

Machine Learning

A. Supervised Learning

A.3. Regularization

Lars Schmidt-Thieme

Information Systems and Machine Learning Lab (ISMLL)
Institute for Computer Science
University of Hildesheim, Germany

Syllabus

- Fri. 21.10. (1) 0. Introduction
- A. Supervised Learning**
- Fri. 28.10. (2) A.1 Linear Regression
- Fri. 4.11. (3) A.2 Linear Classification
- Fri. 11.11. (4) A.3 Regularization
- Fri. 18.11. (5) A.4 High-dimensional Data
- Fri. 25.11. (6) A.5 Nearest-Neighbor Models
- Fri. 2.12. (7) A.6 Support Vector Machines
- Fri. 9.12. (8) A.7 Decision Trees
- Fri. 16.12. (9) A.8 A First Look at Bayesian and Markov Networks
- B. Unsupervised Learning**
- Fri. 13.1. (10) B.1 Clustering
- Fri. 20.1. (11) B.2 Dimensionality Reduction
- Fri. 27.1. (12) B.3 Frequent Pattern Mining
- C. Reinforcement Learning**
- Fri. 3.2. (13) C.1 State Space Models
- (14) C.2 Markov Decision Processes

Outline


1. The Problem of Overfitting
2. Model Selection
3. Regularization
4. Hyperparameter Optimization

Outline


1. The Problem of Overfitting
2. Model Selection
3. Regularization
4. Hyperparameter Optimization

Fitting of models


Linear model (RSS= 11353.52)



Quadratic model (RSS= 10824.72)



Polynomial model (RSS= 7029.37)



Underfitting/Overfitting

Underfitting:

- ▶ the model is not complex enough to explain the data well.
- ▶ results in poor predictive performance.

Overfitting:

- ▶ the model is too complex, it describes the
 - ▶ noise instead of the
 - ▶ underlying relationship between target and predictors.
- ▶ results in poor predictive performance as well.

Remark: Given N points (x_n, y_n) without repeated measurements (i.e. $x_n \neq x_m, n \neq m$), there exists a polynomial of degree $N - 1$ with RSS equal to 0.

Outline

1. The Problem of Overfitting
2. Model Selection
3. Regularization
4. Hyperparameter Optimization

Model Selection Measures

Model selection means: we have a set of models, e.g.,

$$Y = \sum_{i=0}^{p-1} \beta_i X_i$$

indexed by p (i.e., one model for each value of p), make a choice which model **describes** the data best.

If we just look at **losses** / **fit measures** such as RSS, then

the larger p , the better the fit

or equivalently

the larger p , the lower the loss

as the model with p parameters can be **reparametrized** in a model with $p' > p$ parameters by setting

$$\beta'_i = \begin{cases} \beta_i, & \text{for } i \leq p \\ 0, & \text{for } i > p \end{cases}$$

Model Selection Measures

One uses **model selection measures** of type

$$\text{model selection measure} = \text{fit} - \text{complexity}$$

or equivalently

$$\text{model selection measure} = \text{loss} + \text{complexity}$$

The smaller the loss (= lack of fit), the better the model.

The smaller the complexity, the simpler and thus better the model.

The model selection measure tries to find a trade-off between fit/loss and complexity.

Model Selection Measures

Akaike Information Criterion (AIC): (maximize)

$$\text{AIC} := \log L - p$$

or (minimize)

$$\text{AIC} := -2 \log L + 2p$$

Bayes Information Criterion (BIC) /

Bayes-Schwarz Information Criterion: (maximize)


$$\text{BIC} := \log L - \frac{p}{2} \log N$$

where L denotes the likelihood, N the number of samples.


Variable Backward Selection

{ A, F, H, I, J, L, P }
AIC = 63.01


Variable Backward Selection



Variable Backward Selection



Variable Backward Selection



X removed variable

Outline

1. The Problem of Overfitting
2. Model Selection
- 3. Regularization**
4. Hyperparameter Optimization

Shrinkage

Model selection operates by

- ▶ fitting models for a set of models with varying complexity and then picking the "best one" ex post,
- ▶ omitting some parameters completely (i.e. forcing them to be 0).

