

Introduction to Machine Learning

NPFL 054

<http://ufal.mff.cuni.cz/course/npfl054>

Barbora Hladká
hladka@ufal.mff.cuni.cz

Martin Holub
holub@ufal.mff.cuni.cz

Charles University,
Faculty of Mathematics and Physics,
Institute of Formal and Applied Linguistics

Outline

- Model complexity, overfitting, bias and variance
- Regularization
 - Ridge regression
 - Lasso
 - Linear regression
 - Logistic regression
 - SVM
- Principal Component Analysis

- Suppose features A_1, \dots, A_m and a set of possible target values Y
- Suppose development data as a set of instances


$$Data = \{(\mathbf{x}_i, y_i), \mathbf{x}_i = \langle x_i^1, \dots, x_i^m \rangle, y_i \in Y\}$$

where \mathbf{x}_i is a feature vector and y_i is its true target value

Let h^* be a best approximation of c trained on $Data$.

Model complexity and overfitting

Finding a model that minimizes generalization error
... is one of central goals of the machine learning process




Model complexity and overfitting

No universal definition

Here ... **model complexity** is the number of hypothesis parameters

$$\Theta = \langle \Theta_0, \dots, \Theta_m \rangle$$

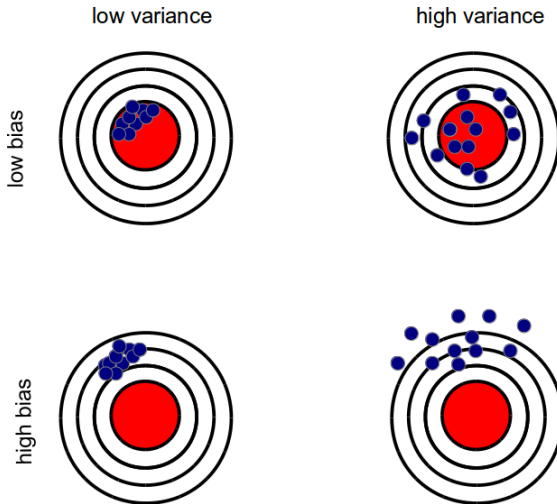
Model complexity and overfitting



Bias and variance

- 1 Select a machine learning algorithm
 - 2 Get k different training sets
 - 3 Get k predictors h_1^*, \dots, h_k^*
- **Bias** measures error that originates from the learning algorithm
 - how far off in general the predictions by k predictors are from the true output value
 - **Variance** measures error that originates from the training data
 - how much the predictions for a test instance vary between k predictors

Bias and variance



Generalization error $\text{error}_{\mathcal{D}}(h)$ measures how well a hypothesis h generalizes beyond the used training data set, to unseen data with distribution \mathcal{D} . Usually it is defined as follows

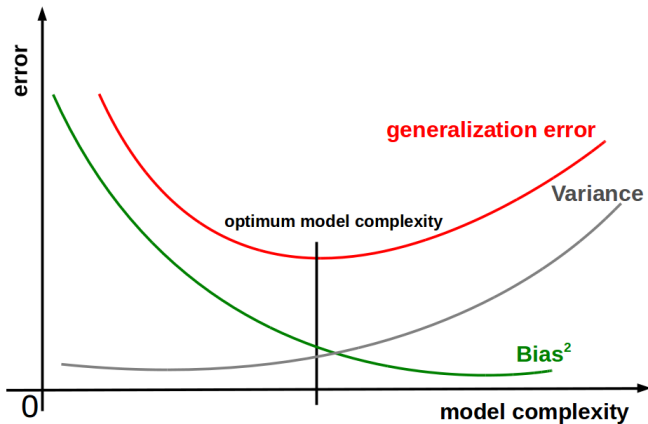
- for **regression**: $\text{error}_{\mathcal{D}}(h) = \mathbb{E}(\hat{y}_i - y_i)^2$
- for **classification**: $\text{error}_{\mathcal{D}}(h) = \Pr(\hat{y}_i \neq y_i)$

Decomposition of $\text{error}_{\mathcal{D}}(h)$

$$\text{error}_{\mathcal{D}}(h) = \text{Bias}^2 + \text{Variance}$$

Bias and variance

- underfitting = high bias
- overfitting = high variance




Bias and variance


k-Nearest Neighbor

- $\uparrow k \rightarrow \downarrow$ variance and \uparrow bias
- $\downarrow k \rightarrow \uparrow$ variance and \downarrow bias
- Increasing k "simplifies" decision boundary (averaging more instances)

