



Knowledge Discovery in Databases with Exercises Summer Semester 2025

Submission 3: Clustering

About this Assignment

In this assignment, your task is to implement the algorithms for **K-means** and **DBSCAN**. For this purpose, you have access to a basic code skeleton, some helper classes, and several test cases.

Key Data

- **Max. Group Size:** 3
- **Max. Points:** 30
- **Estimated Workload:** 3 - 4 hours

How to Work on the Assignment

To start working on the assignment, you'll need to accept the assignment via GitHub Classroom by clicking the provided link. This will set up a new GitHub repository for your group, packed with all the necessary files for the assignment. If you're joining an existing group, it'll add you to that group's repository.¹

Once that's done, you have two main options for working on your assignment. You can clone the repository² to your local machine by navigating to **Code** → **Local**, which allows you to work directly from your computer. Alternatively, you might prefer using GitHub Codespaces by selecting **Code** → **Codespaces** for a virtual online environment, complete with the ability to run Python through the **Terminal** provided.

Whichever method you choose, it's crucial to commit and push your changes back to the repository to submit your solution². After your submission, GitHub Actions takes over to automatically grade your solution and provide feedback. You'll find this feedback in the **Actions** tab of your repository. If you didn't receive full points, you can improve your solution and push the changes back to the repository to trigger a reevaluation.

¹Each student must join individually. You can join groups while accepting an assignment.

²If you're unfamiliar with Git or GitHub, check out this helpful guide: <https://github.com/git-guides/>

How to Prepare the Transfer the Points to StudOn

In addition to joining the GitHub Classroom, you also need to register your GitHub username on StudOn. This is necessary to transfer the points you've earned on GitHub to StudOn. To do this, enter your GitHub username in **Submission 3 - GitHub Username**. Make sure to enter your username correctly, as otherwise, the points cannot be transferred.

After submission deadline, the points you've earned on GitHub will be transferred to StudOn. This process is not immediate and may take a few days. If you have any questions or issues, please contact us via the StudOn forum.

Restrictions

Within the scope of your implementation, you are not permitted to modify the helper classes, the test cases, or the provided GitHub Actions.

This will be checked on a random basis, and any attempt to do so will result in zero points for the involved group, similar to the consequences of plagiarism.

Task 1: K-means

(14 Points)

K-means is a simple and widely used clustering algorithm. It partitions a dataset into k clusters by iteratively assigning each data point to the cluster with the nearest centroid and updating the centroids based on the mean of the data points in the cluster.

Task 1.1

The first step of the K-means algorithm is to distribute the data points to the k partitions. This partition can be done randomly or by using a more sophisticated method. All partitions have to be non-empty after the initialization and each data point has to be assigned to exactly one partition.

Open `kmeans.py` and implement `_initialize_partitions`, which initializes the partitions for the K-means algorithm:

```
1 def _initialize_partitions(self, points: List[Point]):
2     """
3     Initializes the partitions (self.partitions) by assigning each point
4     to a cluster/partition.
5     All clusters/partitions are non-empty after this method is called.
6
7     Parameters:
8     points (List[Point]): The points to cluster
9     """
10    # Check if the number of points is greater or equal to the number of clusters
11    if len(points) < self.k:
12        raise ValueError(
13            "Number of points has to be greater or equal to the number of clusters."
14        )
15
16    # TODO
```

The method expects a list of `Points` and does not return anything. The method should put each point into one of the partitions available at `self.partitions`. The number of partitions is given by the variable `self.k`.

You can test whether your implementation is correct by executing the following command in the console:

```
1 pytest tests/kmeans/test_initialize_partitions.py
```

Task 1.2

A second important part of the K-means algorithm is to update the centroids of the partitions. The centroid of a partition is the mean of all points in the partition.

Open `kmeans.py` and implement `_update_centroids`, which updates the centroids of the partitions:

```

1 def _update_centroids(self):
2     """
3     Updates the centroids of the partitions and writes the new centroids
4     into self.centroids.
5     """
6     # Make sure that all partitions are non-empty
7     for partition in self.partitions:
8         if len(partition) == 0:
9             # Throw an error if a partition is empty
10            raise ValueError(
11                "All partitions have to be non-empty before updating the centroids."
12            )
13
14     # TODO

```

The method does not expect any parameters and does not return anything. The method should update the `self.centroids` based on the mean of all points in the partition. The i -th centroid in `self.centroids` should refer to the i -th partition in `self.partitions`.

You can test whether your implementation is correct by executing the following command in the console:

```

1 pytest tests/kmeans/test_update_centroids.py

```

Task 1.3

The last step of the K-means algorithm is to assign each data point to the cluster with the nearest centroid.

