

Updates

The updated version of this homework is maintained at <http://www.cs.cornell.edu/Courses/cs4786/2015sp/assignments/hw2.pdf>. Announcements of changes will be made on the course assignments webpage. Additionally, announcements of any *crucial* updates will be made via email, using the CMS system.

The changes mentioned below are marked in orange in the following pages as well.

- Feb 19, 7:13pm:
 - second Q1.1: number has been corrected to Q1.3
 - Q1.3: clarify that your spectral clustering should use k-means
 - Deliverables: clarify that we only need scatter plots for the final clustering, not for every intermediate iteration of k-means or single link.
 - Deliverables: clarification that each line (i.e., each row, *not* each column) of `means.csv` should be a cluster center.
 - Deliverables: add the requirement to submit the initial cluster centroids for spectral clustering.

Prefatory note We've purposely made this assignment short to make it easy to finish before spring break. Other material on clustering not covered in this assignment may be included in future assignments.

Instructions Due at 5pm Friday March 27 on CMS. Submit what you have at least once by an hour before that deadline, even if you haven't quite added all the finishing touches — CMS allows resubmissions up to, but not after, the deadline. If there is an emergency such that you need an extension, contact the professors.

You may work in groups of one up to four¹. Each group of two or more people must create a group on CMS well before the deadline (there is both an invitation step and an accept process; make sure both sides of the handshake occur), and submits 1 submission per group. **We will not be automatically transferring the groups from A1, so that you can change groups if you want, but this means you should form groups anew for A2.** You may choose different groups for different assignments. Please ensure that each member of the group can individually defend or explain your group's submission equally well.

You will submit both a writeup and some datafiles you create. The writeup can be handwritten or typeset, but please make sure it is easily readable either way.

Keep an eye on the course webpage for any announcements or updates.

Academic integrity policy We distinguish between “merely” violating the rules for a given assignment and violating *academic integrity*. To violate the latter is to commit *fraud* by claiming credit for someone else's work. For this assignment, an example of the former would be getting an answer from person X who is not in your CMS-declared group but stating in your homework that X was the source of that particular answer. You would cross the line into fraud if you did not mention X. The worst-case outcome for the former is a grade penalty; the worst-case scenario in the latter is academic-integrity hearing procedures.

The way to avoid violating academic integrity is to always document any portions of work you submit that are due to or influenced by other sources, even if those sources weren't permitted by the rules.²

¹The choice of the number “four” is intended to reflect the idea of allowing collaboration, but requiring that all group members be able to fit “all together at the whiteboard”, and thus all be participating equally at all times. (Admittedly, it will be a tight squeeze around a laptop, but please try.)

²We make an exception for sources that can be taken for granted in the instructional setting, namely, the course materials. To minimize documentation effort, we also do not expect you to credit the course staff for ideas you get from them, although it's nice to do so anyway.

Q1 (Clustering Sensitivity). In class, we covered three clustering methods: K-means, single link and spectral clustering (normalized-cut version). The goal of this assignment is for you to explore the sensitivity of the clustering methods we’ve introduced by showing that small perturbations of the initial data can lead to a quite different clustering, even for binary clusterings.

You may use code packages provided by other people or sources — be sure to credit these sources appropriately. But, *if you use external code, it is your responsibility to make sure that the resulting clusterings are the same as would be produced if you were to reimplement precisely what was presented in class.* For example, if you use k-means clustering code that makes multiple runs and then averages over the runs in some way³, then your clustering result may differ from what our testing harness comes up with. Thus, carefully read the documentation of any external code. Specifically for k-means, if you are using external code, make sure to specify explicitly initial centroids (usually this can be added as extra parameter).

The three problems below have roughly equal weight. In grading, we care about your explanations at least as much as the datasets you provide.

In this assignment, K , the number of clusters per clustering, is fixed at 2, and n , the number of data points each initial dataset should contain, is fixed at 30. When you are asked to provide a vector c of cluster assignments, in such vectors, the t^{th} entry c_t is 1 if the t^{th} datapoint is in the first cluster, 0 otherwise.⁴

Q 1.1 K-means:

- Create an initial data matrix $X^{\text{kmeans},I}$ with 30 points each in \mathbb{R}^2 . Also create the vector $c^{\text{kmeans},I} \in \mathbb{R}^{30}$ of cluster assignments you get by running the K-means algorithm on this data along with the initial two cluster centers $\mu_1, \mu_2 \in \mathbb{R}^2$ you chose to use. $c^{\text{kmeans},I}$ **should have an equal number of 1’s and 0’s, that is, clusters of equal size.**
- Add anywhere between 1 to 3 points to $X^{\text{kmeans},I}$ to create a new data matrix $X^{\text{kmeans},II}$. **These 1 to 3 points must be within the smallest rectangle bounding the points in $X^{\text{kmeans},I}$, and must be the last vectors in your matrix.** Run the K-means algorithm on this modified dataset with the **same initial cluster centers** μ_1, μ_2 you used for $X^{\text{kmeans},I}$ and produce the new cluster assignment vector $c^{\text{kmeans},II}$.
- **Goal: $c^{\text{kmeans},II}$ and $c^{\text{kmeans},I}$ must vary by over 30%. That is⁵,**

$$\min_{C=c^{\text{kmeans},II}, C=1-c^{\text{kmeans},II}} \frac{1}{30} \sum_{t=1}^{30} \mathbb{1}_{\{c_t^{\text{kmeans},I} \neq C_t\}} \geq 0.3$$

Q 1.2 Single Link:

- Create an initial data matrix $X^{\text{s-link},I}$ with 30 points each in \mathbb{R}^2 . Also create the vector $c^{\text{s-link},I}$ of cluster assignments you get by running the single link clustering algorithm on it. $c^{\text{s-link},I}$ **should have an equal number of 1’s and 0’s.**

³Hint: we didn’t just make this up.

