# Introduction to Machine Learning (67577) Lecture 11 

## Shai Shalev-Shwartz

School of CS and Engineering, The Hebrew University of Jerusalem

Dimensionality Reduction

## Dimensionality Reduction

- Dimensionality Reduction $=$ taking data in high dimensional space and mapping it into a low dimensional space


## Dimensionality Reduction

- Dimensionality Reduction $=$ taking data in high dimensional space and mapping it into a low dimensional space
- Why?


## Dimensionality Reduction

- Dimensionality Reduction $=$ taking data in high dimensional space and mapping it into a low dimensional space
- Why?
- Reduces training (and testing) time


## Dimensionality Reduction

- Dimensionality Reduction $=$ taking data in high dimensional space and mapping it into a low dimensional space
- Why?
- Reduces training (and testing) time
- Reduces estimation error


## Dimensionality Reduction

- Dimensionality Reduction $=$ taking data in high dimensional space and mapping it into a low dimensional space
- Why?
- Reduces training (and testing) time
- Reduces estimation error
- Interpretability of the data, finding meaningful structure in data, illustration


## Dimensionality Reduction

- Dimensionality Reduction $=$ taking data in high dimensional space and mapping it into a low dimensional space
- Why?
- Reduces training (and testing) time
- Reduces estimation error
- Interpretability of the data, finding meaningful structure in data, illustration
- Linear dimensionality reduction: $\mathbf{x} \mapsto W \mathbf{x}$ where $W \in \mathbb{R}^{n, d}$ and $n<d$


## Outline

(1) Principal Component Analysis (PCA)
(2) Random Projections
(3) Compressed Sensing

## Principal Component Analysis (PCA)

$$
\mathbf{x} \mapsto W \mathbf{x}
$$

- What makes $W$ a good matrix for dimensionality reduction ?


## Principal Component Analysis (PCA)

$$
\mathbf{x} \mapsto W \mathbf{x}
$$

- What makes $W$ a good matrix for dimensionality reduction ?
- Natural criterion: we want to be able to approximately recover $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$


## Principal Component Analysis (PCA)

$$
\mathbf{x} \mapsto W \mathbf{x}
$$

- What makes $W$ a good matrix for dimensionality reduction ?
- Natural criterion: we want to be able to approximately recover $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$
- PCA:


## Principal Component Analysis (PCA)

$$
\mathbf{x} \mapsto W \mathbf{x}
$$

- What makes $W$ a good matrix for dimensionality reduction ?
- Natural criterion: we want to be able to approximately recover $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$
- PCA:
- Linear recovery: $\tilde{\mathbf{x}}=U \mathbf{y}=U W \mathbf{x}$


## Principal Component Analysis (PCA)

$$
\mathbf{x} \mapsto W \mathbf{x}
$$

- What makes $W$ a good matrix for dimensionality reduction ?
- Natural criterion: we want to be able to approximately recover x from $\mathbf{y}=W \mathbf{x}$
- PCA:
- Linear recovery: $\tilde{\mathbf{x}}=U \mathbf{y}=U W \mathbf{x}$
- Measures "approximate recovery" by averaged squared norm: given examples $\mathbf{x}_{1}, \ldots, \mathbf{x}_{m}$, solve

$$
\underset{W \in \mathbb{R}^{n, d}, U \in \mathbb{R}^{d, n}}{\operatorname{argmin}} \sum_{i=1}^{m}\left\|\mathbf{x}_{i}-U W \mathbf{x}_{i}\right\|^{2}
$$

## Solving the PCA Problem

$$
\underset{W \in \mathbb{R}^{n, d}, U \in \mathbb{R}^{d, n}}{\operatorname{argmin}} \sum_{i=1}^{m}\left\|\mathbf{x}_{i}-U W \mathbf{x}_{i}\right\|^{2}
$$

## Solving the PCA Problem

$$
\underset{W \in \mathbb{R}^{n, d}, U \in \mathbb{R}^{d, n}}{\operatorname{argmin}} \sum_{i=1}^{m}\left\|\mathbf{x}_{i}-U W \mathbf{x}_{i}\right\|^{2}
$$

