

Mathematics of Image and Data Analysis

Math 5467

Lecture 21: Graph-based embeddings

Instructor: Jeff Calder

Email: jcalder@umn.edu

<http://www-users.math.umn.edu/~jwcalder/5467S21>

Last time

- Graph-based semi-supervised learning.

Today

- Graph-based embeddings (spectral and t-SNE)

Embeddings of high dimensional data

High dimensional data is hard to visualize and work with. Embeddings to low dimensional spaces help us visualize data and improve the performance of data analysis algorithms.

- PCA (linear dimension reduction)
- Spectral embedding (today)
- t-Distributed Stochastic Neighbor Embedding (t-SNE) (also today)

The key is to embed the data while still preserving important structures.

Spectral embeddings

Let v_1, v_2, v_3, \dots be the normalized eigenvectors of L , in order of increasing eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \dots$. The spectral embedding corresponding to L is the map $\Phi : I_m \rightarrow \mathbb{R}^k$ (recall $I_m = \{1, 2, \dots, m\}$ are the indices of our datapoints) given by


$$(1) \quad \Phi(i) = (v_1(i), v_2(i), \dots, v_k(i)).$$

Since the first eigenvector v_1 is the trivial constant eigenvector, it is also common to omit this to obtain the embedding

$$\Phi(i) = (v_2(i), v_3(i), \dots, v_{k+1}(i)).$$

There are other normalizations of the graph Laplacian that are commonly used, such as the symmetric normalization $L = D^{-1/2}(D - W)D^{-1/2}$, and the spectral embedding for a normalized Laplacian is defined analogously.

Spectral embeddings



The intuition behind the spectral embedding is encapsulated in the following simple result.

Proposition 1. If $A \subset I_m$ is a disconnected component of the graph, which means that $W(i, j) = 0$ for all $i \in A$ and $j \in I_m \setminus A$, then the indicator function of A , denoted u_A , satisfies

$$Lu_A = (D - W)u_A = 0. \quad \leftarrow \text{eigenvector}$$

$$Lu_A(i) = \sum_{j=1}^m w(i, j) (u_A(i) - u_A(j))$$

$$\textcircled{1} \quad \underline{i \in A} \quad = \sum_{j \in A} w(i, j) \underbrace{(u_A(i) - u_A(j))}_{1 - 1 = 0}, \quad \text{since } w(i, j) = 0 \text{ for } j \notin A$$

$$\textcircled{2} \quad \underline{i \notin A} \quad = \sum_{j \notin A} w(i, j) (u_A(i) - u_A(j)) = 0$$

since $u_A(i) = u_A(j) = 0$



If we have k disconnected

clusters A_1, A_2, \dots, A_k , then

$u_{A_1}, u_{A_2}, \dots, u_{A_k}$ are the first
 k eigen vectors of L_f and the
spectral embedding is

