# Least Squares 

Pontus Giselsson

## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Machine learning

- Machine learning can very roughly be divided into:
- Supervised learning
- Unsupervised learning
- Semisupervised learning (between supervised and unsupervised)
- Reinforcement learning
- We will focus on supervised learning


## Supervised learning

- Let $(x, y)$ represent object and label pairs
- Object $x \in \mathcal{X} \subseteq \mathbb{R}^{n}$
- Label $y \in \mathcal{Y} \subseteq \mathbb{R}^{K}$
- Available: Labeled training data (training set) $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{N}$
- Data $x_{i} \in \mathbb{R}^{n}$, or examples (often $n$ large)
- Labels $y_{i} \in \mathbb{R}^{K}$, or response variables (often $K=1$ )

Objective: Find a model (function) $m(x)$ :

- that takes data (example, object) $x$ as input
- and predicts corresponding label (response variable) $y$

How?:

- learn $m$ from training data, but should generalize to all $(x, y)$


## Relation to optimization

Training the "machine" $m$ consists in solving optimization problem

## Regression vs Classification

There are two main types of supervised learning tasks:

- Regression:
- Predicts quantities
- Real-valued labels $y \in \mathcal{Y}=\mathbb{R}^{K}$ (will mainly consider $K=1$ )
- Classification:
- Predicts class belonging
- Finite number of class labels, e.g., $y \in \mathcal{Y}=\{1,2, \ldots, k\}$


## Examples of data and label pairs

| Data | Label | R/C |
| :--- | :--- | :--- |
| text in email | spam? | C |
| dna | blood cell concentration | R |
| dna | cancer? | C |
| image | cat or dog | C |
| advertisement display | click? | C |
| image of handwritten digit | digit | C |
| house address | selling cost | R |
| stock | price | R |
| sport analytics | winner | C |
| speech representation | spoken word | C |

$R / C$ is for regression or classification

## In this course

Lectures will cover different supervised learning methods:

- Classical methods with convex training problems
- Least squares (this lecture)
- Logistic regression
- Support vector machines
- Deep learning methods with nonconvex training problem

Highlight difference:

- Deep learning (specific) nonlinear model instead of linear


## Notation

- (Primal) Optimization variable notation:
- Optimization literature: $x, y, z$ (as in first part of course)
- Statistics literature: $\beta$
- Machine learning literature: $\theta, w, b$
- Reason: data, labels in statistics and machine learning are $x, y$
- Will use machine learning notation in these lectures
- We collect training data in matrices (one example per row)

$$
X=\left[\begin{array}{c}
x_{1}^{T} \\
\vdots \\
x_{N}^{T}
\end{array}\right] \quad Y=\left[\begin{array}{c}
y_{1}^{T} \\
\vdots \\
y_{N}^{T}
\end{array}\right]
$$

- Columns $X_{j}$ of data matrix $X=\left[X_{1}, \ldots, X_{n}\right]$ are called features


## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Regression training problem

- Objective: Find data model $m$ such that for all $(x, y)$ :

$$
m(x)-y \approx 0
$$

- Let model output $u=m(x)$; Examples of data misfit losses

$$
\begin{aligned}
& L(u, y)=\frac{1}{2}(u-y)^{2} \\
& L(u, y)=|u-y| \\
& L(u, y)= \begin{cases}\frac{1}{2}(u-y)^{2} & \text { if }|u-v| \leq c \\
c(|u-y|-c / 2) & \text { else }\end{cases}
\end{aligned}
$$



Square


1-norm


- Training: find model $m$ that minimizes sum of training set losses

$$
\underset{m}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i}\right), y_{i}\right)
$$

## Supervised learning - Least squares

- Parameterize model $m$ and set a linear (affine) structure

$$
m(x ; \theta)=w^{T} x+b
$$

where $\theta=(w, b)$ are parameters (also called weights)

- Training: find model parameters that minimize training cost

$$
\underset{\theta}{\operatorname{minimize}} \sum_{i=1}^{N} L\left(m\left(x_{i} ; \theta\right), y_{i}\right)=\frac{1}{2} \sum_{i=1}^{N}\left(w^{T} x_{i}+b-y_{i}\right)^{2}
$$

(note: optimization over model parameters $\theta$ )

- Once trained, predict response of new input $x$ as $\hat{y}=w^{T} x+b$


## Example - Least squares

- Find affine function parameters that fit data:



## Example - Least squares

- Find affine function parameters that fit data:

- Data points $(x, y)$ marked with (*), LS model $w x+b(-)$


## Example - Least squares

- Find affine function parameters that fit data:

- Data points $(x, y)$ marked with (*), LS model $w x+b(-)$
- Least squares finds affine function that minimizes squared distance 13


