


Advanced Algorithms

Dimension Reduction

尹一通 Nanjing University, 2022 Fall

Metric Embedding

- Two **metric spaces**: (X, d_X) and (Y, d_Y)



low-distortion: for small $\alpha \geq 1$

$$\forall x_1, x_2 \in X : \quad \frac{1}{\alpha} d_X(x_1, x_2) \leq d_Y(\phi(x_1), \phi(x_2)) \leq \alpha d_X(x_1, x_2)$$

Dimension Reduction

Input: n points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$

Output: n points $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n \in \mathbb{R}^k$ s.t. $\forall 1 \leq i, j \leq n :$

$$(1 - \epsilon) \|\mathbf{x}_i - \mathbf{x}_j\| \leq \|\mathbf{y}_i - \mathbf{y}_j\| \leq (1 + \epsilon) \|\mathbf{x}_i - \mathbf{x}_j\|$$

- Usually we want $k \ll d$.
- How small can k be?
- For what distance $\|\cdot\|$?
- The embedding should be efficiently constructible.

Johnson-Lindenstrauss Theorem/Transformation (JLT)

Johanson-Linenstrauss Theorem

(Johnson-Lindenstrauss 1984)

“In **Euclidian** space, it is always possible to embed a set of n points in *arbitrary* dimension to $O(\log n)$ dimension with constant distortion.”

Theorem (Johnson-Lindenstrauss 1984):

$\forall 0 < \epsilon < 1$, for any set S of n points from \mathbb{R}^d , there is a $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^k$ with $k = O(\epsilon^{-2} \log n)$, such that $\forall \mathbf{x}, \mathbf{y} \in S$:

$$(1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2$$

Johanson-Linenstrauss Theorem

(Johnson-Lindenstrauss 1984)

“In **Euclidian** space, it is always possible to embed a set of n points in *arbitrary* dimension to $O(\log n)$ dimension with constant distortion.”

Theorem (Johnson-Lindenstrauss 1984):

$\forall 0 < \epsilon < 1$, for any set S of n points from \mathbb{R}^d , there is a $A \in \mathbb{R}^{k \times d}$ with $k = O(\epsilon^{-2} \log n)$, such that $\forall \mathbf{x}, \mathbf{y} \in S$:

$$(1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2$$

- The **probabilistic method**: for random $A \in \mathbb{R}^{k \times d}$

$$\Pr \left[\forall \mathbf{x}, \mathbf{y} \in S : (1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \right] = 1 - O\left(\frac{1}{n}\right)$$

w.h.p.

Theorem (Johnson-Lindenstrauss 1984):

$\forall 0 < \epsilon < 1$, for any set S of n points from \mathbb{R}^d , there is a $A \in \mathbb{R}^{k \times d}$ with $k = O(\epsilon^{-2} \log n)$, such that $\forall \mathbf{x}, \mathbf{y} \in S$:

$$(1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2$$

- The **probabilistic method**: for random $A \in \mathbb{R}^{k \times d}$

$$\Pr \left[\forall \mathbf{x}, \mathbf{y} \in S : (1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \right] = 1 - O\left(\frac{1}{n}\right)$$

- **Efficient construction** of random $A \in \mathbb{R}^{k \times d}$:
 - projection onto uniform random k -dimensional subspace; (Johnson-Lindenstrauss; Dasgupta-Gupta)
 - independent Gaussian entries; (Indyk-Motwani)
 - i.i.d. ± 1 entries; (Achlioptas)

Dimension Reduction

Input: n points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$

Output: n points $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n \in \mathbb{R}^k$ s.t. $\forall 1 \leq i, j \leq n :$

$$(1 - \epsilon) \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \leq \|\mathbf{y}_i - \mathbf{y}_j\|_2^2 \leq (1 + \epsilon) \|\mathbf{x}_i - \mathbf{x}_j\|_2^2$$

- for some suitable $k = O(\epsilon^{-2} \log n)$:

J-L Transformation (i.i.d. Gaussian entries):

Entries of $A \in \mathbb{R}^{k \times d}$ are chosen i.i.d. from $\mathcal{N}(0, 1/k)$;
(Gaussian distribution with mean 0 and variance $1/k$)

For $i = 1, 2, \dots, n$: let $\mathbf{y}_i = A\mathbf{x}_i$;

- Gaussian random variable $X \sim \mathcal{N}(\mu, \sigma^2)$:

$$\Pr[X \leq t] = \int_{-\infty}^t \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx$$

$$\begin{aligned} \mathbb{E}[X] &= \mu \\ \mathbf{Var}[X] &= \sigma^2 \end{aligned}$$

Norm Preservation

- $\forall 0 \leq \epsilon \leq 1, \forall$ set S of n points from \mathbb{R}^d
- Random matrix $A \in \mathbb{R}^{k \times d}$ with $k = (\epsilon^{-2} \log n)$:

Johnson-Lindenstrauss Theorem:

With high probability ($\geq 1 - O(1/n)$): $\forall x, y \in S$,

$$(1 - \epsilon) \|x - y\|_2^2 \leq \|Ax - Ay\|_2^2 \leq (1 + \epsilon) \|x - y\|_2^2$$

$$1 - \epsilon \leq \left\| A \frac{(x-y)}{\|x-y\|_2} \right\|_2^2 \leq 1 + \epsilon$$

union bound
over all $O(n^2)$
pairs of $x, y \in S$

unit vector!

for any unit vector $u \in \mathbb{R}^d$:

$$\Pr \left[\left| \|Au\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$

$A \in \mathbb{R}^{k \times d}$: each entry of A is chosen **i.i.d. from** $\mathcal{N}\left(0, \frac{1}{k}\right)$

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$

$$\|A\mathbf{u}\|_2^2 = \sum_{i=1}^k (A\mathbf{u})_i^2 \xrightarrow{\text{linearity of expectation}} \mathbb{E} [\|A\mathbf{u}\|_2^2] = \sum_{i=1}^k \mathbb{E} [(A\mathbf{u})_i^2]$$

$$(A\mathbf{u})_i^2 = \left(\sum_{j=1}^d A_{ij} u_j \right)^2 \quad \text{each } A_{ij} \sim \mathcal{N}\left(0, \frac{1}{k}\right) \text{ i.i.d.}$$

recall:

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

$$\Rightarrow (A\mathbf{u})_i = \sum_{j=1}^d A_{ij} u_j \sim \mathcal{N}\left(0, \frac{\sum_{j=1}^d u_j^2}{k}\right) = \mathcal{N}\left(0, \frac{1}{k}\right)$$

independently!

$A \in \mathbb{R}^{k \times d}$: each entry of A is chosen **i.i.d. from** $\mathcal{N}\left(0, \frac{1}{k}\right)$

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left| \|\mathbf{A}\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$

$$\|\mathbf{A}\mathbf{u}\|_2^2 = \sum_{i=1}^k (\mathbf{A}\mathbf{u})_i^2 \xrightarrow{\text{linearity of expectation}} \mathbb{E} [\|\mathbf{A}\mathbf{u}\|_2^2] = \sum_{i=1}^k \mathbb{E} [(\mathbf{A}\mathbf{u})_i^2]$$

$$(\mathbf{A}\mathbf{u})_i \sim \mathcal{N}\left(0, \frac{1}{k}\right) \text{ i.i.d.}$$

$$\Rightarrow \mathbb{E} [(\mathbf{A}\mathbf{u})_i^2] = \mathbf{Var}[(\mathbf{A}\mathbf{u})_i] + \mathbb{E} [(\mathbf{A}\mathbf{u})_i]^2 = \frac{1}{k}$$

$$\Rightarrow \mathbb{E} [\|\mathbf{A}\mathbf{u}\|_2^2] = \sum_{i=1}^k \mathbb{E} [(\mathbf{A}\mathbf{u})_i^2] = 1$$

