(R) = C_r \cdot R^r$ where C_r is a constant depending only on r.

- **Q2. [15 Points] Low-Rank Approximation via JL Transformation.**

  Consider the low-rank approximation method via JL transformation (Lecture 7). We construct $\mathbf{B} = \mathbf{\Pi A}$ where $\mathbf{\Pi} \in \mathbb{R}^{c \times m}$ satisfies JLT$(\varepsilon/2, \delta, 2k)$, and compute the SVD $\mathbf{B} = \mathbf{U\Sigma V}^T$ with right singular vectors $\mathbf{V} = [\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_n]$. Define $\mathbf{V}_k = [\mathbf{b}_1, \ldots, \mathbf{b}_k]$ (the first k columns) and $\hat{\mathbf{A}} = \mathbf{A}\mathbf{V}_k\mathbf{V}_k^T$.

  Prove that:

  $$\|\mathbf{A} - \hat{\mathbf{A}}\|_F^2 \leq \|\mathbf{A} - \mathbf{A}_k\|_F^2 + 2\sqrt{k}\,\|\mathbf{A}^T\mathbf{A} - \mathbf{B}^T\mathbf{B}\|_F,$$

  where $\mathbf{A}_k$ is the best rank-k approximation of $\mathbf{A}$.

  Hint: The proof follows the same structure as Theorem 3 in Lecture 7 (for column-sampling), but adapted to the JL setting. The key steps are: (1) show $\|\mathbf{A} - \hat{\mathbf{A}}\|_F^2 = \|\mathbf{A}\|_F^2 - \|\mathbf{A}\mathbf{V}_k\|_F^2$ using the orthonormality of $\mathbf{V}_k$; (2) bound $|\|\mathbf{A}\mathbf{V}_k\|_F^2 - \sum_{i=1}^k \sigma_i^2(\mathbf{B})|$; (3) bound $|\sum_{i=1}^k \sigma_i^2(\mathbf{B}) - \sum_{i=1}^k \sigma_i^2(\mathbf{A})|$. For step (2), if $\mathrm{Tr}(\mathbf{X}^T\mathbf{X})$ does not work, try $\mathrm{Tr}(\mathbf{X}\mathbf{X}^T)$.

- **Q3. [20 Points] Yet Another η-net.**

  In Lecture 7, we showed that an internal η-net of the unit sphere can be constructed greedily. Here we explore an alternative construction using a uniform grid $\mathcal{G}$. Note that the points in $\mathcal{G}$ do not necessarily lie on $\mathcal{S}$. We call such a set an external η-net.

  Let $\mathcal{U}$ be a d-dimensional linear subspace of $\mathbb{R}^n$. Without loss of generality, assume $\mathcal{U}$ is spanned by the first d standard basis vectors. Define the grid $\mathcal{G}$ over $[-1, 1]^d$ with spacing ε/d in each coordinate:

  $$\mathcal{G} = \left\{ \left(-1 + \frac{j_1\varepsilon}{d}, \ldots, -1 + \frac{j_d\varepsilon}{d}, 0, \ldots, 0 \right) \in \mathbb{R}^n \;:\; j_i \in \left\{ 0, 1, \ldots, \frac{2d}{\varepsilon} \right\} \right\}.$$

  Assume $\varepsilon < 1/3$ and $\frac{1}{\varepsilon}$ is a positive integer.

  (a) **[5 Points]** An external η-net of $\mathcal{S}$ is defined the same way as an internal one, except that the points in the net need not lie on $\mathcal{S}$.

  (i) Show that $|\mathcal{G}| = \left(\frac{2d}{\varepsilon} + 1\right)^d$.

(ii) Prove that $\mathcal{G}$ is an external $\frac{\varepsilon}{2\sqrt{d}}$-net of $\mathcal{S} = \{\mathbf{u} \in \mathcal{U} : \|\mathbf{u}\|_2 = 1\}$. That is, for every $\mathbf{z} \in \mathcal{S}$, there exists $\mathbf{g} \in \mathcal{G}$ with $\|\mathbf{z} - \mathbf{g}\|_2 \leq \frac{\varepsilon}{2\sqrt{d}}$.

Hint: Since $\|\mathbf{z}\|_2 = 1$, each coordinate satisfies $|z_i| \leq 1$. How close is the nearest grid point in the L^∞ norm? Then convert to L^2.

(b) **[7 Points]** Prove the following structural properties of the grid approximation. For any $\mathbf{z} \in \mathcal{S}$:

(i) There exists a grid vertex $\mathbf{g}^-$ obtained by rounding each coordinate of $\mathbf{z}$ *toward zero*, satisfying $\|\mathbf{g}^-\|_2 \leq 1$.

(ii) There exists a grid vertex $\mathbf{g}^+$ obtained by rounding each coordinate of $\mathbf{z}$ *away from zero*, satisfying $\|\mathbf{g}^+\|_2 \geq 1$.

(iii) The displacement $\Delta = \mathbf{z} - \mathbf{g}$ (for either choice $\mathbf{g} = \mathbf{g}^-$ or $\mathbf{g} = \mathbf{g}^+$) can be decomposed as

$$\Delta = \sum_{i=1}^{d} \alpha_i \cdot \frac{\varepsilon}{d}\, s_i\, \mathbf{e}_i,$$

where $\alpha_i \in [0, 1]$, $s_i \in \{-1, +1\}$, and each vector $\frac{\varepsilon}{d}\, s_i\, \mathbf{e}_i$ is the difference between two adjacent grid points.

Hint: For (i) and (ii), argue coordinate-wise: since 0 is a grid point and $z_i \in [-1, 1]$, what happens when you round toward/away from zero?

(c) **[8 Points]** Let $\mathbf{\Pi} \in \mathbb{R}^{k \times n}$ be a matrix satisfying $\mathrm{JLT}\Big(\varepsilon, \delta, \big(\frac{2d}{\varepsilon} + 1\big)^d\Big)$. Using the results of parts (a) and (b), prove that for any $\mathbf{z} \in \mathcal{S}$:

$$(1 - 3\varepsilon)\|\mathbf{z}\|_2 \leq \|\mathbf{\Pi}\mathbf{z}\|_2 \leq (1 + 3\varepsilon)\|\mathbf{z}\|_2.$$

Hint: Use JLT to bound $\|\mathbf{\Pi}\mathbf{g}\|_2$ for grid points and $\|\mathbf{\Pi}(\frac{\varepsilon}{d} s_i \mathbf{e}_i)\|_2$ for grid steps. Bound $\|\mathbf{\Pi}\Delta\|_2$ via the triangle inequality. Then combine with $\mathbf{g}^-$ for the upper bound and $\mathbf{g}^+$ for the lower bound.