



Machine Learning: Unsupervised Learning

Konstantin Tretyakov

<http://kt.era.ee>

IFI Summer
School 2014

STACC

Software Technology and
Applications Competence Center



So far...



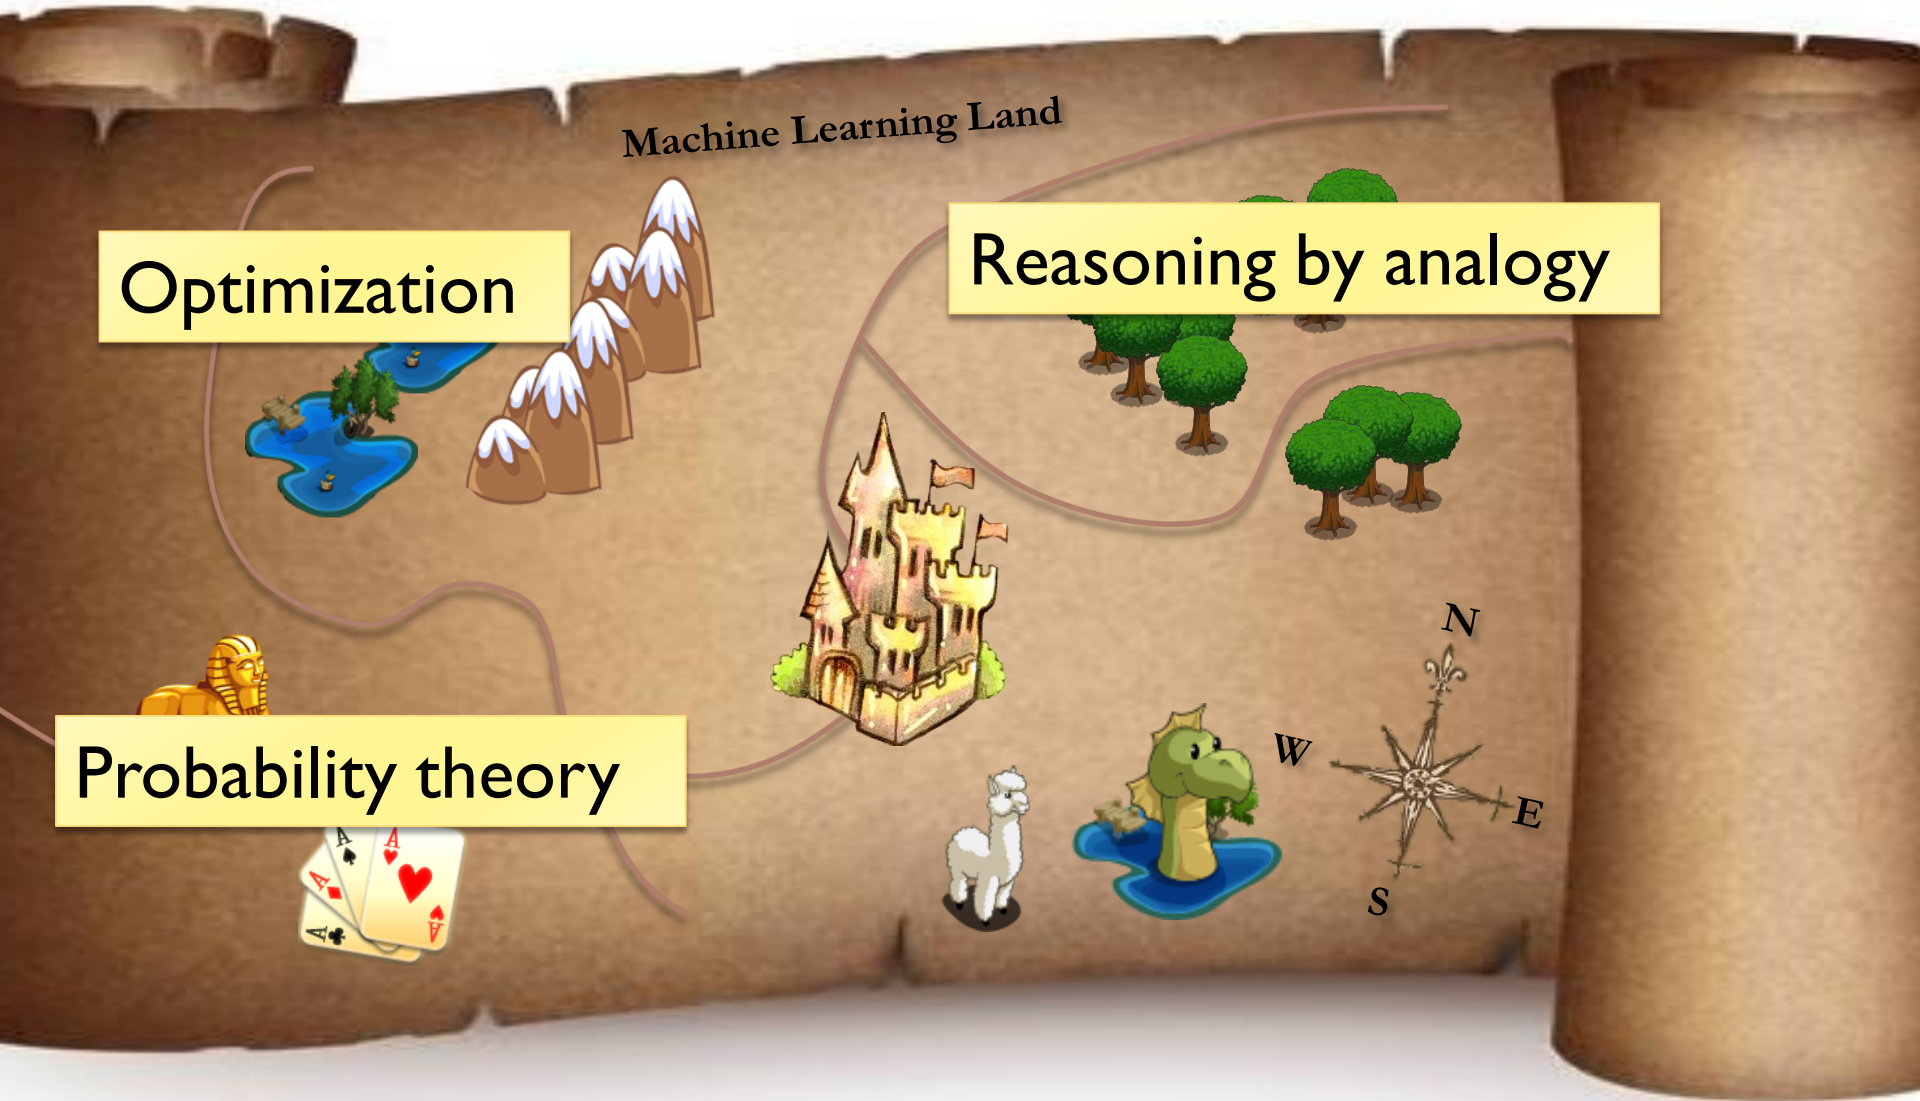
So far...