$A \in \mathbb{R}^{k \times d}$: each entry of A is chosen **i.i.d. from** $\mathcal{N}\left(0, \frac{1}{k}\right)$


for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$

$$\|A\mathbf{u}\|_2^2 = \sum_{i=1}^k (A\mathbf{u})_i^2$$

$$(A\mathbf{u})_i \sim \mathcal{N}\left(0, \frac{1}{k}\right) \text{ i.i.d.}$$

$$\mathbb{E} \left[\|A\mathbf{u}\|_2^2 \right] = 1$$



for i.i.d. $Y_1, Y_2, \dots, Y_k \sim \mathcal{N}\left(0, \frac{1}{k}\right)$

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] = \Pr \left[\left| \sum_{i=1}^k Y_i^2 - \mathbb{E} \left[\sum_{i=1}^k Y_i^2 \right] \right| > \epsilon \right]$$

$A \in \mathbb{R}^{k \times d}$: each entry of A is chosen **i.i.d. from** $\mathcal{N}\left(0, \frac{1}{k}\right)$

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$

for i.i.d. $Y_1, Y_2, \dots, Y_k \sim \mathcal{N}\left(0, \frac{1}{k}\right)$ **consider** $X_i = \sqrt{k} \cdot Y_i$

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] = \Pr \left[\left| \sum_{i=1}^k Y_i^2 - \mathbb{E} \left[\sum_{i=1}^k Y_i^2 \right] \right| > \epsilon \right]$$

Chernoff bound for χ^2 -distributions:

For independent $X_1, \dots, X_k \in \mathcal{N}(0,1) \implies \begin{cases} \Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \right] < e^{-\epsilon^2 k/8} \\ \Pr \left[\sum_{i=1}^k X_i^2 < (1 - \epsilon)k \right] < e^{-\epsilon^2 k/8} \end{cases}$

$A \in \mathbb{R}^{k \times d}$: each entry of A is chosen **i.i.d. from** $\mathcal{N}\left(0, \frac{1}{k}\right)$

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$

for i.i.d. $X_1, X_2, \dots, X_k \sim \mathcal{N}(0, 1)$

$$\Pr \left[\left| \|A\mathbf{u}\|_2^2 - 1 \right| > \epsilon \right] = \Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \text{ or } \sum_{i=1}^k X_i^2 < (1 - \epsilon)k \right] \\ < \frac{1}{n^3} \text{ for suitable } k = O(\epsilon^{-2} \log n)$$

Chernoff bound for χ^2 -distributions:

$$\text{For independent } X_1, \dots, X_k \in \mathcal{N}(0, 1) \implies \begin{cases} \Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \right] < e^{-\epsilon^2 k / 8} \\ \Pr \left[\sum_{i=1}^k X_i^2 < (1 - \epsilon)k \right] < e^{-\epsilon^2 k / 8} \end{cases}$$

Chernoff bound for χ^2 -distributions:

For independent $X_1, \dots, X_k \in \mathcal{N}(0,1) \Rightarrow \begin{cases} \Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \right] < e^{-\epsilon^2 k/8} \\ \Pr \left[\sum_{i=1}^k X_i^2 < (1 - \epsilon)k \right] < e^{-\epsilon^2 k/8} \end{cases}$

for all $\lambda > 0$: $\Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \right] \leq e^{-(1+\epsilon)\lambda k} \cdot \mathbb{E} \left[e^{\lambda \sum_{i=1}^k X_i^2} \right]$

$X \sim \mathcal{N}(0, 1) \rightarrow$

$$\mathbb{E} \left[e^{sX^2} \right] = \frac{1}{\sqrt{1-2s}}$$

$$\mathbb{E}[e^{sX^2}]$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{sx^2} \cdot e^{-x^2/2} dx$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-(1-2s)x^2/2} dx$$

$$\text{set } y = \sqrt{1-2s} x$$

$$= \frac{1}{\sqrt{1-2s}} \underbrace{\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-y^2/2} dy}_{=1} = \frac{1}{\sqrt{1-2s}}$$

Chernoff bound for χ^2 -distributions:

$$\text{For independent } X_1, \dots, X_k \in \mathcal{N}(0,1) \Rightarrow \begin{cases} \Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \right] < e^{-\epsilon^2 k/8} \\ \Pr \left[\sum_{i=1}^k X_i^2 < (1 - \epsilon)k \right] < e^{-\epsilon^2 k/8} \end{cases}$$

$$\text{for all } \lambda > 0: \Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon)k \right] = \Pr \left[e^{\lambda \sum_{i=1}^k X_i^2} > e^{(1+\epsilon)\lambda k} \right]$$

$$\leq e^{-(1+\epsilon)\lambda k} \cdot \mathbb{E} \left[e^{\lambda \sum_{i=1}^k X_i^2} \right] = e^{-(1+\epsilon)\lambda k} \cdot \prod_{i=1}^k \mathbb{E} \left[e^{\lambda X_i^2} \right]$$

$$X \sim \mathcal{N}(0,1) \rightarrow$$

$$\mathbb{E} \left[e^{sX^2} \right] = \frac{1}{\sqrt{1-2s}}$$

$$\leq e^{-\epsilon\lambda k} \left(\frac{e^{-\lambda}}{\sqrt{1-2\lambda}} \right)^k$$

$$\leq e^{-\epsilon\lambda k + 2\lambda^2 k} \quad \text{for } \lambda < 1/4$$

$$= e^{-\epsilon^2 k/8} \quad \text{choosing } \lambda = \epsilon/4$$

Johanson-Linenstrauss Theorem

(Johnson-Lindenstrauss 1984)

“In **Euclidian** space, it is always possible to embed a set of n points in *arbitrary* dimension to $O(\log n)$ dimension with constant distortion.”

Theorem (Johnson-Lindenstrauss 1984):

$\forall 0 < \epsilon < 1$, for any set S of n points from \mathbb{R}^d , there is a $A \in \mathbb{R}^{k \times d}$ with $k = O(\epsilon^{-2} \log n)$, such that $\forall \mathbf{x}, \mathbf{y} \in S$:

$$(1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2$$

- The **probabilistic method**: for random $A \in \mathbb{R}^{k \times d}$

$$\Pr \left[\forall \mathbf{x}, \mathbf{y} \in S : (1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \right] = 1 - O\left(\frac{1}{n}\right)$$

w.h.p.

Theorem (Johnson-Lindenstrauss 1984):

$\forall 0 < \epsilon < 1$, for any set S of n points from \mathbb{R}^d , there is a $A \in \mathbb{R}^{k \times d}$ with $k = O(\epsilon^{-2} \log n)$, such that $\forall \mathbf{x}, \mathbf{y} \in S$:

$$(1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2$$

- The **probabilistic method**: for random $A \in \mathbb{R}^{k \times d}$

$$\Pr \left[\forall \mathbf{x}, \mathbf{y} \in S : (1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \right] = 1 - O\left(\frac{1}{n}\right)$$

- **Efficient construction** of random $A \in \mathbb{R}^{k \times d}$:

- projection onto uniform random k -dimensional subspace; (Johnson-Lindenstrauss; Dasgupta-Gupta)
- independent Gaussian entries; (Indyk-Motwani)
- i.i.d. $-1/+1$ entries; (Achlioptas)

Dimension Reduction

Input: n points $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d$

Output: n points $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n \in \mathbb{R}^k$ s.t. $\forall 1 \leq i, j \leq n :$

$$(1 - \epsilon) \|\mathbf{x}_i - \mathbf{x}_j\|_2^2 \leq \|\mathbf{y}_i - \mathbf{y}_j\|_2^2 \leq (1 + \epsilon) \|\mathbf{x}_i - \mathbf{x}_j\|_2^2$$


- for some suitable $k = O(\epsilon^{-2} \log n)$:

J-L Transformation (uniform k -dim subspace):