## Theorem

Let $A=\sum_{i=1}^{m} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}$ and let $\mathbf{u}_{1}, \ldots, u_{n}$ be the $n$ leading eigenvectors of $A$. Then, the solution to the PCA problem is to set the columns of $U$ to be $\mathbf{u}_{1}, \ldots, \mathbf{u}_{n}$ and to set $W=U^{\top}$

## Proof main ideas

- $U W$ is of rank $n$, therefore its range is $n$ dimensional subspace, denoted $S$


## Proof main ideas

- $U W$ is of rank $n$, therefore its range is $n$ dimensional subspace, denoted $S$
- The transformation $\mathbf{x} \mapsto U W \mathbf{x}$ moves $\mathbf{x}$ to this subspace


## Proof main ideas

- $U W$ is of rank $n$, therefore its range is $n$ dimensional subspace, denoted $S$
- The transformation $\mathbf{x} \mapsto U W \mathbf{x}$ moves $\mathbf{x}$ to this subspace
- The point in $S$ which is closest to $\mathbf{x}$ is $V V^{\top} \mathbf{x}$, where columns of $V$ are orthonormal basis of $S$


## Proof main ideas

- $U W$ is of rank $n$, therefore its range is $n$ dimensional subspace, denoted $S$
- The transformation $\mathbf{x} \mapsto U W \mathbf{x}$ moves $\mathbf{x}$ to this subspace
- The point in $S$ which is closest to $\mathbf{x}$ is $V V^{\top} \mathbf{x}$, where columns of $V$ are orthonormal basis of $S$
- Therefore, we can assume w.l.o.g. that $W=U^{\top}$ and that columns of $U$ are orthonormal


## Proof main ideas

Observe:

$$
\begin{aligned}
\left\|\mathbf{x}-U U^{\top} \mathbf{x}\right\|^{2} & =\|\mathbf{x}\|^{2}-2 \mathbf{x}^{\top} U U^{\top} \mathbf{x}+\mathbf{x}^{\top} U U^{\top} U U^{\top} \mathbf{x} \\
& =\|\mathbf{x}\|^{2}-\mathbf{x}^{\top} U U^{\top} \mathbf{x} \\
& =\|\mathbf{x}\|^{2}-\operatorname{trace}\left(U^{\top} \mathbf{x} \mathbf{x}^{\top} U\right)
\end{aligned}
$$

## Proof main ideas

Observe:

$$
\begin{aligned}
\left\|\mathbf{x}-U U^{\top} \mathbf{x}\right\|^{2} & =\|\mathbf{x}\|^{2}-2 \mathbf{x}^{\top} U U^{\top} \mathbf{x}+\mathbf{x}^{\top} U U^{\top} U U^{\top} \mathbf{x} \\
& =\|\mathbf{x}\|^{2}-\mathbf{x}^{\top} U U^{\top} \mathbf{x} \\
& =\|\mathbf{x}\|^{2}-\operatorname{trace}\left(U^{\top} \mathbf{x} \mathbf{x}^{\top} U\right)
\end{aligned}
$$

Therefore, an equivalent PCA problem is

$$
\underset{U \in \mathbb{R}^{d, n}: U^{\top} U=I}{\operatorname{argmax}} \operatorname{trace}\left(U^{\top}\left(\sum_{i=1}^{m} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right) U\right)
$$

## Proof main ideas

Observe:

$$
\begin{aligned}
\left\|\mathbf{x}-U U^{\top} \mathbf{x}\right\|^{2} & =\|\mathbf{x}\|^{2}-2 \mathbf{x}^{\top} U U^{\top} \mathbf{x}+\mathbf{x}^{\top} U U^{\top} U U^{\top} \mathbf{x} \\
& =\|\mathbf{x}\|^{2}-\mathbf{x}^{\top} U U^{\top} \mathbf{x} \\
& =\|\mathbf{x}\|^{2}-\operatorname{trace}\left(U^{\top} \mathbf{x} \mathbf{x}^{\top} U\right)
\end{aligned}
$$

Therefore, an equivalent PCA problem is

$$
\underset{U \in \mathbb{R}^{d, n}: U^{\top} U=I}{\operatorname{argmax}} \operatorname{trace}\left(U^{\top}\left(\sum_{i=1}^{m} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}\right) U\right)
$$

The solution is to set $U$ to be the leading eigenvectors of $A=\sum_{i=1}^{m} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}$.