## Solving for constant term

- Constant term $b$ also called bias term or intercept
- What is optimal $b$ ?

$$
\underset{w, b}{\operatorname{minimize}} \frac{1}{2} \sum_{i=1}^{N}\left(w^{T} x_{i}+b-y_{i}\right)^{2}
$$

- Optimality condition w.r.t. $b$ (gradient w.r.t. $b$ is 0 ):

$$
0=N b+\sum_{i=1}^{N}\left(w^{T} x_{i}-y_{i}\right) \quad \Leftrightarrow \quad b=\bar{y}-w^{T} \bar{x}
$$

where $\bar{x}=\frac{1}{N} \sum_{i=1}^{N} x_{i}$ and $\bar{y}=\frac{1}{N} \sum_{i=1}^{N} y_{i}$ are mean values

## Equivalent problem

- Plugging in optimal $b=\bar{y}-w^{T} \bar{x}$ in least squares estimate gives

$$
\underset{w, b}{\operatorname{minimize}} \frac{1}{2} \sum_{i=1}^{N}\left(w^{T} x_{i}+b-y_{i}\right)^{2}=\frac{1}{2} \sum_{i=1}^{N}\left(w^{T}\left(x_{i}-\bar{x}\right)-\left(y_{i}-\bar{y}\right)\right)^{2}
$$

- Let $\tilde{x}_{i}=x_{i}-\bar{x}$ and $\tilde{y}_{i}=y_{i}-\bar{y}$, then it is equivalent to solve

$$
\underset{w}{\operatorname{minimize}} \frac{1}{2} \sum_{i=1}^{N}\left(w^{T} \tilde{x}_{i}-\tilde{y}_{i}\right)^{2}=\frac{1}{2}\|X w-Y\|_{2}^{2}
$$

where $X$ and $Y$ now contain all $\tilde{x}_{i}$ and $\tilde{y}_{i}$ respectively

- Obviously $\tilde{x}_{i}$ and $\tilde{y}_{i}$ have zero averages (by construction)
- Will often assume averages subtracted from data and responses


## Least squares - Solution

- Training problem

$$
\underset{w}{\operatorname{minimize}} \frac{1}{2}\|X w-Y\|_{2}^{2}
$$

- Strongly convex if $X$ full column rank
- Features linearly independent and more examples than features
- Consequences: $X^{T} X$ is invertible and solution exists and is unique
- Optimal $w$ satisfies (set gradient to zero)

$$
0=X^{T} X w-X^{T} Y
$$

if $X$ full column rank, then unique solution $w=\left(X^{T} X\right)^{-1} X^{T} Y$

## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Nonaffine example

- What if data that cannot be well approximated by affine mapping?



## Nonaffine example

- What if data that cannot be well approximated by affine mapping?



## Nonaffine example

- What if data that cannot be well approximated by affine mapping?



## Adding nonlinear features

- A linear model is not rich enough to model relationship
- Try, e.g., a quadratic model

$$
m(x ; \theta)=b+\sum_{i=1}^{n} w_{i} x_{i}+\sum_{i=1}^{n} \sum_{j=1}^{i} q_{i j} x_{i} x_{j}
$$

where $x=\left(x_{1}, \ldots, x_{n}\right)$ and parameters $\theta=(b, w, q)$

- For $x \in \mathbb{R}^{2}$, the model is

$$
\begin{aligned}
& m(x ; \theta)=b+w_{1} x_{1}+w_{2} x_{2}+q_{11} x_{1}^{2}+q_{12} x_{1} x_{2}+q_{22} x_{2}^{2}=\theta^{T} \phi(x) \\
& \text { where } x=\left(x_{1}, x_{2}\right) \text { and }
\end{aligned}
$$

$$
\begin{aligned}
\theta & =\left(b, w_{1}, w_{2}, q_{11}, q_{12}, q_{22}\right) \\
\phi(x) & =\left(1, x_{1}, x_{2}, x_{1}^{2}, x_{1} x_{2}, x_{2}^{2}\right)
\end{aligned}
$$

- Add nonlinear features $\phi(x)$, but model still linear in parameter $\theta$


## Least squares with nonlinear features

- Can, of course, use other nonlinear feature maps $\phi$
- Gives models $m(x ; \theta)=\theta^{T} \phi(x)$ with increased fitting capacity
- Use least squares estimate with new model

$$
\underset{\theta}{\operatorname{minimize}} \frac{1}{2} \sum_{i=1}^{N}\left(m\left(x_{i} ; \theta\right)-y_{i}\right)^{2}=\frac{1}{2} \sum_{i=1}^{N}\left(\theta^{T} \phi\left(x_{i}\right)-y_{i}\right)^{2}
$$

which is still convex since $\phi$ does not depend on $\theta$ !

