

Due: Wednesday, May 8 at 11:59 pm

Deliverables:

1. Submit a PDF of your homework, **with an appendix listing all your code**, to the Gradescope assignment entitled “Homework 7 Write-Up”. In addition, please include, as your solutions to each coding problem, the specific subset of code relevant to that part of the problem. You may typeset your homework in LaTeX or Word (submit PDF format, **not** .doc/.docx format) or submit neatly handwritten and scanned solutions. **Please start each question on a new page**. If there are graphs, include those graphs in the correct sections. **Do not** put them in an appendix. We need each solution to be self-contained on pages of its own.
 - In your write-up, please state with whom you worked on the homework.
 - In your write-up, please copy the following statement and sign your signature next to it. (Mac Preview and FoxIt PDF Reader, among others, have tools to let you sign a PDF file.) We want to make it *extra* clear so that no one inadvertently cheats.

“I certify that all solutions are entirely in my own words and that I have not looked at another student’s solutions. I have given credit to all external sources I consulted.”
2. Submit all the code needed to reproduce your results to the Gradescope assignment entitled “Homework 7 Code”. Yes, you must submit your code twice: in your PDF write-up following the directions as described above so the readers can easily read it, and once in compilable/interpretable form so the readers can easily run it. Do **NOT** include any data files we provided. Please include a short file named README listing your name, student ID, and instructions on how to reproduce your results. Please take care that your code doesn’t take up inordinate amounts of time or memory. If your code cannot be executed, your solution cannot be verified.

1 Regularized and Kernel k-Means

Recall that in k -means clustering we attempt to minimize the objective

$$\min_{C_1, C_2, \dots, C_k} \sum_{i=1}^k \sum_{x_j \in C_i} \|x_j - \mu_i\|_2^2, \text{ where}$$
$$\mu_i = \operatorname{argmin}_{\mu_i \in \mathbb{R}^d} \sum_{x_j \in C_i} \|x_j - \mu_i\|_2^2 = \frac{1}{|C_i|} \sum_{x_j \in C_i} x_j, \quad i = 1, 2, \dots, k.$$

The samples are $\{x_1, \dots, x_n\}$, where $x_j \in \mathbb{R}^d$. C_i is the set of sample points assigned to cluster i and $|C_i|$ is its cardinality. Each sample point is assigned to exactly one cluster.

- (a) What is the minimum value of the objective when $k = n$ (the number of clusters equals the number of sample points)?
- (b) (Regularized k -means) Suppose we add a regularization term to the above objective. The objective is now

$$\sum_{i=1}^k \left(\lambda \|\mu_i\|_2^2 + \sum_{x_j \in C_i} \|x_j - \mu_i\|_2^2 \right).$$

Show that the optimum of

$$\min_{\mu_i \in \mathbb{R}^d} \lambda \|\mu_i\|_2^2 + \sum_{x_j \in C_i} \|x_j - \mu_i\|_2^2$$

is obtained at $\mu_i = \frac{1}{|C_i| + \lambda} \sum_{x_j \in C_i} x_j$.

- (c) Here is an example where we would want to regularize clusters. Suppose there are n students who live in a \mathbb{R}^2 Euclidean world and who wish to share rides efficiently to Berkeley for their final exam in CS189. The university permits k vehicles which may be used for shuttling students to the exam location. The students need to figure out k good locations to meet up. The students will then walk to the closest meet up point and then the shuttles will deliver them to the exam location. Let x_j be the location of student j , and let the exam location be at $(0, 0)$. Assume that we can drive as the crow flies, i.e., by taking the shortest path between two points. Write down an appropriate objective function to minimize the total distance that the students and vehicles need to travel. Hint: your result should be similar to the regularized k -means objective.
- (d) (Kernel k -means) Suppose we have a dataset $\{x_i\}_{i=1}^n, x_i \in \mathbb{R}^\ell$ that we want to split into k clusters, i.e., finding the best k -means clustering *without the regularization*. Furthermore, suppose we know *a priori* that this data is best clustered in an impractically high-dimensional feature space \mathbb{R}^m with an appropriate metric. Fortunately, instead of having to deal with the (implicit) feature map $\Phi : \mathbb{R}^\ell \rightarrow \mathbb{R}^m$ and (implicit) distance metric¹, we have a kernel function $\kappa(x_1, x_2) = \Phi(x_1) \cdot \Phi(x_2)$ that we can compute easily on the raw samples. How should we perform the kernelized counterpart of k -means clustering?