## Value of the objective

It is easy to see that

$$
\min _{W \in \mathbb{R}^{n, d}, U \in \mathbb{R}^{d, n}} \sum_{i=1}^{m}\left\|\mathbf{x}_{i}-U W \mathbf{x}_{i}\right\|^{2}=\sum_{i=n+1}^{d} \lambda_{i}(A)
$$

## Centering

- It is a common practice to "center" the examples before applying PCA, namely:
- First calculate $\boldsymbol{\mu}=\frac{1}{m} \sum_{i=1}^{m} \mathbf{x}_{i}$
- Then apply PCA on the vectors $\left(\mathbf{x}_{1}-\boldsymbol{\mu}\right), \ldots,\left(\mathbf{x}_{m}-\boldsymbol{\mu}\right)$
- This is also related to the interpretation of PCA as variance maximization (will be given in exercise)


## Efficient implementation for $d \gg m$ and kernel PCA

- Recall: $A=\sum_{i=1}^{m} \mathbf{x}_{i} \mathbf{x}_{i}^{\top}=X^{\top} X$ where $X \in \mathbb{R}^{m, d}$ is a matrix whose $i$ 'th row is $\mathbf{x}_{i}^{\top}$.
- Let $B=X X^{\top}$. That is, $B_{i, j}=\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle$
- If $B \mathbf{u}=\lambda \mathbf{u}$ then

$$
A\left(X^{\top} \mathbf{u}\right)=X^{\top} X X^{\top} \mathbf{u}=X^{\top} B \mathbf{u}=\lambda\left(X^{\top} \mathbf{u}\right)
$$

- So, $\frac{X^{\top} \mathbf{u}}{\left\|X^{\top} \mathbf{u}\right\|}$ is an eigenvector of $A$ with eigenvalue $\lambda$
- We can therefore calculate the PCA solution by calculating the eigenvalues of $B$ instead of $A$
- The complexity is $O\left(m^{3}+m^{2} d\right)$
- And, it can be computed using a kernel function


## Pseudo code

## PCA

## input

A matrix of $m$ examples $X \in \mathbb{R}^{m, d}$
number of components $n$
if $(m>d)$
$A=X^{\top} X$
Let $\mathbf{u}_{1}, \ldots, \mathbf{u}_{n}$ be the eigenvectors of $A$ with largest eigenvalues else
$B=X X^{\top}$
Let $\mathbf{v}_{1}, \ldots, \mathbf{v}_{n}$ be the eigenvectors of $B$ with largest eigenvalues for $i=1, \ldots, n$ set $\mathbf{u}_{i}=\frac{1}{\left\|X^{\top} \mathbf{v}_{i}\right\|} X^{\top} \mathbf{v}_{i}$
output: $\mathbf{u}_{1}, \ldots, \mathbf{u}_{n}$

## Demonstration



## Demonstration

- $50 \times 50$ images from Yale dataset
- Before (left) and after reconstruction (right) to 10 dimensions



## Demonstration

- Before and after



## Demonstration

- Images after dim reduction to $\mathbb{R}^{2}$
- Different marks indicate different individuals



## Outline

## (1) Principal Component Analysis (PCA)

(2) Random Projections

## (3) Compressed Sensing

## What is a successful dimensionality reduction?

- In PCA, we measured succes as squared distance between $\mathbf{x}$ and a reconstruction of $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$


## What is a successful dimensionality reduction?

- In PCA, we measured succes as squared distance between $\mathbf{x}$ and a reconstruction of $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$
- In some cases, we don't care about reconstruction, all we care is that $\mathbf{y}_{1}, \ldots, \mathbf{y}_{m}$ will retain certain properties of $\mathbf{x}_{1}, \ldots, \mathbf{x}_{m}$


