

Lecture 9

Dimensionality Reduction: Johnson-Lindenstrauss Lemma

Course: Algorithms for Big Data

Instructor: Hossein Jowhari

Department of Computer Science and Statistics
Faculty of Mathematics
K. N. Toosi University of Technology

Spring 2021

Distance-Preserving Dimensionality Reduction

Given n vectors $A = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ in \mathbb{R}^d ,

we want a fast transformation $f : \mathbb{R}^d \rightarrow \mathbb{R}^t$ so that

- ▶ (Distances are approximately preserved):

For all pairs $\mathbf{x}, \mathbf{y} \in A$, we have $\|f(\mathbf{x}) - f(\mathbf{y})\| \approx \|\mathbf{x} - \mathbf{y}\|$

- ▶ (The dimension is reduced considerably)

$$t \ll d$$

Such a transformation will be very useful in practice (when the input dimension is large)

Two applications

Faster clustering algorithms: As an example, recall that for the k -center problem we had a 2-factor approximation algorithm with running time $O(nkd)$ (when the points lie in \mathbb{R}^d .)

Using the transformation $f : \mathbb{R}^d \rightarrow \mathbb{R}^t$, we first compute $f(x)$ for all $x \in A$

Then we run the k -center alg. on $A' = \{f(x_1), \dots, f(x_n)\}$

Running time = $\underbrace{\text{time of transformation}} + \underbrace{\text{time of clustering } A'}_{O(nkt)}$

Approximation quality ≈ 2

Approximate nearest neighbor queries: Most exact nearest neighbor data structures have time complexity $n^{O(d)}$

How to reduce the dimension?

- ▶ **A bad idea:** Randomly select a small subset of dimensions. Restrict every \mathbf{x} to the selected dimensions. In other words, choose $S \subseteq \{1, \dots, d\}$ randomly. $f(\mathbf{x}) = \mathbf{x}_S$

Example:

$$\mathbf{x}_1 = \overbrace{(1, 0, \dots, 0)}^d, \mathbf{x}_2 = \overbrace{(0, 1, 0, \dots, 0)}^d, \|\mathbf{x} - \mathbf{y}\| = 1$$

$$S = \{3, 5, 20, 25\}$$

$$f(\mathbf{x}_1) = (0, 0, 0, 0), f(\mathbf{x}_2) = (0, 0, 0, 0), \|f(\mathbf{x}) - f(\mathbf{y})\| = 0$$

- ▶ Dimensionality reduction methods such as PCA do not preserve the pairwise distances.

Johnson-Lindenstrauss Lemma

JL Lemma (existential formulation): Let $\epsilon \in (0, \frac{1}{2})$. Given a set of n vectors $A = \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathbb{R}^d$, there is a mapping $f: \mathbb{R}^d \rightarrow \mathbb{R}^t$ where $t = O(\frac{\log n}{\epsilon^2})$ where

$$\forall x, y \in A, (1 - \epsilon)\|\mathbf{x} - \mathbf{y}\| \leq \|f(\mathbf{x}) - f(\mathbf{y})\| \leq (1 + \epsilon)\|\mathbf{x} - \mathbf{y}\|$$

JL lemma is essentially **tight** with respect to the target dimension t :

Noga Alon has shown that any such mapping requires $t = \Omega(\frac{\log n}{\epsilon^2 \log \frac{1}{\epsilon}})$

Larsen and Nelson have shown that any linear mapping requires $t = \Omega(\frac{\log n}{\epsilon^2})$

Construction of the mapping f : Let \mathbf{M} be a $t \times d$ matrix where every entry M_{ij} is an independent random sample from the normal standard distribution $N(0, 1)$. In other words, each $M_{ij} \sim N(0, 1)$

$$\frac{1}{\sqrt{t}} \underbrace{\begin{bmatrix} M_{11} & \dots & M_{1d} \\ M_{21} & \dots & M_{2d} \\ \vdots & \vdots & \vdots \\ M_{t1} & \dots & M_{td} \end{bmatrix}}_{\mathbf{M}} \underbrace{\begin{bmatrix} x_{11} & x_{12} & \dots & x_{1n} \\ x_{21} & x_{22} & \dots & x_{2n} \\ \vdots & \vdots & \dots & \vdots \\ x_{d1} & x_{d2} & \dots & x_{dn} \end{bmatrix}}_{\mathbf{A}} = \begin{bmatrix} \vdots & \vdots & \vdots \\ f(\mathbf{x}_1) & \vdots & f(\mathbf{x}_n) \\ \vdots & \vdots & \vdots \end{bmatrix}$$