Machine Learning Land

Optimization

Reasoning by analogy

Probability theory



So far...

Optimization

Fermat's theorem
Gradient methods
Batch & On-line

MLE,
MAP,
Bayesian Estimation
Risk Optimization

Probability theory

Machine Learning

Reasoning by analogy

k-NN,
Kernel methods



So far...

Optimization

Fermat's theorem

Gradient

Batch & C

MLE,

MAP,

Bayesian Estimation

Risk Optimization

Probability theory

Machine Learning

Reasoning by analogy

k-NN,

k-NN, Decision tree learning


Linear regression (OLS, RR),

Linear classification (SVM, Perceptron),

Bayes & Naïve Bayes

Kernel-<xxx>

ods



Next

Optimization

Fermat's theorem

Gradient

Batch & C

MLE,

MAP,

Bayesian Estimation

Risk Optimization

Probability theory

Machine Learning

Reasoning by analogy

k-NN,

k-NN, Decision tree learning

Linear regression (OLS, RR),

Linear classification (SVM, Perceptron),


Bayes & Naïve Bayes


Kernel-<xxx>


ods

**Unsupervised
Learning**


N
E
S








AATAACGGCCCCGATGAGGAAACGAACGGTCGCACT
AAGATGAGACATGTCCCGAAAGGTGCATAAGTTAT
GGACGAAAAACTTTCTTCGCCCTTTGATGTGCCCC
AGCGCGGGATGAGGATCAGCCCCCGCATTAGTTCA
ATATGCGAGCTTTCGCGCTCGGAAAGGGCAATAAA
GCGACGGCCCCGATGAGGGGTGTTACTAGATTGGA
TGGGTGGTTCAGATCTCGGCTTACCCCCCTTTATCA
ACCCTGCTACAGACTCGTTGAGAATGCTACGGATC




Data Mining




Unsupervised learning patterns

Decomposition



Clustering



Quiz




- ▶ Why would one need clustering?




Hierarchical clustering




Hierarchical clustering



Hierarchical clustering




Hierarchical clustering




Hierarchical clustering

Complete linkage




Hierarchical clustering

Single linkage




Hierarchical clustering

Average linkage




Hierarchical clustering


Ward linkage




Hierarchical clustering




Hierarchical clustering




Hierarchical clustering



Partitional vs Hierarchical




Partitional clustering finds
a fixed number of clusters




Hierarchical clustering creates a
series of clusterings contained in
each other


K-means




K-means




K-means




K-means




K-means



K-means



K-means



$$\operatorname{argmin}_{c_1, \dots, c_K} \sum_i \|x_i - c_{\text{closest_to}(i)}\|^2$$

K-means

$$\operatorname{argmin}_{c_1, \dots, c_K} \sum_i \|x_i - c_{\text{closest_to}(i)}\|^2$$

- **Need to find cluster centers c_k .**
 $c_1 = ? , c_2 = ? , \dots , c_K = ?$

K-means

$$\operatorname{argmin}_{c_1, \dots, c_K} \sum_i \|x_i - c_{\text{closest_to}(i)}\|^2$$

- **Need to find cluster centers c_k .**
 $c_1 = ? , c_2 = ? , \dots , c_K = ?$
- **Introduce *latent variables* (one for each x_i)**
 $a_i = \text{closest_cluster_center}(i)$
 $a_1 = ? , a_2 = ? , a_3 = ? , \dots , a_n = ?$


K-means

$$\operatorname{argmin}_{c_1, \dots, c_K} \sum_i \|x_i - c_{\text{closest_to}(i)}\|^2$$


- **For fixed c_k we can find optimal a_i**
- **For fixed a_i we can find optimal c_k .**
- **Iterate to convergence.**



Fuzzy vs Hard



Each object belongs to each cluster with some weight (the weight can be zero)



Each object belongs to exactly one cluster

Gaussian Mixture Modeling

$$\mathbf{X} \sim [N(\boldsymbol{\mu}_1, \sigma_1^2) \text{ or } N(\boldsymbol{\mu}_2, \sigma_2^2)]$$

Given \mathbf{X} , estimate $\boldsymbol{\mu}_i, \sigma_i^2$

Gaussian Mixture Modeling

$$X \sim [N(\boldsymbol{\mu}_1, \sigma_1^2) \text{ or } N(\boldsymbol{\mu}_2, \sigma_2^2)]$$


Given X , estimate μ_i, σ_i^2

 **MLE**


Gaussian Mixture Modeling

$$X \sim [N(\boldsymbol{\mu}_1, \sigma_1^2) \text{ or } N(\boldsymbol{\mu}_2, \sigma_2^2)]$$

Given X , estimate μ_i, σ_i^2



MLE



Expectation-Maximization (EM)

SKLearn's Clustering

```
from sklearn.cluster
import
    Ward,
    KMeans,
    DBScan,
    MeanShift,
    SpectralClustering,
    AffinityPropagation
```

SKLearn's Clustering

```
from sklearn.cluster
```

```
import
```

```
Ward,  
KMeans,  
DBScan,  
MeanShift,
```

Use feature vectors

```
SpectralClustering,  
AffinityPropagation
```


Use distance matrix

Quiz


- ▶ Fuzzy clustering means that _____
- ▶ K-means finds a set of cluster centers, which have the smallest _____
- ▶ K-means can get stuck in a local minimum (Y/N)?

Unsupervised learning patterns


Decomposition



Clustering




Canonical basis




$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \alpha \begin{pmatrix} 1 \\ 0 \end{pmatrix} + \beta \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

Alternative basis



$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \alpha \begin{pmatrix} 0.9 \\ -0.1 \end{pmatrix} + \beta \begin{pmatrix} 0.1 \\ 0.9 \end{pmatrix}$$

Alternative basis







$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \alpha \begin{pmatrix} 0.4 \\ -0.6 \end{pmatrix} + \beta \begin{pmatrix} 0.6 \\ 0.4 \end{pmatrix}$$

Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \cdots + \alpha_m \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix}$$





Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \cdots + \alpha_m \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix}$$




Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 0.0 \\ 0.1 \\ 0.1 \\ 0.2 \\ \vdots \\ 0.0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0.3 \\ 0.2 \\ 0.2 \\ 0.1 \\ \vdots \\ 0.3 \end{pmatrix} + \alpha_m \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ 0.0 \\ \vdots \\ 0.0 \end{pmatrix}$$



Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 0.0 \\ 0.1 \\ 0.1 \\ 0.2 \\ \vdots \\ 0.0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0.3 \\ 0.2 \\ 0.2 \\ 0.1 \\ \vdots \\ 0.3 \end{pmatrix} + \alpha_m \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ 0.0 \\ \vdots \\ 0.0 \end{pmatrix}$$



Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \dots + \alpha_m \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 1 \end{pmatrix}$$




Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 0.0 \\ 0.1 \\ 0.1 \\ 0.2 \\ \vdots \\ 0.0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0.3 \\ 0.2 \\ 0.2 \\ 0.1 \\ \vdots \\ 0.3 \end{pmatrix} + \alpha_m \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ 0.0 \\ \vdots \\ 0.0 \end{pmatrix}$$



Linear Decomposition

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \\ x_{100000} \end{pmatrix} = \alpha_1 \begin{pmatrix} 0.0 \\ 0.1 \\ 0.1 \\ 0.2 \\ \vdots \\ 0.0 \end{pmatrix} + \alpha_2 \begin{pmatrix} 0.3 \\ 0.2 \\ 0.2 \\ 0.1 \\ \vdots \\ 0.3 \end{pmatrix} + \alpha_m \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ 0.0 \\ \vdots \\ 0.0 \end{pmatrix}$$



Linear Decomposition

$$\mathbf{x} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \cdots + \alpha_m \mathbf{v}_m$$

Linear Decomposition

$$\mathbf{x} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \cdots + \alpha_m \mathbf{v}_m$$

$$\mathbf{x} = \begin{pmatrix} \vdots & \vdots & \cdots & \vdots \\ \mathbf{v}_1 & \mathbf{v}_2 & \cdots & \mathbf{v}_m \\ \vdots & \vdots & \cdots & \vdots \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \vdots \\ \alpha_m \end{pmatrix}$$

Linear Decomposition

$$\mathbf{x} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \cdots + \alpha_m \mathbf{v}_m$$

$$\mathbf{x} = \mathbf{V} \boldsymbol{\alpha}$$

Linear Decomposition

$$\mathbf{x} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \cdots + \alpha_m \mathbf{v}_m$$

$$\mathbf{x} = \mathbf{V} \boldsymbol{\alpha}$$

$$\boldsymbol{\alpha} = ?$$


Linear Decomposition

$$\mathbf{x} = \alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2 + \cdots + \alpha_m \mathbf{v}_m$$


$$\mathbf{x} = \mathbf{V} \boldsymbol{\alpha}$$

$$\boldsymbol{\alpha} = \mathbf{V}^+ \mathbf{x}$$


How do we find a good basis?