Open `kmeans.py` and implement `_reassign_points(self)`, which assigns each data point to the partition with the nearest centroid:

```

1 def _reassign_points(self) -> bool:
2     """
3     Reassigns each point to the partition with the closest centroid.
4
5     Ensures that each partition is non-empty after reassigning the points,
6     by randomly reassigning a single point from a random non-empty partition
7     with more than one element into each empty partition.
8
9     This is necessary to avoid empty partitions, which would lead to k-means
10    not producing k clusters, but k - n clusters (n is the number of empty partitions).
11
12    Returns:
13    bool: True if the reassignment changed the partitions, False otherwise
14    """
15    # TODO

```

The method does not expect any parameters and returns a boolean. The method should remove all points from their previous partition and add them to the partition with the nearest centroid if there is a closer centroid. If the reassignment changed the partitions, the method should return `True`, otherwise `False`.

If there are empty partitions after the reassignment, the method should randomly reassign a single point from a random non-empty partition with more than one element into each empty partition. This is necessary to avoid empty partitions, which would lead to K-means not producing k clusters, but $k - n$ clusters (n being the number of empty partitions).

You can test whether your implementation is correct by executing the following command in the console:

```
1 pytest tests/kmeans/test_reassign_points.py
```

Task 1.4

The K-means algorithm is iterative. The algorithm stops if the partitions do not change anymore or if a maximum number of iterations is reached.

Open `kmeans.py` and implement `fit`, which combines the previous steps to implement the K-means algorithm:

```
1 def fit(self, points: List[Point]):  
2     """  
3     Fit the K-Means clustering instance to the given points.  
4  
5     Parameters:  
6     points (List[Point]): The points to cluster  
7     """  
8     # TODO
```

The method expects a list of `Points` and does not return anything. The method should implement the K-means algorithm by calling the methods `_initialize_partitions`, `_update_centroids`, and `_reassign_points`.

You can test whether your implementation is correct by executing the following command in the console:

```
1 pytest tests/kmeans/test_fit.py
```

Task 2: DBSCAN

DBSCAN is a density-based clustering algorithm that groups together points that are closely packed together. It is based on two parameters: ε and *MinPts*.

Task 2.1

(2 Points)

A core part of the DBSCAN algorithm is to find all points that are within a distance of ε of a given point, the so-called ε -Neighborhood.

Open `dbscan.py` and implement `_epsilon_neighborhood`, which returns all points that are within a distance of ε of a given point:

```
1 def _get_neighborhood(self, point: Point, points: List[Point]) -> List[Point]:
2     """
3     Get the neighborhood of a point.
4
5     Parameters:
6     point (Point): The point to get the neighborhood of.
7     points (List[Point]): The points to consider.
8
9     Returns:
10    List[Point]: The points in the neighborhood.
11    """
12    # TODO
```

The method expects the `Point` for which the neighborhood should be determined and a list of all `Points` that should be considered. The method should return all points that are within a distance of ε of the given point as a list of `Points`. ε is given by the variable `self.epsilon`.

You can test whether your implementation is correct by executing the following command in the console:

```
1 pytest tests/dbscan/test_epsilon_neighborhood.py
```

Task 2.2

(14 Points)

The function `_epsilon_neighborhood` can be used to implement the DBSCAN algorithm. The algorithm can be implemented either iteratively or recursively.

Open `dbscan.py` and implement `fit`, which implements the DBSCAN algorithm:

```
1 def fit(self, points: List[Point]) -> None:
2     """
3     Fit the DBSCAN clustering instance to the given points.
4
5     Parameters:
6     points (List[Point]): The points to cluster.
7     """
8     # TODO
```

The method expects a list of `Points` and does store the found clusters in the variable `self.clusters` and the noise points in the variable `self.noise`.

The method can be implemented either iteratively or recursively. Additional helper methods can be implemented if necessary, but they will not be tested and therefore award no points.

You can test whether your implementation is correct by executing the following command in the console:

```
1  pytest tests/dbscan/test_fit.py
```

Appendices


In [Task 1](#) and [Task 2](#) test cases are provided and used to grade the submission.

The most test cases are based on the following data sets:

Small Point Dataset

All test cases starting with the prefix `test_with_small_point_dataset` are based on the small dataset of 2D Points known from Exercise Sheet 5 - Task 1/Task 2.

The dataset is structured as follows:



Bigger Point Dataset

All test cases starting with the prefix `test_with_bigger_point_dataset` are based on a bigger dataset of one-hundred 2D Points.

The dataset is structured as follows:

