Gabriel Stoltz

# An Introduction to Machine Learning

April 14, 2025

# Preface

# Machine learning is not a spectators' sport

# Bibliographical guide

The main reference for this course is the introductory book [40] by Kevin Murphy, which focuses on algorithms and numerical methods. It covers a wide range of methods and techniques, and can be found online for free at the url

# https://probml.github.io/pml-book/book1.html

Another important reference, for the more theoretical parts of the course, is the book on learning theory from Francis Bach [4], which can be found at the page

#### https://www.di.ens.fr/%7Efbach/ltfp\_book.pdf

Other nice references mentioned in the course include [].

S. Shalev-Shwartz and S. Ben-David, Understanding Machine Learning M. Mohri, A. Rostamizadeh and A. Talwalkar, Foundations of Machine Learning D. Barber, Bayesian Reasoning and Machine Learning C. Bishop, Pattern Recognition and Machine Learning E. Alpaydin, Introduction to Machine Learning P. Mehta, M. Bukov, C.-H. Wang, A.G.R. Day, C. Richardson, C.K. Fisher, D.J. Schwab, A high-bias, low-variance introduction to Machine Learning for physicists, Physics Reports 810, 1-124 (2019) [arXiv reference]

# Exercises

More difficult questions are indicated by this symbol in the margin. These questions would typically not be asked at the exam.

# Some practical comments

This set of lecture notes is the first typed version I have on this material. It therefore certainly contains various typos. Do not hesitate to point me out anything you spot!

Paris, 18 July 2023

 $Gabriel\ Stoltz$ 

Ś

# Contents

1	A first experience in machine learning				
	1.1	What is machine learning?	1		
	1.2	Supervised learning	7		
	1.3	Local averaging methods	15		
<b>2</b>	Least square regression				
	2.1	General framework of linear least square regression	26		
	2.2	Ordinary least square regression	27		
	2.3	Ridge regression	35		
	2.4	LASSO regression	43		
3	Classification with logistic regression				
	3.1	Convexification of the loss	53		
	3.2	Logistic regression	61		
4	(Stochastic) Gradients methods				
	4.1	Minimization problems to solve and general strategy	67		
	4.2	Simple gradient descent	73		
	4.3	Deterministic methods beyond simple gradient descent	76		
	4.4	Stochastic gradient descent and its extensions	78		
<b>5</b>	Pri	ncipal component analysis	83		
	5.1	Motivating the need for dimensionality reduction	83		
	5.2	Deriving PCA from reconstruction error	84		
	5.3	PCA in practice	88		
	5.4	Interpretation of PCA	89		
	5.5	Extensions of PCA	91		
6	Support vector machines				
	6.1	Linear classification for separable data sets	99		
	6.2	Linear classification for non separable data sets	105		
	6.3	Multiclass SVM	108		
	6.4	Kernel SVM	109		
7	Tre	es and ensemble methods	115		
	7.1	Regression and classifications trees	115		
	7.2	Bagging and averaging techniques	121		
	7.3	Boosting	123		

# VIII Contents

8	Neu	ral networks	129			
	8.1	Feed-forward neural networks	129			
	8.2	Training neural networks	133			
	8.3	Unsupervised learning with neural networks	138			
	8.4	Other types of neural networks	144			
9	Clu	Clustering methods				
	9.1	Aims and scope of clustering	145			
	9.2	K-means and center-based clustering methods	146			
	9.3	Hierarchical clustering	149			
	9.4	Clustering with mixture models	150			
	9.5	Density based clustering	155			
	9.6	Spectral clustering	155			
10	Con	aplements: Elements of statistical learning	161			
	10.1	Empirical risk minimization and statistical learning theory	161			
	10.2	Model selection	166			
	10.3	Statistical learning theory	168			
11	Imp	lementing and debuging machine learning programs	171			
<b>References</b>						

1.1	What	is mach	ine learning?	1
	1.1.1	The (big	) data paradigm	2
	1.1.2	Definitio	ns	2
		1.1.2.1	Machine learning tasks	2
		1.1.2.2	Relationship with other scientific fields	3
	1.1.3	Types of	elearning	3
	1.1.4	Learning	stages	4
1.2	Super	vised lea	arning	7
	1.2.1	Mathem	atical framework	7
		1.2.1.1	Classification	8
		1.2.1.2	Regression	9
	1.2.2	Decision	theory and statistical learning	10
		1.2.2.1	Decision theory	10
		1.2.2.2	Statistical learning	14
		1.2.2.3	Learning from data	15
1.3	Local	averagir	ng methods	15
	1.3.1	Presenta	tion of the K-nearest neighbor algorithm	16
	1.3.2	Choosing $K$ by (cross) validation		
		1.3.2.1	Large training data sets	17
		1.3.2.2	Small training data sets	17
	1.3.3	Theoretical analysis of the K-nearest neighbor algorithm		
		1.3.3.1	Local methods	18
		1.3.3.2	Consistency analysis for local methods	19
		1.3.3.3	Consistency analysis for <i>K</i> -nearest neighbors	21

Not everything can be learned. The question which underlies machine learning at large is: What can be learned, and under which conditions? This chapter gives an introduction to machine learning, by first discussing in Section 1.1 what machine learning is about through some overview of the field. We next turn in Section 1.2 to supervised learning, which is the type of learning we will mostly consider in these lecture notes. We conclude this introductory chapter with our first examples of machine learning methods, based on local averaging, in particular K-nearest neighbors (see Section 1.3).

# 1.1 What is machine learning?

We give here an overview of machine learning based on [40, Chapter 1], [3, Chapter 1], [39, Chapter 1], [6, Chapter 13] and [49, Chapter 1].

# 1.1.1 The (big) data paradigm

Although machine learning as a scientific field has been present since decades, it gained a lot of popularity in the past decade, thanks to a new paradigm where

- data can be abundant and easily obtained, and therefore comes first in many settings (before the system is even being modeled);
- efficient computing devices are available.

Examples of situations where data is easily obtained include: internet trafic (creation of data when visiting webpages), music or movie ratings (when liking songs on some apps, writing a review for a movie, ...), customer behaviors (recording information on products bought online), etc.

Gathering data is not a goal in itself. What one ultimately aims at is to predict new events or behaviors based on the gathered information. For instance, a supermarket chain may want to be able to say which customer is likely to buy which product, when, how much, ... while the customer may want to find products that fit his/her needs. Another example is spam classification: the user may want to tailor the classifier so that it better sorts the incoming emails. A last example is character recognition in images, for instance for automatic recognition of amounts when depositing bank checks. In all these situations, one may lack a good knowledge of the underlying process, and be unable to model it convincingly. The idea is then to make up for this by relying on accumulated experience in the form of data. This can be coined as "re-cognizing" events or data, relying on the definition of "cognize" as "perceive, become aware of".

# 1.1.2 Definitions

Machine learning is a subfield of artificial intelligence,<sup>1</sup> and has strong links with data mining, data analysis and data visualization. It can be defined in various ways, including

- automatically detecting meaningful patterns in data;
- transforming past information/experience (data) into predictions as accurate as possible;
- improving the performance at certain tasks when having more experience.

All these statements explicitly or implicitly suggest that machine learning is about designing accurate and efficient prediction *algorithms* (to be emphasized as a clear list of instructions/procedures to follow in order to obtain a result). The so-obtained algorithm should be able to adapt to modifications in behaviors. Illustrative examples to this end are spam classification and fraud detection.

# 1.1.2.1 Machine learning tasks

The application domains of machine learning are ubiquitous:

- science (processing of measurement data in astronomy, physics; DNA analysis in biology, ...)
- medicine (for instance medical diagnosis)
- image recognition/segmentation (for medical applications, social media, etc)
- text analysis/classification
- speech processing/natural language processing (for instance translations)
- finance/banking (credit applications, fraud detection, prediction of stock market prices, ...)
- playing games (chess, go)

These application domains require machine learning tools for various tasks:

<sup>&</sup>lt;sup>1</sup> There are many definitions of artificial intelligence (AI)... We use here elements from the Wikipedia page of the topic. In these notes, we consider AI as a scientific field studying and developing "the intelligence of machines or software, as opposed to the intelligence of humans or other animals [...] The various subfields of IA research are centered around particular goals and the use of particular tools. The traditional goals of AI research include reasoning, knowledge representation, planning, learning, natural language processing, perception, and support for robotics. General intelligence (the ability to complete any task performable by a human) is among the field's long-term goals."

- classification (many possible applications: credit scoring, digit classification, face recognition and biometrics, ...), see for instance Chapters 3 and 6;
- regression, see for instance Chapter 2;
- ranking (which corresponds to ordering items given a criterion, e.g. webpages when performing a query on google);
- clustering (for example to find communities in social networks), see Chapter 9;
- dimensionality reduction (e.g. for image compression), see Chapter 5 and Section 8.3;
- data generation (for instance with tools such as ChatGPT, Code-Llama, Dall-E, ...).

**Exercise 1.1.** Find other examples of machine learning tasks.

#### 1.1.2.2 Relationship with other scientific fields

Machine learning is a proteiform scientific field, at the crossroad of various other fields, in particular:

- probability and statistics. Distinctive features of machine learning are (i) the fact that the distribution of the data is unknown (distribution free setting), so that theoretical results are established over (large) classes of distributions; (ii) non asymptotic results are of prime interest, as the number of data points may not be large in comparison with the number of degrees of freedom/unknowns of the machine learning model; (iii) the key ouput of machine learning algorithms are predictions on unseen data, so that metrics of success are based on this criterion, and not on the quality of the estimation of the parameters of the model as in statistics. Somehow, machine learning cares less about checking statistical hypotheses than finding causes for the observed data. It also has a more computer science oriented mindset than traditional statistics.
- *optimization theory*, as many models of machine learning require some form of parameter optimization with respect to some loss/cost function. This optimization is performed in (very) high dimension, often for non convex targets.
- *computer science*, as the algorithmic part of machine learning is pivotal. A particular attention is paid to equilibrating the prediction performance and the computational cost of the methods which are being implemented.
- *artificial intelligence*: machine learning can be considered as a specific component, which allows to transform "experience" (in the form of data) into "actions" (based on the prediction provided by the algorithm). There is however no attempt to emulate or reproduce human intelligence in machine learning.

# 1.1.3 Types of learning

In these lecture notes, we will mostly perform *supervised learning*, although *unsupervised learning* will also be considered to some extent. This will be done in a *statistical framework*, using *batch learning* with a *passive teacher*. In order to make sense of these sentences, let us comment on the various alternative options:

• Supervised learning is described more precisely in Section 1.2. In short, this means that data points come with labels, as opposed to unsupervised learning where data is unlabeled. Sometimes, a semi-supervised setting is considered, where only a fraction of the data points are labeled. A typical example of supervised learning is spam detection, for which the training database needs to include examples of spams and legitimate emails in order to make predictions on new incoming emails. An example of unsupervised learning for emails would be anomaly detection, *i.e.* marking a new incoming email as deviating from the usual email content. Supervised learning is intuitively more efficient, as some extra information is available compared to unsupervised learning. However, attaching a label to the data requires some (possibly costly and time consuming) preprocessing of the data – think for instance about medical diagnosis, where a medical doctor needs to go over various images and pinpoint whether there

is a pathology or not. This motivates working with unlabeled data, or only with a fraction of labeled data.

- Active learning corresponds to settings where the algorithm can request new data to be generated in some regions of data space; whereas, in *passive learning*, the data is given to the algorithm without the possibility to give some feedback on the next data location.
- Online learning arises in situations where data continuously flows in, while batch learning corresponds to settings where the data is handed out at the beginning of the procedure. One example of online learning is finance, where stockbrokers aim at predicting stock market prices in real time for their daily decisions.
- *Passive* vs. *adversarial training* depends on the helpfulness of the teacher. A passive teacher is encountered when the data is random, which corresponds to the framework of statistical learning. This should be contrasted with the adversarial scenario, where data is generated in order to trick the learner into making mistakes. This is useful for algorithms targeted at fraud detection for instance, or more generally to obtain guarantees in the worst case scenario (in opposition to typical scenarios, as in the statistical learning framework).

**Unsupervised learning.** From a mathematical viewpoint, and anticipating to some extent on the presentation of supervised learning in Section 1.2, unsupervised learning aims at fitting an unconditional model  $p_{data}(x)$  which can generate new data x, as opposed to supervised learning which considers data points x and labels y, and fits a conditional model  $p_{data}(x|y)$  in order to make predictions. In essence, unsupervised learning aims at finding compact descriptions of data. The main interest of this learning framework is that it avoids to collect labeled data. It can also be considered in situations where the labeling is ambiguous (as for text or image classifications when some categories are close to each other). More generally, it allows to find patterns or "explanations" in high dimensional data instead of focusing on low dimensional inputs only.

One difficulty with unsupervised learning is the absence of ground truth provided by labels. Evaluating the quality of the model therefore requires some work. Possible approaches to this end are (i) to check the likelihood of the generated data (using some test data and density estimations); (ii) to use the learned unsupervised (reduced) representation of the data for classification/regression tasks and hope to be more efficient than with the full data; (iii) sometimes, a qualitative gain in the understanding of the model ("interpretability") can be considered as a good enough motivation for unsupervised learning.

**Reinforcement learning.** Another learning framework, which we will not touch upon, is provided by reinforcement learning, where an agent learns to interact with an environment, through some policy (list of actions), which is updated using some reward function. In this framework, data is also unlabeled, but the output of the algorithm is a balance between exploration of new data (which can be seen as taking actions in a dynamic environment) and exploitation of the data collected until now (to update the parameters of the method in order to make better predictions, the metric for the quality of the prediction being some reward function). A prominent example is provided by learning to play chess or go. In essence, reinforcement learning amounts to "learning with a critic" (with some occasional thumbs up or down) as opposed to "learning with a teacher".

# 1.1.4 Learning stages

We briefly discuss in this section the overall workflow associated with machine learning techniques. Only the first step is made precise below, the other ones being discussed in the following chapters:

- (1) Preparation of the data: The first step is to collect data/examples, to preprocess them by curating the dataset, and next apply some featurization procedure and label the data points for supervised learning. The dataset can then be decomposed into three subsets: a training dataset, a validation dataset (which will be used to find the best parameters of the model) and a test set. The latter set is never used to find the parameters of the model; its sole purpose is to assess the quality of the prediction.
- (2) Choose a loss function, which measures the performance of the prediction.

- (3) Choose the hypothesis set (i.e. the class of models which is considered) and the hyperparameters for the algorithm.
- (4) Find the parameters of the model leading to the best prediction by (i) for a given value of the hyperparameters, finding the parameters of the model giving the smallest loss for the training dataset; (ii) selecting the hyperparameters to obtain the smallest loss on the validation set; (iii) assessing the quality of the so-obtained prediction on the test set.

A focus on data preparation. The nature and quality of the data is key to the success of the learning process. Although this is of paramount importance in practice and requires quite some care, we consider in these lecture notes that the data has been correctly prepared, and therefore focus on the algorithmic and theoretical aspects of learning. In order to test methods, there are nowadays various benchmark datasets, of increasing complexity. For instance, for image classification, one can start with the rather simple MNIST dataset for which the task is to classify handwritten digits, and then progressively complexify the classification task by considering fashion-MNIST (recognizing clothing items), the various CIFAR datasets, the ImageNet dataset, etc. Some websites such as Kaggle gather standard datasets (see also the references on the course webpage).

After collecting raw data, the first task is to remove duplicates, and check for missing data. Various rules exist in the latter case to make up for missing fields (mean value imputation, use of generative methods to fill in fields, etc). Categorical data, such as colors (for a wine, an item of clothing, an animal, etc), genders, etc., require some care in the way they are handled, as algorithms and computers work with numbers. The customary way of proceeding is to use one-hot encoding to transform a categorical variable with K possible values into a vector  $\{0,1\}^K$  with only one non-zero entry; more precisely, the categorical component  $x_i$  is transformed into the vector  $(\mathbf{1}_{x_i=1},\ldots,\mathbf{1}_{x_i=K})$  (the interest of this transformation will become more clear in contexts such as the one of Section 3.2.4 on multiclass logistic regression). This is one instance of featurization, the process of transforming the raw data into a more relevant input. This featurization step can be much more involved for certain applications. For instance, in order to predict forces in atomistic models of materials, the raw input, which is the atomistic configuration of a system, is featurized using various functions computing radial and angular moments of the distribution of neighboring atoms around each other atom. Another example is provided by data inputs which are not of the same sizes, such as sequences of words or chunks of temporal series. Let us emphasize here that coming up with (or learning) a good representation of the data is still an active research field in various application domains of machine learning. In any case, some exploratory data analysis needs to be conducted in order to understand the salient properties of the dataset; and possibly remove some attributes/features (for instance through pairwise scatter plots to identify redundant or irrelevant features).

In the remainder of these lecture notes, we always consider that the data is organized in the following design matrix, where each row represents one of the *n* data point  $x_i = (x_{i,1}, \ldots, x_{i,d}) \in \mathbb{R}^{1 \times d}$ , and each of the *d* columns represents a feature (also called attribute):

$$X = \begin{pmatrix} x_{1,1} & x_{1,2} \dots & x_{1,d} \\ x_{2,1} & x_{2,2} \dots & x_{2,d} \\ \vdots & \vdots & \ddots & \vdots \\ x_{n,1} & x_{n,2} \dots & x_{n,d} \end{pmatrix} \in \mathbb{R}^{n \times d}.$$

**Normalization.** In many methods, it is convenient to have data features of similar sizes, and in fact of order 1. This can be achieved through linear transformations. The three main options to this end are

• standardization: this corresponding to a normalization of the data component by component, *i.e.* centering the data in each column and rescaling it so that the empirical variance is 1. More precisely, this amounts to setting  $\tilde{x}_{i,k} = (x_{i,k} - b_k)/a_k$  (alternatively,  $x_{i,k} = a_k \tilde{x}_{i,k} + b_k$ ) with  $a_k, b_k$  chosen such that

$$\sum_{i=1}^{n} \widetilde{x}_{i,k} = 0, \qquad \frac{1}{n} \sum_{i=1}^{n} \widetilde{x}_{i,k}^{2} = 1.$$
(1.1)

The values  $a_k, b_k$  can in fact be analytically computed, see Exercise 1.2 below.

• *whitening:* this corresponds to a global normalization of the data, *i.e.* applying an affine transformation so that the data is centered and with a unit empirical covariance. The empirical covariance of the data inputs is

$$\Sigma_n = \frac{1}{n} \sum_{i=1}^n \left( x_i - \overline{x}_n \right)^\top \left( x_i - \overline{x}_n \right) \in \mathbb{R}^{d \times d}.$$
(1.2)

Note that the matrix  $\Sigma_n$  is positive symmetric. Assuming moreover that it is definite, one can consider  $\Sigma_n^{-1/2}$  by spectral calculus, and introduce

$$\widetilde{x}_i = (x_i - \overline{x}_n) \, \Sigma_n^{-1/2}. \tag{1.3}$$

A simple computation shows that the transformed data set  $\{\tilde{x}_1, \ldots, \tilde{x}_n\}$  is by construction centered and with unit covariance; see Exercise 1.3. The original data inputs can be recovered by undoing the transformation as  $x_i = \bar{x}_n + \tilde{x}_i \Sigma_n^{1/2}$ .

• *scaling:* this corresponds to a linear transformation mapping the data to the interval [0, 1], with minimal value 0 and maximal value 1 in each column. More precisely,

$$\widetilde{x}_{i,k} = \frac{x_{i,k} - m_k}{M_k - m_k}, \qquad m_k = \min_{1 \le j \le n} x_{j,k}, \quad M_k = \max_{1 \le j \le n} x_{j,k}.$$

In all cases, when the machine learning method is trained on the transformed data set  $\{\tilde{x}_1, \ldots, \tilde{x}_n\}$ , some preprocessing is necessary to make predictions for a new (test) data point x', as one first needs to transform it the same way the data points were transformed. When normalization is considered, this means that the prediction is performed on  $\tilde{x}'$  whose components are  $\tilde{x}'_k = (x' - b_k)/a_k$ ; when whitening is applied, prediction is performed on  $\tilde{x}' = \sum_n^{-1/2} (x' - \bar{x}_n)$ ; for scaling, prediction is done on  $\tilde{x}' = (x' - m_k)/(M_k - m_k)$ . Let us emphasize that, in all cases, the parameters of the transformation are those computed on the training data set.

**Exercise 1.2 (Standardization).** Prove that the values  $a_k, b_k$  which allow to satisfy (1.1) are the empirical standard deviation and the empirical average:

$$a_k = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_{i,k} - \overline{x}_n)^2}, \qquad b_k = \overline{x}_n = \frac{1}{n} \sum_{i=1}^n x_{i,k}.$$

**Correction.** One first determines  $b_k$  from the centering condition (first equality in (1.1)), and then  $a_k$  from the moment condition (second equality in (1.1)).

**Exercise 1.3 (Whitening).** Prove that the transformed data points (1.3) are such that

$$\frac{1}{n}\sum_{i=1}^{n}\widetilde{x}_{i}=0\in\mathbb{R}^{d},\qquad\frac{1}{n}\sum_{i=1}^{n}\widetilde{x}_{i}^{\top}\widetilde{x}_{i}=\mathrm{Id}_{d}.$$

Correction. The first equality is easily obtained as

$$\sum_{i=1}^{n} \widetilde{x}_i = \left[\sum_{i=1}^{n} (x_i - \overline{x}_n)\right] \Sigma_n^{-1/2},$$

the sum on the right hand side being 0. For the second equality, one writes

$$\frac{1}{n}\sum_{i=1}^{n}\widetilde{x}_{i}^{\top}\widetilde{x}_{i} = \frac{1}{n}\sum_{n}^{-1/2}\left(\sum_{i=1}^{n}\left(x_{i}-\overline{x}_{n}\right)^{\top}\left(x_{i}-\overline{x}_{n}\right)\right)\Sigma_{n}^{-1/2} = \Sigma_{n}^{-1/2}\Sigma_{n}\Sigma_{n}^{-1/2} = \mathrm{Id}_{d},$$

from which the conclusion follows.

# 1.2 Supervised learning

We make precise the mathematical framework behind the subclass of machine learning problems corresponding to supervised learning, as this will be the core topic considered in these lectures. The presentation is based on [4, Sections 2.1 and 2.2] and [39, Section 2.4.2].

#### **1.2.1** Mathematical framework

Given a dataset of n elements  $\mathcal{D} = \{(x_1, y_1), \ldots, (x_n, y_n)\} \subset \mathcal{X} \times \mathcal{Y}$  (couples of inputs/outputs) or features/labels), the aim of supervised learning is to predict the output or label  $y \in \mathcal{Y}$  on an unseen data point  $x \in \mathcal{X}$  ("test data"). Equivalently, this amounts to learning a mapping  $\mathcal{X} \to \mathcal{Y}$ .

In many practical cases,  $\mathcal{X}$  is (isomorphic to)  $\mathbb{R}^d$ , either explicitly or implicitly (for instance when using kernels, see Section 5.5.2 and Chapter 6). Typical examples include physical measurements of various quantities at different locations (e.g. pressure, temperture, wind speed, humidity, etc. at various places in France for weather forecasting), or pixel values for images.

The label or output space  $\mathcal{Y}$  depends on the type of problem which is considered. For classification (see Section 1.2.1.1) it is a discrete space:  $\mathcal{Y} = \{0, 1\}$  or  $\{-1, 1\}$  for binary classification, and  $\mathcal{Y} = \{1, \ldots, K\}$  when there are K classes. For regression problems,  $\mathcal{Y} = \mathbb{R}$  or a higher dimensional space  $\mathbb{R}^{K}$ .

The measure of performance, *i.e.* the quality of the mapping  $f: \mathcal{X} \to \mathcal{Y}$  which is learnt, depends on the task (classification vs. regression). Let us however warn the reader that the criterion for performance is in practice not always defined as precisely as in these lecture notes. Moreover, it can be the case that the ouput y cannot be a deterministic function of the input x, for instance when there is some measurement noise, or when there are extra features which are unobserved. A typical example would be the determination of the gender of a person depending on the height, weight and shoe size. In any case, the function f can be quite complicated, and nonlinear.

There are two (related) key challenges, which we will constantly encounter, when trying to learn the function  $f: \mathcal{X} \to \mathcal{Y}$ :

- the number of data points n which are observed can be rather small in view of the complexity of the function to learn. This requires to carefully assess the interpolation and extrapolation capabilities of the mapping f. Typically, interpolation is much easier than extrapolation: loosely speaking, prediction at new inputs in between data points is more reliable than predictions in regions not covered by data points.
- the input space can be of very large dimension d (the so-called "curse of dimensionality"). This raises scalability issues for the algorithms, and also limits the capacity of the model to make correct predictions on new data – the so-called generalization ability.

Formalizing the learning framework. We consider the paradigm of statistical learning, where elements in the dataset are assumed to be identically and independently distributed (i.i.d.) according to some unknown probability measure  $p_{\text{data}}(dx dy)$ . Expectations with respect to the latter probability measure are denoted by  $\mathbb{E}$ . The mapping  $\mathcal{X} \to \mathcal{Y}$  to be learned is parametrized by  $\theta \in \Theta$ , so that the model is sought in the parametric class of functions  $\{f_{\theta}, \theta \in \Theta\}$ . The ideal learning problem can then be formulated as the following minimization of the expected risk:

$$\min_{\theta \in \Theta} \mathcal{R}(\theta), \qquad \mathcal{R}(\theta) = \mathbb{E}\left[\ell\left(y, f_{\theta}(x)\right)\right], \qquad (1.4)$$

for some loss function  $\ell: \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$  (typically having nonnegative values). In practice, only a finite number of data points are available through a dataset  $\mathcal{D} = \{(x_1, y_1), \dots, (x_n, y_n)\} \subset (\mathcal{X} \times \mathcal{Y})^n$ . In this context, a learning algorithm is formally an application which associates a function  $f \in$  $\mathscr{F}(\mathcal{X},\mathcal{Y})$  (the set of measurable functions from  $\mathcal{X}$  to  $\mathcal{Y}$ ) to a dataset  $\mathcal{D}$  of size n. This is done in practice by relying on the so-called empirical risk minimization (see Definition 1.3 for a more formal presentation), where the expectation in (1.4) is approximated by a finite sum. More precisely, the algorithm returns  $f_{\hat{\theta}_n}$  with

7

$$\widehat{\theta}_n \in \operatorname*{argmin}_{\theta \in \Theta} \widehat{\mathcal{R}}_n(\theta), \qquad \widehat{\mathcal{R}}_n(\theta) = \frac{1}{n} \sum_{i=1}^n \ell\left(y_i, f_{\theta}(x_i)\right).$$

**Exercise 1.4.** For  $\theta$  fixed, give sufficient conditions on  $\ell$  and  $p_{\text{data}}$  ensuring that  $\widehat{\mathcal{R}}_n(\theta)$  converges to  $\mathcal{R}(\theta)$  as  $n \to +\infty$ , and make precise in which sense convergence holds.

**Correction.** The convergence result is based on the Law of Large Numbers. We assume to this end that

$$\mathbb{E}\left(\left|\ell(y_1, f_{\theta}(x_1))\right|\right) = \int_{\mathcal{X} \times \mathcal{Y}} \left|\ell(y, f_{\theta}(x))\right| p_{\text{data}}(dx \, dy) < +\infty,$$

and therefore obtain that, for any  $\theta \in \Theta$ , the empirical risk  $\widehat{\mathcal{R}}_n(\theta)$  almost surely converges to the population risk  $\mathcal{R}(\theta)$  as  $n \to +\infty$ .

Let us however already emphasize that the goal of machine learning is not to minimize the loss on the training data, but on future data not yet seen, to guarantee some good prediction ability. This generalization performance can be discussed within the framework of decision theory, see Section 1.2.2 below and Chapter 10.

#### 1.2.1.1 Classification

In classification problems, the aim is to predict a label  $y \in \mathcal{Y}$  in a discrete set  $\mathcal{Y}$ . For multiclass classification,  $\mathcal{Y} = \{1, \ldots, K\}$ , with K the number of classes considered in the classification problem. For binary classification,  $\mathcal{Y} = \{0, 1\}$  for certain algorithms (for instance logistic regression, see Chapter 3), while  $\mathcal{Y} = \{-1, 1\}$  for others (for instance support vector machines, see Chapter 6). One example besides the MNIST dataset, and which we will consider for the hands-on, is the Iris data set (see for instance [40, Section 1.2.1.1]), for which there are d = 4 features and K = 3 classes.

There are various choices for the loss functions. A natural one is  $\ell(z, y) = \mathbf{1}_{y \neq z}$ , for which the associated risk function  $\mathcal{R}(\theta) = \mathbb{P}(y \neq f_{\theta}(x))$  in (1.4) is called the misclassification error. However, as discussed in Section 3.1, the corresponding risk function is difficult to optimize. The use of gradient-like optimization algorithms, as those presented in Chapter 4, suggests to resort to smooth surrogate loss functions. Appropriate choices for these surrogate functions are discussed in Section 3.1.

Let us conclude this section by presenting a typical way of predicting classes in multiclass classification, relying on the softmax function, which takes as argument  $(a_1, \ldots, a_K) \in \mathbb{R}^K$  and returns

$$S_K(a_1,\ldots,a_K) = \left(\frac{e^{a_1}}{\sum_{k=1}^K e^{a_k}}, \frac{e^{a_2}}{\sum_{k=1}^K e^{a_k}}, \ldots, \frac{e^{a_K}}{\sum_{k=1}^K e^{a_k}}\right) \in [0,1]^K.$$

Note that the denominators in the arguments of  $S_K$  ensure that the components of  $S_K$  sum to 1, so that  $S_K$  returns a discrete probability. The choice of an exponential function is convenient to transform real numbers into positive ones.

The prediction is based on some function  $f_{\theta} : \mathcal{X} \to \mathbb{R}^{K}$ , so that the vector  $S_{K}(f_{\theta}(x))$  gives, for a new input  $x \in \mathcal{X}$ , the probabilities of the label y to be in one of the classes; more precisely, the probability that y = k is the k-th component of  $S_{K}(f_{\theta}(x))$ , denoted by  $S_{K}(f_{\theta}(x))_{k}$ . A natural classification rule is therefore

$$y \in \operatorname*{argmax}_{1 \leq k \leq K} \left\{ S_K(f_\theta(x))_k \right\}.$$

In fact, the choice  $f_{\theta}(x) = W^{\top}x + b$ , with  $\theta = (W, b)$  where  $W \in \mathbb{R}^{d \times K}$  and  $b \in \mathbb{R}^{K}$  (we interpret here x as a column vector), corresponds to multinomial logistic regression (see Section 3.2.4). The part W of the parameter  $\theta$  is called weight, while the part b is called bias (let us however emphasize that this has nothing to do with the notion of bias for estimators, as encountered in statistics).

#### 1.2 Supervised learning

**Remark 1.1.** For K = 2 classes (binary classification), a simple computation shows that

$$S_2(a_1, a_2) = \left(\frac{e^{a_1}}{e^{a_1} + e^{a_2}}, 1 - \frac{e^{a_1}}{e^{a_1} + e^{a_2}}\right) = (\sigma(a_1 - a_2), 1 - \sigma(a_1 - a_2)),$$

where we introduced the sigmoid function

$$\sigma(z) = \frac{1}{1 + e^{-z}} = \frac{e^z}{1 + e^z}.$$

This shows that it suffices to learn  $a_1 - a_2$ , and not  $a_1, a_2$  separately. The model can then be parametrized using only  $\sigma(f_{\theta}(x))$  with  $f_{\theta}(x) = W^{\top}x + b \in \mathbb{R}$  where  $W \in \mathbb{R}^{K}$  and  $b \in \mathbb{R}$ .

#### 1.2.1.2 Regression

Regression problems correspond to situations where the space  $\mathcal{Y}$  is continuous, typically  $\mathcal{Y} = \mathbb{R}^{K}$ . The most famous choice of loss function is the square loss  $\ell(y,z) = \|y-z\|_2^2$  (with  $\|\cdot\|_2$  the standard Euclidean norm on  $\mathbb{R}^{K}$ ), for which the associated empirical risk minimization reads

$$\widehat{\mathcal{R}}_{n}(\theta) = \frac{1}{n} \sum_{i=1}^{n} \|y_{i} - f_{\theta}(x_{i})\|_{2}^{2}.$$
(1.5)

Other loss functions can be considered, for instance the mean absolute loss (with  $y = (y_1, \ldots, y_K)$ )

$$\ell(y,z) = \|y - z\|_1 = \sum_{k=1}^{K} |y_k - z_k|,$$

or generalizations of the latter loss. The interest of these functions is that large deviations are less penalized, which is important to limit the impact of outliers in the dataset and make regression more robust.

**Exercise 1.5.** Prove that (1.5) is proportional to the log-likelihood of the sample  $\mathcal{D}$  in a statistical model where y is distributed according to a Gaussian distribution centered on  $f_{\theta}(x)$  with given variance  $\sigma^2 > 0$ .

**Correction.** The log-likelihood writes  

$$L_n(\theta) = \prod_{i=1}^n \sqrt{\frac{1}{2\pi\sigma^2}} \exp\left(-\frac{\|f_\theta(x_i) - y_i\|_2^2}{2\sigma^2}\right),$$
so that  

$$\widehat{\mathfrak{T}}_n(\theta) = \frac{2\sigma^2}{2\sigma^2} \left(1 - L_n(\theta) - \frac{n}{2}L_n(\theta) - \frac{2}{2\sigma^2}\right)$$

$$\widehat{\mathcal{R}}_n(\theta) = -\frac{2\sigma^2}{n} \left( \log L_n(\theta) + \frac{n}{2} \log(2\pi\sigma^2) \right).$$

Typical instances of models used in regression are:

- linear regression  $f_{\theta}(x) = b + \sum_{j=1}^{d} w_j x_j$  for  $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ , with  $\theta = (b, w_1, \dots, w_d) \in \mathbb{R}^d$  $\mathbb{R}^{d+1}$ :
- polynomial regression, where some data point  $x \in \mathbb{R}$  is first featurized as  $\phi(x) = (1, x, \dots, x^p)$ (with p the degree of the polynomial), and then a prediction is made as  $f_{\theta}(x) = \theta^{\top} \phi(x)$  for some vector  $\theta \in \mathbb{R}^{p+1}$ . Note that this regression is still linear in the parameter  $\theta$  to estimate;

9

- 10 1 A first experience in machine learning
- regression using (deep) neural networks, presented in Chapter 8, and various other models generically relies on the prediction function  $f_{\theta}(x, \theta) = w^{\top} \phi(x, v)$ , where the parameters  $\theta = (w, v)$  are decomposed in two parts: a part w which is used to linearly combine the output of a (nonlinear) featurization function, and a part v which characterizes the featurization function  $\phi(x, v)$ . For neural networks for instance, this featurization function is of the form

$$\phi(x,v) = \rho_L \left( \rho_{L-1} \left( \rho_{L-2} \left( \dots \rho_1(x, W_1, b_1) \right), W_{L-1}, b_{L-1} \right), W_L, b_L \right)$$

where L is the number of layers, and  $v = (W_1, b_1, \ldots, W_L, b_L)$  gathers the parameters to be learnt at each layer.

#### 1.2.2 Decision theory and statistical learning

Decision theory (see Section 1.2.2.1) is a mathematical framework which allows to construct the best predictor when the distribution of the data  $p_{data}$  is known. However, as mentioned earlier, this distribution is typically not known in practice. One should therefore see the predictor provided by decision theory as some ideal choice which provides some baseline lower bound to the loss/risk which can be actually achieved. Statistical learning (see Section 1.2.2.2) aims at obtaining quantitative bounds/guarantees on how much the model which is learnt departs from the optimal choice provided by decision theory.

#### 1.2.2.1 Decision theory

**Definition 1.1 (Expected risk).** Given a measurable function  $f : \mathcal{X} \to \mathcal{Y}$ , a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$  and a probability distribution  $p_{data}$  on  $\mathcal{X} \times \mathcal{Y}$ , the expected risk of a prediction function is

$$\mathcal{R}(f) = \mathbb{E}\left[\ell(Y, f(X))\right] = \int_{\mathcal{X} \times \mathcal{Y}} \ell(y, f(x)) \, p_{\text{data}}(dx \, dy).$$
(1.6)

Note that, in this definition, we distinguished the random variables  $(X, Y) \sim p_{\text{data}}$ , and their realizations  $(x, y) \in \mathcal{X} \times \mathcal{Y}$ . In the remainder of these lectures notes, we will often denote the random variables with small letters. It should be clear from the context whether (x, y) denotes a random variable or one of its realizations.

**Example 1.1.** For classification, the loss function  $\ell(y, z) = \mathbf{1}_{y \neq z}$  leads to the expected risk  $\mathcal{R}(f) = \mathbb{P}(f(x) \neq y)$ , which is the error rate; alternatively,  $\mathcal{R}(f) = 1 - \mathbb{P}(f(x) = y)$ , where  $\mathbb{P}(f(x) = y)$  is the accuracy.

**Example 1.2.** For regression, a common choice is  $\ell(y, z) = ||y - z||_2^2$ , in which case the expected risk is the mean square error between the random variables f(X) and Y.

**Remark 1.2.** For some methods, the prediction provided by the function f is itself random, and depends on extra random variables (think of a binary classification example where one would guess the label at random). In this case, there is no simple outcome f(x), but rather a distribution of possible values for f(x). The definition (1.6) then still makes sense, upon further taking the expectation over the extra randomness.

In order to find the function  $f^*$  which minimizes the expected risk, we condition the values of the expected risk on the realizations of the input<sup>2</sup> (distinguishing here for clarity the random variable X and its realizations):

$$\mathcal{R}(f) = \mathbb{E}\left[\mathbb{E}\left(\ell(Y, f(X)) \mid X\right)\right] = \mathbb{E}\left[r(f(X) \mid X)\right] = \int_{\mathcal{X}} r\left(f(x') \mid x'\right) \, p_{\text{data}}(dx'),$$

 $<sup>^2</sup>$  Here and in the sequel, we abuse notation and denote  $p_{\rm data}$  and its marginal distributions by the same symbol.

where

$$r(z \mid x') = \mathbb{E}\left(\ell(Y, z) \mid X = x'\right)$$

is the conditional expectation of the loss. This reformulation makes it clear that, for any realization x' of X, the associated prediction f(x') should minimize  $r(\cdot | x')$ . This observation is summarized in the following result.

**Proposition 1.1.** The expected risk is minimized at a Bayes predictor  $f^* : \mathcal{X} \to \mathcal{Y}$  defined as<sup>3</sup>

$$f^{\star}(x') \in \operatorname*{argmin}_{z \in \mathcal{Y}} r(z \,|\, x'). \tag{1.7}$$

The Bayes risk is the risk of all Bayes predictors:

$$\mathcal{R}^{\star} = \mathbb{E}\left[\ell\left(Y, f^{\star}(X)\right)\right] = \mathbb{E}\left[\inf_{z \in \mathcal{Y}} \mathbb{E}\left(\ell(Y, z) \mid X\right)\right].$$

Let us emphasize that the Bayes predictor defined in (1.7) may not be unique as the set of minimizers for x' given may contain several elements. However, all Bayes predictors lead to the same Bayes risk. The Bayes risk is usually positive, unless the dependence between X and Y is deterministic.

**Definition 1.2 (Excess risk).** The excess risk of a measurable function  $f: \mathcal{X} \to \mathcal{Y}$  is  $\mathcal{R}(f) - \mathcal{R}^*$ .

Machine learning would be trivial if a Bayes predictor  $f^*$  was known! This would however require the knowledge of the distribution of outcomes y conditionally on the value of the input x, which is typically not the case as the distribution of the data is unknown.

**Exercise 1.6.** Compute the Bayes predictor and the Bayes risk for the loss function considered in Examples 1.1 (with  $\mathcal{Y} = \{0, 1\}$ ) and 1.2 (for one-dimensional outputs, i.e.  $\mathcal{Y} = \mathbb{R}$ ).

**Correction.** We first consider binary classification. We compute the conditional expectation of the loss by distinguishing the realizations of the label Y:

$$r(z \mid x') = \ell(1, z)\mathbb{P}(Y = 1 \mid X = x') + \ell(0, z)\mathbb{P}(Y = 0 \mid X = x')$$
  
=  $\ell(1, z)\eta(x') + \ell(0, z)(1 - \eta(x')),$  (1.8)

where  $\eta(x') = \mathbb{P}(Y = 1 | X = x')$ . Therefore, for x' fixed, the conditional expectation of the loss r(z | x') is minimized by retaining the minimum of the terms  $\eta(x')$  and  $1 - \eta(x')$ , *i.e.* 

$$f^{\star}(x') = \begin{vmatrix} 0 & \text{if } \eta(x') < 1/2, \\ 1 & \text{if } \eta(x') > 1/2, \end{vmatrix}$$

which can be summarized as  $f^*(x') = \mathbf{1}_{\eta(x')>1/2}$ . A specific decision should be made to break ties when  $\eta(x') = 1/2$ , for instance choosing the label at random; or, when  $\mathbb{P}(\eta(X) = 1/2) = 0$ , deciding that the label is always 0 or 1. The corresponding Bayes risk is

$$\mathcal{R}^{\star} = \mathbb{E}[\min\{\eta(X), 1 - \eta(X)\}].$$

Let us next consider regression. We also start by computing the conditional expectation of the loss:

$$r(z \mid x') = \mathbb{E}\left[ (Y - z)^2 \mid X = x' \right] = \mathbb{E}\left[ Y^2 \mid X = x' \right] - 2z\mathbb{E}\left[ Y \mid X = x' \right] + z^2$$

<sup>&</sup>lt;sup>3</sup> We assume here that there is a well defined minimizer; otherwise the argmin should be replaced by an arginf.

The right hand side is a second order polynomial with a nonnegative prefactor for  $z^2$ . It is therefore minimized for  $z = \mathbb{E}[Y | X = x']$ . This shows the Bayes predictor is the conditional expectation

$$f^{\star}(x') = \mathbb{E}\left[Y \mid X = x'\right].$$

The associated Bayes risk is the expectation of the conditional variance:

$$\mathcal{R}^{\star} = \mathbb{E}\left[\left(Y - \mathbb{E}\left[Y \mid X\right]\right)^{2}\right] = \mathbb{E}\left[Y^{2}\right] - \mathbb{E}\left[\left(\mathbb{E}\left[Y \mid X\right]\right)^{2}\right] = \mathbb{E}\left(\mathbb{E}\left[Y^{2} \mid X\right]\right) - \mathbb{E}\left[\left(\mathbb{E}\left[Y \mid X\right]\right)^{2}\right] = \mathbb{E}\left[\operatorname{Var}\left(Y \mid X\right)\right].$$

**Exercise 1.7.** For binary classification on  $\mathcal{Y} = \{-1, 1\}$ , compute the Bayes predictor and the Bayes risk for the loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$  defined by its values  $\ell(-1, 1) = c_-$  (cost of false positive),  $\ell(1, -1) = c_+$  (cost of false negative), and  $\ell(-1, -1) = \ell(1, 1) = 0$ . Interpret the so-obtained Bayes predictor.

**Correction.** We restart from (1.8), adapted to the setting  $\mathcal{Y} = \{-1, 1\}$ :

$$r(z \mid x') = \ell(1, z)\eta(x') + \ell(-1, z)(1 - \eta(x')),$$

so that  $r(1 | x') = c_{-}(1 - \eta(x'))$  and  $r(-1 | x') = c_{+}\eta(x')$ . Therefore,

$$f^{\star}(x') = \begin{vmatrix} 1 & \text{if } c_{-}(1 - \eta(x')) < c_{+}\eta(x'), \\ -1 & \text{if } c_{-}(1 - \eta(x')) > c_{+}\eta(x'), \end{vmatrix}$$

which can be summarized as  $f^{\star}(x') = \mathbf{1}_{\eta(x') > c_-/(c_++c_-)} - \mathbf{1}_{\eta(x') < c_-/(c_++c_-)}$ . The associated Bayes risk is

$$\mathcal{R}^{\star} = \mathbb{E}\left[\min\left\{c_{+}\eta(X), c_{-}(1-\eta(X))\right\}\right].$$

When  $c_{-} = c_{+}$ , one recovers the result of Exercise 1.6. When  $c_{-} < c_{+}$  (false negatives are more penalized then false positives), the prediction function is biased towards predicting positive outputs 1 compared to the situation when  $c_{+} = c_{-}$ . Conversely, when  $c_{-} > c_{+}$  (false positives are more penalized then false negatives), the prediction function is biased towards predicting -1 compared to the situation when  $c_{+} = c_{-}$ .

An example of situation where asymmetric cost functions are useful is spam classification. If the label 1 is for spams, then  $c_{-}$  should be (much) larger than  $c_{+}$ , as not classifying an email as spam (false negative) can be considered as a smaller nuisance than classifying a genuine and potentially important email as spam (false positive).

**Exercise 1.8 (Bayes predictor for robust regression).** Robust regression amounts to considering the loss function  $\ell(y, z) = |y - z|$  for regression problems with  $\mathcal{Y} = \mathbb{R}$ , instead of the square loss  $(y - z)^2$ . The interest of the loss function based on absolute values and not their square is that outliers have a smaller impact on the prediction.

(a) Consider a random variable Z which admits a positive continuous density  $\rho(z)$  with respect to the Lebesgue measure, and is integrable. Prove that

$$m = \mathop{argmin}_{c \in \mathbb{R}} \mathbb{E}|Z - c|$$

is the median of Z, namely the (here unique) value such that  $\mathbb{P}(Z \leq m) = \mathbb{P}(Z \geq m)$ .

(b) Assume for simplicity that the (unknown) distribution  $p_{data}(dx dy)$  of the data points has a positive continuous density with respect to the Lebesgue measure dx dy. Compute the Bayes predictor and the associated Bayes risk.

(c) We no longer assume that the (unknown) distribution  $p_{data}(dx dy)$  of the data points has a positive continuous density with respect to the Lebesgue measure dx dy. For a general random variable, the median is defined as a value  $m \in \mathbb{R}$  such that

$$\mathbb{P}(Z\leqslant m)\geqslant \frac{1}{2}, \qquad \mathbb{P}(Z\geqslant m)\geqslant \frac{1}{2}.$$

Note that the inequalities can be strict (in particular when there is a Dirac mass at m). Compute the Bayes predictor and the associated Bayes risk.

# Correction.

(a) It holds

$$F(c) := \mathbb{E}|Z - c| = \int_{-\infty}^{c} (c - z)\rho(z) \, dz + \int_{c}^{+\infty} (z - c)\rho(z) \, dz.$$

Note first that  $F(c) \to +\infty$  as  $c \to \pm\infty$  since  $F(c) \ge \max(c - \mathbb{E}(Z), \mathbb{E}(Z) - c)$ , and F is continuous, so it admits a minimizer. Moreover, the function F is  $C^1$ , and

$$F'(c) = \int_{-\infty}^{c} \rho(z) \, dz - \int_{c}^{+\infty} \rho(z) \, dz.$$
 (1.9)

To establish the differentiability of F, we consider

$$\frac{F(c+h) - F(c)}{h} = \int_{-\infty}^{c} \rho(z) \, dz - \int_{c}^{+\infty} \rho(z) \, dz + \frac{1}{h} \int_{c}^{c+h} (c+h-z)\rho(z) \, dz - \frac{1}{h} \int_{c}^{c+h} (z-c-h)\rho(z) \, dz.$$

The terms on the second line of the previous equality are equal to

$$\frac{1}{h} \int_0^h (h-t)\rho(c+t) \, dt - \frac{1}{h} \int_0^h (t-h)\rho(c+t) \, dt = 2 \int_0^h \left(1 - \frac{t}{h}\right)\rho(c+t) \, dt.$$

The latter term goes to 0 as  $h \to 0$ , as can be seen for instance by dominated convergence (the integral can be rewritten as the integral over  $\mathbb{R}$  of the function  $f_h(t) = \mathbf{1}_{[0,h]}(t)(1 - t/h)\rho(c+t)$ , where  $f_h$  is uniformly bounded by an integrable function since  $|f_h(t)| \leq \rho(c+t)$ , and  $f_h(t)$  converges pointwise to 0 as  $h \to 0$  for almost all  $t \in \mathbb{R}$ ). This shows that (1.9) holds.

Note that F' is continuous and increasing, tending to -1 as  $c \to -\infty$ , and to 1 as  $c \to +\infty$ . There is therefore a unique critical point m for F, which is necessarily a minimizer – hence the uniqueness of the minimizer. This minimizer is characterized by the equality

$$\mathbb{P}(Z \leqslant m) = \int_{-\infty}^{m} \rho(z) \, dz = \int_{m}^{+\infty} \rho(z) \, dz = \mathbb{P}(Z \ge m).$$

(b) The Bayes predictor is obtained as

$$f^*(x) = \operatorname*{argmin}_{z \in \mathbb{R}} \mathbb{E}\left( |Y - z| \mid X = x \right)$$

The result of the previous question shows that  $f^*(x) = m(x)$  is the median of the conditional distribution  $p_{\text{data}}(dy|x)$ :

$$\int_{-\infty}^{m(x)} p_{\text{data}}(dy|x) = \int_{m(x)}^{+\infty} p_{\text{data}}(dy|x).$$

The associated Bayes risk is  $\mathcal{R}^* = \mathbb{E}[|Y - m(X)|].$ 

- 14 1 A first experience in machine learning
  - (c) The Bayes predictor and risk are the same as in the previous question. To see this (and upon replacing Z by  $Z m_Z$  with  $m_Z$  a median of Z), we can assume without loss of generality that  $m_Z = 0$  is a median of Z. It suffices to prove that

$$\forall c \in \mathbb{R}, \qquad \mathbb{E}|Z - c| \leqslant E|Z|$$

We prove the inequality for  $c \ge 0$ , a similar reasoning holding for  $c \le 0$ . Note that |Z - c| - |Z| = c for  $Z \le 0$ , while  $|Z - c| - |Z| \ge -c$  for  $Z \ge 0$ . Therefore,

$$\mathbb{E}\left(|Z-c|-|Z|\right) = \mathbb{E}\left[\left(|Z-c|-|Z|\right)\mathbf{1}_{Z\leqslant 0}\right] + \mathbb{E}\left[\left(|Z-c|-|Z|\right)\mathbf{1}_{Z>0}\right]$$
$$\geq c\mathbb{P}(Z\leqslant 0) - c\mathbb{P}(Z>0) = c\left(2\mathbb{P}(Z\leqslant 0)-1\right) \geq 0,$$

where the last inequality follows from the definition of the median.

**Exercise 1.9.** In a binary classification context on  $\mathcal{Y} = \{0,1\}$  with  $\ell(0,0) = \ell(1,1) = 0$ and  $\ell(1,0) = \ell(0,1)$ , compute the so-called "chance level", which is the risk associated with the random prediction function  $f(x) = \mathbf{1}_{U \leq 1/2}$  with  $U \sim \mathcal{U}[0,1]$  a random variable uniformly distributed on [0,1], independent of the sample points.

**Correction.** The prediction function  $f(x) = \mathbf{1}_{U \leq 1/2}$  has a random output. The associated loss is (denoting for clarity as subscripts with respect to each random variable expectations are taken)

$$\mathcal{R}(f) = \mathbb{E}_X \left( \mathbb{E}_{Y,U} \left[ \ell(Y, \mathbf{1}_{U \leqslant 1/2}) | X \right] \right).$$

Now, recalling the notation  $\eta(x') = \mathbb{P}(Y = 1 | X = x')$ ,

$$\mathbb{E}_{Y,U}\left[\ell(Y, \mathbf{1}_{U\leqslant 1/2}) \middle| X = x'\right] = \mathbb{E}_U\left[\ell(1, \mathbf{1}_{U\leqslant 1/2})\eta(x') + \ell(0, \mathbf{1}_{U\leqslant 1/2})(1 - \eta(x'))\right] \\ = \frac{1}{2}\left[\ell(1, 0)\eta(x') + \ell(0, 1)(1 - \eta(x'))\right] = \frac{1}{2}\ell(1, 0).$$

 $= \frac{1}{2} \left[ \ell(1,0)\eta(x) + \ell(0,1)(1-\eta(x)) \right]$ For the loss function  $\ell(y,z) = \mathbf{1}_{y\neq z}$ , one therefore obtains  $\mathcal{R}(f) = 1/2$ .

#### 1.2.2.2 Statistical learning

Decision theory tells us what to do if the distribution  $p_{data}$  of the data is known, in which case optimal performance can be achieved. Now, in practice, this distribution is unknown. Statistical learning is a framework to obtain bounds or guarantees quantifying how much the the learned model departs from the (unknown) optimal choice. A key concept in this endeavor is empirical risk.

**Definition 1.3 (Empirical risk).** Given a measurable function  $f : \mathcal{X} \to \mathcal{Y}$ , a loss function  $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}_+$  and a data set  $\{(x_1, y_1), \ldots, (x_n, y_n)\}$ , the empirical risk of a prediction function is

$$\widehat{\mathcal{R}}_n(f) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f(x_i)).$$

The empirical risk can in principle be successfully minimized: upon considering a sufficiently flexible class of predictor functions, one can achieve zero training loss. A successful empirical risk minimization does however not imply successful predictions, as the prediction function may fit too closely the training data and learn irrelevant features. This is the celebrated issue of *overfitting*, which we will come back to in Chapter 2. The discrepancy between the expected risk and the empirical risk is measured by the generalization gap

$$\mathcal{R}(f) - \widehat{\mathcal{R}}_n(f)$$

The population risk  $\mathcal{R}(f)$  is approximated in practice by the test risk, which corresponds to some empirical risk based on data points which have not been seen during training, and can therefore be considered as new data points sampled from  $p_{\text{data}}$ .

