1.4 k-means Clustering

k-means clustering aims to partition an unlabeled dataset into clusters. The algorithm is easy:

- Randomly initialize k cluster centroids $\mu_1, \mu_2, ..., \mu_k$.

- Repeat:

    - **Cluster assignment.** Assign each data point to its closest centroid.

    - **Move centroid.** Move each centroid to the average of the points assigned to it.

Essentially, it finds $\mu_1, \mu_2, ..., \mu_k$ that minimizes

$$J(\mu_1, \mu_2, ..., \mu_k) = \frac{1}{m} \sum_{i=1}^{m} ||x^{(i)} - \mu_{c^{(i)}}||^2 \ . \tag{11}$$

Typically, you can initialize the centroids as some subset of your training data. Because you might find a local optima with some unlucky random initialization, you can run k-means 50-1000 times and pick the one with the lowest cost. To choose the number of clusters k, you can do it manually using the "elbow method" (lol). More importantly, you should consider the downstream purpose that you're using the clusters for.