⁴It’s up to you which cluster is the “first” one, so in this sense the cluster labels are arbitrary; we just need to know which points are in different clusters and which points are in the same cluster.

⁵ $\mathbf{1}$ is the vector with all 30 coordinates being 1. We pick C this way because labeling clusters as $1 - 0$ or $0 - 1$ leads to the same groupings but potentially swapped labels, so just looking at label differences isn’t the right way to measure the degree of perturbation. So this measure checks both one labeling and then the “flip” of that labeling.

- Add anywhere between 1 to 3 points to $X^{s\text{-link},I}$ to create the data matrix $X^{s\text{-link},II}$. **These 1 to 3 points must be within the smallest rectangle bounding the points in $X^{s\text{-link},I}$, and must be the last vectors in your matrix.** Run the single link clustering algorithm on this modified dataset and produce the new cluster assignment $c^{s\text{-link},II}$.
- **Goal: $c^{s\text{-link},II}$ and $c^{s\text{-link},I}$ must vary by over 30%. That is,**

$$\min_{C=c^{s\text{-link},II}, C=1-c^{s\text{-link},II}} \frac{1}{30} \sum_{t=1}^{30} \mathbb{1}_{\{c_t^{s\text{-link},I} \neq C_t\}} \geq 0.3$$

Q 1.43 Spectral clustering:

- Provide an initial 30×30 adjacency matrix $A^{\text{spectral},I}$ of an undirected graph (symmetric matrix with 0 – 1 entries). Perform spectral clustering with the normalized Laplacian matrix on the data points and provide the vector $c^{\text{spectral},I}$ of cluster assignments you obtained. $c^{\text{spectral},I}$ **should have an equal number of 1's and 0's.**
- Add anywhere between 1 to 3 edges to $A^{\text{spectral},I}$ to create the new adjacency matrix $A^{\text{spectral},II}$. Run spectral clustering on this modified dataset, **using k-means on the embeddings** to get new cluster assignment vector $c^{\text{spectral},II}$.
- **Goal: $c^{\text{spectral},II}$ and $c^{\text{spectral},I}$ must vary by over 30%. That is,**

$$\min_{C=c^{\text{spectral},II}, C=1-c^{\text{spectral},II}} \frac{1}{30} \sum_{t=1}^{30} \mathbb{1}_{\{c_t^{\text{spectral},I} \neq C_t\}} \geq 0.3$$

Deliverables: Submit a **writeup** explaining the way you generated the data points and the corresponding modifications, and why you expected the new datasets to result in significantly different clusterings. In your write-up, for K-means and single-link clustering, for every cluster **in the final (output) clustering** produced, **include scatter plots** of the points where the points are color-coded according to their corresponding cluster assignments. For K-means, also include your **initial cluster centroids (as larger points or otherwise clearly visible and distinguished from the data points)** in the scatter plots. For spectral clustering, **include a depiction** of both the initial and the modified **graphs**. Color code the nodes according to the spectral cluster assignments.

For each method, also submit the initial data points (adjacency matrix in the spectral clustering case); the modified dataset/adjacency matrix produced by adding the extra 1 to 3 points (or edges); and the cluster assignments you obtained by running the algorithms over the initial and modified datasets. For the K-means algorithm provide the initial cluster means μ_1, μ_2 you started with.

Specifically, submit your datasets as csv files obeying the following requirements. `XkmeansI.csv` and `XslinkI.csv` must each consist of exactly 30 lines, each consisting of 2 comma-separated values. `XkmeansII.csv` and `XslinkII.csv` must each be between 31 and 33 lines, where each line contains 2 comma-separated values. `AspectralI.csv` and `AspectralII.csv` must each be 30 lines long, each line consisting of 30 comma-separated values. Finally, `ckmeanI.csv`, `ckmeanII.csv`, `cslinkI.csv`, `cslinkII.csv`, `cspectralI.csv` and `cspectralII.csv` are each 30 lines, where each line contains one value that is either 0 or 1, indicating the cluster assignment of the corresponding original point. Also submit cluster centers $\mu_1, \mu_2 \in \mathbb{R}^2$ for Q1.1 in file `means.csv` containing 2 lines, **representing μ_1 and μ_2 , respectively**, each of which consists of 2 comma-separated values. **And, submit in file `spectralmeans.csv` the cluster centers $\mu_1^{\text{spectral}}, \mu_2^{\text{spectral}} \in \mathbb{R}^2$ that you used to initialize k-means on the embeddings for Q1.3.** File

`spectralmeans.csv` should contain 2 lines, representing μ_1^{spectral} and μ_2^{spectral} , respectively, each of which consists of 2 comma-separated values.

Note: in this assignment, points will be deducted for submissions of dataset that do not conform precisely to our instructions. (Last time, we altered our grading code to handle transposes and the like).