# add picture here

Let us make precise the various steps involved in training a model and then assessing its quality. Given a dataset  $\mathcal{D}$ , we first separate it into a training set  $\mathcal{D}_{\text{train}}$  and a test set  $\mathcal{D}_{\text{test}}$ . A typical ratio is to keep 80% of the data points in the training set, and 20% in the test set. This allows to balance the need to have many data points in order for the training procedure to see as many instances as possible, while still keeping enough data points in order to have a sufficiently reliable estimation of the expected risk. The next step is to find a predictor function  $\hat{f}_n$  which minimizes the empirical risk over the data set  $\mathcal{D}_{\text{train}}$ , among a prescribed class  $\mathcal{F}$  of functions  $\mathcal{X} \to \mathcal{Y}$ :

$$\widehat{f_n} \in \underset{f \in \mathcal{F}}{\operatorname{argmin}} \left\{ \frac{1}{|\mathcal{D}_{\operatorname{train}}|} \sum_{(x_i, y_i) \in \mathcal{D}_{\operatorname{train}}} \ell(y_i, f(x_i)) \right\}.$$

The final step is to approximate the expected risk as

$$\frac{1}{|\mathcal{D}_{\text{test}}|} \sum_{(x_j, y_j) \in \mathcal{D}_{\text{test}}} \ell\left(y_j, \widehat{f}_n(x_j)\right).$$

In fact, we will see later on that one needs to further divide the training set in an actual training set, and a validation set, used to fix the hyperparameters of the method. The important point here is that the test set is never used to determine the predictor function  $\hat{f}_n$  which minimizes the empirical risk. It is only used at the very end, to assess the quality of the prediction.

#### 1.2.2.3 Learning from data

Let us conclude this introduction to machine learning by mentioning three important ways to obtain a (good) prediction from the data at hand rather than from the full knowledge of (the unknown distribution)  $p_{\text{data}}$ :

- in local averaging methods, one relies on some interpolation procedure with the data at hand to provide predictions for new data points. One famous example is the K-nearest neighbor method, presented in Section 1.3;
- the most common situation we will consider is to obtain prediction functions from the minimization of the empirical risk;
- we will only briefly consider other ways to find predictors, in particular using the boosting paradigm (see Section 7.3). One could also rely on probabilistic methods, which we however do not consider in these lecture notes.

# 1.3 Local averaging methods

The material in this section is based on [40, Chapter 16] and [4, Chapter 6], as well as [6, Chapter 14] and [49, Chapter 19]. The main idea behind local averaging methods is that "things that look alike must be alike", *i.e.* close-by points should have similar labels. A successful prediction relies in this context on the smoothness of the data, in the sense that the labels/values y should depend smoothly on the inputs x (if the labels/values y were truly random, there would be no way to learn meaningful patterns in the data and generalize this knowledge to new data points...).

#### 1.3.1 Presentation of the K-nearest neighbor algorithm

Nearest neighbor methods are based on a majority vote among the closest neighbors for classification, and some form of averaging over the closest neighbors for regression. They provide a useful starting method in machine learning since they are easy to code and to understand. In essence, they amount to memorizing a training set, and predicting the label of a new instance x' based on the labels of the closest points  $x_i$  in the training set.

To make this precise, we first need to choose some metric  $\mathbf{d} : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$  to measure distances between data points. This is a key element in the method. A typical choice is the Euclidean distance, which however behaves badly in high dimensions (where all points are far away from each other). A good practice to alleviate these issues is to perform some rescaling of the data, for instance through standardization or whitening (recall the discussion at the end of Section 1.1.4), or to rely on the Mahalanobis distance

$$\mathbf{d}(x,x') = \sqrt{(x-x')^{\top} \Sigma_n^{-1} (x-x')},$$

where the covariance  $\Sigma_n$  of the data inputs is defined in (1.2) (in fact this corresponds to considering a simple Euclidean distance on the whitehed data). Note that the distance function actually used can be considered as a hyperparameter of the method, to be set using some (cross) validation procedure.

Additionally, it may be beneficial to first perform some dimensionality reduction (using for instance principal component analysis, see Chapter 5, or its nonlinear generalization based on autoencoders, see Section 8.3) – both in order to reduce the computational cost of the method, and to possibly improve predictions since only the large scale features of the data are retained. This heuristic consideration is backed up by the theoretical estimates discussed around (1.23).

**Algorithm 1.1.** Fix a number of neighbors  $K \ge 1$  and a training data set  $\{(x_1, y_1), \ldots, (x_n, y_n)\} \subset \mathcal{X} \times \mathcal{Y}$ . For a new input  $x' \in \mathcal{X}$ , sort the data points by increasing distance<sup>4</sup>

$$d\left(x_{i_1(x')}, x'\right) \leq d\left(x_{i_2(x')}, x'\right) \leq \ldots \leq d\left(x_{i_n(x')}, x'\right), \tag{1.10}$$

where  $\{i_1(x'), \ldots, i_n(x')\} = \{1, \ldots, n\}$ . The prediction is then performed as

- (classification) y' is the majority of labels among the labels of the K-closest neighbors  $\{y_{i_1(x')}, \ldots, y_{i_K(x')}\};$
- (regression) y' is the average of the values for the K-closest neighbors:  $y' = \frac{1}{K} \sum_{k=1}^{K} y_{i_k(x')}$ .

The value K = 1 corresponds to predictions based on a Voronoi tesselation of the input space. This means in practice that a new input inherits the label or value of the closest input in the data set. Intuitively, this seems however not very robust a choice, as a small variation in the inputs can lead to abrupt changes in the predictions. As we will see in Section 1.3.3.3, the variance in the prediction is indeed large when K is small. On the other hand, when K is large, decision boundaries separating various values for the predictions will be smoother, and the predictors more robust. However, the predictions may be quite biased when K is large. Overall, this calls for some optimal choice of K based on a bias/variance trade-off. We will see both theoretical guidelines which give firm foundations to this choice (see Section 1.3.3), as well as a practical manner of fixing the value of K based on a key procedure in machine learning, cross validation (see Section 1.3.2).

Let us finally discuss the algorithmic cost of finding the closest neighbr. A priori, this scales as O(nd) (when computing the distance to all elements in the data set), but the cost can be reduced by clever algorithmic strategies based on a tree search (see the discussions in [40, Section 16.1.3] and [49, Section 19.3]). This also suggests that it makes sense to sparsify to some extent the training set and keep only the most relevant points.

confirm

<sup>&</sup>lt;sup>4</sup> with some rule to break ties, for instance  $i_k(x') \leq i_{k+1}(x')$  (*i.e.* always choose the smallest index), or choosing at random among equidistant points.

#### 1.3.2 Choosing K by (cross) validation

We give here a brief introduction to (cross) validation, referring to Section 10.2 for complements. We distinguish two cases, depending on whether the training data set is large or not.

#### 1.3.2.1 Large training data sets

When enough data is available, the training data set  $\mathcal{D}$  is decomposed into an actual training data set  $\mathcal{D}_{train}$  and a validation set  $\mathcal{D}_{val}$ . The validation set should be large enough so that the population risk can be well approximated by computing the empirical risk over it; while most of the data should still be used for training. A common rule of thumb is to devote 20% of the initial data  $\mathcal{D}$  for validation, and keep 80% for training. The assignment to  $\mathcal{D}_{train}, \mathcal{D}_{val}$  is done at random, typically by shuffling the data.

Once these sets are obtained, the validation procedure amounts to iterating over the possible values of the hyperparameters of the method, here the number of neighbors  $K \ge 1$ , training the model for these hyperparameters (which need not be done for K-nearest neighbors as no training is required in the method), and then computing the empirical risk on the validation set:

$$\widehat{\mathcal{R}}_{n}^{\mathrm{val}}(K) = \frac{1}{|\mathcal{D}_{\mathrm{val}}|} \sum_{(x_{j}, y_{j}) \in \mathcal{D}_{\mathrm{val}}} \ell\left(y_{j}, \widehat{y}_{j}\right),$$

where  $n = |\mathcal{D}_{\text{train}}|$  is here the size of the actual training set  $\mathcal{D}_{\text{train}}$ , and  $\hat{y}_j$  is the prediction for the input  $x_j$  of the method trained on  $\mathcal{D}_{\text{train}}$ . One then chooses the value of K for which  $\widehat{\mathcal{R}}_n^{\text{val}}(K)$  is minimal.

Let us emphasize once again that the procedure crucially relies on  $\mathcal{D}_{\text{train}} \cap \mathcal{D}_{\text{val}} = \emptyset$ , *i.e.* no validation data is used in the training procedure. The statistical consistency of validation is discussed in Section 10.2.

#### 1.3.2.2 Small training data sets

When the data set  $\mathcal{D}$  is small, there are two obstructions to the validation procedure described in Section 1.3.2.1: first, the validation set would be small, and so the approximation of the population risk by the empirical risk computed on the validation set would be unreliable; second, one would prefer not to loose data on validation, and use as much data as possible for training. Crossvalidation is a procedure to address these two issues.

Concretely, the data set  $\mathcal{D}$  is first shuffled, then decomposed into L folds of same sizes  $|\mathcal{D}|/L$  (assuming that this number is an integer). Among these folds, one is kept for validation, while the remaining L-1 ones are used for training. The empirical risk is next computed with the data of the fold which was chosen for validation. This is then repeated for all possible L choices for the validation set, the approximation of the expected risk being obtained as the average of the empirical risks computed for each validation fold. More precisely, denoting by  $\mathcal{D}_k$  for  $1 \leq k \leq L$  the various folds and by  $n = |\mathcal{D}|$  the size of the initial data set, the expected risk is approximated as

$$\widehat{\mathcal{R}}_{n,L}^{\text{CV}} = \frac{1}{L} \sum_{k=1}^{L} \widehat{\mathcal{R}} \left( \mathcal{D} \setminus \mathcal{D}_k, \mathcal{D}_k \right), \qquad (1.11)$$

where

$$\widehat{\mathcal{R}}\left(\mathcal{D}_{\text{train}}, \mathcal{D}_{\text{val}}\right) = \frac{1}{\left|\mathcal{D}_{\text{val}}\right|} \sum_{(x_j, y_j) \in \mathcal{D}_{\text{val}}} \ell\left(y_j, \widehat{y}_j^{\mathcal{D}_{\text{train}}}\right),$$

with  $\hat{y}_j^{\mathcal{D}_{\text{train}}}$  the prediction for the input  $x_j$  and the method trained on  $\mathcal{D}_{\text{train}}$ .