1-nearest neighbour




5-nearest neighbour




Bias and variance

k-Nearest Neighbor

5-nearest neighbour




8-nearest neighbour




Bias and variance

k-Nearest Neighbor

8-nearest neighbour




10-nearest neighbour




Bias and variance

k-Nearest Neighbor

10-nearest neighbour



15-nearest neighbour



Regularization

We want a model in between which is

- powerful enough to model the underlying structure of data
- not so powerful to model the structure of the training data

Let's prevent overfitting by **complexity regularization**, a technique that regularizes the parameter estimates, or equivalently, shrinks the parameter estimates towards zero.

- A machine learning algorithm estimates hypothesis parameters

$$\Theta = \langle \Theta_0, \Theta_1, \dots, \Theta_m \rangle$$

using Θ^* that minimizes loss function for the data D

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \operatorname{loss}(\Theta)$$

- **Regularization**

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \operatorname{loss}(\Theta) + \lambda * \mathbf{penalty}(\Theta)$$

where $\lambda \geq 0$ is a **tuning parameter**

Regularization – Ridge regression

$$\text{penalty}(\boldsymbol{\Theta}) = \Theta_1^2 + \dots + \Theta_m^2$$

$\Theta_1^2 + \dots + \Theta_m^2$ is the ℓ_2 norm

$$\boldsymbol{\Theta}^* = \underset{\boldsymbol{\Theta}}{\text{argmin}} \text{loss}(\boldsymbol{\Theta}) + \lambda * (\Theta_1^2 + \dots + \Theta_m^2)$$


The penalty is applied to $\Theta_1, \dots, \Theta_m$, but not to Θ_0 since the goal is to regularize the estimated association between each feature and the target value.

Ridge regression

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \operatorname{loss}(\Theta) + \lambda * (\Theta_1^2 + \dots + \Theta_m^2)$$

- **Let** $\Theta_{\lambda 1}^*, \dots, \Theta_{\lambda m}^*$ be ridge regression parameter estimates for a particular value of λ
- **Let** $\Theta_1^*, \dots, \Theta_m^*$ be unregularized parameter estimates
- $0 \leq \frac{\Theta_{\lambda 1}^{*2} + \dots + \Theta_{\lambda m}^{*2}}{\Theta_1^{*2} + \dots + \Theta_m^{*2}} \leq 1$
- **When** $\lambda = 0$, **then** $\Theta_{\lambda i}^* = \Theta_i^*$ for $i = 1, \dots, m$
- **When** λ is extremely large, **then** $\Theta_{\lambda i}^*$ is very small for $i = 1, \dots, m$
- **When** λ between, we are fitting a model and shrinking the parameters

Ridge regression



Regularization – Lasso


$$\text{penalty}(\Theta) = |\Theta_1| + \dots + |\Theta_m|$$

$|\Theta_1| + \dots + |\Theta_m|$ is the ℓ_1 norm

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \operatorname{loss}(\Theta) + \lambda * (|\Theta_1| + \dots + |\Theta_m|)$$

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \operatorname{loss}(\Theta) + \lambda * (|\Theta_1| + \dots + |\Theta_m|)$$

- **Let** $\Theta_{\lambda 1}^*, \dots, \Theta_{\lambda m}^*$ be lasso regression parameter estimates
- **Let** $\Theta_1^*, \dots, \Theta_m^*$ be unregularized parameter estimates
- **When** $\lambda = 0$, **then** $\Theta_{\lambda i}^* = \Theta_i^*$ for $i = 1, \dots, m$
- **When** λ grows, **then** the impact of penalty grows
- **When** λ is extremely large, **then** $\Theta_{\lambda i}^* = 0$ for $i = 1, \dots, m$



Difference between Ridge regression and Lasso

Ridge regression shrinks all the parameters but eliminates none, while the Lasso can shrink some parameters to zero.

