Mathematics of Image and Data Analysis
Math 5467

Graph-based embeddings

Instructor: Jeff Calder
Email: jcalder@umn.edu

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

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

$$\Phi(i) = (v_1(i), v_2(i), \ldots, v_k(i)). \tag{1}$$

$k << m$

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

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

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 analagously.

Spectral embeddings

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

Proposition 1. *If $A \subset I_m$ is a disconected 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*

$$Lu_A = (D - W)u_A = 0.$$

$$u_A(i) = \begin{cases} 1, & i \in A \\ 0, & i \in I_m \setminus A \end{cases}$$

<u>Proof</u>: $Lu_A(i) = \sum_{j=1}^{m} W(i,j)\left(u_A(i) - u_A(j)\right)$

$$= \sum_{j \in A} W(i,j)\left(u_A(i) - u_A(j)\right) + \sum_{j \in I_m \setminus A} W(i,j)\left(u_A(i) - u_A(j)\right).$$

$= S_1 + S_2$

If $i \in A$ then $S_1 = 0$, since $i, j \in A$

$S_2 = 0$, since $i \in A, j \notin A$ so $w(i,j) = 0$.

If $i \notin A$ then $S_1 = 0$, $i \in A, j \in A \hookrightarrow w(i,j) = 0$

$S_2 = 0$, $\hookrightarrow i \notin A, j \notin A$

so $u_A(i) = 0$

$u_A(j) = 0$

$\boxed{\checkmark}$

Spectral embedding of MNIST

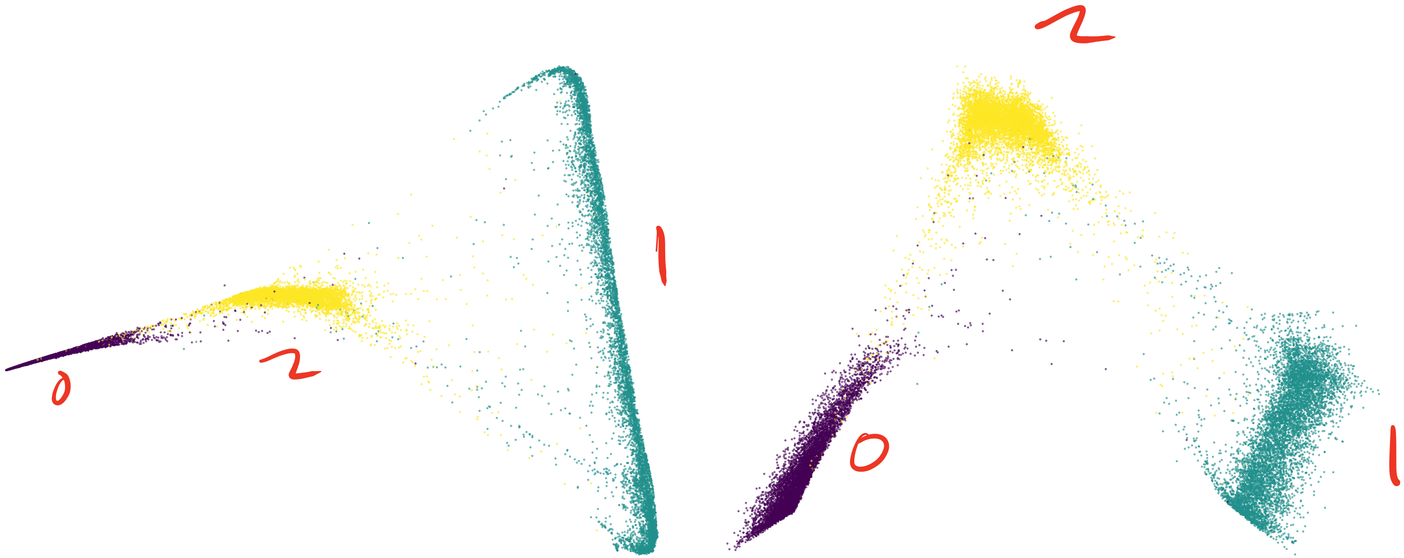

(a) Unnormalized (b) Normalized

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

$W(i,j) \geq 0$

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

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

$P^T = P$

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., $\mathbf{1}^T P \mathbf{1} = 1$.

