April 9, 2025, Lecture Note: Fairness in Clustering

Nishanth Nagendran, Aryan Esmailpour

CS 516: Responsible Data Science and Algorithmic Fairness; Spring 2025
Abolfazl Asudeh; www.cs.uic.edu/∼ asudeh/teaching/archive/cs516spring25/

1 Introduction to Clustering

Clustering is a fundamental concept with wide-ranging applications across machine learning, data mining, and other fields such as sociology and biology. At its core, clustering involves grouping a set of elements based on some notion of similarity. However, its definition and implementation vary depending on the context.

Definition 1 (Clustering)**.** Given a universe of elements $U = \{e_1, e_2, \ldots, e_n\}$, a *clustering* is a partitioning of U into k subsets, $C = \{C_1, C_2, \ldots, C_k\}$, such that $\bigcup_{j=1}^{k} C_j = U$ and $C_i \cap C_j = \emptyset$ for all $i \neq j$.

This definition assumes a *hard partition*, where each element belongs to exactly one cluster. However, some applications use a *soft (fuzzy) partition*, where elements have a probability of belonging to each cluster. A well-known example is the Gaussian Mixture Model (GMM), which assigns membership probabilities based on a mixture of Gaussian distributions. In this scribe, we focus on hard partitioning, particularly center-based clustering.

1.1 Types of Hard Partitioning

Hard clustering can be categorized into three main types:

1. **Center-Based Clustering**: k cluster centers are selected, and each element is assigned to the nearest center based on a distance metric, typically Euclidean distance.
2. **Hierarchical Clustering**: U is partitioned using a tree-like structure, either bottom-up (agglomerative) or top-down (divisive).
3. **Density-Based Clustering**: Clusters are formed by identifying connected regions of high density, e.g., DBSCAN.

Here, we narrow our focus to center-based clustering, specifically the k-means algorithm, and explore fairness within this framework.

1.2 Center-Based Clustering

In center-based clustering, we select k centers $c_1, c_2, \ldots, c_k \in \mathbb{R}^d$ and assign each element $e_i \in U$ to the cluster with the nearest center. Let $\delta(e_i, c_j) = \|e_i - c_j\|_2$ denote the Euclidean distance (i.e., $p = 2$ in the L_p-norm). The objective is to minimize an aggregate function over these distances, such as the sum in k-means:

$$\text{Objective: } \min_{C_1,\ldots,C_k} \sum_{j=1}^{k} \sum_{e_i \in C_j} \delta(e_i, c_j)^2.$$

Variants like k-center (minimize the maximum distance) or k-median (minimize the sum of distances) use different aggregate functions.

2 Fairness in Center-Based Clustering

Fairness in clustering addresses biases or disparities in how elements are grouped or how resources are allocated based on cluster assignments. We consider *group fairness*, where each element $e_i \in U$ belongs to a demographic group $g_i \in G = \{g_1, g_2, \ldots, g_m\}$, and fairness ensures (almost) equitable treatment across groups.

2.1 Demographic Parity (Balance)

Consider a scenario where cluster assignments confer advantages (e.g., promotions) or disadvantages, and group membership is known but outcome labels are not. A fair clustering might aim for *demographic parity*, ensuring each cluster's group proportions mirror the overall population.

Definition 2 (Balance)**.** A clustering $C = \{C_1, \ldots, C_k\}$ satisfies *balance* if, for all groups $g_i \in G$ and clusters $C_j \in C$,

$$\frac{|g_i|}{|U|} = \frac{|g_i \cap C_j|}{|C_j|}.$$

This strict condition is often impractical, so a relaxed version allows an additive error $\alpha \geq 0$:

$$\left| \frac{|g_i|}{|U|} - \frac{|g_i \cap C_j|}{|C_j|} \right| \leq \alpha.$$

Example 1. Suppose U represents employees, clusters determine promotion eligibility, and G denotes demographic groups. Balance ensures promotion opportunities are proportionally distributed across different demographic groups.

2.2 Fair Resource Allocation

In another scenario, cluster centers represent facilities (e.g., MRI devices), and each element accesses only its cluster's center. A fairness issue arises when dense regions receive more centers, forcing sparse regions to travel farther. Additionally, overloading a center (e.g., too many patients per MRI) degrades service quality.

Definition 3 (Socially Fair Clustering)**.** A clustering is *socially fair* if the average distance from elements in each group to their assigned center is equal across groups. Formally, for each group $g_i \in G$ and centers $c_1, \ldots, c_k$,

$$\frac{1}{|g_i|} \sum_{e \in g_i} \delta(e, c_{\text{assign}(e)}) = \text{constant for all } g_i,$$

where $c_{\text{assign}(e)}$ is the center assigned to e.

2.3 Fair Center Selection

Suppose centers represent selected entities, each associated with a demographic group. Fairness in center selection requires the selected center's group distribution to reflect the overall population.

Definition 4 (Fair Center Selection)**.** A set of k centers $\{c_1, \ldots, c_k\}$ satisfies *fair center selection* if, for each group $g_i \in G$, the proportion of centers from g_i approximates its population proportion:

$$\frac{|\{c_j : c_j \in g_i\}|}{k} \approx \frac{|g_i|}{|U|}.$$

Alternative fairness objectives might use the maximum (instead of average) distance or group representation as the aggregate function.