Loss function

A loss function $L(\hat{y}, y)$ measures the cost of predicting \hat{y} when the true value is $y \in \{-1, +1\}$. Commonly used loss functions are

- **Squared** (RSS) $L(\hat{y}, y) = (y - \hat{y})^2$
- **Zero-one** (0/1) $L(\hat{y}, y) = I(y\hat{y} \leq 0)$
indicator variable I is 1 if $y\hat{y} \leq 0$, 0 otherwise
- **Hinge** $L(\hat{y}, y) = \max(0, 1 - y\hat{y})$
- **Logistic** $L(\hat{y}, y) = \max(0, \log(1 + e^{-y\hat{y}}))$
- **Exponential** $L(\hat{y}, y) = e^{-y\hat{y}}$

Recap of linear regression

Linear regression is a regression algorithm

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \sum_{i=1}^n (h(\mathbf{x}_i) - y_i)^2$$

where

- $h(\mathbf{x}) = \Theta_0 + \Theta_1 x_1 + \dots + \Theta_m x_m$
- loss function = residual sum of squares

Intepretation of Θ

- $h(\mathbf{x}) = \Theta_0 + \Theta_1 x_1 + \cdots + \Theta_m x_m$
- Θ_j gives an average change in a target value with one-unit change in feature A_j , holding other features fixed

Regularized linear regression

$$h(\mathbf{x}) = \Theta_0 + \Theta_1 x_1 + \cdots + \Theta_m x_m$$

$$\text{loss}(\Theta) = \text{RSS} = \sum_{i=1}^n (h(\mathbf{x}_i) - y_i)^2$$


$$\Theta^* = \underset{\Theta}{\text{argmin}} \sum_{i=1}^n (h(\mathbf{x}_i) - y_i)^2 + \lambda * \text{penalty}(\Theta)$$

Ridge regression – alternative formulation

$$\Theta^* = \operatorname{argmin}_{\Theta} \sum_{i=1}^n (h(\mathbf{x}_i) - y_i)^2$$


$$\text{subject to } \Theta_1^2 + \dots + \Theta_m^2 \leq s$$

- the gray circle represents the feasible region for Ridge regression
- the contours represent different loss values for the unregularized model



Ridge regression – alternative formulation

- If s is large enough so that the minimum loss value falls into the region of **ridge regression** parameter estimates then the alternative formulation yields the primary solution.




Lasso – alternative formulation

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} \sum_{i=1}^n (h(\mathbf{x}_i) - y_i)^2$$

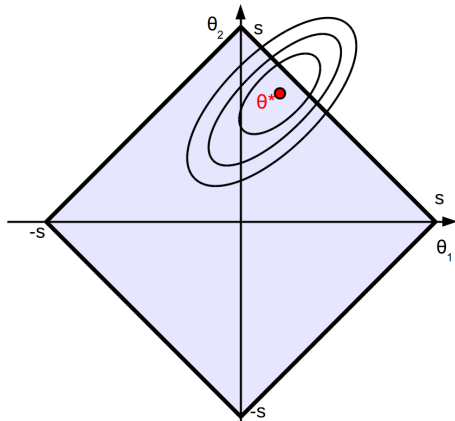
subject to $|\Theta_1| + \dots + |\Theta_m| \leq s$

- the grey square represents the feasible region of the Lasso
- the contours represent different loss values for the unregularized model
- the feasible point that minimizes the loss is more likely to happen on the coordinates on the Lasso graph than on the Ridge regression graph since the Lasso graph is more angular



Lasso – alternative formulation

- If s is large enough so that the minimum loss value falls into the region of **loss** parameter estimates then the alternative formulation yields the primary solution.



Recap of logistic regression

Logistic regression is a classification algorithm

Assume $Y = \{0, 1\}$

- **modeling the probability** $h(\mathbf{x}) = \Pr(Y = 1 | \mathbf{x}; \Theta)$

$$h(\mathbf{x}) = g(\Theta^T \mathbf{x}) = \frac{1}{1 + e^{-\Theta^T \mathbf{x}}}, \text{ where } \Theta = \langle \Theta_0, \dots, \Theta_m \rangle$$

- **prediction function** of \mathbf{x}

$$= \begin{cases} 1 & \text{if } h(\mathbf{x}) \geq 0.5 \\ 0 & \text{if } h(\mathbf{x}) < 0.5 \end{cases}$$