Shrinkage follows a similar idea:

- ▶ smaller parameters mean a simpler hypothesis/less complex model. Hence, small parameters should be preferred in general.
- ▶ a term is added to the model equation to penalize high parameters instead of forcing them to be 0.

Shrinkage

There are various types of shrinkage techniques for different application domains.

L1/Lasso Regularization: $\lambda \sum_{j=1}^p |\hat{\beta}_j| = \lambda \|\hat{\beta}\|_1$

L2/Tikhonov Regularization: $\lambda \sum_{j=1}^p \hat{\beta}_j^2 = \lambda \|\hat{\beta}\|_2^2$

Elastic Net: $\lambda_1 \|\hat{\beta}\|_1 + \lambda_2 \|\hat{\beta}\|_2^2$

Ridge Regression

Ridge regression is a combination of

$$\underbrace{\sum_{i=1}^n (y_i - \hat{y}_i)^2}_{\text{L2 loss}} + \lambda \underbrace{\sum_{j=1}^p \beta_j^2}_{\text{L2 regularization}}$$

= L2 loss + λ L2 regularization

Ridge Regression (Closed Form)

Ridge regression: minimize

$$\begin{aligned}
 \text{RSS}_\lambda(\hat{\beta}) &= \text{RSS}(\hat{\beta}) + \lambda \sum_{j=1}^p \hat{\beta}_j^2 = \langle \mathbf{y} - \mathbf{X}\hat{\beta}, \mathbf{y} - \mathbf{X}\hat{\beta} \rangle + \lambda \sum_{j=1}^p \hat{\beta}_j^2 \\
 \Rightarrow \hat{\beta} &= (\mathbf{X}^T \mathbf{X} + 2\lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}, \quad \mathbf{I} := \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & 1 \end{pmatrix}
 \end{aligned}$$

with $\lambda \geq 0$ a **complexity parameter** / **regularization parameter**.

As solutions of ridge regression are not equivariant under scaling of the predictors, data is normalized before ridge regression:

$$x'_{i,j} := \frac{x_{i,j} - \bar{x}_{\cdot,j}}{\hat{\sigma}(x_{\cdot,j})}$$

Ridge Regression (Gradient Descent)

- 1: **procedure** RIDGE-REGR-
GD($\hat{y} : \mathbb{R}^P \rightarrow \mathbb{R}, \hat{\beta}^{(0)} \in \mathbb{R}^{P+1}, \alpha, t_{\max} \in \mathbb{N}, X \in \mathbb{R}^{N \times P}, y \in \mathbb{R}^N$)
- 2: **for** $t = 1, \dots, t_{\max}$ **do**
- 3: $\hat{\beta}_0^{(t)} := \hat{\beta}_0^{(t-1)} - \alpha \left(2 \sum_{n=1}^N - (y_n - \hat{y}(X_n)) \right)$
- 4: **for** $j = 1, \dots, P$ **do**
- 5: $\hat{\beta}_j^{(t)} := \hat{\beta}_j^{(t-1)} - \alpha \left(2 \sum_{n=1}^N -X_{n,j} (y_n - \hat{y}(X_n)) + 2\lambda \hat{\beta}_j^{(t-1)} \right)$
- 6: **if** converged **then**
- 7: **return** $\hat{\beta}^{(t)}$

L2-Regularized Update Rule

$$\hat{\beta}_j^{(t)} := \underbrace{(1 - 2\alpha\lambda)}_{\text{shrinkage}} \hat{\beta}_j^{(t-1)} - \alpha \left(2 \sum_{n=1}^N -X_{n,j} (y_n - \hat{y}(X_n)) \right)$$

Tikhonov Regularization Derivation (1/2)

Treat the true parameters θ_j as random variables Θ_j with the following distribution (**prior**):

$$\Theta_j \sim \mathcal{N}(0, \sigma_\Theta), \quad j = 1, \dots, p$$

Then the **joint likelihood of the data and the parameters** is

$$L_{\mathcal{D}, \Theta}(\theta) := \left(\prod_{i=1}^n p(x_i, y_i | \theta) \right) \prod_{j=1}^p p(\Theta_j = \theta_j)$$

and the conditional joint log likelihood of the data and the parameters

$$\log L_{\mathcal{D}, \Theta}^{\text{cond}}(\theta) := \left(\sum_{i=1}^n \log p(y_i | x_i, \theta) \right) + \sum_{j=1}^p \log p(\Theta_j = \theta_j)$$

and

$$\log p(\Theta_j = \theta_j) = \log \frac{1}{\sqrt{2\pi\sigma_\Theta}} e^{-\frac{\theta_j^2}{2\sigma_\Theta^2}} = -\log(\sqrt{2\pi\sigma_\Theta}) - \frac{\theta_j^2}{2\sigma_\Theta^2}$$