Linear projection




Linear projection



Linear projection




Linear projection



Idea: Maximize projection variance

- ▶ For a point x_i and a unit basis vector v the length of projection of x_i onto v is given by

$$p = \langle v, x_i \rangle = v^T x_i$$



Projection variance

$$p_i = v^T x_i$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \frac{1}{n} \sum_i (p_i - \bar{p})^2$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \frac{1}{n} \sum_i (p_i - \bar{p})^2$$

$$\mathbf{v} = \operatorname{argmax}_v \sigma_v^2$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \frac{1}{n} \sum_i (p_i - \bar{p})^2$$

Pre-center data, so that $\bar{p} = \mathbf{v}^T \bar{\mathbf{x}} = 0$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \frac{1}{n} \sum_i (p_i)^2$$

Pre-center data, so that $\bar{p} = \mathbf{v}^T \bar{\mathbf{x}} = 0$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \frac{1}{n} \sum_i (p_i)^2 = \frac{1}{n} \|\mathbf{p}\|^2$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\begin{aligned}\sigma_{\mathbf{v}}^2 &= \frac{1}{n} \sum_i (p_i)^2 = \frac{1}{n} \|\mathbf{p}\|^2 \\ &= \frac{1}{n} \|\mathbf{X}\mathbf{v}\|^2\end{aligned}$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \dots$$

$$= \frac{1}{n} \|\mathbf{X}\mathbf{v}\|^2 = \frac{1}{n} (\mathbf{X}\mathbf{v})^T (\mathbf{X}\mathbf{v})$$

$$= \frac{1}{n} \mathbf{v}^T \mathbf{X}^T \mathbf{X} \mathbf{v} = \mathbf{v}^T \boldsymbol{\Sigma} \mathbf{v}$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \mathbf{v}^T \mathbf{\Sigma} \mathbf{v}$$

Projection variance

$$p_i = \mathbf{v}^T \mathbf{x}_i$$

$$\sigma_v^2 = \mathbf{v}^T \Sigma \mathbf{v}$$

Data covariance matrix
 $\mathbf{X}^T \mathbf{X}$

Objective function

$$\operatorname{argmax}_{\boldsymbol{v}} \boldsymbol{v}^T \boldsymbol{\Sigma} \boldsymbol{v}$$

$$s. t. \|\boldsymbol{v}\|^2 = 1$$

Optimization

$$\operatorname{argmax}_{\boldsymbol{v}} \boldsymbol{v}^T \boldsymbol{\Sigma} \boldsymbol{v}$$

$$s. t. \|\boldsymbol{v}\|^2 = 1$$

Method of Lagrange multipliers...

$$\boldsymbol{\Sigma} \boldsymbol{v} = \lambda \boldsymbol{v}$$

Optimization

$$\operatorname{argmax}_v v^T \Sigma v$$

$$s. t. \|v\|^2 = 1$$

Method of Lagrange multipliers...

$$\Sigma \boxed{v} = \boxed{\lambda} v$$

Eigenvector of Σ

Eigenvalue

Example


```
Xc = X - mean(X, axis=0)
```

```
Sigma = Xc.T * Xc / n
```


```
(= cov(Xc, rowvar=0) )
```

```
lambdas, vs = eig(Sigma)
```



Example




Example




Example



Example



Example



Principal Components Analysis

Principal components are the **eigenvectors** of the **covariance matrix**.

$$V, \lambda = \text{eig}(\Sigma)$$

Principal Components Analysis


Principal components are the **eigenvectors** of the **covariance matrix**.

$$V, \lambda = \text{eig}(\Sigma)$$

For each PC, the corresponding eigenvalue λ_i shows the **amount of variance explained** by the component.

