

Lab 4: Clustering

Due date: Thursday, May 3..

Overview

In this assignment you will use unsupervised learning techniques to discover clusters in a number of simple datasets. You will implement two cluster analysis methods, ***k*-means clustering** and **agglomerative hierarchical clustering**.

Assignment Preparation

This is a pair programming assignment. Each student teams up with a partner. While this assignment can be easily split into two parts (one person does *k*-means clustering algorithm, the other — hierarchical clustering algorithm), I **strongly recommend** that both students take active part in implementing both algorithms.

Note: Open-source versions of clustering algorithms are plentiful. It is considered cheating to re-use code for clustering algorithms that was written by someone else. (You may use code from other *explicitly acknowledged sources* for other purposes in the program, e.g., for parsing input).

Data

We provide multiple **simple datasets** for use in this assignment. Each dataset consists of two files: a CSV (comma-separated values) file containing the actual data and a text `.txt` file containing the *header* of the dataset, specifying the names of the attributes.

Most of the datasets were adopted from

<http://people.sc.fsu.edu/~burkardt/datasets/hartigan/hartigan.html>

which contains a list of datasets from [1] and [2]. Some of these datasets contain real data, some — are synthesized. The remaining data files come from the FATAL ACCIDENTS dataset, collected by Peter Oyler. The dataset contains information about automobile accidents that resulted in fatalities. All CSV files for all datasets have exactly the same format.

CSV file format. The first line of each CSV file will be analogous to the restrictions file in the **Lab 3** assignment. It is a **binary vector** identifying which columns of the CSV file shall be used for clustering. In some CSV files, one of the columns is used as a `rowId`, and should be excluded from the analysis, while in some other files, all columns contain data.

0 in position i of the first row means that column i should be ignored in the cluster analysis.

1 in position 1 of the first row means that column i should be used in the cluster analysis.

The remaining rows of the CSV file contain data points.

Text files. Text files need not be processed by your programs. Text files contain the original headers for each dataset, split from the CSV data to simplify parsing. The headers describe the nature of the dataset, contain references to its origins and specify the names of each column in the dataset. If you want to use attribute headers in your program (e.g., for generating output), you may either add a header file name as a parameter to your program, or deduce the header file name (it will always be `header_` followed by the name of the data file w/o extension, followed by `.txt`).

List of Datasets

The following datasets are *officially* provided to you for this assignment. Each file listed below, as well as a zip archive of the entire collection is available at the following URL¹:

<http://wiki.csc.calpoly.edu/csc466-2009/wiki/Lab4Data>

CSV File	Header File	Dataset	Source
<code>4clusters.csv</code>	<code>header_4clusters.txt</code>	4 clusters in 2D space (synthetic)	Spaeth[2]
<code>mammal_milk.csv</code>	<code>header_mammal_milk.txt</code>	constituents of mammal milk	Hartigan[1]
<code>economy.csv</code>	<code>header_economy.txt</code>	profit vs. equity in sectors of economy	Hartigan[1]
<code>planets.csv</code>	<code>header_planets.txt</code>	sightings of minor planets	Hartigan[1]
<code>many_clusters.csv</code>	<code>many_clusters_headers.txt</code>	many clusters in 2D space (synthetic)	Spaeth[2]
<code>birth_death_rate.csv</code>	<code>header_birth_death_rate.txt</code>	birth and death rates of countries	Hartigan[1]
<code>AccidentsSet01.csv</code>	<code>header_AccidentsSet01.csv</code>	fatal automotive accidents	P. Oyler
<code>AccidentsSet02.csv</code>	<code>header_AccidentsSet02.csv</code>	fatal automotive accidents	P. Oyler
<code>AccidentsSet03.csv</code>	<code>header_AccidentsSet03.csv</code>	fatal automotive accidents	P. Oyler

¹I used an already existing page from the previous course's wiki and added data to it. You have no write permissions to the wiki, but should have full read rights.

More datasets may be made available by the instructor, as well as can be manufactured and used by each team, but only the datasets listed above will be used as part of your lab deliverables.

Lab Assignment

You will investigate the behavior of two common clustering procedures, ***k*-means clustering** and **agglomerative hierarchical clustering** on the datasets provided to you. Part of your assignment involves implementation of the two clustering procedures. In addition to this, you will study the datasets using your implementations of the clustering algorithms, and will submit the results you have obtained.

For simplicity, the lab refers to Java programs whenever specific pieces of code are mentioned. However, each team is free to implement clustering algorithms in their programming language of choice (*provided that the instructions on how to run their programs are submitted as part of the lab deliverables*).

k-means Clustering

You will implement the ***k*-means clustering** algorithm as a Java program `kmeans.java`.