The k rows A_1, \dots, A_k of $A \in \mathbb{R}^{k \times d}$ are orthogonal unit vectors $\in \mathbb{R}^d$ chosen uniform at random;

For $i = 1, 2, \dots, n$: let $\mathbf{y}_i = \sqrt{\frac{d}{k}} A \mathbf{x}_i$;


- $A \in \mathbb{R}^{k \times d}$: projection onto a uniform k -dimensional subspace



$A \in \mathbb{R}^{k \times d}$: projection onto uniform random k -dim subspace

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left| \left\| \sqrt{\frac{d}{k}} A \mathbf{u} \right\|_2^2 - 1 \right| > \epsilon \right] < \frac{1}{n^3}$$




$$\Pr \left[\left\| A \mathbf{u} \right\|_2^2 > (1 + \epsilon) \frac{k}{d} \text{ or } \left\| A \mathbf{u} \right\|_2^2 < (1 - \epsilon) \frac{k}{d} \right] < \frac{1}{n^3}$$

$A \in \mathbb{R}^{k \times d}$: projection onto uniform random k -dim subspace

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left\| A\mathbf{u} \right\|_2^2 > (1 + \epsilon) \frac{k}{d} \right] < \frac{1}{2n^3}$$

$$\Pr \left[\left\| A\mathbf{u} \right\|_2^2 < (1 - \epsilon) \frac{k}{d} \right] < \frac{1}{2n^3}$$



“inner-products are invariant under rotations”

$A \in \mathbb{R}^{k \times d}$: projection onto uniform random k -dim subspace

for any **unit vector** $\mathbf{u} \in \mathbb{R}^d$:

$$\Pr \left[\left\| A\mathbf{u} \right\|_2^2 > (1 + \epsilon) \frac{k}{d} \right] < \frac{1}{2n^3}$$


$$\Pr \left[\left\| A\mathbf{u} \right\|_2^2 < (1 - \epsilon) \frac{k}{d} \right] < \frac{1}{2n^3}$$

uniform random **unit vector** $\in \mathbb{R}^d$:

$$Y = (Y_1, \dots, Y_k, Y_{k+1}, \dots, Y_d)$$

fixed k -dimensional subspace:

$$Z = (Y_1, \dots, Y_k)$$



$\|A\mathbf{u}\|$ is **identically distributed** as $\|Z\|$

uniform random unit vector $\in \mathbb{R}^d$: $Y = (Y_1, \dots, Y_d)$

sample $X = (X_1, \dots, X_d) \in \mathbb{R}^d$ where $X_i \sim \mathcal{N}(0,1)$ i.i.d.;

$$\text{let } Y = \frac{X}{\|X\|};$$

density: $\Pr[X = x] = \prod_{i=1}^d \frac{1}{\sqrt{2\pi}} e^{-x_i^2/2} = (2\pi)^{-d/2} e^{-\|x\|_2^2/2}$

Spherically symmetric!

for some suitable $k = O(\varepsilon^{-2} \log n)$:

$$\Pr \left[\sum_{i=1}^k Y_i^2 > (1 + \epsilon) \frac{k}{d} \right] < \frac{1}{2n^3}$$


$$\Pr \left[\sum_{i=1}^k Y_i^2 < (1 - \epsilon) \frac{k}{d} \right] < \frac{1}{2n^3}$$

i.i.d. **Gaussian** random variables $X_1, X_2, \dots, X_d \sim \mathcal{N}(0,1)$

for some suitable $k = O(\varepsilon^{-2} \log n)$:

$$\Pr \left[\sum_{i=1}^k X_i^2 > (1 + \epsilon) \frac{k}{d} \sum_{i=1}^d X_i^2 \right] < \frac{1}{2n^3}$$

$$\Pr \left[\sum_{i=1}^k X_i^2 < (1 - \epsilon) \frac{k}{d} \sum_{i=1}^d X_i^2 \right] < \frac{1}{2n^3}$$



$$\Pr \left[(d - (1 + \epsilon)k) \sum_{i=1}^k X_i^2 - (1 + \epsilon)k \sum_{i=k+1}^d X_i^2 > 0 \right] < \frac{1}{2n^3}$$

$$\Pr \left[(d - (1 - \epsilon)k) \sum_{i=1}^k X_i^2 - (1 - \epsilon)k \sum_{i=k+1}^d X_i^2 < 0 \right] < \frac{1}{2n^3}$$

$$X_1, X_2, \dots, X_d \sim \mathcal{N}(0,1)$$

n):

$$\epsilon)k \sum_{i=k+1}^d X_i^2 > 0 \Big] < \frac{1}{2n^3}$$

$$\epsilon)k \sum_{i=k+1}^d X_i^2 < 0 \Big] < \frac{1}{2n^3}$$

$$\begin{aligned} E[e^{sX^2}] &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{sx^2} \cdot e^{-x^2/2} dx \end{aligned}$$

$$= \frac{1}{\sqrt{2\pi}} \int e^{-(1-2s)x^2/2} dx$$

$$\text{set } y = \sqrt{1-2s} x$$

$$= \frac{1}{\sqrt{1-2s}} \underbrace{\frac{1}{\sqrt{2\pi}} \int e^{-y^2/2} dy}_{=1} = \frac{1}{\sqrt{1-2s}}$$

$$\text{red square} = \Pr \left[\exp \left\{ \lambda \left(((1-\epsilon)k - d) \sum_{i=1}^k X_i^2 + (1-\epsilon)k \sum_{i=k+1}^d X_i^2 \right) \right\} > 1 \right] \quad (\text{arbitrary } \lambda > 0)$$

$$\leq \mathbb{E} \left[\exp \left\{ \lambda \left(((1-\epsilon)k - d) \sum_{i=1}^k X_i^2 + (1-\epsilon)k \sum_{i=k+1}^d X_i^2 \right) \right\} \right] \quad (\text{Markov's inequality})$$

$$= \prod_{i=1}^k \mathbb{E} \left[e^{\lambda((1-\epsilon)k-d)X_i^2} \right] \prod_{i=k+1}^d \mathbb{E} \left[e^{\lambda(1-\epsilon)kX_i^2} \right]$$

i.i.d. **Gaussian** random variables $X_1, X_2, \dots, X_d \sim \mathcal{N}(0,1)$

for some suitable $k = O(\epsilon^{-2} \log n)$:

$$\Pr \left[(d - (1 + \epsilon)k) \sum_{i=1}^k X_i^2 - (1 + \epsilon)k \sum_{i=k+1}^d X_i^2 > 0 \right] < \frac{1}{2n^3}$$

$$\Pr \left[(d - (1 - \epsilon)k) \sum_{i=1}^k X_i^2 - (1 - \epsilon)k \sum_{i=k+1}^d X_i^2 < 0 \right] < \frac{1}{2n^3}$$

$$\blacksquare = \Pr \left[\exp \left\{ \lambda \left(((1 - \epsilon)k - d) \sum_{i=1}^k X_i^2 + (1 - \epsilon)k \sum_{i=k+1}^d X_i^2 \right) \right\} > 1 \right] \quad (\text{arbitrary } \lambda > 0)$$

$$\leq (1 - 2\lambda((1 - \epsilon)k - d))^{-\frac{k}{2}} (1 - 2\lambda(1 - \epsilon)k)^{-\frac{d-k}{2}}$$

$$\leq (1 - \epsilon)^{-\frac{k}{2}} \left(1 - \frac{\epsilon k}{d - k} \right)^{\frac{d-k}{2}} \leq \exp \left(-\frac{\epsilon^2 k}{4} \right) \quad \begin{array}{l} \text{set:} \\ \lambda = \frac{\epsilon}{2(1 - \epsilon)(d - (1 - \epsilon)k)} \end{array}$$

Theorem (Johnson-Lindenstrauss 1984):