$$\Phi(i) = (u_{A_1}(i), u_{A_2}(i), \dots, u_{A_k}(i)) = e_j$$

if $i \in A_j$

Spectral embedding of MNIST

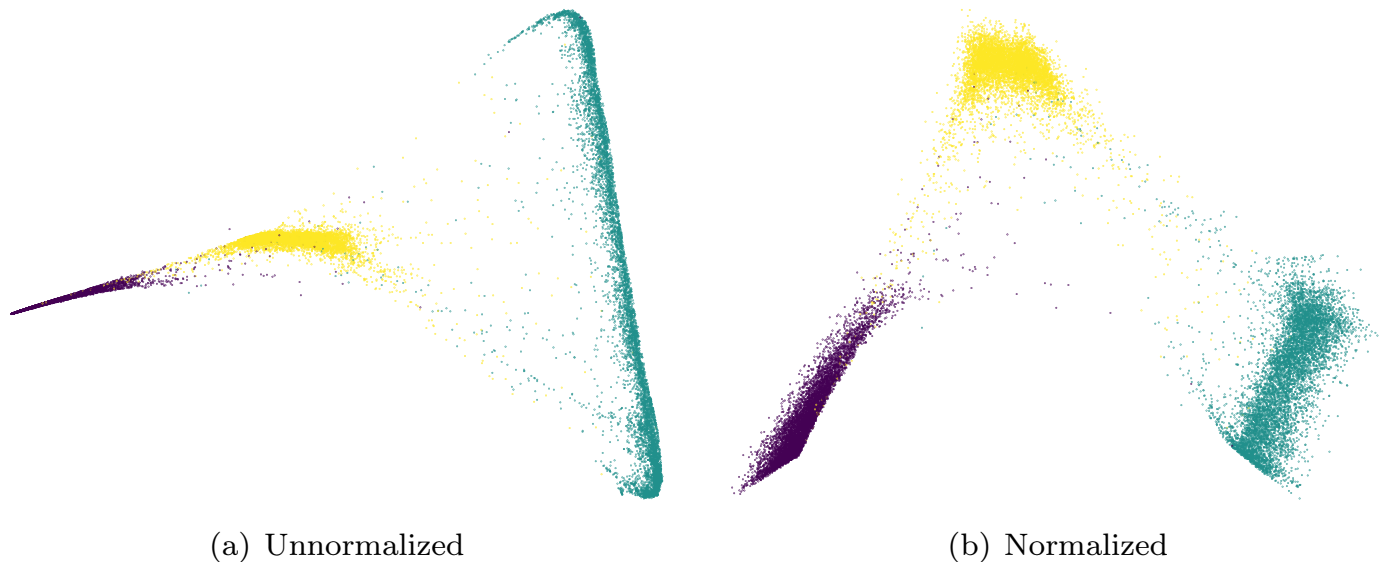


Figure 1: Example of spectral embeddings in the plane $k = 2$ of the 0, 1, and 2 digits of the MNIST dataset using the unnormalized $L = D - W$ and symmetric normalized $L = D^{-1/2}(D - W)D^{-1/2}$ graph Laplacians.

t-SNE

The t-Stochastic Neighbor Embedding (t-SNE) tries to find embedded points whose pairwise similarities match as closely as possible the given weight matrix W for the graph.

From the weight matrix W , t-SNE constructs a probability weight matrix

$$(2) \quad P = \frac{1}{2m}(D^{-1}W + W^T D^{-1}), \quad P(i, i) = 0$$

where D is the diagonal matrix of degrees $d(i) = \sum_{j=1}^m W(i, j)$. We say *probability* since all entries of P sum to one, i.e., $\mathbf{1}^T P \mathbf{1} = 1$.

$$\begin{aligned} \mathbf{1}^T P \mathbf{1} &= \frac{1}{2m} (\mathbf{1}^T D^{-1} W \mathbf{1} + \mathbf{1}^T W^T D^{-1} \mathbf{1}), & \text{diag}(D) &= W \mathbf{1} \\ &= \frac{1}{2m} (\mathbf{1}^T \mathbf{1} + \mathbf{1}^T \mathbf{1}), & D^{-1} W \mathbf{1} &= \mathbf{1} \\ &= \frac{1}{2m} (m + m) = 1, & \mathbf{1}^T W^T D^{-1} &= \mathbf{1}^T \end{aligned}$$

t-SNE

t-SNE aims to find embedded points $y_1, y_2, \dots, y_m \in \mathbb{R}^k$, where usually $k = 2$ or $k = 3$, so that the similarity between y_i and y_j matches $P(i, j)$ as closely as possible. The similarity matrix for the y_i , denoted Q , is the $m \times m$ matrix defined by


$$(3) \quad Q(i, i) = 0 \quad Q(i, j) = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{\ell \neq s} (1 + \|y_s - y_\ell\|^2)^{-1}}.$$

\uparrow all $\ell, s = 1, \dots, m, \ell \neq s$.