3 Socially Fair k-Means Clustering

The standard k-means algorithm, known as Lloyd's heuristic, is a commonly used approach for solving the k-means clustering problem. It proceeds as follows:

1. Select k centers $c_1, \ldots, c_k$ randomly from U.

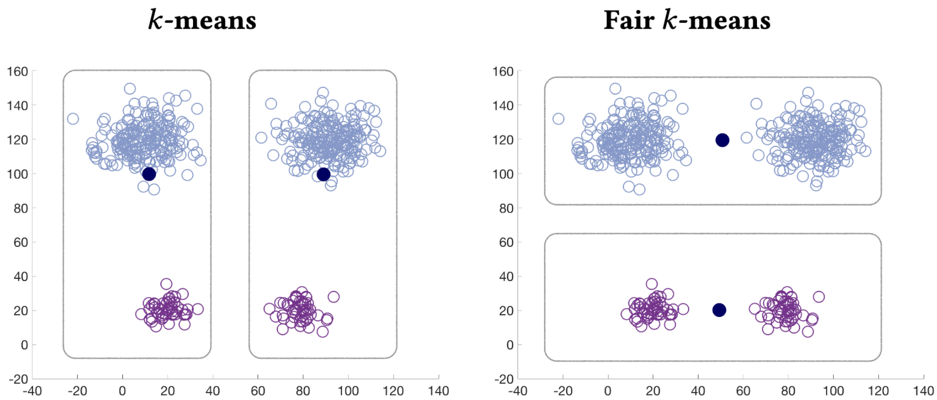

Figure 1: Two demographic groups are shown with blue and purple. The 2-means objective minimizing the average clustering cost prefers the clustering (and centers) shown in the left figure. This clustering incurs a much higher average clustering cost for purple than for blue. The clustering in the right figure has a more equitable clustering cost for the two groups [2]. The figure is from [2].

2. Repeat until convergence:

   (a) Assign each $e_i \in U$ to the nearest center: $C_j = \{e_i : \delta(e_i, c_j) \leq \delta(e_i, c_{j'}) \text{ for all } j'\}$.

   (b) Update each center c_j to the centroid of C_j: $c_j = \frac{1}{|C_j|} \sum_{e_i \in C_j} e_i$.

3.1 Socially Fair Variant

Now, we focus on discussing the socially fair k-means clustering, explaining the approach taken in [2]. To incorporate social fairness, we modify k-means to minimize the *maximum average distance* of any group to its center, rather than the overall average distance. Although this definition is not exactly equivalent to what we defined as "socially fair" in Definition 3, it is intuitively acceptable that in many cases, minimizing the maximum average distance is similar to making the average distances equal for all the groups as it is desired according to Definition 3. An example is shown in Figure 1. Note that this intuition is only true when there is no restriction on the cluster centers, and any point in the plane (continuous) can be chosen as a cluster center. Therefore, we focus on the problem defined below. For each cluster C_j and group g_i, define the group-specific average distance:

$$d(g_i, c_j) = \frac{1}{|g_i \cap C_j|} \sum_{e \in g_i \cap C_j} \delta(e, c_j).$$

The objective becomes:

$$\min_{c_1, \ldots, c_k} \max_{g_i \in G, j \in \{1, \ldots, k\}} d(g_i, c_j).$$

To solve this problem, we can take an approach similar to the Lloyd's algorithm, referred to as *Fair Lloyd Algorithm*. The only essential difference in the Fair version of the algorithm is in *step (b)*, where instead of updating each c_j to be the centroid of the points in C_j, it is updated so that the average distance of points of every color to the new c_j, is equal.

1. Select k centers $c_1, \ldots, c_k$ randomly from U.

2. Repeat until convergence:

   (a) Assign each $e_i \in U$ to the nearest center: $C_j = \{e_i : \delta(e_i, c_j) \leq \delta(e_i, c_{j'}) \text{ for all } j'\}$.

   (b) Update each center c_j such that $\frac{1}{|g_{i_1} \cap C_j|} \sum_{e \in g_{i_1} \cap C_j} \delta(e, c_j) = \frac{1}{|g_{i_2} \cap C_j|} \sum_{e \in g_{i_2} \cap C_j} \delta(e, c_j)$, for every pair of groups $g_{i_1}, g_{i_2} \in G$, where $g_{i_1} \cap C_j \neq \emptyset$, and $g_{i_2} \cap C_j \neq \emptyset$.

The point c_j in *step (b)* almost (under some mild assumptions) always exists [2]. The main problem is to find such a point efficiently. Interestingly, when there are only two demographic groups ($|G| = 2$), this could be done efficiently [2] using a simple approach. The high-level idea is as follows. Assume that there are two demographic groups g_1 and g_2. For each C_j, let p_1 (resp. p_2) be the centroid of the points in $g_1 \cap C_j$ (resp. $g_2 \cap C_j$). It can be shown that the point c_j in *step (b)* of the Fair-Lloyd's algorithm always lies on the line connecting p_1 and p_2 [2]. The points p_1 and p_2 can be found in $O(n)$ by traversing all the data points. To find the point c_j on this line, a binary search can be performed.

For the case where there are more than two groups ($|G| > 2$), a more complicated approach is presented in [2].

4 Conclusion

Fairness in clustering extends traditional objectives to ensure equitable outcomes across demographic groups. From demographic parity to socially fair resource allocation and center selection, these notions address real-world biases in applications like facility placement and decision-making. In this note, we reviewed some of the definitions used in the literature for introducing fairness to some of the clustering problems. Finally, we discussed the algorithm of [2] for solving the socially fair variant of the k-means clustering. An interested reader is referred to [1] for a survey on fairness in clustering problems.

References

[1] Anshuman Chhabra, Karina Masalkovaitė, and Prasant Mohapatra. An overview of fairness in clustering. *IEEE Access*, 9:130698–130720, 2021.

[2] Mehrdad Ghadiri, Samira Samadi, and Santosh Vempala. Socially fair k-means clustering. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 438–448, New York, NY, USA, 2021. Association for Computing Machinery.