$\forall 0 < \epsilon < 1$, for any set S of n points from \mathbb{R}^d , there is a $A \in \mathbb{R}^{k \times d}$ with $k = O(\epsilon^{-2} \log n)$, such that $\forall \mathbf{x}, \mathbf{y} \in S$:

$$(1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2$$

- The **probabilistic method**: for random $A \in \mathbb{R}^{k \times d}$

$$\Pr \left[\forall \mathbf{x}, \mathbf{y} \in S : (1 - \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \leq \|A\mathbf{x} - A\mathbf{y}\|_2^2 \leq (1 + \epsilon) \|\mathbf{x} - \mathbf{y}\|_2^2 \right] = 1 - O\left(\frac{1}{n}\right)$$

- **Efficient construction** of random $A \in \mathbb{R}^{k \times d}$:
 - projection onto uniform random k -dimensional subspace; (Johnson-Lindenstrauss; Dasgupta-Gupta)
 - independent Gaussian entries; (Indyk-Motwani)
 - i.i.d. $-1/+1$ entries; (Achlioptas)

Nearest Neighbor Search (NNS)


Nearest Neighbor Search (NNS)

- Metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$

Query: a point $x \in X$

Find the datapoint y_i that is **closest** to x .



Applications in:

- database systems
- pattern matching
- machine learning
- image processing
- bioinformatics
-


Nearest Neighbor Search (NNS)

Data: n points $y_1, y_2, \dots, y_n \in [N]^d$


Query: a point $x \in [N]^d$

Find the datapoint y_i that is **closest** to x .

when **dimension** d is small:



k-d tree



Voronoi diagram

Nearest Neighbor Search (NNS)

- Hamming space $\{0,1\}^d$:

Data: n points $y_1, y_2, \dots, y_n \in \{0,1\}^d$

Query: a point $x \in \{0,1\}^d$

Find the datapoint y_i that is **closest** to x .

when **dimension** d is high:

say $d \gg \log n$

Curse of dimensionality:

It is conjectured that to solve NNS in high dimension requires either super-polynomial(n) space or super-polynomial(d) time.

Blessing: randomization + approximation

Approximate Near Neighbor (ANN)

- Metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$


Query: a point $x \in X$

c -ANN (*A*pproximate *N*earest *N*eighbor):

Find a y_i such that $\text{dist}(x, y_i) \leq c \cdot \min_{1 \leq j \leq n} \text{dist}(x, y_j)$

(c, r) -ANN (*A*pproximate *N*ear *N*eighbor):

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$
“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$
arbitrary if otherwise



$$r_0 = D_{\min} = \min_{1 \leq i < j \leq n} \text{dist}(y_i, y_j)$$

$$r_k = \sqrt{c} \cdot r_{k-1}$$

$$r_{\log_c(D_{\max}/D_{\min})} = D_{\max} = \max_{1 \leq i < j \leq n} \text{dist}(y_i, y_j)$$

- Metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$

Query: a point $x \in X$

c -ANN (**A**pproximate **N**earest **N**eighbor):


Find a y_i such that $\text{dist}(x, y_i) \leq c \cdot \min_{1 \leq j \leq n} \text{dist}(x, y_j)$

(c, r) -ANN (**A**pproximate **N**ear **N**eighbor):

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$
 “no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$
 arbitrary if otherwise

$$\text{let } R = \frac{D_{\max}}{D_{\min}} \quad \begin{aligned} D_{\max} &= \max_{1 \leq i < j \leq n} \text{dist}(y_i, y_j) \\ D_{\min} &= \min_{1 \leq i < j \leq n} \text{dist}(y_i, y_j) \end{aligned}$$

$\forall r$: (\sqrt{c}, r) -ANN
 can be solved with **space** s
 and **query time** t



c -ANN can be solved
 within space $O(s \log_c R)$
 and query time $O(t \log \log_c R)$

- **Hamming space** $\{0,1\}^d$:

Data: n points $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_n \in \{0,1\}^d$

Query: a point $\mathbf{x} \in \{0,1\}^d$

(c, r) -ANN (*A*pproximate *N*ear *N*eighbor):

return a \mathbf{y}_i that $\text{dist}(\mathbf{x}, \mathbf{y}_i) \leq c \cdot r$ if $\exists \mathbf{y}_j$ s.t. $\text{dist}(\mathbf{x}, \mathbf{y}_j) \leq r$

answer “no” if $\forall \mathbf{y}_i, \text{dist}(\mathbf{x}, \mathbf{y}_i) > c \cdot r$

arbitrary if otherwise

- **High dimension:** $d \gg \log n$

Dimension Reduction:

Let k, p and s to be fixed later;

sample $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

for $i = 1, 2, \dots, n$: let $\mathbf{z}_i = A\mathbf{y}_i \in \{0,1\}^k$ on **finite field** $\text{GF}(2)$;

store all **s-balls** $B_s(\mathbf{u}) = \{\mathbf{y}_i \mid \text{dist}(\mathbf{u}, \mathbf{z}_i) \leq s\}$ for all $\mathbf{u} \in \{0,1\}^k$;

$$z_i(j) = (A\mathbf{y}_i)_j = \left(\sum_{\ell=1}^d A_{j\ell} \mathbf{y}_i(\ell) \right) \bmod 2$$

- **Hamming space** $\{0,1\}^d$:

Data: n points $y_1, y_2, \dots, y_n \in \{0,1\}^d$

Query: a point $x \in \{0,1\}^d$

- **High dimension:** $d \gg \log n$

Dimension Reduction:

Let k, p and s to be fixed later;

sample $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

for $i = 1, 2, \dots, n$: let $z_i = Ay_i \in \{0,1\}^k$ on **finite field** $\text{GF}(2)$;

store all **s-balls** $B_s(u) = \{y_i \mid \text{dist}(u, z_i) \leq s\}$ for all $u \in \{0,1\}^k$;

To answer **query** $x \in \{0,1\}^d$:
 retrieve $B_s(Ax)$;
 if $B_s(Ax) = \emptyset$ return “no”
 else return any $y_i \in B_s(Ax)$

$$z_i(j) = (Ay_i)_j = \left(\sum_{\ell=1}^d A_{j\ell} y_i(\ell) \right) \bmod 2$$

space: $O(n2^k)$ **query time:** $O(kd)$ computation + $O(1)$ memory access

- **Hamming space** $\{0,1\}^d$:

Data: n points $y_1, y_2, \dots, y_n \in \{0,1\}^d$

Query: a point $x \in \{0,1\}^d$

- **High dimension:** $d \gg \log n$

Dimension Reduction:

Let k, p and s to be fixed later;

sample $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

for $i = 1, 2, \dots, n$: let $z_i = Ay_i \in \{0,1\}^k$ on **finite field** GF(2);

$$z_i(j) = (Ay_i)_j = \left(\sum_{\ell=1}^d A_{j\ell} y_i(\ell) \right) \bmod 2$$

for suitable $k = O(\log n)$, p and s :

$$\begin{aligned} \forall x, y \in \{0,1\}^d: \quad \text{dist}(x, y) \leq r &\quad \Rightarrow \quad \Pr[\text{dist}(Ax, Ay) > s] < 1/n^2 \\ \text{dist}(x, y) > c \cdot r &\quad \Rightarrow \quad \Pr[\text{dist}(Ax, Ay) \leq s] < 1/n^2 \end{aligned}$$

union bound

(c, r) -**ANN** is solved w.h.p.

random $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

computation on $\text{GF}(2)$: $(A\mathbf{x})_i = \left(\sum_{j=1}^d A_{ij} x_j \right) \bmod 2$

for suitable p and s :

$$\begin{aligned} \forall \mathbf{x}, \mathbf{y} \in \{0,1\}^d: \quad \text{dist}(\mathbf{x}, \mathbf{y}) \leq r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] < e^{-\Omega(k)} \\ \text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] < e^{-\Omega(k)} \end{aligned}$$

row vector $A_{i\cdot}$: *i.i.d.* entries $\in \text{Bernoulli}(p)$

$$\Pr[(A\mathbf{x})_i \neq (A\mathbf{y})_i] = \Pr[\langle A_{i\cdot}, \mathbf{x} \rangle \neq \langle A_{i\cdot}, \mathbf{y} \rangle] = \frac{1}{2} \left(1 - (1 - 2p)^{\text{dist}(\mathbf{x}, \mathbf{y})} \right)$$

Why?