Tikhonov Regularization Derivation (2/2)

Dropping the terms that do not depend on θ_j yields:

$$\begin{aligned} \log L_{\mathcal{D}, \Theta}^{\text{cond}}(\theta) &:= \left(\sum_{i=1}^n \log p(y_i | x_i, \theta) \right) + \sum_{j=1}^p \log p(\Theta_j = \theta_j) \\ &\propto \left(\sum_{i=1}^n \log p(y_i | x_i, \theta) \right) - \frac{1}{2\sigma_{\Theta}^2} \sum_{j=1}^p \theta_j^2 \end{aligned}$$

This also gives a semantics to the complexity / regularization parameter λ :

$$\lambda = \frac{1}{2\sigma_{\Theta}^2}$$

but σ_{Θ}^2 is unknown. (We will see methods to estimate λ soon.)

The parameters θ that maximize the joint likelihood of the data and the parameters are called **Maximum A posteriori Estimators (MAP estimators)**.

L2-Regularized Logistic Regression (Gradient Descent)

$$\log L_{\mathcal{D}}^{\text{cond}}(\hat{\beta}) = \sum_{i=1}^n y_i \langle x_i, \hat{\beta} \rangle - \log(1 + e^{\langle x_i, \hat{\beta} \rangle}) - 2\lambda \sum_{j=1}^P \hat{\beta}_j^2$$

1: **procedure** LOG-REGR-

GA($L_{\mathcal{D}}^{\text{cond}} : \mathbb{R}^{P+1} \rightarrow \mathbb{R}, \hat{\beta}^{(0)} \in \mathbb{R}^{P+1}, \alpha, t_{\max} \in \mathbb{N}, \epsilon \in \mathbb{R}^+$)

2: **for** $t = 1, \dots, t_{\max}$ **do**

3: $\hat{\beta}_0^{(t)} := \hat{\beta}_0^{(t-1)} + \alpha \sum_{i=1}^n \left(y_i - p \left(Y = 1 | X = x_i; \hat{\beta}^{(t-1)} \right) \right)$

4: **for** $j = 1, \dots, P$ **do**

5: $\hat{\beta}_j^{(t)} :=$

$\hat{\beta}_j^{(t-1)} + \alpha \left(\sum_{i=1}^n x_{i,j} \left(y_i - p \left(Y = 1 | X = x_i; \hat{\beta}^{(t-1)} \right) \right) - 2\lambda \hat{\beta}_j^{(t-1)} \right)$

6: **if** $L_{\mathcal{D}}^{\text{cond}}(\hat{\beta}^{(t-1)}) - L_{\mathcal{D}}^{\text{cond}}(\hat{\beta}^{(t)}) < \epsilon$ **then**

7: **return** $\hat{\beta}^{(t)}$

8: **error** "not converged in t_{\max} iterations"

L2-Regularized Logistic Regression (Newton)

Newton update rule:

$$\hat{\beta}^{(t)} := \hat{\beta}^{(t-1)} + \alpha H^{-1} \nabla_{\hat{\beta}} p \left(Y = 1 | X = x_i; \hat{\beta}^{(t-1)} \right)$$

$$p_i = p \left(Y = 1 | X = x_i; \hat{\beta}^{(t-1)} \right)$$

$$\nabla_{\hat{\beta}} L_{\mathcal{D}}^{\text{cond}} = \begin{pmatrix} \sum_{i=1}^n -(y_i - p_i) \\ \sum_{i=1}^n -x_{i,1} (y_i - p_i) - 2\lambda \hat{\beta}_1 \\ \vdots \\ \sum_{i=1}^n -x_{i,P} (y_i - p_i) - 2\lambda \hat{\beta}_P \end{pmatrix}$$

$$H = \sum_{i=1}^n -p_i (1 - p_i) x_i x_i^T - 2\lambda I$$

Outline

1. The Problem of Overfitting
2. Model Selection
3. Regularization
- 4. Hyperparameter Optimization**

What is Hyperparameter Optimization?