Recap of logistic regression

- $\frac{h(\mathbf{x})}{1 - h(\mathbf{x})} = \text{odds ratio}$
- **log odds is linear**

$$\log \frac{h(\mathbf{x})}{1 - h(\mathbf{x})} = \Theta^T \mathbf{x}$$

- **recall linear regression**

$$h(\mathbf{x}) = \Theta^T \mathbf{x}$$

Recap of logistic regression

Interpretation of Θ

Suppose $\Theta = \langle \Theta_0, \Theta_1 \rangle$

- linear regression $h(\mathbf{x}) = \Theta_0 + \Theta_1 x_1$: Θ_1 gives an average change in a target value with one-unit change in A_1
- logistic regression $\log \frac{h(\mathbf{x})}{1-h(\mathbf{x})} = \Theta_0 + \Theta_1 x_1$: Θ_1 gives an average change in logit $h(\mathbf{x})$ with one-unit change in A_1

Recap of logistic regression

Estimating Θ by maximizing the likelihood

- loss function

$$L(\Theta) = \sum_{i=1}^n y_i \log P(y_i | \mathbf{x}_i; \Theta) + (1 - y_i) \log(1 - P(y_i | \mathbf{x}_i; \Theta))$$

- optimization task

$$\begin{aligned}\Theta^* &= \operatorname{argmax}_{\Theta} L(\Theta) \\ &= \operatorname{argmin}_{\Theta} -L(\Theta) \\ &= \operatorname{argmin}_{\Theta} \sum_{i=1}^n -y_i \log P(y_i | \mathbf{x}_i; \Theta) - (1 - y_i) \log(1 - P(y_i | \mathbf{x}_i; \Theta))\end{aligned}$$

Recap of logistic regression

Multinomial logistic regression $Y = \{y_1, \dots, y_k\}$

- train k one-versus-all binary classifiers h_i^* , $i = 1, \dots, k$
- classify \mathbf{x} into the class K that maximizes $h_K^*(\mathbf{x})$

Regularized logistic regression

$$\Theta^* = \underset{\Theta}{\operatorname{argmin}} -L(\Theta) + \lambda * \operatorname{penalty}(\Theta)$$

Logistic regression with Ridge regression

$$L(\Theta) = -\left[\sum_{i=1}^n y_i \log(h(\mathbf{x}_i)) + (1 - y_i) \log(1 - h(\mathbf{x}_i))\right] + \lambda \sum_{j=1}^m \Theta_j^2$$

$$\Theta^* = \operatorname{argmin}_{\Theta} L(\Theta)$$

SVM and Logistic regression

Logistic regression with Ridge regression

$$\begin{aligned}L(\Theta) &= -\left[\sum_{i=1}^n y_i \log(h(\mathbf{x}_i)) + (1 - y_i) \log(1 - h(\mathbf{x}_i))\right] + \lambda \sum_{j=1}^m \Theta_j^2 = \\&= \sum_{i=1}^n y_i (-\log(h(\mathbf{x}_i))) + (1 - y_i) (-\log(1 - h(\mathbf{x}_i))) + \lambda \sum_{j=1}^m \Theta_j^2 = \\&= \sum_{i=1}^n y_i L_1(\Theta) + (1 - y_i) L_0(\Theta) + \lambda \sum_{j=1}^m \Theta_j^2\end{aligned}$$

$$\mathbf{A} + \lambda \mathbf{B} \equiv \mathbf{C} \mathbf{A} + \mathbf{B}, \mathbf{C} = \frac{1}{\lambda}$$

$$\operatorname{argmin}_{\Theta} L(\Theta) = \operatorname{argmin}_{\Theta} \sum_{j=1}^m \Theta_j^2 + \mathbf{C} \left[\sum_{i=1}^n y_i L_1(\Theta) + (1 - y_i) L_0(\Theta) \right]$$

where $L_1(\Theta) = -\log \frac{1}{1+e^{-\Theta^T \mathbf{x}}}$ and $L_0(\Theta) = -\log(1 - \frac{1}{1+e^{-\Theta^T \mathbf{x}}})$