Input. The program shall take as input two parameters:

```
java kmeans <Filename> <k>
```

`<Filename>` is the name of the CSV file containing the input dataset.

`<k>` is the number of clusters the program has to produce.

Internals. There is a variety of versions of the ***k*-means clustering algorithm**. Each team is encouraged to (a) select the specific version and, possibly (b) experiment with different versions ² to discover the one that produces better clusters (or works faster).

While the datasets you are working with are small, you should implement a version of the **disk-based *k*-means clustering algorithm** in a sense that your implementation **access each data point exactly once per iteration of the algorithm**. You are allowed to store all data in main memory, since the datasets are very small.

Among the different decisions that each team can make are:

- Choice of **initial cluster centroids**. Random, SelectCentroids procedure, pre-selected, other variations. . . .

²This is where two people can work in parallel on the same program.

- Choice of **centroid recomputation method**. Typically, the **mean** point of the cluster (hence the "*k-means*" in the name of the method), but you can use **medians** or **modes** if so desired.
- **Stopping criterion**. Minimum point reassignment, minimum centroid change, threshold on **sum of squared error (SSE)** (for the latter — you need to find a good threshold).
- **Distance measure**. While Euclidean distance makes sense for some of the datasets (the 2D synthetic ones, e.g.), you may choose to try other distance measures to see if your cluster organization changes/becomes more pronounced.
- **Dealing with outliers**. Attempt to detect outliers vs. no outlier detection.

Output. The output of your program shall consist of two things:

1. Description of the detected clusters.
2. Evaluation of the computed clustering of the data.

Since our datasets are small, describe detected clusters **by listing all points belonging to them** (this also simplifies grading/comparison). If data points come with **rowIds**, you are allowed to output just them.

Because we lack well-established (and suppliable to your program) **ground truth** on clusters, only **internal evaluation measures** for clusters need to be computed and reported.

For each cluster, compute and report:

1. Number of points in the cluster.
2. Coordinates of its centroid.
3. Maximum, minimum and the average distance from a point to cluster centroid.
4. Sum of Squared Errors (SSE) for the points in the cluster.

You may *additionally* choose to report any other internal measures, e.g., **inter-cluster distances**.

Here is an example of an output for a single cluster (w/o the SSE):

```
Cluster 0:
Center: 26.25,28.0,
Max Dist. to Center: 10.077822185373186
Min Dist. to Center: 3.010398644698074
Avg Dist. to Center: 6.569312544990502
4 Points:
25.0,38.0,
```

32.0,27.0,
26.0,25.0,
22.0,22.0,

(Your output does not have to match this, but should report the information in an easy-to-read way).

Hierarchical Clustering

You will implement the **agglomerative hierarchical clustering** method as a `hclustering.java` program.

Input. Your program shall take two parameters as input:

```
java hclustering <Filename> [<threshold>]
```

<Filename> is the name of the CSV file containing the input dataset.

<threshold> is the *optional* threshold at which your program will "cut" the cluster hierarchy to report the clusters.

If <threshold> parameter is **specified** in the input, your program shall produce both the cluster hierarchy, and the appropriate list of clusters cut at the specified threshold.

If <threshold> parameter is **not specified** in the input, your program shall produce the cluster hierarchy **alone**.

Internals. Commonly, **hierarchical clustering algorithms** take as input the distance matrix for the points in the input dataset. Your algorithm will work with the same data as the ***k*-means clustering algorithm**, i.e., the actual dataset. Your implementation of the hierarchical clustering method will be responsible for computing the distances between all points in the dataset.

Additionally, as with ***k*-means clustering**, in your implementation of the **agglomerative hierarchical clustering** algorithm, you can select the specific details of implementation that you prefer, or experiment with multiple options.

Among the features you have the select are:

- **Distance measure for points.** See ***k*-means clustering algorithm** discussion.
- **Distance measure for clusters.** Select between **single-link**, **complete-link** and **average-link** methods (you can also try **centroid** or **Ward's** methods).

Output. `hclustering.java` shall produce as output the following information:

- The **dendrogram of the cluster hierarchy** constructed by the algorithm. This hierarchy is independent of the threshold and shall be produced each time the program is run.
- The **actual clusters** obtained by applying the input threshold to the computed **cluster hierarchy**.
- The **evaluation** measures specifying the quality of the obtained clusters.

The latter two types of output should be produced **only if** the threshold value is specified in the input. If it is not specified, only the tree is produced and returned.

Dendrogram. Create an XML version of the cluster hierarchy dendrogram and output it. You need to output the dendrogram XML to `stdout`, but you may, if you want, also create an XML file to store it.