Let $\mathbf{M}^{(i)}$ be the i -th row of \mathbf{M} .

$$f(\mathbf{x}) = \frac{1}{\sqrt{t}} \mathbf{M} \mathbf{x} = \left(\frac{1}{\sqrt{t}} \mathbf{M}^{(1)} \cdot \mathbf{x}, \dots, \frac{1}{\sqrt{t}} \mathbf{M}^{(t)} \cdot \mathbf{x} \right)$$

An example: $n = 5, d = 7, t = 4$

$$\frac{1}{\sqrt{4}} \underbrace{\begin{bmatrix} -0.23 & -0.02 & -0.22 & -0.68 & +0.39 & +0.24 & +0.36 \\ +0.08 & +0.46 & +0.68 & +0.47 & -0.28 & +1.90 & +1.13 \\ +0.89 & -0.24 & +0.83 & +1.92 & -0.47 & +0.10 & +0.33 \\ +0.47 & -1.42 & -1.09 & +2.27 & -0.90 & -0.99 & -0.11 \end{bmatrix}}_M \underbrace{\begin{bmatrix} 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_A = \begin{bmatrix} 0 & -0.11 & -0.24 & -0.38 & -0.08 \\ 0 & +0.04 & +0.62 & +0.71 & +2.23 \\ 0 & +0.44 & +0.74 & +1.46 & +1.68 \\ 0 & +0.23 & -1.01 & -0.33 & -0.88 \end{bmatrix}$$

Lemma 1: If $t \geq \frac{c \log n}{\epsilon^2}$ for a large enough constant c then with probability at least $3/4$ for all pairs $\mathbf{x}, \mathbf{y} \in A$, we have

$$\|f(\mathbf{x}) - f(\mathbf{y})\| \in [(1 - \epsilon)\|\mathbf{x} - \mathbf{y}\|, (1 + \epsilon)\|\mathbf{x} - \mathbf{y}\|]$$

$$\|f(\mathbf{x}) - f(\mathbf{y})\| \approx_{\epsilon} \|\mathbf{x} - \mathbf{y}\|$$

Lets consider a special case:

$A = \{\mathbf{0}, \mathbf{x}\}$ where \mathbf{x} is a unit vector in \mathbb{R}^n . $\|\mathbf{x}\| = 1$.

We want $\|f(\mathbf{x}) - f(\mathbf{0})\| = \|f(\mathbf{x})\| \approx_{\epsilon} 1$.

Lemma 2: Given $\mathbf{x} \in \mathbb{R}^d$ where $\|\mathbf{x}\| = 1$, assuming $t \geq \frac{c \log(\frac{1}{\delta})}{\epsilon^2}$ when c is a large enough constant then we have

$$Pr(\|f(\mathbf{x})\|^2 \in [1 - \epsilon, 1 + \epsilon]) \geq 1 - \delta$$

Before proving Lemma 2, we show Lemma 1 is a consequence of Lemma 2.

Observation 1: Since $f(\mathbf{x}) - f(\mathbf{y}) = f(\mathbf{x} - \mathbf{y})$ (the mapping f is linear) then it is enough to show that for any arbitrary vector $\mathbf{z} \in \mathbb{R}^d$ we have

$$\|f(\mathbf{z})\| \approx_{\epsilon} \|\mathbf{z}\|$$

Observation 2: Let $\mathbf{z}' = \frac{\mathbf{z}}{\|\mathbf{z}\|}$. The vector \mathbf{z}' is a unit vector. If we have $\|f(\mathbf{z}')\| \approx_{\epsilon} 1$ then (by linearity of f) we have

$$\|f(\mathbf{z})\| = \|f(\|\mathbf{z}\|\mathbf{z}')\| = \|\|\mathbf{z}\|f(\mathbf{z}')\| = \|\mathbf{z}\|\|f(\mathbf{z}')\| \approx_{\epsilon} \|\mathbf{z}\|$$

Observation 3: There are $\binom{n}{2}$ pair of vectors in A . From Lemma 2 and the above observations we have for a pair $\mathbf{x}, \mathbf{y} \in A$, if $t \geq \frac{c \log(\frac{1}{\delta})}{\epsilon^2}$ then $\|f(\mathbf{x}) - f(\mathbf{y})\| \approx_{\epsilon} \|\mathbf{x} - \mathbf{y}\|$ with probability $1 - \delta$.