SVM and regularized logistic regression

- Regularized logistic regression

$$\operatorname{argmin}_{\Theta} \sum_{j=1}^m \Theta_j^2 + C \sum_{i=1}^n \log(1 + e^{-\bar{y}_i \Theta^T \mathbf{x}_i})$$

where

$$\bar{y}_i = \begin{cases} -1 & \text{if } y_i = 0 \\ 1 & \text{if } y_i = 1 \end{cases}$$

- SVM

$$\operatorname{argmin}_{\Theta} \sum_{j=1}^m \Theta_j^2 + C \sum_{i=1}^n \max(0, 1 - y_i \Theta^T \mathbf{x}_i)$$

Soft-margin is equivalent to the regularization problem

Hinge loss: $\max(0, 1 - y_i \Theta^T \mathbf{x})$

- 1 $y_i \Theta^T \mathbf{x}_i > 1$: no contribution to loss
- 2 $y_i \Theta^T \mathbf{x}_i = 1$: no contribution to loss
- 3 $y_i \Theta^T \mathbf{x}_i < 1$: contribution to loss

SVM and regularized logistic regression

- $\xi_i \geq 0$ is equivalent to $\xi_i = \max(0, 1 - y_i \Theta^T \mathbf{x}_i)$
- $\operatorname{argmin}_{\Theta} L(\Theta) = \operatorname{argmin}_{\Theta} C \sum_{i=1}^n \max(0, 1 - y_i \Theta^T \mathbf{x}_i) + \sum_{j=1}^m \Theta_j^2 =$


$$= \operatorname{argmin}_{\Theta} C \sum_{i=1}^n \xi_i + \sum_{j=1}^m \Theta_j^2$$

s.t. $\Theta^T \mathbf{x}_i \geq 1 - \xi_i$ if $y_i = 1$ and $\Theta^T \mathbf{x}_i \leq -1 + \xi_i$ if $y_i = -1$

Principal Component Analysis (PCA)

- a tool to analyze the data
- a tool to do dimensionality reduction

Auto data set



Basic concepts needed

- data analysis: measures of center and spread, covariance and correlation
- linear algebra: eigenvectors, eigenvalues, dot product, basis

How two features are related

Both covariance and correlation indicate how closely two features relationship follows a straight line.

- **Covariance** measures the degree of the linear relationship between two features

$$\text{cov}(X, Y) = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})$$

- > 0 both features increase or decrease together
- < 0 while one feature increases the other decreases
- $= 0$ features are linearly independent of each other

- **Covariance matrix** of features A_1, \dots, A_m represents covariance among them

$$\text{COV}(A_1, \dots, A_m) = \begin{pmatrix} \text{var}(A_1) & \text{cov}(A_1, A_2) & \dots & \text{cov}(A_1, A_m) \\ \text{cov}(A_2, A_1) & \text{var}(A_2) & \dots & \text{cov}(A_2, A_m) \\ \dots & \dots & \dots & \dots \\ \text{cov}(A_m, A_1) & \text{cov}(A_m, A_2) & \dots & \text{var}(A_m) \end{pmatrix}$$

How two features are related

- **Correlation** measures the degree to which the features tend to move together.

$$-1 \leq \text{cor}(X, Y) = \frac{\text{cov}(X, Y)}{s_X s_Y} \leq 1$$

Data analysis

Auto data set

```
> cov(Auto[c("mpg", "cylinders", "horsepower", "weight")])  
  
#           mpg  cylinders horsepower    weight  
# mpg          60.91814 -10.352928 -233.85793 -5517.441  
# cylinders    -10.35293   2.909696   55.34824  1300.424  
# horsepower  -233.85793   55.348244 1481.56939 28265.620  
# weight      -5517.44070 1300.424363 28265.62023 721484.709  
  
> cor(Auto[c("mpg", "cylinders", "horsepower", "weight")])  
  
#           mpg  cylinders horsepower    weight  
# mpg          1.0000000 -0.7776175 -0.7784268 -0.8322442  
# cylinders    -0.7776175  1.0000000  0.8429834  0.8975273  
# horsepower  -0.7784268  0.8429834  1.0000000  0.8645377  
# weight      -0.8322442  0.8975273  0.8645377  1.0000000
```