Principal Components Analysis

Eigenvalue spectrum of Σ



Principal Components Analysis

Data projection onto PC i :

$$\mathbf{p} = \mathbf{X}\mathbf{v}_i$$

Data projection onto multiple PCs:

$$\mathbf{X}_{\text{proj}} = \mathbf{X}\mathbf{V}_*$$

Data reconstruction from PC coordinates:

$$\mathbf{X}_{\text{proj}}\mathbf{V}_*^T = \mathbf{X}$$

SKLearn's PCA

```
from sklearn.decomposition
                                import PCA


model = PCA(n_components=2)
model.fit(X)
X_t = model.transform(X)

model.components_[1, :]
```

SKLearn's PCA


```
from sklearn.decomposition
import
    PCA,
    SparsePCA,
    ProbabilisticPCA,
    KernelPCA,
    FastICA,
    NMF,
    DictionaryLearning,
    . . .
```

PCA: Geometric intuition



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


PCA: Geometric intuition



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

$$\mathbf{X}' = \mathbf{X} \begin{pmatrix} 5 & 0 \\ 0 & 0.9 \end{pmatrix}$$

PCA: Geometric intuition




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

$$\mathbf{X}' = \mathbf{X} \begin{pmatrix} 5 & 0 \\ 0 & 0.9 \end{pmatrix}$$

$$\begin{aligned} \mathbf{X}'' &= \mathbf{X}' \begin{pmatrix} \cos 0.8 & -\sin 0.8 \\ \sin 0.8 & \cos 0.8 \end{pmatrix} \end{aligned}$$


PCA: Geometric intuition



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

$$\mathbf{X}'' = \mathbf{X} \cdot \mathbf{D} \cdot \mathbf{R}$$

PCA: Geometric intuition




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

$$\mathbf{X}'' = \mathbf{X} \cdot \mathbf{D} \cdot \mathbf{R}$$

$$\begin{aligned} & (\mathbf{X}'')^T (\mathbf{X}'') \\ &= (\mathbf{XDR})^T (\mathbf{XDR}) \end{aligned}$$

PCA: Geometric intuition




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

$$\mathbf{X}'' = \mathbf{X} \cdot \mathbf{D} \cdot \mathbf{R}$$

$$\begin{aligned} & (\mathbf{X}'')^T (\mathbf{X}'') \\ &= (\mathbf{XDR})^T (\mathbf{XDR}) \\ &= \mathbf{R}^T \mathbf{D}^T \mathbf{X}^T \mathbf{XDR} \end{aligned}$$

PCA: Geometric intuition




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

$$\mathbf{X}'' = \mathbf{X} \cdot \mathbf{D} \cdot \mathbf{R}$$

$$\begin{aligned} & (\mathbf{X}'')^T (\mathbf{X}'') \\ &= (\mathbf{XDR})^T (\mathbf{XDR}) \\ &= \mathbf{R}^T \mathbf{D}^T \mathbf{X}^T \mathbf{X} \mathbf{D} \mathbf{R} \\ &= \mathbf{R}^T \mathbf{D}^2 \mathbf{R} \end{aligned}$$

PCA: Geometric intuition




$$X \sim N(0,1)$$

$$X'' = X \cdot D \cdot R$$

$$\begin{aligned} & (X'')^T (X'') \\ &= (XDR)^T (XDR) \\ &= R^T D^T X^T X D R \\ &= R^T D^2 R \end{aligned}$$


PCA: Geometric intuition




$$X \sim N(0,1)$$

$$X'' = X \cdot D \cdot R$$

$$\begin{aligned} (X'')^T (X'') &= (X D R)^T (X D R) \\ &= R^T D^T X^T X D R \\ &= R^T D^2 R \end{aligned}$$





Quiz

- ▶ Principal components are _____ of the _____ matrix.
- ▶ Eigenvalue spectrum shows how much _____ is explained by each _____.
- ▶ If $\Sigma = V\Lambda V^T$, then
$$X_{\text{proj}} = \underline{\hspace{2cm}}$$