The XML representing the dendrogram is simple and straightforward. The dendrogram is a binary labeled tree. You will use three XML elements:


- `<tree>` to indicate the root node of the dendrogram.
- `<node>` for all inner nodes.
- `<leaf>` for all leaf nodes containing individual data points. Leaf elements are `EMPTY`.

The `<tree>` and `<node>` elements will contain an attribute `height` representing the height label of the node (the distance between the two clusters that are merged).

The `<leaf>` elements will contain an attribute `data` whose contents shall be the data point from the dataset represented by the leaf node (or its `rowId`). For the sake of completeness, feel free to add a `height` attribute with the value of 0 to the `<leaf>` elements.

The `<tree>` and all `<node>` elements shall have exactly two children: either `<node>` or `<leaf>` elements.

Here is an example of a small XML output and the dendrogram it represents.



Clusters and measures. Whenever a **threshold** is specified, your program shall, in addition to producing the full dendrogram, also report the clusters that are formed when the dendrogram is cut at the specified threshold. The clusters and their parameters shall be reported in the same way as your ***k*-means clustering algorithm** implementation does it.

Study of the datasets

You will use your implementations of the ***k*-means clustering** and **agglomerative hierarchical clustering** methods to study the six datasets provided to you and to find the *best* (in your opinion) clusters in the data.

The results of your study will be compiled in the report that each team will prepare and submit.

Use both methods. For each dataset provided to you, you need to use both methods (programs) to discover the best clusters. Your report will indicate the best results you achieved using each of the two programs.

Best clustering. Each implementation comes with one *predefined* parameter (k , number of clusters, for *k*-means clustering; **threshold** for hierarchical clustering), which will affect the output of the respective program. In addition, each team may choose to implement multiple choices for various parts of each algorithm (e.g., you may implement all three *linking* techniques for the hierarchical clustering, or a number of different ways to select the initial centroids for the *k*-means clustering).

You will run your programs on each dataset with different settings of the input parameters (and any choices that you have implemented) to determine which runs produce the best clustering result.

The evaluation measures computed by your programs will help you assess the quality of the output.

Additionally, you may choose to determine the **ground truth** for some (or all) of the datasets and compare the outputs to the results you obtain. Ground truth can be determined in some cases by investigating/visualizing the dataset using any tools available to you. 2D and some 3D datasets, for example, can be plotted as scatterplots using Excel or MatLab and visual observation may help you determine the clusters that should be detected by your program. (For some datasets, like the **economy** dataset, which has 10 variables, simple visualization may be infeasible).

Report. Each team will prepare a report of discoveries. The report will consist of the following information:

1. **Study Design.** Provide brief description of your implementations of the two clustering methods. Specify which features (such as distance

measures, centroid choices, etc.) you have implemented and used in the study.

2. **Results.** For each dataset, and for each algorithm provide the description of the best result achieved. Specify the input parameters (and features) of the program that lead to the best result, and provide a snapshot of the obtained output (feel free to paste the output verbatim, or to use a more succinct representation of the clusters).
3. **Discussion.** May be embedded into reporting of results. Mention any interesting observations about specific datasets/methods that you have made during the study.
4. **Analysis.** Present an overall conclusion of your results. Did both methods perform equally well? Did one method work better on some data and the other - on different data? When did each of the algorithms perform well? What features (if any) of each algorithm lead to better performance in what cases?

Deliverables and submission instructions

This lab has both electronic and hardcopy deliverables.

Electronic deliverables.

- `kmeans.java` and all supporting files.
- `hclustering.java` and all supporting files.
- `README` file specifying what has been implemented, and providing any instructions for compiling and running the code. (must contain the names of all students on the team).
- **The report.** The electronic version of the report can be submitted in any readable format (pdf, Word, HTML). In addition, each team should put the report on the course wiki page (a page for the reports will be created).

Submit all electronic deliverables as a single zip or gzipped tar archive (`lab04.zip` or `lab04.tar.gz`). Use the following command

```
$ handin dekhtyar-grader lab04 lab04.<ext>
```

Hardcopy Deliverables.

- **The report.** Submit a hard copy of your report during the *May 12* lab period. (if you need more time to complete the report, submit the hardcopy to my office on Friday. No late penalty will be assessed if the electronic submission was on time.)

References

- [1] John Hartigan, *Clustering Algorithms*, Wiley, 1975. ISBN 0-471-35645-X.
- [2] Helmut Spaeth, *Cluster Dissection and Analysis, Theory, FORTRAN Programs, Examples*, Ellis Horwood, 1985.