

1. Local Optimum Problem in K-Means

Now that you have learned the k-means clustering algorithm, let's take a look at how the initial choice of cluster locations can affect the result of the algorithm. The k-means algorithm is an optimization problem that iteratively makes steps towards the "best" placement for the clusters. However, the algorithm can get stuck in a placement for the clusters that appears "best" locally, but a better solution could exist.




Figure 1: An arbitrary function with two local minimum values

Let's start by considering a simplified optimization problem in which we want to find the value x that gives us the smallest value for a function $f(x)$. Let $f(x)$ be the function shown in Figure 1. Assuming that we don't know anything about the behavior of $f(x)$, we need to design an algorithm that can step through the function to find the smallest value. Let's suppose that our algorithm works as follows:

- Starting at some point x , find the slope of $f(x)$ at the point x
 - Determine which direction (left or right) the function slopes down
 - Move a little bit in that direction and update x to be the new position
 - Repeat this process until we can't find a direction that $f(x)$ decreases
- (a) Using the algorithm described above, which minimum point will it find when we start at points A, B, and C, and which starting point gives us the smallest value for $f(x)$?
- (b) Now, let's extend what we learned in this 1-dimensional example to the k-means clustering algorithm. First, we need to identify a function that takes cluster locations as input and outputs a single value that tells us how well the clusters fit the data. What might this function look like?
- (c) Recall that the k-means algorithm has two parts:
- Assign each point to the nearest cluster
 - Update the location of the clusters by moving them to the centers of their assigned points

Show that both steps decreases (or at the very least does not increase) the value of the function we defined in the previous part.

- (d) We have now just shown that each step of the k-means algorithm is similar to our optimization algorithm from the first part. At each step, we go in a direction that decreases the value of our function. We observed that always going in a direction that decreases the function can cause us to get stuck in a local optimum. This also holds true for the k-means algorithm. Consider the set of four points shown in Figure 2. We want to fit two clusters to this data, and we will see how our choice for the starting location of the two clusters affects the output. Let's suppose that for the two clusters, we assign the initial locations to points A and B, and for our second example, let's assign the locations to points A and C. After running k-means for both starting locations, where do the clusters end up, and which starting configuration gives us better cluster positions?




Figure 2: Set of 4 data points that we will use k-means to find the two clusters

Contributors:

- Stephen Bailey.
- Baruch Sterin.