- Build new data matrix (with one column per feature in $\phi$ )

$$
X=\left[\begin{array}{c}
\phi\left(x_{1}\right)^{T} \\
\vdots \\
\phi\left(x_{N}\right)^{T}
\end{array}\right]
$$

to arrive at least squares formulation

$$
\underset{\theta}{\operatorname{minimize}} \frac{1}{2}\|X \theta-Y\|_{2}^{2}
$$

- The more features, the more parameters $\theta$ to optimize (lifting)


## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Nonaffine example

- Fit polynomial of degree $k$ to data using LS ( $J$ is cost):



## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Generalization and overfitting

- Generalization: How well does model perform on unseen data
- Overfitting: Model explains training data, but not unseen data
- How to reduce overfitting/improve generalization?


## Tikhonov Regularization

- Example indicates: Reducing $\|\theta\|_{2}$ seems to reduce overfitting
- Least squares with Tikhonov regularization:

$$
\underset{\theta}{\operatorname{minimize}} \frac{1}{2}\|X \theta-Y\|_{2}^{2}+\frac{\lambda}{2}\|\theta\|_{2}^{2}
$$

- Regularization parameter $\lambda \geq 0$ controls fit vs model expressivity
- Optimization problem called ridge regression in statistics
- (Could regularize with $\|\theta\|_{2}$, but square easier to solve)
- (Don't regularize $b$ - constant data offset gives different solution)


## Ridge Regression - Solution

- Recall ridge regression problem for given $\lambda$ :

$$
\underset{\theta}{\operatorname{minimize}} \frac{1}{2}\|X \theta-Y\|_{2}^{2}+\frac{\lambda}{2}\|\theta\|_{2}^{2}
$$

- Objective $\lambda$-strongly convex for all $\lambda>0$, hence unique solution
- Objective is differentiable, Fermat's rule:

$$
\begin{aligned}
0=X^{T}(X \theta-Y)+\lambda \theta & \Longleftrightarrow \quad\left(X^{T} X+\lambda I\right) \theta=X^{T} Y \\
& \Longleftrightarrow \quad \theta=\left(X^{T} X+\lambda I\right)^{-1} X^{T} Y
\end{aligned}
$$

## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Ridge Regression - Example

- Same problem data as before
- Fit 10-degree polynomial with Tikhonov regularization
- $\lambda$ : regularization parameter, $J$ LS cost, $\|\theta\|_{2}$ norm of weights



## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Selecting model hyperparameters

- Parameters in machine learning models are called hyperparameters
- Ridge model has polynomial order and $\lambda$ as hyperparameters
- How to select hyperparameters?


## Holdout

- Randomize data and assign to train, validate, or test set



## Training set:

- Solve training problems with different hyperparameters


## Validation set:

- Estimate generalization performance of all trained models
- Use this to select model that seems to generalize best


## Test set:

- Final assessment on how chosen model generalizes to unseen data
- Not for model selection, then final assessment too optimistic


## Holdout - Comments

- Typical division between sets $50 / 25 / 25$ (or $70 / 20 / 10$ )
- Sometimes no test set (then no assessment of final model)
- If no test set, then validation set often called test set
- Can work well if lots of data, if less, use ( $k$-fold) cross validation


## $k$-fold cross validation

- Similar to hold out - divide first into training/validate and test set
- Divide training/validate set into $k$ data chunks
- Train $k$ models with $k-1$ chunks, use $k$ :th chunk for validation
- Loop

1. Set hyperparameters and train all $k$ models
2. Evaluate generalization score on its validation data
3. Sum scores to get model performance

- Select final model hyperparameters based on best score
- Simpler model with slightly worse score may generalize better
- Estimate generalization performance via test set


## 4-fold cross validation - Graphics



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Evaluate generalization score/performance

- Ridge regression example generalization, validation data ( $\diamond$ )
- $\lambda$ : regularization parameter, $J_{t}$ train cost, $J_{v}$ validation cost



## Selecting model

- Average training and test error vs model complexity
- Average training error smaller than average test error
- Large $\lambda$ (left) model not rich enough
- Small $\lambda$ (right) model too rich (overfitting)



## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Feature selection

- Assume $X \in \mathbb{R}^{m \times n}$ with $m<n$ (fewer examples than features)
- Want to find a subset of features that explains data well
- Example: Which genes in genome control eyecolor


## Lasso

- Feature selection by regularizing least squares with 1-norm:

$$
\underset{w}{\operatorname{minimize}} \frac{1}{2}\|X w-Y\|_{2}^{2}+\lambda\|w\|_{1}
$$