## What is a successful dimensionality reduction?

- In PCA, we measured succes as squared distance between $\mathbf{x}$ and a reconstruction of $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$
- In some cases, we don't care about reconstruction, all we care is that $\mathbf{y}_{1}, \ldots, \mathbf{y}_{m}$ will retain certain properties of $\mathbf{x}_{1}, \ldots, \mathbf{x}_{m}$
- One option: do not distort distances. That is, we'd like that for all $i, j,\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\| \approx\left\|\mathbf{y}_{i}-\mathbf{y}_{j}\right\|$


## What is a successful dimensionality reduction?

- In PCA, we measured succes as squared distance between $\mathbf{x}$ and a reconstruction of $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$
- In some cases, we don't care about reconstruction, all we care is that $\mathbf{y}_{1}, \ldots, \mathbf{y}_{m}$ will retain certain properties of $\mathbf{x}_{1}, \ldots, \mathbf{x}_{m}$
- One option: do not distort distances. That is, we'd like that for all $i, j,\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\| \approx\left\|\mathbf{y}_{i}-\mathbf{y}_{j}\right\|$
- Equivalently, we'd like that for all $i, j, \frac{\left\|W \mathbf{x}_{i}-W \mathbf{x}_{j}\right\|}{\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\|} \approx 1$


## What is a successful dimensionality reduction?

- In PCA, we measured succes as squared distance between $\mathbf{x}$ and a reconstruction of $\mathbf{x}$ from $\mathbf{y}=W \mathbf{x}$
- In some cases, we don't care about reconstruction, all we care is that $\mathbf{y}_{1}, \ldots, \mathbf{y}_{m}$ will retain certain properties of $\mathbf{x}_{1}, \ldots, \mathbf{x}_{m}$
- One option: do not distort distances. That is, we'd like that for all $i, j,\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\| \approx\left\|\mathbf{y}_{i}-\mathbf{y}_{j}\right\|$
- Equivalently, we'd like that for all $i, j, \frac{\left\|W \mathbf{x}_{i}-W \mathbf{x}_{j}\right\|}{\left\|\mathbf{x}_{i}-\mathbf{x}_{j}\right\|} \approx 1$
- Equivalently, we'd like that for all $\mathbf{x} \in Q$, where $Q=\left\{\mathbf{x}_{i}-\mathbf{x}_{j}: i, j \in[m]\right\}$, we'll have $\frac{\|W \mathbf{x}\|}{\|x\|} \approx 1$


## Random Projections do not distort norms

- Random projection: The transformation $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a random matrix


## Random Projections do not distort norms

- Random projection: The transformation $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a random matrix
- We'll analyze the distortion due to $W$ s.t. $W_{i, j} \sim N(0,1 / n)$


## Random Projections do not distort norms

- Random projection: The transformation $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a random matrix
- We'll analyze the distortion due to $W$ s.t. $W_{i, j} \sim N(0,1 / n)$
- Let $\mathbf{w}_{i}$ be the $i$ 'th row of $W$. Then:

$$
\begin{aligned}
\mathbb{E}\left[\|W \mathbf{x}\|^{2}\right] & =\sum_{i=1}^{n} \mathbb{E}\left[\left(\left\langle\mathbf{w}_{i}, \mathbf{x}\right\rangle\right)^{2}\right]=\sum_{i=1}^{n} \mathbf{x}^{\top} \mathbb{E}\left[\mathbf{w}_{i} \mathbf{w}_{i}^{\top}\right] \mathbf{x} \\
& =n \mathbf{x}^{\top}\left(\frac{1}{n} I\right) \mathbf{x}=\|\mathbf{x}\|^{2}
\end{aligned}
$$