For **uniform** $\mathbf{u} \in \{0,1\}^d$: $\Pr[\langle \mathbf{u}, \mathbf{x} \rangle \neq \langle \mathbf{u}, \mathbf{y} \rangle] = \frac{1}{2}$

generate $A_{i\cdot}$ as:

1. each $j \in [d]$ joins $D \subseteq [d]$ independently with probability $2p$;
2. for each $j \in D$: samples a uniform and independent $A_{ij} \in \{0,1\}$;
3. for each $j \notin D$: $A_{ij} = 0$;

$A_{i\cdot}$ restricted on D is a uniform Boolean vector!

random $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

computation on GF(2): $(A\mathbf{x})_i = \left(\sum_{j=1}^d A_{ij} x_j \right) \bmod 2$

for suitable p and s :

$$\begin{aligned} \forall \mathbf{x}, \mathbf{y} \in \{0,1\}^d: \quad \text{dist}(\mathbf{x}, \mathbf{y}) \leq r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] < e^{-\Omega(k)} \\ \text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] < e^{-\Omega(k)} \end{aligned}$$

row vector $A_{i\cdot}$: *i.i.d.* entries $\in \text{Bernoulli}(p)$

$$\Pr[(A\mathbf{x})_i \neq (A\mathbf{y})_i] = \Pr[\langle A_{i\cdot}, \mathbf{x} \rangle \neq \langle A_{i\cdot}, \mathbf{y} \rangle] = \frac{1}{2} \left(1 - (1 - 2p)^{\text{dist}(\mathbf{x}, \mathbf{y})} \right)$$

choose p to satisfy $(1-2p) = 2^{-1/r}$

$$\text{dist}(\mathbf{x}, \mathbf{y}) \leq r \Rightarrow \Pr[(A\mathbf{x})_i \neq (A\mathbf{y})_i] \leq 1/4$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r \Rightarrow \Pr[(A\mathbf{x})_i \neq (A\mathbf{y})_i] > 1/2 - 2^{-(c+1)}$$

random $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

computation on GF(2): $(A\mathbf{x})_i = \left(\sum_{j=1}^d A_{ij} x_j \right) \bmod 2$

for suitable p and s :

$$\begin{aligned} \forall \mathbf{x}, \mathbf{y} \in \{0,1\}^d: \quad \text{dist}(\mathbf{x}, \mathbf{y}) \leq r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] < e^{-\Omega(k)} \\ \text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] < e^{-\Omega(k)} \end{aligned}$$

choose p to satisfy $(1-2p) = 2^{-1/r}$

$$\text{dist}(\mathbf{x}, \mathbf{y}) \leq r \Rightarrow \Pr[(A\mathbf{x})_i \neq (A\mathbf{y})_i] \leq 1/4$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r \Rightarrow \Pr[(A\mathbf{x})_i \neq (A\mathbf{y})_i] > 1/2 - 2^{-(c+1)}$$

$$\text{dist}(A\mathbf{x}, A\mathbf{y}) = X = \sum_{i=1}^k X_i \quad \text{where } X_i = \begin{cases} 1 & \text{if } (A\mathbf{x})_i \neq (A\mathbf{y})_i \\ 0 & \text{otherwise} \end{cases}$$

independent trials

random $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

computation on GF(2): $(A\mathbf{x})_i = \left(\sum_{j=1}^d A_{ij} x_j \right) \bmod 2$

for suitable p and s :

$$\begin{aligned} \forall \mathbf{x}, \mathbf{y} \in \{0,1\}^d: \quad \text{dist}(\mathbf{x}, \mathbf{y}) \leq r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] < e^{-\Omega(k)} \\ \text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] < e^{-\Omega(k)} \end{aligned}$$

choose p to satisfy $(1-2p) = 2^{-1/r}$

$$\text{dist}(\mathbf{x}, \mathbf{y}) \leq r \Rightarrow \Pr[X_i = 1] \leq 1/4 \Rightarrow \mathbf{E}[X] \leq k/4$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r \Rightarrow \Pr[X_i = 1] > 1/2 - 2^{-(c+1)} \Rightarrow \mathbf{E}[X] \leq (1/2 - 2^{-(c+1)})k$$

$$\text{dist}(A\mathbf{x}, A\mathbf{y}) = X = \sum_{i=1}^k X_i \quad \text{where} \quad X_i = \begin{cases} 1 & \text{if } (A\mathbf{x})_i \neq (A\mathbf{y})_i \\ 0 & \text{otherwise} \end{cases}$$

$$\text{choose } s = (1/4 + 1/2 - 2^{-(c+1)})k/2 = (3/8 - 2^{-(c+2)})k$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) \leq r \Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] \leq \Pr[X > \mathbf{E}X + (1/8 - 2^{-(c+2)})k]$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r \Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] \leq \Pr[X < \mathbf{E}X - (1/8 - 2^{-(c+2)})k]$$

Chernoff-Hoeffding Bound

Chernoff Bound:

For $X = \sum_{i=1}^n X_i$, where $X_1, \dots, X_n \in \{0,1\}$ are *independent*
(or *negatively associated*),

for any $t > 0$:

$$\Pr [X \geq \mathbb{E}[X] + t] \leq \exp \left(-\frac{2t^2}{n} \right)$$

$$\Pr [X \leq \mathbb{E}[X] - t] \leq \exp \left(-\frac{2t^2}{n} \right)$$

random $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

computation on GF(2): $(A\mathbf{x})_i = \left(\sum_{j=1}^d A_{ij} x_j \right) \bmod 2$

for suitable p and s :

$$\begin{aligned} \forall \mathbf{x}, \mathbf{y} \in \{0,1\}^d: \quad \text{dist}(\mathbf{x}, \mathbf{y}) \leq r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] < e^{-\Omega(k)} \\ \text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r &\Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] < e^{-\Omega(k)} \end{aligned}$$

choose p to satisfy $(1-2p) = 2^{-1/r}$

$$\text{dist}(\mathbf{x}, \mathbf{y}) \leq r \Rightarrow \Pr[X_i = 1] \leq 1/4 \Rightarrow \mathbf{E}[X] \leq k/4$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r \Rightarrow \Pr[X_i = 1] > 1/2 - 2^{-(c+1)} \Rightarrow \mathbf{E}[X] \leq (1/2 - 2^{-(c+1)})k$$

$$\text{dist}(A\mathbf{x}, A\mathbf{y}) = X = \sum_{i=1}^k X_i \quad \text{where} \quad X_i = \begin{cases} 1 & \text{if } (A\mathbf{x})_i \neq (A\mathbf{y})_i \\ 0 & \text{otherwise} \end{cases}$$

$$\text{choose } s = (1/4 + 1/2 - 2^{-(c+1)})k/2 = (3/8 - 2^{-(c+2)})k$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) \leq r \Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) > s] \leq \Pr[X > \mathbf{E}X + (1/8 - 2^{-(c+2)})k] < \exp(-2(1/8 - 2^{-(c+2)})^2 k)$$

$$\text{dist}(\mathbf{x}, \mathbf{y}) > c \cdot r \Rightarrow \Pr[\text{dist}(A\mathbf{x}, A\mathbf{y}) \leq s] \leq \Pr[X < \mathbf{E}X - (1/8 - 2^{-(c+2)})k] < \exp(-2(1/8 - 2^{-(c+2)})^2 k)$$