$$\sum_{i,j=1}^{m} P(i,j) = 1$$

$$\mathbf{1}^T P \mathbf{1} = \frac{1}{2m}\Big(\underbrace{\mathbf{1}^T D^{-1} W \mathbf{1}}_{= m} + \underbrace{\mathbf{1}^T W^T D^{-1} \mathbf{1}}_{m = \mathbf{1}^T D^{-1} W \mathbf{1}}\Big) = 1$$

$$m = \mathbf{1}^T D^{-1} \underbrace{W \mathbf{1}}_{\text{vector of degrees}}$$

= ones vector $\mathbb{1}$

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

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

$\sum_{i,j} Q(i,j) = 1$

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

y_m

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

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

Why Kullbck-Leibler?

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

y_m

- 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

3

The gradient of the Kullback-Leibler divergence

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

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

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

where h is the time-step.

t-SNE embedding of MNIST

1 1

4 7

4 9

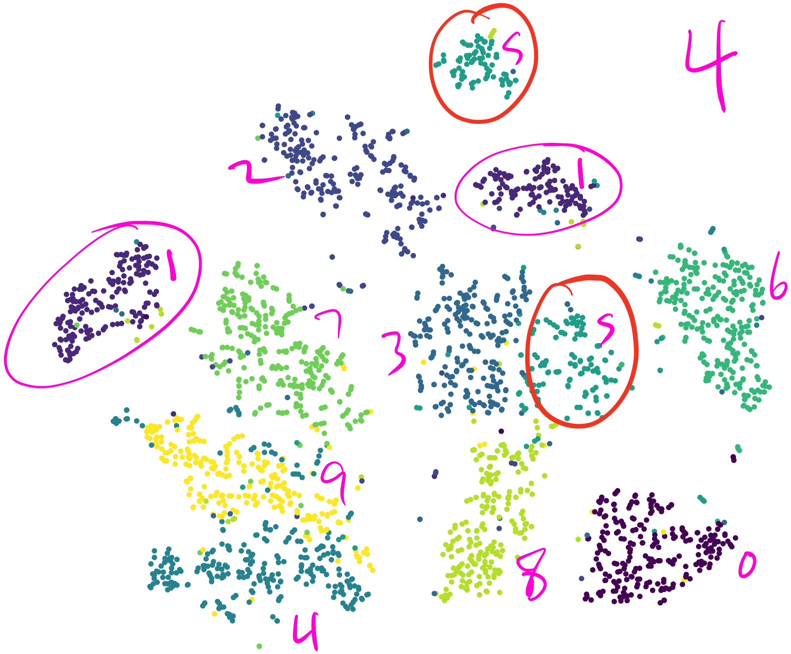

Figure 2: A t-SNE embedding of 2500 images from the MNIST dataset, with colors corresponding to the digit labels of each image.

Gradient computation:

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

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

Can reduce to

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

$$+ \overbrace{\left[\sum_{i \neq j} P(i,j)\right]}^{=1} \log\left(\sum_{\ell \neq s} (1 + \|y_s - y_\ell\|^2)^{-1}\right)$$

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

$$= A + B$$

$\underline{B}$

$$\nabla_{y_\ell} B = \sum_{i \neq j} P(i,j) \nabla_{y_\ell} \log\left(1 + \|y_i - y_j\|^2\right)$$

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

If $\ell = i$, $\nabla_{y_\ell} \|y_i - y_j\|^2 = 2(y_i - y_j)$

If $\ell = j$, $\nabla_{y_\ell} \|y_i - y_j\|^2 = -2(y_i - y_j)$

If $\ell \neq i, j$ $\nabla_{y_\ell} \|y_i - y_j\|^2 = 0$

Kronecker Delta $\delta_{ij} = \begin{cases} 1, & i = j \\ 0, & i \neq j \end{cases}$

$$\nabla_{y_\ell} \|y_i - y_j\|^2 = 2(\delta_{\ell i} - \delta_{\ell j})(y_i - y_j)$$

$$\nabla_{y_\ell} B = \sum_{i \neq j} P(i,j) \underbrace{(1 + \|y_i - y_j\|^2)^{-1}}_{Z Q(i,j)} 2(\delta_{\ell i} - \delta_{\ell j})(y_i - y_j)$$