## Random Projections do not distort norms

- Random projection: The transformation $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a random matrix
- We'll analyze the distortion due to $W$ s.t. $W_{i, j} \sim N(0,1 / n)$
- Let $\mathbf{w}_{i}$ be the $i$ 'th row of $W$. Then:

$$
\begin{aligned}
\mathbb{E}\left[\|W \mathbf{x}\|^{2}\right] & =\sum_{i=1}^{n} \mathbb{E}\left[\left(\left\langle\mathbf{w}_{i}, \mathbf{x}\right\rangle\right)^{2}\right]=\sum_{i=1}^{n} \mathbf{x}^{\top} \mathbb{E}\left[\mathbf{w}_{i} \mathbf{w}_{i}^{\top}\right] \mathbf{x} \\
& =n \mathbf{x}^{\top}\left(\frac{1}{n} I\right) \mathbf{x}=\|\mathbf{x}\|^{2}
\end{aligned}
$$

- In fact, $\|W \mathbf{x}\|^{2}$ has a $\chi_{n}^{2}$ distribution, and using a measure concentration inequality it can be shown that

$$
\mathbb{P}\left[\left|\frac{\|W \mathbf{x}\|^{2}}{\|\mathbf{x}\|^{2}}-1\right|>\epsilon\right] \leq 2 e^{-\epsilon^{2} n / 6}
$$

## Random Projections do not distort norms

- Applying the union bound over all vectors in $Q$ we obtain:


## Lemma (Johnson-Lindenstrauss lemma)

Let $Q$ be a finite set of vectors in $\mathbb{R}^{d}$. Let $\delta \in(0,1)$ and $n$ be an integer such that

$$
\epsilon=\sqrt{\frac{6 \log (2|Q| / \delta)}{n}} \leq 3 .
$$

Then, with probability of at least $1-\delta$ over a choice of a random matrix $W \in \mathbb{R}^{n, d}$ with $W_{i, j} \sim N(0,1 / n)$, we have

$$
\max _{\mathbf{x} \in Q}\left|\frac{\|W \mathbf{x}\|^{2}}{\|\mathbf{x}\|^{2}}-1\right|<\epsilon
$$

## Outline

## (1) Principal Component Analysis (PCA)

## (2) Random Projections

## (3) Compressed Sensing

## Compressed Sensing

- Prior assumption: $\mathbf{x} \approx U \boldsymbol{\alpha}$ where $U$ is orthonormal and $\|\boldsymbol{\alpha}\|_{0} \stackrel{\text { def }}{=}\left|\left\{i: \alpha_{i} \neq 0\right\}\right| \leq s$ for some $s \ll d$


## Compressed Sensing

- Prior assumption: $\mathbf{x} \approx U \boldsymbol{\alpha}$ where $U$ is orthonormal and $\|\boldsymbol{\alpha}\|_{0} \stackrel{\text { def }}{=}\left|\left\{i: \alpha_{i} \neq 0\right\}\right| \leq s$ for some $s \ll d$
- E.g.: natural images are approximately sparse in a wavelet basis


## Compressed Sensing

- Prior assumption: $\mathbf{x} \approx U \boldsymbol{\alpha}$ where $U$ is orthonormal and $\|\boldsymbol{\alpha}\|_{0} \stackrel{\text { def }}{=}\left|\left\{i: \alpha_{i} \neq 0\right\}\right| \leq s$ for some $s \ll d$
- E.g.: natural images are approximately sparse in a wavelet basis
- How to "store" x ?


## Compressed Sensing

- Prior assumption: $\mathbf{x} \approx U \boldsymbol{\alpha}$ where $U$ is orthonormal and $\|\boldsymbol{\alpha}\|_{0} \stackrel{\text { def }}{=}\left|\left\{i: \alpha_{i} \neq 0\right\}\right| \leq s$ for some $s \ll d$
- E.g.: natural images are approximately sparse in a wavelet basis
- How to "store" x ?
- We can find $\boldsymbol{\alpha}=U^{\top} \mathbf{x}$ and then save the non-zero elements of $\boldsymbol{\alpha}$


## Compressed Sensing

- Prior assumption: $\mathbf{x} \approx U \boldsymbol{\alpha}$ where $U$ is orthonormal and $\|\boldsymbol{\alpha}\|_{0} \stackrel{\text { def }}{=}\left|\left\{i: \alpha_{i} \neq 0\right\}\right| \leq s$ for some $s \ll d$
- E.g.: natural images are approximately sparse in a wavelet basis
- How to "store" x ?
- We can find $\boldsymbol{\alpha}=U^{\top} \mathbf{x}$ and then save the non-zero elements of $\boldsymbol{\alpha}$
- Requires order of $s \log (d)$ storage