- **Eigenvector \mathbf{u} , eigenvalue λ : $\mathbf{A}\mathbf{u} = \lambda\mathbf{u}$**
 - \mathbf{u} does not change its direction under the transformation
 - $\lambda\mathbf{u}$ scales a vector \mathbf{u} by λ ; it changes its length, not its direction
- ① The covariance matrix of an $n \times m$ matrix \mathbf{X} is an $m \times m$ symmetric matrix given by $\frac{1}{n-1}\mathbf{X}\mathbf{X}^T$
- ② Any symmetric matrix $m \times m$ has a set of orthonormal eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m$ and associated eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_m$
 - for any i , $\mathbf{A}\mathbf{v}_i = \lambda_i\mathbf{v}_i$
 - $\|\mathbf{v}_i\| = 1$
 - $\mathbf{v}_i\mathbf{v}_j = 0$ if $i \neq j$
- ③ \mathbf{A} is a symmetric $m \times m$ matrix and \mathbf{E} is an $m \times m$ matrix whose i -th column is the i -th eigenvector of \mathbf{A} . The eigenvectors are ordered in terms of decreasing values of their associated eigenvalues. Then there is a diagonal matrix \mathbf{D} such that $\mathbf{A} = \mathbf{E}\mathbf{D}\mathbf{E}^T$
- ④ If the rows of \mathbf{E} are orthogonal, then $\mathbf{E}^{-1} = \mathbf{E}^T$

- **Dot product** of $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{R}^m$: $\mathbf{x}_1 \mathbf{x}_2 = \sum_{i=1}^m x_{1_i} x_{2_i}$
- **Basis** of \mathcal{R}^m is a set of linearly independent vectors $\mathbf{u}_1, \dots, \mathbf{u}_m$
 - none of them is a linear combination of other vectors
 - $\mathbf{u}_i \mathbf{u}_j = 0, i, j = 1, \dots, m, i \neq j$
 - any $\mathbf{u} = c_1 \mathbf{u}_1 + \dots + c_m \mathbf{u}_m$
 - for example, the standard basis of the 3-dimensional Euclidean space \mathcal{R}^3 consists of $\mathbf{x} = \langle 1, 0, 0 \rangle, \mathbf{y} = \langle 0, 1, 0 \rangle, \mathbf{z} = \langle 0, 0, 1 \rangle$. It is an example of orthonormal basis, so called *naive* basis **!**

Principal Component Analysis

- **instances** $Data = \{\mathbf{x}_i; \mathbf{x}_i \in \mathcal{R}^m\}, |Data| = n$
- **features** $Attr = \{A_1, \dots, A_m\}$
- **representation of $Data$ for PCA derivation**

$$\mathbf{X} = \begin{pmatrix} x_{11} & \dots & x_{1n} \\ x_{21} & \dots & x_{2n} \\ \dots & \dots & \dots \\ x_{m1} & \dots & x_{mn} \end{pmatrix}$$



Which features to keep?

- features that change a lot, i.e. high variance
- features that do not depend on others, i.e. low covariance

Which features to ignore?

- features with some noise, i.e. low variance

PCA principles

- ① high correlation \sim high redundancy
- ② the most important feature has the largest variance

- **Question**

Is there any other representation of \mathbf{X} to extract the most important features?

- **Answer**

Another basis

$$\mathbf{P}^T \mathbf{X} = \mathbf{Z}$$

where \mathbf{P} transforms \mathbf{X} into \mathbf{Z}

Heading for the \mathbf{P} matrix


$$\mathbf{P} = \begin{pmatrix} \mathbf{p}_{11} & \dots & \dots & \mathbf{p}_{1m} \\ \mathbf{p}_{21} & \dots & \dots & \mathbf{p}_{2m} \\ \dots & \dots & \dots & \dots \\ \mathbf{p}_{m1} & \dots & \dots & \mathbf{p}_{mm} \end{pmatrix}$$

- **Principal components** of \mathbf{X} are the vectors $\mathbf{p}_i = \langle p_{1i}, \dots, p_{mi} \rangle$
- **Principal component loadings** of \mathbf{p}_i are the elements p_{i1}, \dots, p_{im}

Heading for P

$$\mathbf{Z} = \begin{pmatrix} \mathbf{p}_1\mathbf{x}_1 & \dots & \dots & \mathbf{p}_1\mathbf{x}_n \\ \mathbf{p}_2\mathbf{x}_1 & \dots & \dots & \mathbf{p}_2\mathbf{x}_n \\ \dots & \dots & \dots & \dots \\ \mathbf{p}_m\mathbf{x}_1 & \dots & \dots & \mathbf{p}_m\mathbf{x}_n \end{pmatrix}$$

i -principal component scores of n instances are $\mathbf{p}_i\mathbf{x}_1, \mathbf{p}_i\mathbf{x}_2, \dots, \mathbf{p}_i\mathbf{x}_n$



Heading for \mathbf{P}

- What is a good choice of \mathbf{P} ?
- What features we would like \mathbf{Z} to exhibit?