- **Hamming space** $\{0,1\}^d$:

Data: n points $y_1, y_2, \dots, y_n \in \{0,1\}^d$

Query: a point $x \in \{0,1\}^d$

- **High dimension:** $d \gg \log n$

Dimension Reduction:

Let $k = \frac{\ln n}{(1/8 - 2^{-(c+2)})^2}$, $p = \frac{1}{2} - 2^{-1-1/r}$, $s = (\frac{3}{8} - 2^{-(c+2)})k$;

sample $k \times d$ Boolean matrix A with *i.i.d.* entries $\in \text{Bernoulli}(p)$;

for $i = 1, 2, \dots, n$: let $z_i = Ay_i \in \{0,1\}^k$ on **finite field** $\text{GF}(2)$;

store all **s-balls** $B_s(u) = \{y_i \mid \text{dist}(u, z_i) \leq s\}$ for all $u \in \{0,1\}^k$;

To answer **query** $x \in \{0,1\}^d$:
 retrieve $B_s(Ax)$;
 if $B_s(Ax) = \emptyset$ return “no”
 else return any $y_i \in B_s(Ax)$

space: $n^{O(1)}$

query time: $O(d \ln n)$

solve (c,r) -**ANN** w.h.p.

Locality Sensitive Hashing (LSH)

Locality-Sensitive Hashing (LSH)

(Indyk-Motwani 1998)

- Metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$

Query: a point $x \in X$

(c, r) -ANN (*A*pproximate *N*ear *N*eighbor):

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$

“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$

arbitrary if otherwise

Locality-sensitive hashing (LSH):

A random $h : X \rightarrow U$ is an (r, cr, p, q) -LSH if $\forall x, y \in X$:

$$\text{dist}(x, y) \leq r \implies \Pr[h(x) = h(y)] \geq p$$

$$\text{dist}(x, y) > c \cdot r \implies \Pr[h(x) = h(y)] \leq q$$

Locality-Sensitive Hashing (LSH)

- Metric space (X, dist) :

Locality-sensitive hashing (LSH):

A random $h : X \rightarrow U$ is an (r, cr, p, q) -LSH if $\forall x, y \in X$:

$$\text{dist}(x, y) \leq r \implies \Pr[h(x) = h(y)] \geq p$$

$$\text{dist}(x, y) > c \cdot r \implies \Pr[h(x) = h(y)] \leq q$$

Proposition:

$$\begin{array}{ccc} \exists \text{ an } (r, cr, p, q)\text{-LSH} & \implies & \exists \text{ an } (r, cr, p^k, q^k)\text{-LSH} \\ h : X \rightarrow U & & h : X \rightarrow U^k \end{array}$$

draw **independent** h_1, \dots, h_k according to distribution of h

$$g(x) = (h_1(x), h_2(x), \dots, h_k(x)) \in U^k$$

- Metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$

Query: a point $x \in X$

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$

“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$

suppose we have $(r, cr, p^*, 1/n)$ -LSH $g: X \rightarrow U$

$\forall x, y \in X: \text{dist}(x, y) \leq r \implies \Pr[g(x) = g(y)] \geq p^*$

$\text{dist}(x, y) > c \cdot r \implies \Pr[g(x) = g(y)] \leq 1/n$

Data structure: sort y_1, y_2, \dots, y_n according to $g(y_i)$;

Upon **query** $x \in X$: find all y_i that $g(x) = g(y_i)$ by binary search;

if encounter a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ return this y_i ;

else return “no”;

if real answer is “no”: always correct

if real answer is not “no”: correct with probability $\geq p^*$

space: $O(n)$ time: $O(\log n) + O(1)$ in expectation

- Metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$

Query: a point $x \in X$

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$

“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$

suppose we have $(r, cr, p^*, 1/n)$ -LSH $g: X \rightarrow U$

Hash functions: *i.i.d.* instances g_1, \dots, g_ℓ of g , where $\ell = 1/p^*$;

Data structure: ℓ sorted tables;

in table- j : sort y_1, y_2, \dots, y_n according to $g_j(y_i)$

Upon **query** $x \in X$:

find $\leq 10\ell$ such y_i that $\exists j, g_j(x) = g_j(y_i)$ by binary search;

if encounter a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ return this y_i ;

else return “no”;

metric space (X, dist) $(r, cr, p^*, 1/n)$ -LSH $g: X \rightarrow U$

Data: n points $y_1, y_2, \dots, y_n \in X$ **Query:** $x \in X$

Hash functions: *i.i.d.* instances g_1, \dots, g_ℓ of g , where $\ell = 1/p^*$;

Data structure: ℓ sorted tables;

in table- j : sort y_1, y_2, \dots, y_n according to $g_j(y_i)$

Upon **query** $x \in X$:

find $\leq 10\ell$ such y_i that $\exists j, g_j(x) = g_j(y_i)$ by binary search;

if encounter a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ return this y_i ;

else return “no”;

space: $O(n\ell) = O(n/p^*)$ time: $O(\ell \cdot \log n) = O((\log n) / p^*)$

if real answer is “no” : $\forall y_i, \text{dist}(x, y_i) > c \cdot r \Rightarrow$ always correct

if $\exists y_s$ s.t. $\text{dist}(x, y_s) \leq r$

$\Pr[\text{answer “no”}] \leq ?$

metric space (X, dist) $(r, cr, p^*, 1/n)$ -LSH $g: X \rightarrow U$

Data: n points $y_1, y_2, \dots, y_n \in X$ **Query:** $x \in X$

Hash functions: *i.i.d.* instances g_1, \dots, g_ℓ of g , where $\ell = 1/p^*$;

Data structure: ℓ sorted tables;

in table- j : sort y_1, y_2, \dots, y_n according to $g_j(y_i)$

Upon **query** $x \in X$:

find $\leq 10\ell$ such y_i that $\exists j, g_j(x) = g_j(y_i)$ by binary search;

if encounter a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ return this y_i ;

else return “no”;

if $\exists y_s$ s.t. $\text{dist}(x, y_s) \leq r$

$$\Pr[\text{answer “no”}] \leq \Pr[\forall j, g_j(x) \neq g_j(y_s)] \leq (1 - p^*)^\ell \leq 1/e$$
$$+ \Pr[>10\ell \text{ bad } y_i \text{ that } \text{dist}(x, y_i) > c \cdot r \text{ but } \exists j \text{ s.t. } g_j(x) = g_j(y_i)]$$

$$\text{Markov inequality} \leq \mathbf{E}[\# \text{ of such bad } y_i] / 10\ell \leq \frac{\ell \cdot n \cdot (1/n)}{10\ell} \leq 0.1$$

linearity of expectation

metric space (X, dist) $(r, cr, p^*, 1/n)$ -LSH $g: X \rightarrow U$

Data: n points $y_1, y_2, \dots, y_n \in X$ **Query:** $x \in X$

Hash functions: *i.i.d.* instances g_1, \dots, g_ℓ of g , where $\ell = 1/p^*$;

Data structure: ℓ sorted tables;

in table- j : sort y_1, y_2, \dots, y_n according to $g_j(y_i)$

Upon **query** $x \in X$:

find $\leq 10\ell$ such y_i that $\exists j, g_j(x) = g_j(y_i)$ by binary search;

if encounter a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ return this y_i ;

else return “no”;

space: $O(nl) = O(n/p^*)$ time: $O(l \cdot \log n) = O((\log n) / p^*)$

if real answer is “no” : $\forall y_i, \text{dist}(x, y_i) > c \cdot r \Rightarrow$ always correct

if $\exists y_s$ s.t. $\text{dist}(x, y_s) \leq r$

$\Rightarrow \Pr[\text{answer “no”}] \leq 1/e + 0.1 < 0.5$

- (c, r) -ANN in metric space (X, dist) :