## Compressed Sensing

- Prior assumption: $\mathbf{x} \approx U \boldsymbol{\alpha}$ where $U$ is orthonormal and $\|\boldsymbol{\alpha}\|_{0} \stackrel{\text { def }}{=}\left|\left\{i: \alpha_{i} \neq 0\right\}\right| \leq s$ for some $s \ll d$
- E.g.: natural images are approximately sparse in a wavelet basis
- How to "store" x ?
- We can find $\boldsymbol{\alpha}=U^{\top} \mathbf{x}$ and then save the non-zero elements of $\boldsymbol{\alpha}$
- Requires order of $s \log (d)$ storage
- Why go to so much effort to acquire all the $d$ coordinates of $\mathbf{x}$ when most of what we get will be thrown away? Can't we just directly measure the part that won't end up being thrown away?


## Compressed Sensing

Informally, the main premise of compressed sensing is the following three "surprising" results:
(1) It is possible to fully reconstruct any sparse signal if it was compressed by $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a matrix which satisfies a condition called Restricted Isoperimetric Property (RIP). A matrix that satisfies this property is guaranteed to have a low distortion of the norm of any sparse representable vector.

## Compressed Sensing

Informally, the main premise of compressed sensing is the following three "surprising" results:
(1) It is possible to fully reconstruct any sparse signal if it was compressed by $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a matrix which satisfies a condition called Restricted Isoperimetric Property (RIP). A matrix that satisfies this property is guaranteed to have a low distortion of the norm of any sparse representable vector.
(2) The reconstruction can be calculated in polynomial time by solving a linear program.

## Compressed Sensing

Informally, the main premise of compressed sensing is the following three "surprising" results:
(1) It is possible to fully reconstruct any sparse signal if it was compressed by $\mathbf{x} \mapsto W \mathbf{x}$, where $W$ is a matrix which satisfies a condition called Restricted Isoperimetric Property (RIP). A matrix that satisfies this property is guaranteed to have a low distortion of the norm of any sparse representable vector.
(2) The reconstruction can be calculated in polynomial time by solving a linear program.
(3) A random $n \times d$ matrix is likely to satisfy the RIP condition provided that $n$ is greater than order of $s \log (d)$.

## Restricted Isoperimetric Property (RIP)

A matrix $W \in \mathbb{R}^{n, d}$ is $(\epsilon, s)$-RIP if for all $\mathbf{x} \neq 0$ s.t. $\|\mathbf{x}\|_{0} \leq s$ we have

$$
\left|\frac{\|W \mathbf{x}\|_{2}^{2}}{\|\mathbf{x}\|_{2}^{2}}-1\right| \leq \epsilon .
$$

## RIP matrices yield lossless compression for sparse vectors

Theorem
Let $\epsilon<1$ and let $W$ be a $(\epsilon, 2 s)$-RIP matrix. Let $\mathbf{x}$ be a vector s.t. $\|\mathbf{x}\|_{0} \leq s$, let $\mathbf{y}=W \mathbf{x}$ and let $\tilde{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{v}: W \mathbf{v}=\mathbf{y}}\|\mathbf{v}\|_{0}$. Then, $\tilde{\mathbf{x}}=\mathbf{x}$.

## RIP matrices yield lossless compression for sparse vectors

## Theorem

Let $\epsilon<1$ and let $W$ be a $(\epsilon, 2 s)$-RIP matrix. Let $\mathbf{x}$ be a vector s.t. $\|\mathbf{x}\|_{0} \leq s$, let $\mathbf{y}=W \mathbf{x}$ and let $\tilde{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{v}: W \mathbf{v}=\mathbf{y}}\|\mathbf{v}\|_{0}$. Then, $\tilde{\mathbf{x}}=\mathbf{x}$.

## Proof.

- Assume, by way of contradiction, that $\tilde{\mathbf{x}} \neq \mathbf{x}$.