Usual values for L are in the range 5 - 10. Another option is to consider L = n, which corresponds to the so-called "leave-one-out" cross-validation procedure. Similarly to the validation procedure described in Section 1.3.2.1, the cross-validation risk (1.11) is computed for various values of the hyperparameters of the method (here, the number K of nearest neighbors), choosing in the end the hyperparameter leading to the smallest cross-validation risk.

#### 1.3.3 Theoretical analysis of the K-nearest neighbor algorithm

Recall that the aim of machine learning, as discussed in Section 1.2.2, is to estimate the Bayes predictor  $f^*$ , knowing only training data points  $\mathcal{D}_{\text{train}} = \{(x_i, y_i)_{1 \leq i \leq n}\}$ , in order to minimize the excess risk  $\mathcal{R}(f) - \mathcal{R}^* = \mathbb{E}[\ell(Y, f(X))] - \mathcal{R}^*$ . We first describe in Section 1.3.3.1 the general framework allowing to study the theoretical properties of a broader class of estimators known as local methods; then perform a consistency analysis for this class of estimators in Section 1.3.3.2, finally specifying to K-nearest neighbors in Section 1.3.3.3.

#### 1.3.3.1 Local methods

Local methods, such as K-nearest neighbors, approximate  $f^*$  without any form of optimization. Conceptually, this is done by approximating the conditional distribution  $p_{\text{data}}(dy | x)$  of labels y given inputs x by some distribution  $\hat{p}(dy | x)$ , so that the minimization problem (1.7)

$$f^{\star}(x) \in \operatorname*{argmin}_{z \in \mathcal{Y}} \int_{\mathcal{Y}} \ell(y, z) \, p_{\mathrm{data}}(dy \,|\, x)$$

is approximated by

$$\widehat{f}(x) \in \operatorname*{argmin}_{z \in \mathcal{Y}} \int_{\mathcal{Y}} \ell(y, z) \, \widehat{p}(dy \,|\, x)$$

This corresponds to some form of plug-in estimator, as encountered in Statistics. Let us illustrate this general strategy for two prominent examples:

• for multi-class classification over C classes, with the 0-1 loss  $\ell(y, z) = \mathbf{1}_{y \neq z}$ ,

$$\widehat{f}(x) \in \operatorname*{argmin}_{z \in \{1, \dots, C\}} \left\{ \sum_{k=1}^{C} \mathbf{1}_{z \neq k} \, \widehat{p}(y = k \,|\, x) \right\} = \operatorname*{argmin}_{z \in \{1, \dots, C\}} \left\{ 1 - \sum_{k=1}^{C} \mathbf{1}_{z = k} \, \widehat{p}(y = k \,|\, x) \right\},$$

so that

$$\widehat{f}(x) \in \operatorname*{argmax}_{k \in \{1,\dots,C\}} \widehat{p}(y=k \,|\, x); \tag{1.12}$$

 for regression with the square loss, the predictor is the conditional expectation associated with p:

$$\widehat{f}(x) = \int_{\mathcal{Y}} y \, \widehat{p}(dy \,|\, x). \tag{1.13}$$

We further consider here linear estimators, for which the approximation to the conditional distribution  $p_{\text{data}}(dy | x)$  is a convex combination of Dirac masses at  $y_i$ , namely

$$\widehat{p}(dy \,|\, x) = \sum_{i=1}^{n} \widehat{w}_i(x) \delta_{y_i}(dy),$$

where the weights depend only on x and  $x_1, \ldots, x_n$  (and not on the labels) and are such that

$$\forall x \in \mathcal{X}, \qquad \forall i \in \{1, \dots, n\}, \quad \widehat{w}_i(x) \ge 0, \qquad \sum_{i=1}^n \widehat{w}_i(x) = 1. \tag{1.14}$$

In addition to K-nearest neighbors, partition estimators and kernel density estimators (known as Nadaraya–Watson estimators) also belong to this class; see [4, Section 6.2]. In this setting, the estimators (1.12) and (1.13) respectively read

$$\widehat{f}(x) \in \operatorname*{argmax}_{k \in \{1,...,C\}} \left\{ \sum_{i=1}^{n} \widehat{w}_i(x) \mathbf{1}_{y_i=k} \right\},$$

and

$$\widehat{f}(x) = \sum_{i=1}^{n} \widehat{w}_i(x) y_i.$$

The K-nearest neighbors methods corresponds to the choice

$$\widehat{w}_j(x) = \begin{vmatrix} \frac{1}{K} & \text{if } j \in \{i_1(x), \dots, i_K(x)\}, \\ 0 & \text{otherwise.} \end{vmatrix}$$
(1.15)

#### 1.3.3.2 Consistency analysis for local methods

We discuss the consistency of local methods for regression; see the introduction of [4, Section 6.3] for comments and references on the classification case. Consistency is understood here as the excess risk converging to 0 in some sense as the number of training data points goes to infinity. This notion of consistency (rather than, say, quantifying the convergence of predictors to the Bayes predictor) is adapted to machine learning problems as the aim is to devise good predictors, as measured by the fact that the prediction error is small.

We make the following assumptions, to simplify the presentation.

**Assumption 1.1.** The noise is bounded: There exists  $\sigma \in \mathbb{R}_+$  such that  $|Y - \mathbb{E}(Y | X)| \leq \sigma$  almost surely.

**Assumption 1.2.** The Bayes predictor  $f^*(x) = \mathbb{E}[Y | X = x]$  is B-Lipschitz for some continuous distance function  $d: \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$ :

$$\forall (x_1, x_2) \in \mathcal{X}^2, \qquad |f^{\star}(x_1) - f^{\star}(x_2)| \leqslant B \mathsf{d}(x_1, x_2).$$

A careful inspection of the proof shows that it would be sufficient to assume that  $\mathbb{E}\left[|Y - \mathbb{E}(Y \mid X)|^2\right] \leq \sigma^2$  instead of Assumption 1.1. Weaker conditions than Assumption 1.2 are discussed in [4, Section 6.4] (based on the notion of universal consistency).

The first step is to decompose the prediction error as the sum of two contributions, one depending on the realization of the labels  $y_i$  (which therefore measures some variance) and one deterministic term for  $x_1, \ldots, x_n$  given (akin to some bias). More precisely, recalling the abuse of notation where  $y_i$  is a random variable, and using the second condition in (1.14),

$$\begin{aligned} \widehat{f}(x) - f^{\star}(x) &= \sum_{i=1}^{n} y_{i} \widehat{w}_{i}(x) - \mathbb{E}(Y \mid X = x) \\ &= \sum_{i=1}^{n} \widehat{w}_{i}(x) \left[ y_{i} - \mathbb{E}(y_{i} \mid x_{i}) \right] + \sum_{i=1}^{n} \widehat{w}_{i}(x) \left[ \mathbb{E}(y_{i} \mid x_{i}) - \mathbb{E}(Y \mid X = x) \right] \\ &= \sum_{i=1}^{n} \widehat{w}_{i}(x) \left[ y_{i} - \mathbb{E}(y_{i} \mid x_{i}) \right] + \sum_{i=1}^{n} \widehat{w}_{i}(x) \left[ f^{\star}(x_{i}) - f^{\star}(x) \right]. \end{aligned}$$

We can then consider the square error, conditionally on the inputs  $x_1, \ldots, x_n$ :

$$\mathbb{E}\left[\left|\widehat{f}(x) - f^{\star}(x)\right|^{2} \middle| x_{1}, \dots, x_{n}\right]$$

$$= \mathbb{E}\left[\left(\sum_{i=1}^{n} \widehat{w}_{i}(x) \left[y_{i} - \mathbb{E}(y_{i} \mid x_{i})\right]\right)^{2} \middle| x_{1}, \dots, x_{n}\right] + \left(\sum_{i=1}^{n} \widehat{w}_{i}(x) \left[f^{\star}(x_{i}) - f^{\star}(x)\right]\right)^{2}$$

$$= \sum_{i=1}^{n} \widehat{w}_{i}(x)^{2} \mathbb{E}\left[\left(y_{i} - \mathbb{E}(y_{i} \mid x_{i})\right)^{2} \middle| x_{i}\right] + \left(\sum_{i=1}^{n} \widehat{w}_{i}(x) \left[f^{\star}(x_{i}) - f^{\star}(x)\right]\right)^{2}.$$

The right hand side can be seen as the sum of a variance term and a squared bias, as for standard mean square errors encountered in the mathematical analysis of estimators in parametric models of Statistics. In view of Assumptions 1.1 and 1.2, and using a discrete Cauchy–Schwarz inequality first for the second term,

$$\mathbb{E}\left[\left|\widehat{f}(x) - f^{\star}(x)\right|^{2} \left|x_{1}, \dots, x_{n}\right] \leqslant \sigma^{2} \sum_{i=1}^{n} \widehat{w}_{i}(x)^{2} + \left(\sum_{i=1}^{n} \widehat{w}_{i}(x)\right) \left(\sum_{i=1}^{n} \widehat{w}_{i}(x) \left|f^{\star}(x_{i}) - f^{\star}(x)\right|^{2}\right) \\ \leqslant \sigma^{2} \sum_{i=1}^{n} \widehat{w}_{i}(x)^{2} + B^{2} \sum_{i=1}^{n} \widehat{w}_{i}(x) \mathsf{d}(x_{i}, x)^{2}.$$
(1.16)

Let us next use these estimates to quantify the excess risk  $\mathcal{R}(\hat{f}) - \mathcal{R}^*$ , and in fact the expectation of the excess risk with respect to realizations of the training data. Let us emphasize that we consider here only the expected excess risk for simplicity. In some frameworks, such as the probably approximately correct (PAC) setting (see Chapter 10) one may prefer bounds which hold stronger than in average, but this requires more work/assumptions on the distribution of the data. For a given data set, the excess risk can be rewritten as

$$\begin{aligned} \mathcal{R}(\widehat{f}) - \mathcal{R}^{\star} &= \mathbb{E}\left[\left(Y' - \widehat{f}(X')\right)^{2}\right] - \mathcal{R}^{\star} = \mathbb{E}\left[\left(Y' - f^{\star}(X') + f^{\star}(X') - \widehat{f}(X')\right)^{2}\right] - \mathcal{R}^{\star} \\ &= \mathbb{E}\left[\left(f^{\star}(X') - \widehat{f}(X')\right)^{2}\right] + 2\mathbb{E}\left[\mathbb{E}\left(Y' - f^{\star}(X') \mid X'\right)\left(f^{\star}(X') - \widehat{f}(X')\right)\right] \\ &= \mathbb{E}\left[\left(\widehat{f}(X') - f^{\star}(X')\right)^{2}\right] = \int_{\mathcal{X}} \left(\widehat{f}(x') - f^{\star}(x')\right)^{2} p_{\text{data}}(dx'), \end{aligned}$$

where, as indicated by the last integral, the expectation is over the realizations of the test data (x', y'). Therefore, taking now the expectation with respect to the training set (indicated by  $\mathbb{E}_{\mathcal{D}}$  for clarity; keeing the integral notation for expectations with respect to the test data; note that  $\hat{f}$  depends on  $\mathcal{D}$  through the weights  $\hat{w}_i$ , but does not depend on the test data point), and next using (1.16),

$$\mathbb{E}_{\mathcal{D}}\left[\mathcal{R}(\widehat{f})\right] - \mathcal{R}^{\star} = \int_{\mathcal{X}} \mathbb{E}_{\mathcal{D}}\left[\left(\widehat{f}(x') - f^{\star}(x')\right)^{2}\right] p_{\text{data}}(dx')$$

$$\leq \sigma^{2} \sum_{i=1}^{n} \int_{\mathcal{X}} \mathbb{E}_{\mathcal{D}}\left(\widehat{w}_{i}(x')^{2}\right) p_{\text{data}}(dx') + B^{2} \int_{\mathcal{X}} \mathbb{E}_{\mathcal{D}}\left[\sum_{i=1}^{n} \widehat{w}_{i}(x') \mathbf{d}(x_{i}, x')^{2}\right] p_{\text{data}}(dx'). \quad (1.17)$$