Goal: \mathbf{Z} is a new representation of \mathbf{X}

The new features are linear combinations of the original features whose weights are given by \mathbf{P} .

The covariance matrix of \mathbf{Z} is diagonal and the entries on the diagonal are in descending order, i.e. the covariance of any pair of distinct features is zero, and the variance of each of our new features is listed along the diagonal.

Heading for P

- principal components are new basis vectors to represent \mathbf{x}_j , $j = 1, \dots, n$
- $\mathbf{p}_i \mathbf{x}_j$ is a projection of \mathbf{x}_j on \mathbf{p}_i
- changing the basis does not change data, it changes their representation

The covariance matrix $\text{cov}(A_1, A_2, \dots, A_m)$:

- on the diagonal, large values correspond to interesting structure
- off the diagonal, large values correspond to high redundancy

Derivation of PCA

- 1 preprocessing *Data*
mean normalization to get centered data $\rightarrow \mathbf{X}$
- 2 $\text{cov}(\mathbf{X}) = \mathbf{A} = \frac{1}{n-1} \mathbf{X}\mathbf{X}^T$
- 3 Compute eigenvectors $\mathbf{v}_1, \dots, \mathbf{v}_m$ and eigenvalues $\lambda_1, \dots, \lambda_m$ of \mathbf{A}
- 4 Take the eigenvectors, order them by eigenvalues, i.e. by significance, highest to lowest: $\mathbf{p}_1, \dots, \mathbf{p}_m, \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_m$
- 5 The principal components $\mathbf{p}_1, \dots, \mathbf{p}_m$ become columns of \mathbf{P}

$$\mathbf{p}_i = \begin{pmatrix} p_{1i} \\ \dots \\ p_{mi} \end{pmatrix}$$

Properties of PCA

$$\mathbf{P}^T \mathbf{X} = \mathbf{Z}$$

The i -th diagonal value of $\text{cov}(\mathbf{Z})$ is the variance of \mathbf{X} along \mathbf{p}_i .

$$\mathbf{Z} = \begin{pmatrix} \mathbf{p}_1 \mathbf{x}_1 & \dots & \dots & \mathbf{p}_1 \mathbf{x}_n \\ \mathbf{p}_2 \mathbf{x}_1 & \dots & \dots & \mathbf{p}_2 \mathbf{x}_n \\ \dots & \dots & \dots & \dots \\ \mathbf{p}_m \mathbf{x}_1 & \dots & \dots & \mathbf{p}_m \mathbf{x}_n \end{pmatrix}$$

- We calculate a rotation of the original coordinate system such that all non-diagonal elements of the new covariance matrix become zero.
- The eigenvectors (principal components) define the basis of the new coordinate axes and the eigenvalues correspond to the diagonal elements of the new covariance matrix.
- So the eigenvalues, by definition, define the variance along the corresponding principal components.

Properties of PCA

$$\text{cov}(\mathbf{P}^T \mathbf{X}) \stackrel{\text{see p.46.1}}{=} \frac{1}{n-1} (\mathbf{P}^T \mathbf{X})(\mathbf{P}^T \mathbf{X})^T =$$

$$\frac{1}{n-1} \mathbf{P}^T \mathbf{X} \mathbf{X}^T \mathbf{P} \stackrel{\text{let } \mathbf{A} = \mathbf{X} \mathbf{X}^T}{=} \frac{1}{n-1} \mathbf{P}^T \mathbf{A} \mathbf{P} =$$

$$\stackrel{\text{see p.46.3}}{=} \frac{1}{n-1} \mathbf{P}^T (\mathbf{P} \mathbf{D} \mathbf{P}^T) \mathbf{P} \stackrel{\text{see p.46.4}}{=} \frac{1}{n-1} \mathbf{P}^T (\mathbf{P}^T)^{-1} \mathbf{D} \mathbf{P}^T (\mathbf{P}^T)^{-1} = \frac{1}{n-1} \mathbf{D}$$

A geometric interpretation for the first principal component p_1

It defines a direction in feature space along which the data vary the most. If we project the n instances $\mathbf{x}_1, \dots, \mathbf{x}_n$ onto this direction, the projected values are the principal component scores z_{11}, \dots, z_{n1} themselves.