The discrepancy between P and Q is measured with the Kullback-Leibler divergence

$$(4) \quad E(y_1, y_2, \dots, y_k) = D(P \| Q) := \sum_{i \neq j} P(i, j) \log \left(\frac{P(i, j)}{Q(i, j)} \right).$$

t-SNE finds the embedded points y_i by minimizing E with gradient descent.



Why Kullback-Leibler?

$$E(y_1, y_2, \dots, y_k) = D(P||Q) := \sum_{i \neq j} P(i, j) \log \left(\frac{P(i, j)}{Q(i, j)} \right).$$

- When $P(i, j) \gg 0$, forces $Q(i, j) \sim P(i, j)$; i.e., preserve local structure.
- When $P(i, j) \sim 0$ we don't care what $Q(i, j)$ does; i.e., allow global structure to change.
- We cannot preserve all information in a dimension reduction.

Gradient descent

$$\log(P) - \log(Q)$$

The gradient of the Kullback-Leibler divergence

$$E(y_1, y_2, \dots, y_k) = D(P\|Q) := \sum_{i \neq j} P(i, j) \log \left(\frac{P(i, j)}{Q(i, j)} \right)$$

in the variable y_i is given by

$$\nabla_{y_i} E = \underbrace{4Z \sum_{j:j \neq i} P(i, j) Q(i, j) (y_i - y_j)}_{\text{Attraction}} - \underbrace{4Z \sum_{j:j \neq i} Q(i, j)^2 (y_i - y_j)}_{\text{Repulsion}}.$$

where $Z = \sum_{i \neq j} (1 + \|y_i - y_j\|^2)^{-1}$.

Gradient descent is

$$y_i^{k+1} = y_i^k - h \nabla_{y_i} E(y_1^k, y_2^k, \dots, y_m^k),$$

where h is the time-step.

t-SNE embedding of MNIST




Figure 2: A t-SNE embedding of 2500 images from the MNIST dataset, with colors corresponding to the digit labels of each image.

Gradient Computation Can write

$$E = - \sum_{i \neq j} P(i, j) \log(Q(i, j))$$

$$Q(i, j) = \frac{(1 + \|y_i - y_j\|^2)^{-1}}{\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}}$$

$$\log(Q(i, j)) = - \log(1 + \|y_i - y_j\|^2)$$

$$- \log\left(\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}\right)$$

$$E = \underbrace{\sum_{i \neq j} P(i, j) \log(1 + \|y_i - y_j\|^2)}_A + \underbrace{\log\left(\sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}\right)}_B$$

$$\nabla_{y_l} A = \sum_{i \neq j} P(i, j) \nabla_{y_l} \log(1 + \|y_i - y_j\|^2)$$

$$= \sum_{i \neq j} P(i, j) (1 + \|y_i - y_j\|^2)^{-1} \nabla_{y_l} \|y_i - y_j\|^2$$

If $l = i$ then $\nabla_{y_l} \|y_i - y_j\|^2 = 2(y_i - y_j)$

If $l = j$ then $\nabla_{y_l} \|y_i - y_j\|^2 = \nabla_{y_l} \|y_j - y_i\|^2 = 2(y_j - y_i)$

If $l \neq i$ and $l \neq j$ then $\nabla_{y_l} \|y_i - y_j\|^2 = 0$.

$$\text{Let } d_{ij} = \begin{cases} 1, & i=j \\ 0, & i \neq j \end{cases}$$

$$\begin{aligned} \nabla_{y_l} \|y_i - y_j\|^2 &= d_{il} 2(y_i - y_j) + d_{jl} 2(y_j - y_i) \\ &= 2(y_i - y_j)(d_{il} - d_{jl}). \end{aligned}$$

$$\begin{aligned} \nabla_{y_l} A &= \sum_{i \neq j} P(i, j) (1 + \|y_i - y_j\|^2)^{-1} 2(y_i - y_j)(d_{il} - d_{jl}) \\ &= 2 \sum_{j: j \neq l} P(l, j) (1 + \|y_l - y_j\|^2)^{-1} (y_l - y_j) \end{aligned}$$

$$-2 \sum_{i:i \neq l} P(l,i) (1 + \|y_i - y_l\|^2)^{-1} (y_i - y_l)$$

$i \leftrightarrow j$

$$= 2 \sum_{j:j \neq l} P(l,j) (1 + \|y_l - y_j\|^2)^{-1} (y_l - y_j)$$

$Z Q(l,j)$

$$\nabla_{y_l} A = 4 \sum_{j:j \neq l} P(l,j) (1 + \|y_l - y_j\|^2)^{-1} (y_l - y_j)$$