Many learning algorithms \mathcal{A}_λ have hyperparameters λ (learning rate, regularization). After choosing them, \mathcal{A}_λ can be used to map the training data D_{train} to a function \hat{y} by minimizing some loss $\mathcal{L}(x; \hat{y})$.

Identifying good values for the hyperparameters λ is called **hyperparameter optimization**.


Hence, hyperparameter optimization is a second level optimization

$$\operatorname{argmin}_{\lambda \in \Lambda} \frac{1}{|D_{\text{calib}}|} \sum_{x \in D_{\text{calib}}} \mathcal{L}(x; \mathcal{A}_\lambda(D_{\text{train}})) = \operatorname{argmin}_{\lambda \in \Lambda} \Psi(\lambda)$$

where Ψ is the **hyperparameter response function** and D_{calib} a **calibration set**.


Why Hyperparameter Optimization

- ▶ So far only model parameters were optimized.
- ▶ Hyperparameters (such as regularization λ and learning rate α) came “out of the blue”.
- ▶ Hyperparameters can have a big impact on the prediction quality.




Grid Search

- ▶ Choose for each hyperparameter a set of values $\Lambda_1, \dots, \Lambda_q$.
- ▶ $\Lambda = \prod_{i=1}^q \Lambda_i$ is then the combination of all hyperparameters in all Λ_i s.
- ▶ Then choose the hyperparameter $\lambda \in \Lambda$ with best performance on D_{calib} .



Random Search

- ▶ Instead of choosing hyperparameters on a grid, choose random hyperparameters λ for Λ (within a reasonable space).
- ▶ Provides better results than grid search in cases of insensitive hyperparameters.



What is the Calibration Data?

Whenever a learning process depends on a hyperparameter, the hyperparameter can be estimated by picking the value with the lowest error.

If this is done on test data, one actually uses test data in the training process (“train on test”), thereby lessening its usefulness for estimating the test error.

Therefore, one splits the training data again in

- ▶ (proper) training data and
- ▶ **calibration data.**

The calibration data figures as test data during the training process.

Cross Validation

Instead of a single split into

training data, (validation data,) and test data

cross validation splits the data in k parts (of roughly equal size)

$$D = D_1 \cup D_2 \cup \dots \cup D_k, \quad D_i \text{ pairwise disjoint}$$

and averages performance over k learning problems

$$D_{\text{train}}^{(i)} = D \setminus D_i, \quad D_{\text{test}}^{(i)} = D_i \quad i = 1, \dots, k$$

Common is 5- and 10-fold cross validation.

n -fold cross validation is also known as **leave one out**.

Cross Validation

How many folds to use in k -fold cross validation?

$k = n$ / leave one out:

- ▶ approximately unbiased for the true prediction error.
- ▶ high variance as the n training sets are very similar.
- ▶ in general computationally costly as n different models have to be learnt.

$k = 5$:

- ▶ lower variance.
- ▶ bias could be a problem,
due to smaller training set size the prediction error could be overestimated.

Summary

- ▶ The problem of underfitting can be overcome by using more complex models, e.g., having
 - ▶ variable interactions as in polynomial models.
- ▶ The problem of overfitting can be overcome by
 - ▶ model / variable selection as well as by
 - ▶ (parameter) shrinkage.
- ▶ Applying L2-regularization to Linear and Logistic Regression requires only few changes in the learning algorithm
- ▶ Shrinkage introduces a hyperparameter λ that cannot be learned by direct loss minimization.
- ▶ Estimating the best hyperparameters can be considered as a meta-learning problem. They can be estimated e.g. by
 - ▶ Grid Search or
 - ▶ Random Search.

using validation data.

Further Readings

- ▶ [JWHT13, chapter 3], [Mur12, chapter 7], [HTFF05, chapter 3].

References



Trevor Hastie, Robert Tibshirani, Jerome Friedman, and James Franklin.

The elements of statistical learning: data mining, inference and prediction, volume 27.
2005.



Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani.

An introduction to statistical learning.
Springer, 2013.



Kevin P. Murphy.

Machine learning: a probabilistic perspective.
The MIT Press, 2012.