Proportion of Variance Explained (PVE)

How much of the information in a given data set is lost by projecting the instances onto the first few principal components?

In other words, how much of the variance in the data is not contained in the first few principal components?

- total variance in \mathbf{X} : $\sum_{j=1}^m \text{var}(A_j) = \sum_{i=1}^m \frac{1}{n} \sum_{i=1}^n x_{ij}^2$ (assuming feature normalization)
- variance expressed by \mathbf{p}_k : $\frac{1}{n} \sum_{i=1}^n z_{ki}^2$
- $\text{PVE}(\mathbf{p}_k) = \frac{\sum_{i=1}^n z_{ki}^2}{\sum_{i=1}^m \sum_{i=1}^n x_{ij}^2}$
- $\text{PVE}(\mathbf{p}_1, \dots, \mathbf{p}_M) = \sum_{i=1}^M \text{PVE}(\mathbf{p}_i)$, $M \leq m$

PCA

Auto data set

```
> a <- Auto[c("mpg", "cylinders", "horsepower", "weight")]
> pca.a <- prcomp(a, scale = TRUE)
> summary(pca.a)


# Importance of components:
#
#              Comp.1    Comp.2    Comp.3    Comp.4
Standard deviation   1.8704 0.49540 0.40390 0.30518
Proportion of Variance 0.8746 0.06135 0.04078 0.02328
Cumulative Proportion 0.8746 0.93593 0.97672 1.00000
```

PCA


Auto data set

Scree plot

Scree plot: Auto data set



Scree plot: Auto data set



PCA

Auto data set

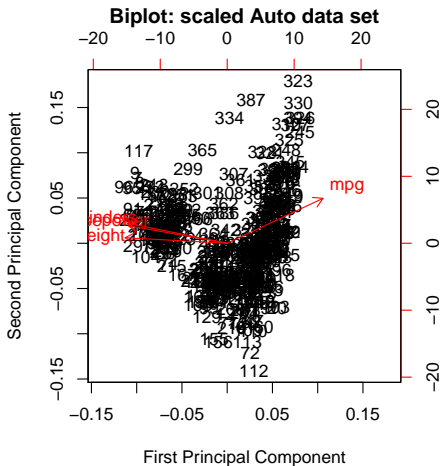
```
> pca.a$rotation
      PC1      PC2      PC3      PC4
mpg      0.4833271 0.8550485 -0.02994982 0.1854453
cylinders -0.5033993 0.3818233 -0.55748381 -0.5385276
horsepower -0.4984381 0.3346173 0.79129092 -0.1159714
weight    -0.5143380 0.1055192 -0.24934614 0.8137252
```

- PC1 places approximately equal weight on cylinders, horsepower, weight with much higher weight on mpg.
- PC2 places most of its weight on mpg and less weight on the other three features.
- Overall, cylinders, horsepower, and weight are located close to each other while mpg is far from the other three. It indicates that the three features are correlated with each other and mpg is less correlated with them.

PCA

Auto data set

A biplot displays both the PC scores and the PC loadings.



The biplot for the Auto data set is showing

- the scores of each example (i.e., cars) on the first two principal components with axes on the top and right
 - see the id cars in black
- the loading of each feature (i.e., mpg, weight, cylinders, horsepower) on the first two principal components with axes on the bottom and left
 - see the red arrows

In general, a $m \times n$ matrix \mathbf{X} has $\min(n - 1, m)$ distinct principal components.

- **Question**

How many principal components are needed?

- **Answer**

Unfortunately, there is no single answer to this question. Study scree plots.

Summary of Examination Requirements

- Bias and variance
- Lasso and Ridge regularization for linear and logistic regression
- SVM and regularized logistic regression
- Principal Component Analysis – derivation, scree plot, biplot