## RIP matrices yield lossless compression for sparse vectors

## Theorem

Let $\epsilon<1$ and let $W$ be a $(\epsilon, 2 s)$-RIP matrix. Let $\mathbf{x}$ be a vector s.t. $\|\mathbf{x}\|_{0} \leq s$, let $\mathbf{y}=W \mathbf{x}$ and let $\tilde{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{v}: W \mathbf{v}=\mathbf{y}}\|\mathbf{v}\|_{0}$. Then, $\tilde{\mathbf{x}}=\mathbf{x}$.

## Proof.

- Assume, by way of contradiction, that $\tilde{\mathbf{x}} \neq \mathbf{x}$.
- Since $\mathbf{x}$ satisfies the constraints in the optimization problem that defines $\tilde{\mathbf{x}}$ we clearly have that $\|\tilde{\mathbf{x}}\|_{0} \leq\|\mathbf{x}\|_{0} \leq s$.


## RIP matrices yield lossless compression for sparse vectors

## Theorem

Let $\epsilon<1$ and let $W$ be a $(\epsilon, 2 s)$-RIP matrix. Let $\mathbf{x}$ be a vector s.t. $\|\mathbf{x}\|_{0} \leq s$, let $\mathbf{y}=W \mathbf{x}$ and let $\tilde{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{v}: W \mathbf{v}=\mathbf{y}}\|\mathbf{v}\|_{0}$. Then, $\tilde{\mathbf{x}}=\mathbf{x}$.

## Proof.

- Assume, by way of contradiction, that $\tilde{\mathbf{x}} \neq \mathbf{x}$.
- Since $\mathbf{x}$ satisfies the constraints in the optimization problem that defines $\tilde{\mathbf{x}}$ we clearly have that $\|\tilde{\mathbf{x}}\|_{0} \leq\|\mathbf{x}\|_{0} \leq s$.
- Therefore, $\|\mathbf{x}-\tilde{\mathbf{x}}\|_{0} \leq 2 s$.


## RIP matrices yield lossless compression for sparse vectors

## Theorem

Let $\epsilon<1$ and let $W$ be a $(\epsilon, 2 s)$-RIP matrix. Let $\mathbf{x}$ be a vector s.t. $\|\mathbf{x}\|_{0} \leq s$, let $\mathbf{y}=W \mathbf{x}$ and let $\tilde{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{v}: W \mathbf{v}=\mathbf{y}}\|\mathbf{v}\|_{0}$. Then, $\tilde{\mathbf{x}}=\mathbf{x}$.

## Proof.

- Assume, by way of contradiction, that $\tilde{\mathbf{x}} \neq \mathbf{x}$.
- Since $\mathbf{x}$ satisfies the constraints in the optimization problem that defines $\tilde{\mathbf{x}}$ we clearly have that $\|\tilde{\mathbf{x}}\|_{0} \leq\|\mathbf{x}\|_{0} \leq s$.
- Therefore, $\|\mathbf{x}-\tilde{\mathbf{x}}\|_{0} \leq 2 s$.
- By RIP on $\mathbf{x}-\tilde{\mathbf{x}}$ we have $\left|\frac{\|W(\mathbf{x}-\tilde{\mathbf{x}})\|^{2}}{\|\mathbf{x}-\tilde{\mathbf{x}}\|^{2}}-1\right| \leq \epsilon$


## RIP matrices yield lossless compression for sparse vectors

## Theorem

Let $\epsilon<1$ and let $W$ be a $(\epsilon, 2 s)$-RIP matrix. Let $\mathbf{x}$ be a vector s.t. $\|\mathbf{x}\|_{0} \leq s$, let $\mathbf{y}=W \mathbf{x}$ and let $\tilde{\mathbf{x}} \in \operatorname{argmin}_{\mathbf{v}: W \mathbf{v}=\mathbf{y}}\|\mathbf{v}\|_{0}$. Then, $\tilde{\mathbf{x}}=\mathbf{x}$.