Derive the underlined portion of this algorithm, and show your work in deriving it. The main issue is that although we define the means μ_i in the usual way, we can't ever compute Φ explicitly because it's

¹Just as how the interpretation of kernels in kernelized ridge regression involves an implicit prior/regularizer as well as an implicit feature space, we can think of kernels as generally inducing an implicit distance metric as well. Think of how you would represent the squared distance between two points in terms of pairwise inner products and operations on them.

way too big. Therefore, in the step where we determine which cluster each sample point is assigned to, we must use the kernel function κ to obtain the right result. (Review the lecture on kernels if you don't remember how that's done.)

Algorithm 1: Kernel k -means

Require: Data matrix $X \in \mathbb{R}^{n \times d}$; Number of clusters K ; kernel function $\kappa(x_1, x_2)$

Ensure: Cluster class $\text{class}(j)$ for each sample x_j .

function KERNEL-K-MEANS(X, c)

 Randomly initialize $\text{class}(j)$ to be an integer in $1, 2, \dots, K$ for each x_j .

while *not converged* **do**

for $i \leftarrow 1$ **to** K **do**

 Set $S_i = \{j \in \{1, 2, \dots, n\} : \text{class}(j) = i\}$.

for $j \leftarrow 1$ **to** n **do**

 Set $\text{class}(j) = \underset{k}{\operatorname{argmin}} \frac{\sum_{i=1}^c \kappa(x_j, x_{S_i})}{|S_i|}$

 Return S_i for $i = 1, 2, \dots, c$.

end function

- (e) The expression you derived may have unnecessary terms or redundant kernel computations. Explain how to eliminate them; that is, how to perform the computation quickly without doing irrelevant computations or redoing computations already done.

2 Low-Rank Approximation

Low-rank approximation tries to find an approximation to a given matrix, where the approximation matrix has a lower rank compared to the original matrix. This is useful for mathematical modeling and data compression. Mathematically, given a matrix M , we try to find \hat{M} in the following optimization problem,

$$\underset{\hat{M}}{\operatorname{argmin}} \|M - \hat{M}\|_F \quad \text{subject to} \quad \operatorname{rank}(\hat{M}) \leq k \quad (1)$$

where $\|A\|_F = \sqrt{\sum_i \sum_j a_{ij}^2}$ is the Frobenius norm, i.e., the sum of squares of all entries in the matrix, followed by a square root.

This problem can be solved using Singular Value Decomposition (SVD). Specifically, let $M = U\Sigma V^T$, where $\Sigma = \operatorname{diag}(\sigma_1, \dots, \sigma_n)$. Then a rank- k approximation of M can be written as $\hat{M} = U\hat{\Sigma}V^T$, where $\hat{\Sigma} = \operatorname{diag}(\sigma_1, \dots, \sigma_k, 0, \dots, 0)$. In this problem, we aim to perform this approximation method on gray-scale images, which can be thought of as a 2D matrix.

- (a) Using the image `low-rank_data/face.jpg`, perform a rank-5, rank-20, and rank-100 approximation on the image. Show both the original image as well as the low-rank images you obtain in your report.
- (b) Now perform the same rank-5, rank-20, and rank-100 approximation on `low-rank_data/sky.jpg`. Show both the original image as well as the low-rank images you obtain in your report.
- (c) In one plot, plot the Mean Squared Error (MSE) between the rank- k approximation and the original image for both `low-rank_data/face.jpg` and `low-rank_data/sky.jpg`, for k ranging from 1 to 100. Be sure to label each curve in the plot. The MSE between two images $I, J \in \mathbb{R}^{w \times h}$ is

$$\operatorname{MSE}(I, J) = \sum_{i,j} (I_{i,j} - J_{i,j})^2. \quad (2)$$

- (d) Find the lowest-rank approximation for which you begin to have a hard time differentiating the original and the approximated images. Compare your results for the face and the sky image. What are the possible reasons for the difference?