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

$$Z = \sum_{\ell \neq s} (1 + \|y_s - y_\ell\|^2)^{-1}$$

$$\boxed{Z\, Q(i,j) = (1 + \|y_i - y_j\|^2)^{-1}}$$

$$D_{y_l} B = 2Z \sum_{i \neq j} P(i,j)\,Q(i,j)\,(\delta_{il} - \delta_{jl})(y_i - y_j)$$

$$= 2Z \sum_{i=1}^{n} \sum_{j \neq i} P(i,j)\,Q(i,j)\,\delta_{il}\,(y_i - y_j)$$

$$- 2Z \sum_{j=1}^{n} \sum_{i \neq j} P(i,j)\,Q(i,j)\,\delta_{jl}\,(y_i - y_j)$$

$$= 2Z \sum_{j \neq l} P(l,j)\,Q(l,j)\,(y_l - y_j)$$

$$- 2Z \sum_{i \neq l} \underbrace{P(i,l)}_{P(l,i)}\,\underbrace{Q(i,l)}_{Q(l,i)}\,\underbrace{(y_i - y_l)}_{-(y_l - y_i)}$$

$$\boxed{D_{y_l} B = 4Z \sum_{j \neq l} P(l,j)\,Q(l,j)\,(y_l - y_j)}$$

Attraction

$\swarrow\ l = i$

$$\underbrace{4Z \sum_{j:j\neq i} P(i,j)Q(i,j)(y_i - y_j)}_{\text{Attraction}} \cdot$$

$$A = \log\left(\sum_{l\neq s}\left(1+\|y_s - y_l\|^2\right)^{-1}\right)$$

$$D_{y_l} A = D_{y_l} \log\underbrace{\left(\sum_{l\neq s}\left(1+\|y_s - y_l\|^2\right)^{-1}\right)}_{z}$$

$$= z^{-1} D_{y_l} \sum_{r\neq s}\left(1+\|y_s - y_r\|^2\right)^{-1}$$

$$= -z^{-1} \sum_{r\neq s}\left(1+\|y_s - y_r\|^2\right)^{-2} D_{y_l}\left(1+\|y_s - y_r\|^2\right)$$

$$= -z^{-1} \sum_{r \neq s} \underbrace{\left(1 + \|y_s - y_r\|^2\right)^{-2}}_{z^2 Q(s,r)^2} 2 \left(\delta_{\ell s} - \delta_{\ell r}\right)(y_s - y_r)$$

$$= -2z \sum_{r \neq s} Q(s,r)^2 \left(\delta_{\ell s} - \delta_{\ell r}\right)(y_s - y_r)$$

$$= -2z \sum_{r \neq \ell} Q(\ell,r)^2 (y_\ell - y_r) + 2z \sum_{s \neq \ell} Q(s,\ell)^2 (y_s - y_\ell)$$

$$\boxed{= -4z \sum_{s \neq \ell} Q(s,\ell)^2 (y_\ell - y_s)}$$

Repulsion

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

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 hunderd iterations:

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

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

Early exaggeration

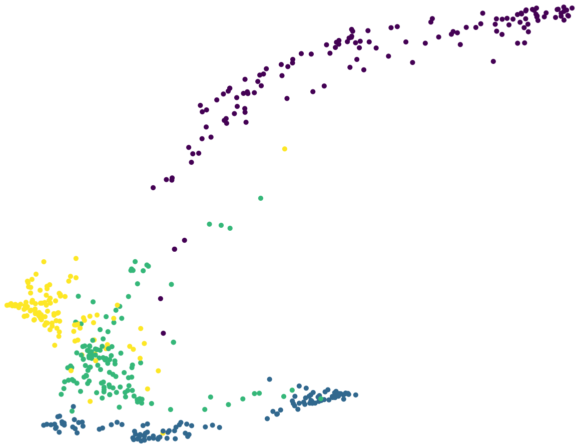

(a) After early exaggeration

(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

MNIST images

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

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

$$f(\sigma_i) = 2^{H(i)} = 10 \quad \forall i$$

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

entropy

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

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

MNIST

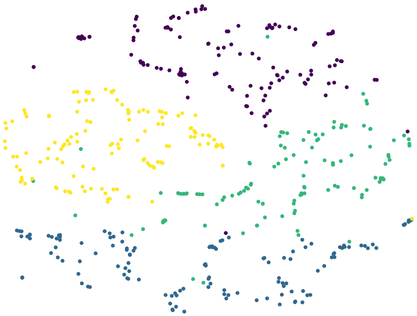

Perplexity 5

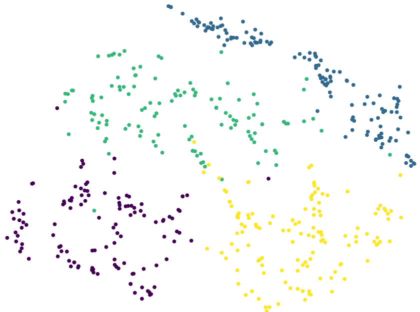

Perplexity 30

Perplexity 50.

Gaussian mixture in 10 dimensions

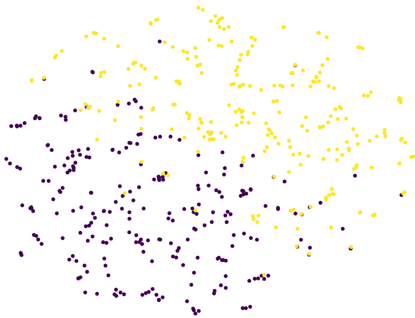

Perplexity 5

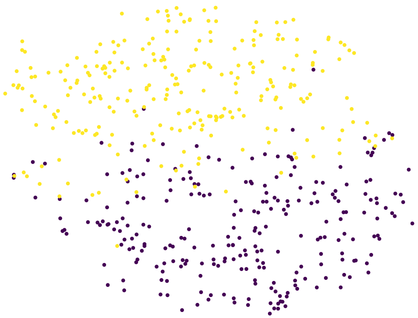

Perplexity 30

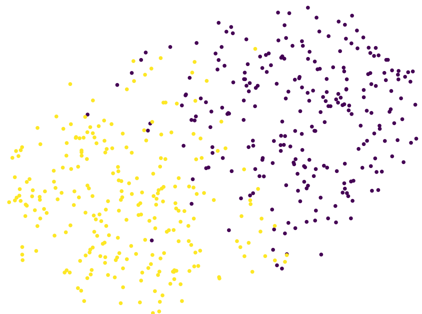

Perplexity 50.

Parabolic curve in 5 dimensions

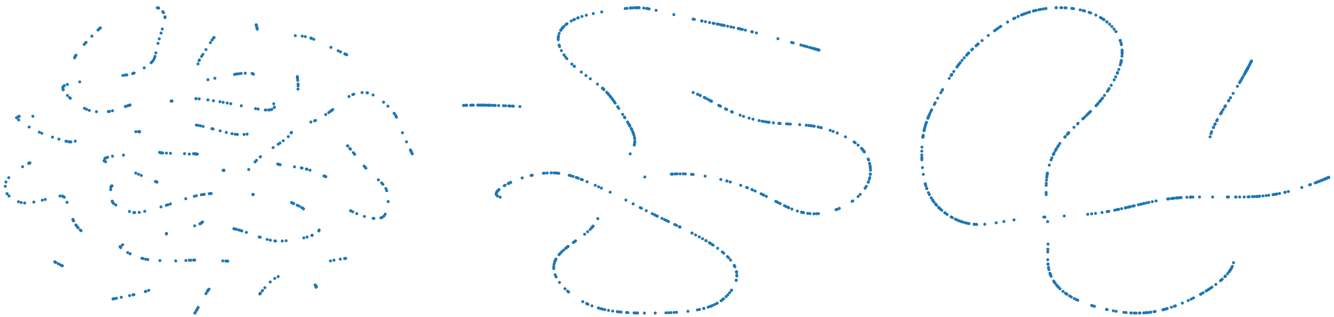

Perplexity 5 Perplexity 30 Perplexity 50.

Graph-based embeddings (.ipynb)