## Proof.

- Assume, by way of contradiction, that $\tilde{\mathbf{x}} \neq \mathbf{x}$.
- Since $\mathbf{x}$ satisfies the constraints in the optimization problem that defines $\tilde{\mathbf{x}}$ we clearly have that $\|\tilde{\mathbf{x}}\|_{0} \leq\|\mathbf{x}\|_{0} \leq s$.
- Therefore, $\|\mathbf{x}-\tilde{\mathbf{x}}\|_{0} \leq 2 s$.
- By RIP on $\mathbf{x}-\tilde{\mathbf{x}}$ we have $\left|\frac{\|W(\mathbf{x}-\tilde{\mathbf{x}})\|^{2}}{\|\mathbf{x}-\tilde{\mathbf{x}}\|^{2}}-1\right| \leq \epsilon$
- But, since $W(\mathbf{x}-\tilde{\mathbf{x}})=\mathbf{0}$ we get that $|0-1| \leq \epsilon$. Contradiction.


## Efficient reconstruction

- If we further assume that $\epsilon<\frac{1}{1+\sqrt{2}}$ then

$$
\mathbf{x}=\underset{\mathbf{v}: W \mathbf{v}=\mathbf{y}}{\operatorname{argmin}}\|\mathbf{v}\|_{0}=\underset{\mathbf{v}: W \mathbf{v}=\mathbf{y}}{\operatorname{argmin}}\|\mathbf{v}\|_{1}
$$

## Efficient reconstruction

- If we further assume that $\epsilon<\frac{1}{1+\sqrt{2}}$ then

$$
\mathbf{x}=\underset{\mathbf{v}: W \mathbf{v}=\mathbf{y}}{\operatorname{argmin}}\|\mathbf{v}\|_{0}=\underset{\mathbf{v}: W \mathbf{v}=\mathbf{y}}{\operatorname{argmin}}\|\mathbf{v}\|_{1}
$$

- The right-hand side is a linear programming problem


## Efficient reconstruction

- If we further assume that $\epsilon<\frac{1}{1+\sqrt{2}}$ then

$$
\mathbf{x}=\underset{\mathbf{v}: W \mathbf{v}=\mathbf{y}}{\operatorname{argmin}}\|\mathbf{v}\|_{0}=\underset{\mathbf{v}: W \mathbf{v}=\mathbf{y}}{\operatorname{argmin}}\|\mathbf{v}\|_{1}
$$

- The right-hand side is a linear programming problem
- Summary: we can reconstruct all sparse vector efficiently based on $O(s \log (d))$ measurements


## PCA vs. Random Projections

- Random projections guarantee perfect recovery for all $O(n / \log (d))$-sparse vectors


## PCA vs. Random Projections

- Random projections guarantee perfect recovery for all $O(n / \log (d))$-sparse vectors
- PCA guarantee perfect recovery if all examples are in an n-dimensional subspace


## PCA vs. Random Projections

- Random projections guarantee perfect recovery for all $O(n / \log (d))$-sparse vectors
- PCA guarantee perfect recovery if all examples are in an n-dimensional subspace
- Different prior knowledge:


## PCA vs. Random Projections

- Random projections guarantee perfect recovery for all $O(n / \log (d))$-sparse vectors
- PCA guarantee perfect recovery if all examples are in an n-dimensional subspace
- Different prior knowledge:
- If the data is $\mathbf{e}_{1}, \ldots, \mathbf{e}_{d}$, random projections will be perfect but PCA will fail


## PCA vs. Random Projections

- Random projections guarantee perfect recovery for all $O(n / \log (d))$-sparse vectors
- PCA guarantee perfect recovery if all examples are in an n-dimensional subspace
- Different prior knowledge:
- If the data is $\mathbf{e}_{1}, \ldots, \mathbf{e}_{d}$, random projections will be perfect but PCA will fail
- If $d$ is very large and data is exactly on an $n$-dim subspace. Then, PCA will be perfect but random projections might fail


## Summary

- Linear dimensionality reduction $\mathbf{x} \mapsto W \mathbf{x}$
- PCA: optimal if reconstruction is linear and error is squared distance
- Random projections: preserves disctances
- Random projections: exact reconstruction for sparse vectors (but with a non-linear reconstruction)
- Not covered: non-linear dimensionality reduction