Data: n points $y_1, y_2, \dots, y_n \in X$

Query: a point $x \in X$

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$

“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$

suppose we have (r, cr, p, q) -LSH $h: X \rightarrow U$




we have $(r, cr, p^k, 1/n)$ -LSH $g: X \rightarrow U^k$

for $k = \log_{(1/q)} n$ so $p^k = p^{\log_{1/q} n} = n^{-\rho}$

where

$$\rho = \frac{\log p}{\log q}$$



solve (c, r) -ANN with space $O(n^{1+\rho})$

query time $O(n^{\rho} \cdot \log n)$ and one-sided error < 0.5

- (c, r) -ANN in metric space (X, dist) :

Data: n points $y_1, y_2, \dots, y_n \in X$


Query: a point $x \in X$

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$

“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$

suppose we have (r, cr, p, q) -LSH $h: X \rightarrow U$

$$\rho = \frac{\log p}{\log q}$$



solve (c, r) -ANN with space $O(n^{1+\rho})$

query time $O(n^\rho \cdot \log n)$ and one-sided error < 0.5

- (c, r) -ANN in **Hamming space** $\{0,1\}^d$:

Data: n points $y_1, y_2, \dots, y_n \in \{0,1\}^d$

Query: a point $x \in \{0,1\}^d$

return a y_i that $\text{dist}(x, y_i) \leq c \cdot r$ if $\exists y_j$ s.t. $\text{dist}(x, y_j) \leq r$

“no” if $\forall y_i, \text{dist}(x, y_i) > c \cdot r$

$\forall x \in \{0,1\}^d$: $h(x) = x_i$ for uniform random $i \in [d]$

$\text{dist}(x, y) \leq r \implies \Pr[h(x) = h(y)] \geq 1 - r/d$

$\text{dist}(x, y) > c \cdot r \implies \Pr[h(x) = h(y)] \leq 1 - cr/d$

$h: \{0,1\}^d \rightarrow \{0,1\}$ is an $(r, cr, 1 - r/d, 1 - cr/d)$ -LSH

$$\rho = \frac{\log(1 - r/d)}{\log(1 - cr/d)} \leq \frac{1}{c}$$

➡ solve (c, r) -ANN in **Hamming space** with space $O(n^{1+1/c})$
 query time $O(n^{1/c} \cdot \log n)$ and one-sided error < 0.5

COMPUTER SCIENCE

A neural algorithm for a fundamental computing problem

Sanjoy Dasgupta,¹ Charles F. Stevens,^{2,3} Saket Navlakha^{4*}

Similarity search—for example, identifying similar images in a database or similar documents on the web—is a fundamental computing problem faced by large-scale information retrieval systems. We discovered that the fruit fly olfactory circuit solves this problem with a variant of a computer science algorithm (called locality-sensitive hashing). The fly circuit assigns similar neural activity patterns to similar odors, so that behaviors learned from one odor can be applied when a similar odor is experienced. The fly algorithm, however, uses three computational strategies that depart from traditional approaches. These strategies can be translated to improve the performance of computational similarity searches. This perspective helps illuminate the logic supporting an important sensory function and provides a conceptually new algorithm for solving a fundamental computational problem.

An essential task of many neural circuits is to generate neural activity patterns in response to input stimuli, so that different inputs can be specifically identified. We studied the circuit used to process odors in the fruit fly olfactory system and uncovered computational strategies for solving a fundamental machine learning problem: approximate similarity (or nearest-neighbors) search.

The fly olfactory circuit generates a “tag” for each odor, which is a set of neurons that fire when that odor is presented (*1*). This tag is critical for learning behavioral responses to different odors (*2*). For example, if a reward (e.g., sugar water) or a punishment (e.g., electric shock) is associated with an odor, that odor becomes attractive (a fly

pendence is removed (*7, 8*); that is, the distribution of firing rates across the 50 PN types is exponential, with close to the same mean for all odors and all odor concentrations (*1*). Thus, the first step in the circuit essentially “centers the mean”—a standard preprocessing step in many computational pipelines—using a technique called divisive normalization (*8*). This step is important so that the fly does not mix up odor intensity with odor type.

The second step, where the main algorithmic insight begins, involves a 40-fold expansion in the number of neurons: Fifty PNs project to 2000 Kenyon cells (KCs), connected by a sparse, binary random connection matrix (*9*). Each KC receives and sums the firing rates from about six randomly

out of the billions of images on the web—that look most similar to your elephant image. This is called the nearest-neighbors search problem, which is of fundamental importance in information retrieval, data compression, and machine learning (*10*). Each image is typically represented as a d -dimensional vector of feature values. (Each odor that a fly processes is a 50-dimensional feature vector of firing rates.) A distance metric is used to compute the similarity between two images (feature vectors), and the goal is to efficiently find the nearest neighbors of any query image. If the web contained only a few images, then brute force linear search could easily be used to find the exact nearest neighbors. If the web contained many images, but each image was represented by a low-dimensional vector (e.g., 10 or 20 features), then space-partitioning methods (*12*) would similarly suffice. However, for large databases with high-dimensional data, neither approach scales (*11*).

In many applications, returning an approximate set of nearest neighbors that are “close enough” to the query is adequate, so long as they can be found quickly. This has motivated an approach for finding approximate nearest neighbors by LSH (*10*). For the fly, as noted, the locality-sensitive property states that two odors that generate similar ORN responses will be represented by two tags that are themselves similar (Fig. 1B). Likewise, for image search, the tag of an elephant image will be more similar to the tag of another elephant image than to the tag of a skyscraper image.

Unlike a traditional (non-LSH) hash function, where the input points are scattered randomly and uniformly over the range, a LSH function provides a distance-preserving embedding of points from d -dimensional space into m -dimensional