- Problem can be written as

$$
\underset{w}{\operatorname{minimize}} \frac{1}{2}\left\|\sum_{i=1}^{n} w_{i} X_{i}-Y\right\|_{2}^{2}+\lambda\|w\|_{1}
$$

if $w_{i}=0$, then feature $X_{i}$ not important

- The 1-norm promotes sparsity (many 0 variables) in solution
- It also reduces size (shrinks) $w$ (like $\|\cdot\|_{2}^{2}$ regularization)
- Problem is called the Lasso problem


## Example - Lasso

- Data $X \in \mathbb{R}^{30 \times 200}$, Lasso solution for different $\lambda$

- For large enough $\lambda$ solution $w=0$
- More nonzero elements in solution as $\lambda$ decreases
- For small $\lambda, 30$ (nbr examples) nonzero $w_{i}$ (i.e., $170 w_{i}=0$ )


## Lasso and correlated features

- Assume two equal features exist, e.g., $X_{1}=X_{2}$, lasso problem is

$$
\operatorname{minimize} \frac{1}{2}\left\|\left(w_{1}+w_{2}\right) X_{1}+\sum_{i=3}^{n} w_{i} X_{i}-Y\right\|_{2}^{2}+\lambda\left(\left|w_{1}\right|+\left|w_{2}\right|+\left\|w_{3: n}\right\|_{1}\right)
$$

- Assume $w^{*}$ solves the problem and let $\Delta:=w_{1}^{*}+w_{2}^{*}>0$ (wlog)
- Then all $w_{1} \in[0, \Delta]$ with $w_{2}=\Delta-w_{1}$ solves problem:
- quadratic cost unchanged since sum $w_{1}+w_{2}$ still $\Delta$
- the remainder of the regularization part reduces to

$$
\min _{w_{1}} \lambda\left(\left|w_{1}\right|+\left|\Delta-w_{1}\right|\right)
$$



- For almost correlated features:
- often only $w_{1}$ or $w_{2}$ nonzero (the one with slightly better fit)
- however, features highly correlated, if $X_{1}$ explains data so does $X_{2}$


## Elastic net

- Add Tikhonov regularization to the Lasso

$$
\operatorname{minimize} \frac{1}{2}\|X w-Y\|^{2}+\lambda_{1}\|w\|_{1}+\frac{\lambda_{2}}{2}\|w\|_{2}^{2}
$$

- This problem is called elastic net in statistics
- Can perform better with correlated features


## Elastic net and correlated features

- Assume equal features $X_{1}=X_{2}$ and that $w^{*}$ solves the elastic net
- Let $\Delta:=w_{1}^{*}+w_{2}^{*}>0(\mathrm{wlog})$, then $w_{1}^{*}=w_{2}^{*}=\frac{\Delta}{2}$
- Data fit cost still unchanged for $w_{2}=\Delta-w_{1}$ with $w_{1} \in[0, \Delta]$
- Remaining (regularization) part is

$$
\min _{w_{1}} \lambda_{1}\left(\left|w_{1}\right|+\left|\Delta-w_{1}\right|\right)+\lambda_{2}\left(w_{1}^{2}+\left(\Delta-w_{1}\right)^{2}\right)
$$

which is minimized in the middle at $w_{1}=w_{2}=\frac{\Delta}{2}$

- For highly correlated features, both (or none) probably selected


## Group lasso

- Sometimes want groups of variables to be 0 or nonzero
- Introduce blocks $w=\left(w_{1}, \ldots, w_{p}\right)$ where $w_{i} \in \mathbb{R}^{n_{i}}$
- The group Lasso problem is

$$
\operatorname{minimize} \frac{1}{2}\|X w-Y\|_{2}^{2}+\lambda \sum_{i=1}^{p}\left\|w_{i}\right\|_{2}
$$

(note $\|\cdot\|_{2}$-norm without square)

- With all $n_{i}=1$, it reduces to the Lasso
- Promotes block sparsity, meaning full block $w_{i} \in \mathbb{R}^{n_{i}}$ would be 0


## Outline

- Supervised learning - Overview
- Least squares - Basics
- Nonlinear features
- Generalization, overfitting, and regularization
- Cross validation
- Feature selection
- Training problem properties


## Composite optimization

- Least squares problems are convex problems of the form

$$
\underset{\theta}{\operatorname{minimize}} f(X \theta)+g(\theta),
$$

where

- $f=\frac{1}{2}\|\cdot-Y\|_{2}^{2}$ is data misfit term
- $X$ is training data matrix (potentially extended with features)
- $g$ is regularization term (1-norm, squared 2-norm, group lasso)
- Function properties
- $f$ is 1 -strongly convex and 1 -smooth and $f \circ X$ is $\|X\|_{2}^{2}$-smooth
- $g$ is convex and possibly nondifferentiable
- Gradient $\nabla(f \circ X)(\theta)=X^{T}(X \theta-Y)$

