
DSC 190 - Homework 05

Due: Wednesday, May 10

Write your solutions to the following problems by either typing them up or handwriting them on another piece of paper. Unless otherwise noted by the problem's instructions, show your work or provide some justification for your answer. Homeworks are due via Gradescope at 11:59 PM.


Problem 1.

In lecture, we saw that one minimizer of the cost function

$$\text{Cost}(\vec{f}) = \sum_i \sum_j w_{ij} (f_i - f_j)^2$$

is the vector $\vec{f} = \frac{1}{\sqrt{n}}(1, 1, \dots, 1)^T$ which embeds all of the nodes of the graph to exactly the same number. The cost of this trivial embedding is zero, and we typically ignore it in favor of the eigenvector of the graph Laplacian with the *next* smallest positive eigenvalue for the embedding.

In the case where the graph has multiple connected components, however, there may be additional embeddings which have a cost of zero. For example, consider the similarity graph G shown below:



The weight of each edge is shown; the weight of non-existing edges is zero.

Find a normalized embedding vector \vec{g} for the above graph such that $\text{Cost}(\vec{g}) = 0$ and $\vec{g} \perp \frac{1}{\sqrt{n}}(1, 1, \dots, 1)^T$; that is, \vec{g} is orthogonal to the vector of all ones.

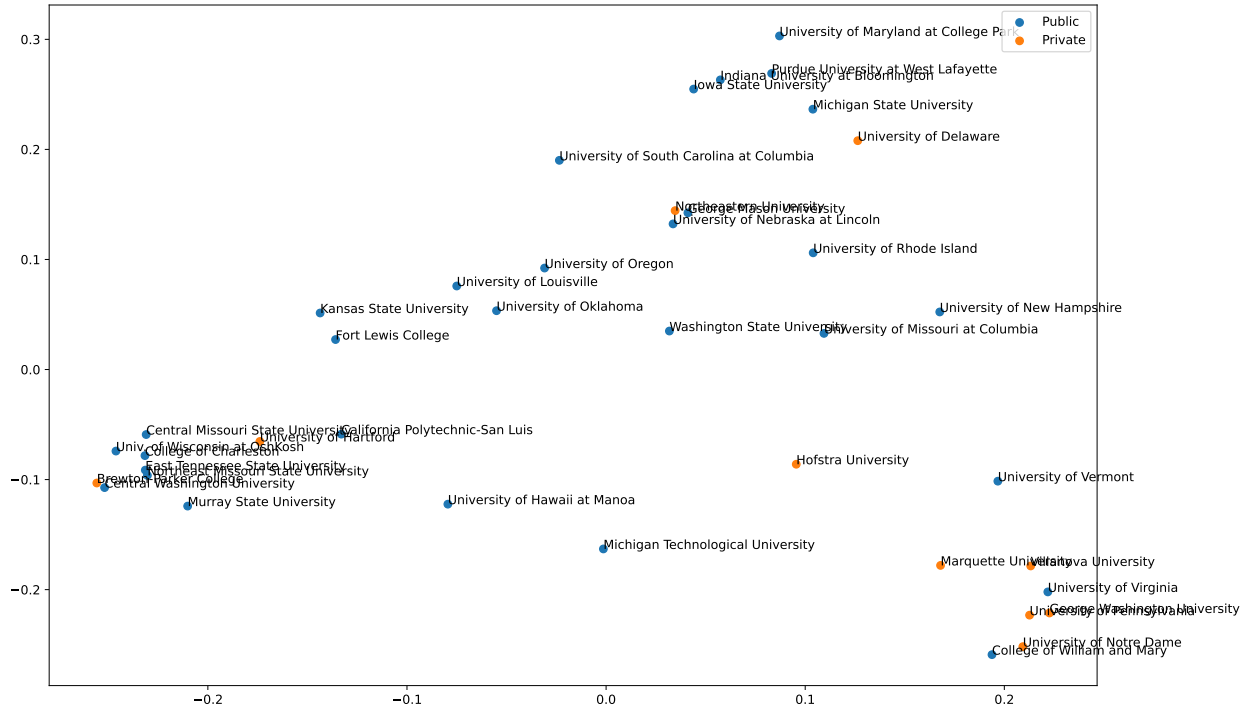
Problem 2.

The data at the link below contains information on 40 colleges in the US:

<https://f000.backblazeb2.com/file/jeldridge-data/006-colleges/colleges-sample.csv>

The data contains 17 numerical features measuring the size of the student body of each college, the graduation rate, etc., and one Boolean feature (“Private”) saying whether the school is private or public. If we consider only the numerical features, each college is a point in \mathbb{R}^{17} .

Using Laplacian Eigenmaps, embed each college as a point in two dimensions and plot the result to obtain a similar figure to that shown below. You should construct a *symmetric* similarity matrix by building a k -neighbors graph (you will need to choose k). You should also consider whether the data needs to be standardized. You should drop the “Private” column for the purpose of computing the embedding – use only the numerical data.



Like the plot above, each point in your plot should be annotated with the name of the college, and the color of the point should show whether the college is public or private. However, your plot may look significantly different from the plot above due to the sign of the eigenvectors your code finds (it is arbitrary), as well as your choice of k in constructing the k -neighbors graph, etc. We will look to see if it shows the same general *patterns* as the plot above. For example, notice how there is a group of large midwestern state schools (including “Michigan State University”, etc.), another group of large private universities (including “Notre Dame”), etc. Your plot should show the same.

You may use packages like `sklearn` to compute the k -neighbors graph, but not to do the actual embedding itself (e.g., do not use anything from `sklearn.manifold`). Hint: the output of some `sklearn` functions is a sparse matrix; it is OK in this case to convert it to a dense `numpy` array. Also note that `sklearn`'s function for building a k -neighbors graph returns a *directed* graph, not an undirected graph, and therefore produces an asymmetric weight matrix. You will need to think about how to “symmetrize” it.