Setting $\delta = \frac{1}{4n^2}$, from the union bound, the statement is true for all pairs in A with probability at least $1 - \binom{n}{2} \frac{1}{4n^2} > 3/4$.

Therefore we get the statement of Lemma 1.

Proof of Lemma 2

Basic facts regarding Gaussian distribution:

- ▶ The Gaussian distribution $N(\mu, \sigma^2)$ with mean μ and variance σ^2 has the following probability density function

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x-\mu}{2\sigma^2}}$$

- ▶ If $X \sim N(\mu_1, \sigma_1^2)$ and $Y \sim N(\mu_2, \sigma_2^2)$ then

$$cX \sim N(c\mu_1, c^2\sigma_1^2)$$

$$X + Y \sim N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$$

Consider $\mathbf{x} \in \mathbb{R}^d$ where $\|\mathbf{x}\| = 1$.

Let Y_i be the i -th coordinate of $\mathbf{M}\mathbf{x}$. Note that

$$f(\mathbf{x}) = \frac{1}{\sqrt{t}}\mathbf{M}\mathbf{x}.$$

Observation: $Y_i = (G_1, \dots, G_d) \cdot \mathbf{x}$ where each G_i is an independent sample from $N(0, 1)$. In other words,

$$Y = G_1x_1 + \dots + G_dx_d$$

Y_i is a linear combination of independent Gaussians. Therefore

$$Y_i \sim N(0, x_1^2 + \dots + x_d^2) = N(0, 1)$$

We need to analyze Y_i^2 since $Y = \|f(\mathbf{x})\|^2 = \frac{1}{t}(Y_1^2 + \dots + Y_t^2)$

$$E[Y_i^2] = \text{Var}[Y_i] + E^2[Y_i] = 1 \Rightarrow E[Y] = 1$$

So, in expectation, $Y = \|f(\mathbf{x})\|$ is exactly 1. Very good but not enough.

We need to bound the probability $Pr(Y > 1 + \epsilon)$.

$Y = \frac{1}{t} \sum_i^t Y_i^2$ is the sum of independent random variables. We could use Chernoff but unfortunately Y_i is not bounded. Still similar ideas that were used in the proof of Chernoff are helpful here. For any $r > 0$

$$Pr(Y > 1 + \epsilon) = Pr(e^{trY} > e^{tr(1+\epsilon)}) \leq \overbrace{\frac{E[e^{trY}]}{e^{tr(1+\epsilon)}}}^{\text{Markov Inequality}} = \prod_{i=1}^t \frac{E[e^{rY_i^2}]}{e^{r(1+\epsilon)}}$$

$$E[e^{rY_i^2}] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{ry^2} e^{-\frac{y^2}{2}} dy = \frac{1}{\sqrt{1-2r}} \quad \text{when } r \leq \frac{1}{2}$$

Therefore for $0 < r < \frac{1}{2}$, we have

$$Pr(Y > 1 + \epsilon) \leq \left(\frac{1}{e^{r(1+\epsilon)} \sqrt{1-2r}} \right)^t$$

Also one can show that $\frac{1}{e^{r(1+\epsilon)} \sqrt{1-2r}} \leq e^{\frac{r^2}{1-2r}}$

Therefore we have

$$Pr(Y > 1 + \epsilon) \leq e^{\frac{tr^2}{1-2r}}$$

We set $r = \frac{\epsilon}{4}$. Using $1 - 2r \geq \frac{1}{2}$ when $\epsilon \leq \frac{1}{2}$, we get

$$Pr(Y > 1 + \epsilon) \leq e^{-\frac{t\epsilon^2}{8}} \leq \frac{\delta}{2} \Rightarrow t \geq \frac{8}{\epsilon^2} \ln\left(\frac{2}{\delta}\right)$$

Similarly we can show

$$Pr(Y < 1 - \epsilon) \leq e^{-\frac{t\epsilon^2}{8}} \leq \frac{\delta}{2}$$

Therefore having $t \geq \frac{8}{\epsilon^2} \ln(\frac{2}{\delta})$ we get

$$Pr(1 - \epsilon \leq Y \leq 1 + \epsilon) \geq 1 - \delta$$

Good news: JL lemma still holds when the Gaussian distribution $N(0, 1)$ is replaced with random $-1, +1$ coefficients.