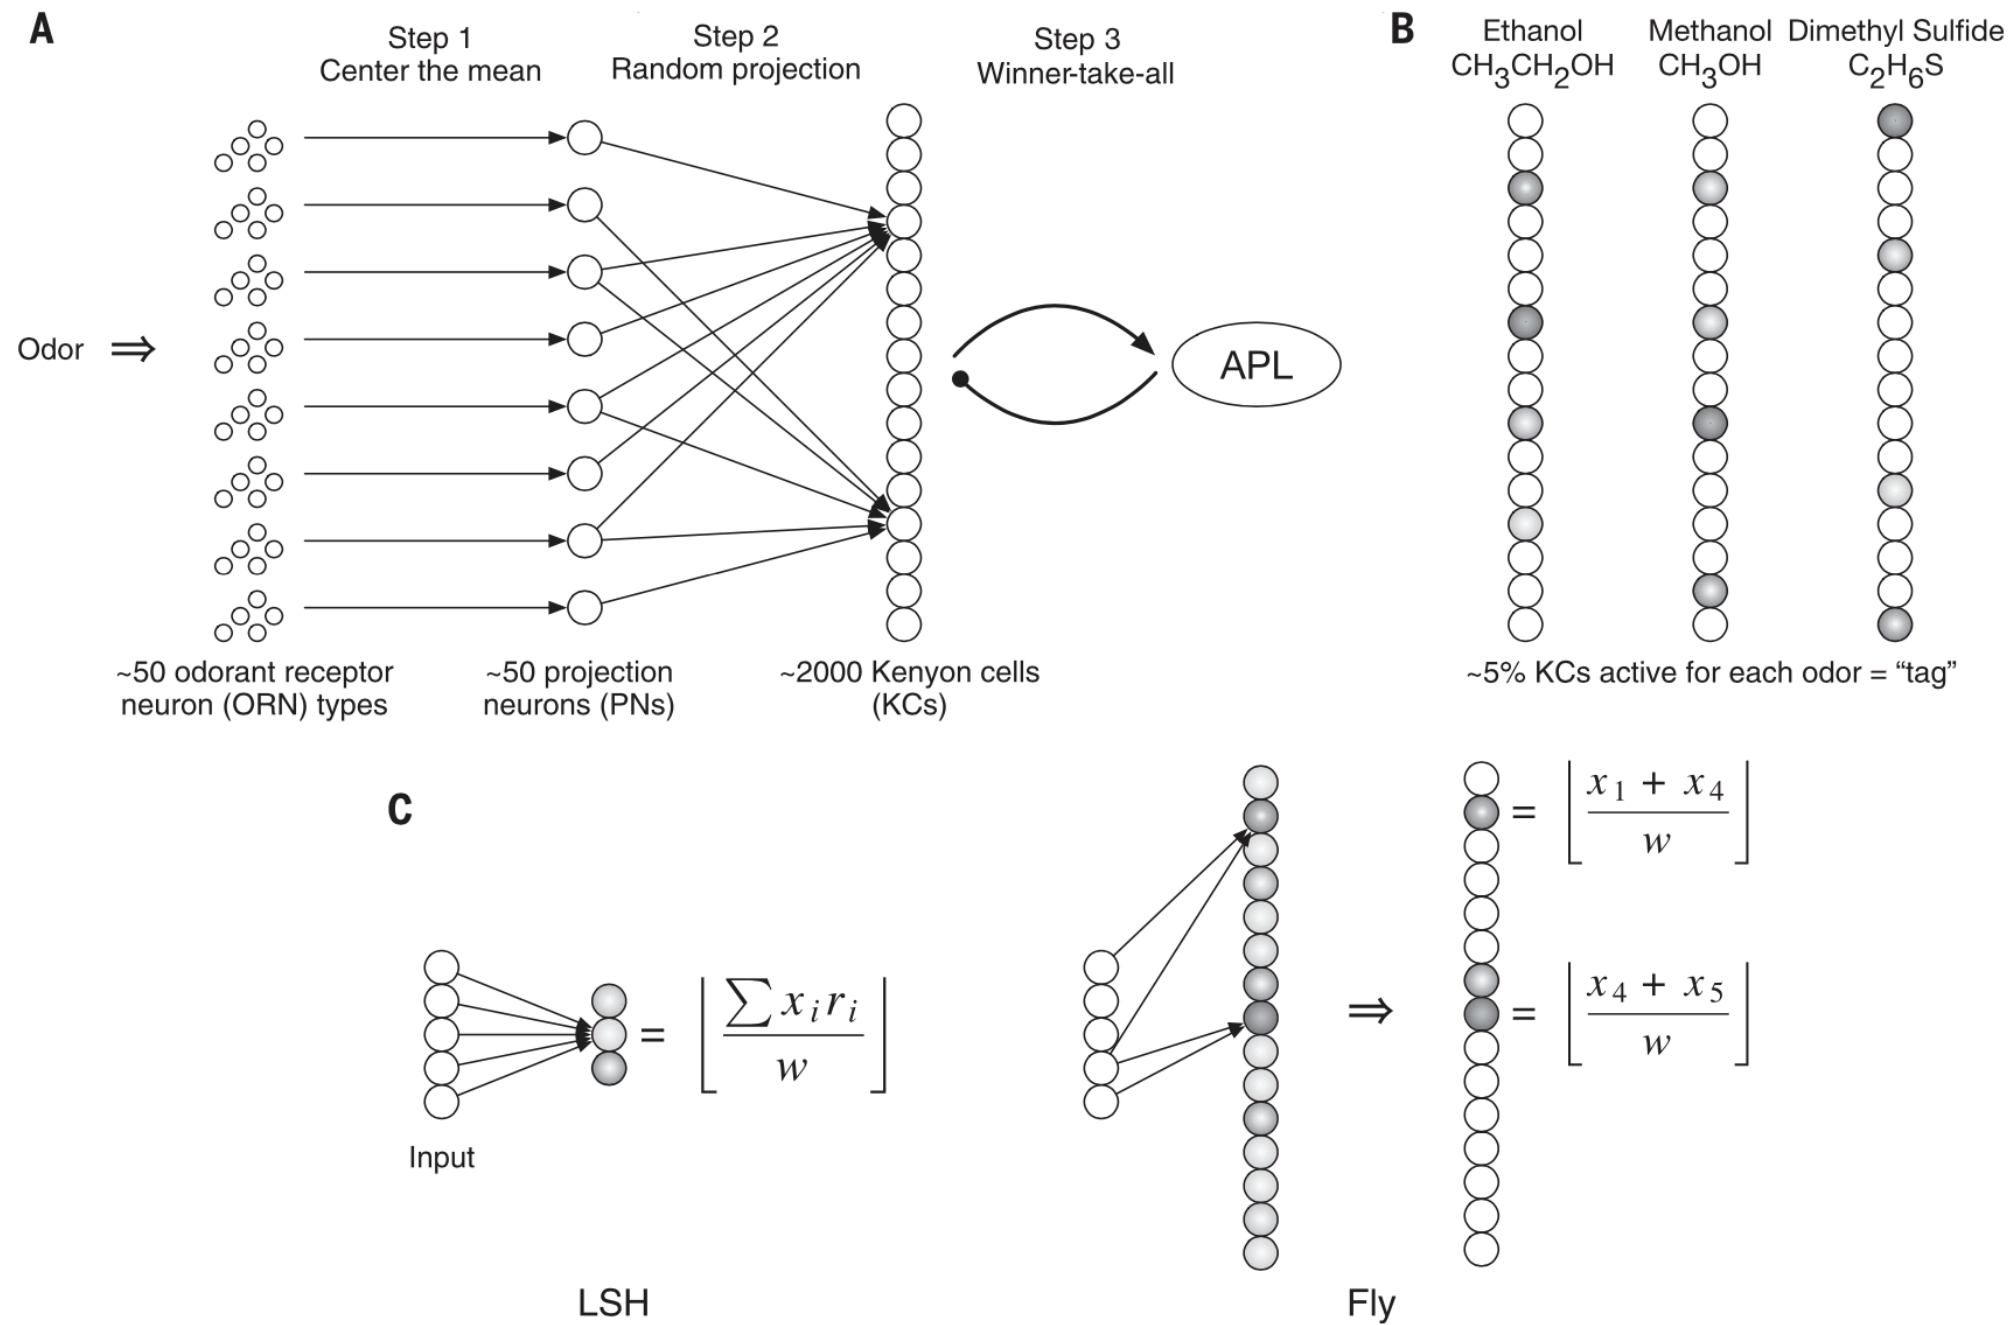


Fig. 1. Mapping between the fly olfactory circuit and locality-sensitive hashing (LSH). (A) Schematic of the fly olfactory circuit. In step 1, 50 ORNs in the fly's nose send axons to 50 PNs in the glomeruli; as a result of this projection, each odor is represented by an exponential distribution of firing rates, with the same mean for all odors and all odor concentrations. In step 2, the PNs expand the dimensionality, projecting to 2000 KCs connected by a sparse, binary random projection matrix. In step 3, the KCs receive feedback inhibition from the anterior paired lateral (APL) neuron, which leaves only the top 5% of KCs to remain firing spikes for the odor. This 5% corresponds to the tag

(hash) for the odor. (B) Illustrative odor responses. Similar pairs of odors (e.g., methanol and ethanol) are assigned more similar tags than are dissimilar odors. Darker shading denotes higher activity. (C) Differences between conventional LSH and the fly algorithm. In the example, the computational complexity for LSH and the fly are the same. The input dimensionality $d = 5$. LSH computes $m = 3$ random projections, each of which requires 10 operations (five multiplications plus five additions). The fly computes $m = 15$ random projections, each of which requires two addition operations. Thus, both require 30 total operations. \mathbf{x} , input feature vector; r , Gaussian random variable; w , bin width constant for discretization.