

CS-233 Theoretical Exercise 13

May 2024

1 K-means Clustering

Question 1: Are the following statements true?

1. The standard way of initializing K-means is by setting $\mu_1 = \dots = \mu_K$ to be equal to a vector of zeros.
2. For some datasets, the "right" or "correct" value of K (the number of clusters) can be ambiguous, and hard even for a human expert looking carefully at the data to determine.
3. One way to reduce the problem of K-means getting stuck in a bad local optimum is to try using multiple random initializations.
4. Since K-means is an unsupervised learning algorithm, it cannot overfit the data, and thus it is always better to have as large a number of clusters as is computationally feasible.

2 Hierarchical Clustering

Very often, clustering algorithms suffer a glaring problem: The number of clusters needs to be fixed before training. One would then have to rerun the entire training process if they were to use a different number of clusters. This imposes high computational costs for application scenarios where the number of clusters is unknown or where one would like to compute the clustering results for different numbers of clusters. In this exercise, we explore a different type of clustering algorithm that is free of such trouble. It is known as hierarchical clustering. More specifically, we are going to study agglomerative clustering, a subtype of hierarchical clustering¹.

Question 2: Let us walk through the process of a hierarchical clustering algorithm with a toy example. Suppose you are given four points A, B, C, and D. Their pairwise distances are denoted in the table on the right.

We start by treating each point as a cluster. If we use a set to denote each cluster, the clusters we have at initialization of this algorithm are: $\{A\}$, $\{B\}$, $\{C\}$, $\{D\}$.

Question 2.1: If we would like to cluster these four points into three clusters, we can merge together the two clusters that have the closest distance. What would be the three clusters after this merging step?

Question 2.2: Now, we are going to merge together two of the three clusters that you obtained in the last question so that we can cluster the four points into two clusters. Again, we are going to merge the two clusters that are closest to each other. We define the distance between two clusters to be *the minimum of the distances between two elements that belong to the two clusters*. What would be the three clusters after this merging step?

	A	B	C	D
A	0	1	4	5
B		0	2	6
C			0	3
D				0

¹The other subtype of hierarchical clustering is divisive clustering.

Question 2.3: Notice that agglomerative clustering is often sensitive to how we compute the distance between clusters. If we define the distance between clusters to be *the maximum of the distances between two elements that belong to the two clusters*, what would be the two clusters we will then get?

With this example, we see how agglomerative clustering effectively clusters data points with a flexible number of clusters. Admittedly, one always needs to compute more fine-grained clusterings before getting coarse-grained ones. Nevertheless, fine-grained clusterings can be of use later. Every step in this algorithm yields a proper clustering.

Question 3: Do the following statements accurately contrast K-means clustering with agglomerative clustering?

1. Agglomerative clustering generally requires more computational resources than K-means clustering.
2. In agglomerative clustering, the "rich get richer" phenomenon can occur, where larger clusters are more likely to merge with other clusters as the process iterates. In K-means, it doesn't.
3. K-means clustering is deterministic, while agglomerative clustering produces different results with each run due to its random initialization process.

3 K-means Clustering Properties

In the K-means clustering algorithm, you are given a set of N points $\mathbf{x}_i \in \mathbb{R}^D$, $i \in \{1, \dots, N\}$ and you want to find the centers of K clusters $\mu = (\mu_1, \dots, \mu_K)$ by minimizing the average distance from the points to the closest cluster center. Formally, you want to minimize the following loss function

$$L(\mu) = \sum_{i=1}^N \min_{j \in \{1, \dots, K\}} \|\mathbf{x}_i - \mu_j\|^2. \quad (1)$$

The K-means algorithm iterates between an assignment step and an update step to solve Eq. 1. In the assignment step, it assigns each point \mathbf{x}_i to the z_i^* -th cluster with

$$z_i^* = \arg \min_{j \in \{1, \dots, K\}} \|\mathbf{x}_i - \mu_j\|^2. \quad (2)$$

In the update step, it updates the cluster centers as $\mu_j^* = \frac{1}{|\{i:z_i=j\}|} \sum_{i:z_i=j} \mathbf{x}_i$, which is the solution of

$$(\mu_1^*, \dots, \mu_K^*) = \arg \min_{(\mu_1, \dots, \mu_K)} \sum_{j=1}^K \left(\sum_{i:z_i=j} \|\mathbf{x}_i - \mu_j\|^2 \right). \quad (3)$$

Question 4: Show that K-means is guaranteed to converge to a local optimum. **Hint:** You need to prove that the loss function is guaranteed to decrease monotonically in each iteration until convergence.

Question 5: What is the time complexity of the K-means clustering algorithm?

Question 6: While the Euclidean distance $\ell_2(\mathbf{x}_i, \mathbf{x}_j)$ is commonly employed in the K-means clustering algorithm and machine learning in general, alternative distance metrics are also available. The $\ell_0(\mathbf{x}_i, \mathbf{x}_j)$ distance computes the count of differing elements between two vectors \mathbf{x}_i and \mathbf{x}_j . The $\ell_1(\mathbf{x}_i, \mathbf{x}_j)$ distance, known as the Manhattan distance, sums the absolute differences between corresponding elements of \mathbf{x}_i and \mathbf{x}_j . The $\ell_\infty(\mathbf{x}_i, \mathbf{x}_j)$ distance represents the maximum absolute difference between corresponding elements of \mathbf{x}_i and \mathbf{x}_j . More generally, the $\ell_p(\mathbf{x}_i, \mathbf{x}_j)$ distance can be defined for all $1 \leq p < \infty$ as

$$\ell_p(\mathbf{x}_i, \mathbf{x}_j) = \|\mathbf{x}_i - \mathbf{x}_j\|_p = \left(\sum_{k=1}^D |\mathbf{x}_i^{(k)} - \mathbf{x}_j^{(k)}|^p \right)^{1/p}.$$

How does $\ell_p(\mathbf{x}_i, \mathbf{x}_j)$ evolve as p increases? Provide a scenario where a low value of p would be advantageous and another where a high value of p would be preferable.