Let $Z = \sum_{k \neq l} (1 + \|y_k - y_l\|^2)^{-1}$ so that

$$Q(l,j) = \frac{(1 + \|y_l - y_j\|^2)^{-1}}{Z}$$

$$\nabla_{y_\ell} A = 4z \sum_{j:j \neq \ell} P(\ell, j) Q(\ell, j) (y_\ell - y_j)$$

↑ Attraction term

$$B = \log \left(\sum_{i \neq j} (1 + \|y_i - y_j\|^2)^{-1} \right) = \log(z)$$

$$\nabla_{y_\ell} B = z^{-1} \nabla_{y_\ell} \sum_{i \neq j} (1 + \|y_i - y_j\|^2)^{-1}$$

$$= -z^{-1} \sum_{i \neq j} \underbrace{(1 + \|y_i - y_j\|^2)^{-2}}_{z^2 \alpha(i, j)^2} \nabla_{y_\ell} \|y_i - y_j\|^2$$

$$= -\tau \sum_{i \neq j} Q(l_{ij})^2 2 (y_i - y_j) (d_{ie} - d_{je})$$

$$= -2\tau \sum_{j: j \neq e} Q(l_{ij})^2 (y_e - y_j)$$

$$+ 2\tau \sum_{i: i \neq e} Q(l_{ie})^2 (y_i - y_e)$$

$$\nabla_{y_e} B = -4\tau \sum_{j: j \neq e} Q(l_{ij})^2 (y_e - y_j)$$

$$\nabla_{y_e} E = \nabla_{y_e} A + \nabla_{y_e} B$$



attraction



repulsion.




Early exaggeration

Gradient descent for t-SNE is *very* slow. To speed it up, *early exaggeration* is used, which amplifies the attraction forces for the first few hundred iterations:


$$\nabla_{y_i} E = 4Z\alpha \sum_{j:j \neq i} P(i,j)Q(i,j)(y_i - y_j) - 4Z \sum_{j:j \neq i} Q(i,j)^2(y_i - y_j),$$

The parameter α is the amplification factor, often $\alpha \approx 10$.

Early exaggeration



(a) After early exaggeration



(b) Final embedding

Figure 3: An example of the t-SNE embedding after the early exaggeration phase and the final embedded, for a small version of MNIST with only 500 images from the digits 0, 1, 2, and 3.

Perplexity

The construction of the weight matrix W is important for the performance of t-SNE. The *perplexity* construction has the form

$$W(i, j) = \exp\left(-\frac{\|x_i - x_j\|^2}{2\sigma_i^2}\right),$$


where σ_i is tuned independently for each x_i depending on a specified *perplexity* level (usually in the range 5 to 50).

The perplexity of the i^{th} row of $W(i, j)$ is $2^{H(i)}$, where

$$H(i) = -\sum_{j=1}^m p(j) \log p(j), \quad p(j) = \frac{W(i, j)}{\sum_{k=1}^m W(i, k)}.$$

The value of σ_i is determined so that the perplexity $2^{H(i)}$ equals a desired user-specified value.

MNIST



Perplexity 5




Perplexity 30




Perplexity 50.


Gaussian mixture in 10 dimensions



Perplexity 5




Perplexity 30




Perplexity 50.


Parabolic curve in 5 dimensions



Perplexity 5




Perplexity 30



Perplexity 50.

$$x_5 = x_4 = x_3 = x_2 = x_1^2$$



$$x_1 \in (0, 1)$$

Graph-based embeddings (.ipynb)