The first term on the previous right hand side, which arises from the inherent label noise (and would therefore be present even for the optimal prediction function  $f^*$ ) involves the factor

$$\mathbb{E}_{\mathcal{D}}\left(\widehat{w}_{i}(x')^{2}\right) = \mathbb{E}_{\mathcal{D}}\left[\left(\widehat{w}_{i}(x') - \frac{1}{n}\right)^{2}\right] + \frac{2}{n} - \frac{1}{n^{2}}$$

which measures the deviation to uniform weights. Note indeed that uniform weights minimize the latter quantity since, by a discrete Cauchy–Schwarz inequality,

$$n\left(\sum_{i=1}^{n}\widehat{w}_{i}(x')^{2}\right)^{2} \geqslant \left(\sum_{i=1}^{n}\widehat{w}_{i}(x')\right)^{2} = 1,$$

with equality if and only if  $\hat{w}_i(x') = 1/n$  for all  $1 \leq i \leq n$ . The expectation of the squares of the weights is therefore related to overfitting issues, as the variance of the weights is largest when all the mass is put in one weight, for a data point which changes from a realization of the data set to another. The second term on the right hand side of (1.17), which depends on the regularity of the optimal prediction function  $f^*$ , measures some bias in the prediction. It is smaller when  $f^*$  varies less, *i.e.* B is small. It is related to underfitting as, in the worst case scenario, all weights are the same and nothing is learned.

It suffices at this stage to ensure that the two terms on the right hand side of (1.17) converge to 0 as  $n \to +\infty$ . Ideally, one would also like to make precise the convergence rate here, and to find guidelines to choose the hyperparameters of the method. This is where one needs to specify the analysis to the algorithm under consideration.

#### 1.3.3.3 Consistency analysis for K-nearest neighbors

We make precise in this section how the two terms on the right hand side of (1.17) depend on various parameters, in particular the number of training points n, the number of neighbors K and the dimension d of the inputs.

The first remark is that the variance term is particularly simple for K-nearest neighbors. Indeed, in view of (1.15),

$$\forall x' \in \mathcal{X}, \qquad \sum_{i=1}^{n} \widehat{w}_i(x')^2 = \frac{1}{K}.$$

The first term on the right hand side of (1.17) can therefore be bounded by  $\sigma^2/K$ . This already shows that one needs to let the number of neighbors K go to infinity when  $n \to +\infty$  for the expected excess risk to converge to 0.

Now, in order to equilibrate between the two contributions in (1.17), we need to more carefully quantify the bias. This requires bounds on the distances to the nearest neighbors; in fact, in view of (1.10) and (1.15),

$$\sum_{i=1}^{n} \widehat{w}_{i}(x') \mathsf{d}(x_{i}, x')^{2} \leq \mathsf{d}\left(x_{i_{K}(x')}, x'\right)^{2}, \qquad (1.18)$$

so controlling the distance to the K-th neighbor is key. We need at this stage to introduce another assumption.

#### **Assumption 1.3.** The distribution $p_{\text{data}}(dx')$ of the inputs has a compact support $\mathcal{X} \subset \mathbb{R}^d$ .

The fact that an additional assumption is needed is related to the so-called *no free lunch theorem*, which roughly says that there is no algorithm that can learn a fully arbitrary distribution (see Section 10.3 for further discussions). In the remainder, we denote by

$$\operatorname{diam}(\mathcal{X}) = \max_{(x,x') \in \mathcal{X}^2} \mathsf{d}(x,x') < +\infty$$

the diameter of the compact set  $\mathcal{X}$  (which is closed and bounded), and choose  $d(x, x') = ||x - x'||_{\infty}$ with

$$\|\xi\|_{\infty} = \max_{1 \leqslant i \leqslant d} |\xi_i|$$

the  $\ell^{\infty}$  norm on  $\mathbb{R}^d$ .

**Lemma 1.1.** Suppose that Assumption 1.3 holds true, and that  $d \ge 2$ . Consider random variables  $(X_1, \ldots, X_n, X_{n+1})$  independently and identically distributed according to p. Then, the expected squared distance between  $X_{n+1}$  and its first nearest neighbor among  $\{X_1, \ldots, X_n\}$  in  $\ell^{\infty}$  norm is upper bounded by  $4 \operatorname{diam}(\mathcal{X})^2 n^{-2/d}$ .

Note that the scaling of the distance with respect to the dimension is not so good, as a very large number of points  $n \sim \varepsilon^{-d}$  is needed in order for the expected squared distance to be of order  $\varepsilon^2$ .

*Proof.* Denote by  $X_{(i)}$  a nearest-neighbor of  $X_i$  among  $\{X_1, \ldots, X_{n+1}\} \setminus \{X_i\}$ . Note that  $R_i^2 = \|X_i - X_{(i)}\|_{\infty}^2$  has the same distribution for all  $1 \leq i \leq n+1$ , so that

$$\mathbb{E}(R_{n+1}^2) = \frac{1}{n+1} \sum_{i=1}^{n+1} \mathbb{E}(R_i^2).$$
(1.19)

Now, the sets  $B_i = \{x \in \mathbb{R}^d \mid ||x - X_i||_{\infty} < R_i/2\}$  are disjoint for  $i \neq j$  because  $||X_i - X_j||_{\infty} \ge \max(R_i, R_j)$ . The union of the sets  $B_1, \ldots, B_{n+1}$  has a diameter smaller than  $2 \operatorname{diam}(\mathcal{X})$  and therefore a volume smaller than  $2^d \operatorname{diam}(\mathcal{X})^d$ . As a consequence,

$$\sum_{i=1}^{n+1} R_i^d \leqslant 2^d \operatorname{diam}(\mathcal{X})^d,$$

and, by convexity (since  $d/2 \ge 1$ ),

$$\left(\frac{1}{n+1}\sum_{i=1}^{n+1}R_i^2\right)^{d/2} \leqslant \frac{1}{n+1}\sum_{i=1}^{n+1}R_i^d \leqslant \frac{2^d \operatorname{diam}(\mathcal{X})^d}{n+1}.$$
(1.20)

Finally,

$$\frac{1}{n+1}\sum_{i=1}^{n+1}R_i^2 \leqslant \frac{4\operatorname{diam}(\mathcal{X})^2}{(n+1)^{2/d}}.$$

The desired result follows by combining the latter bound with (1.19).

The expected square distance between a new data point and its K-th nearest neighbor among n independent data points can then be bounded as follows.

**Lemma 1.2.** Suppose that Assumption 1.3 holds true, and that  $d \ge 2$ . Consider random variables  $(X_1, \ldots, X_n, X_{n+1})$  independently and identically distributed according to p. Then,

$$\mathbb{E}\left(\left\|X_{n+1} - X_{i_K(X_{n+1})}\right\|_{\infty}^2\right) \leqslant 8 \operatorname{diam}(\mathcal{X})^2 \left(\frac{2K}{n}\right)^{2/d}$$

*Proof.* Assume without loss of generality that  $2K \leq n$  (otherwise the bound is trivial). We randomly split the set  $X_1, \ldots, X_n$  into 2K sets of sizes approximately n/(2K) (some sets have sizes  $\lfloor n/(2K) \rfloor$ , others have sizes  $\lceil n/(2K) \rceil$ ), and denote by  $X_{(K)}^j$  a 1-nearest neighbor of  $X_{n+1}$  in the *j*-th set. Then,  $\|X_{n+1} - X_{i_K(X_{n+1})}\|_{\infty}$  is smaller than the *K*-th term among  $\|X_{n+1} - X_{(K)}^j\|_{\infty}$  for  $1 \leq j \leq 2K$  since the minimum to find the *K*-th nearest neighbor is taken over a smaller set of elements  $X_{(K)}^1, \ldots, X_{(K)}^{2K}$  instead of  $X_1, \ldots, X_n$ . Now, for a nonnegative sequence of elements  $a_1, \ldots, a_{2K}$ , the *k*-th smallest term among them

Now, for a nonnegative sequence of elements  $a_1, \ldots, a_{2K}$ , the k-th smallest term among them is upper bounded by  $(a_1 + \cdots + a_{2K})/(2K - k)$  as there are 2K - k terms larger than the k-th smallest term in the sum  $a_1 + \cdots + a_{2K}$ . Applying the latter inequality for  $a_j = \left\| X_{n+1} - X_{(K)}^j \right\|_{\infty}^2$ and k = K gives

$$\left\|X_{n+1} - X_{i_K(X_{n+1})}\right\|_{\infty}^2 \leq \frac{1}{K} \sum_{j=1}^{2K} \left\|X_{n+1} - X_{(K)}^j\right\|_{\infty}^2$$

We next use Lemma 1.1 to upper bound the expectation of  $\left\|X_{n+1} - X_{(K)}^{j}\right\|_{\infty}^{2}$  by  $4 \operatorname{diam}(\mathcal{X})^{2}(n/2K)^{-2/d}$  (in fact, we use (1.20) and replace n by  $\lfloor n/(2K) \rfloor$  to obtain this bound). This gives

$$\mathbb{E}\left(\left\|X_{n+1} - X_{i_{K}(X_{n+1})}\right\|_{\infty}^{2}\right) \leqslant \frac{1}{K} \sum_{j=1}^{2K} \mathbb{E}\left(\left\|X_{n+1} - X_{(K)}^{j}\right\|_{\infty}^{2}\right) \leqslant \frac{8 \operatorname{diam}(\mathcal{X})^{2}}{(n/2K)^{2/d}},$$

which is indeed the desired bound.

The combination of (1.17), (1.18) and Lemma 1.2 finally leads to the following upper bound for the expected excess risk of prediction for the K-th nearest neighbor method:

$$0 \leq \mathbb{E}_{\mathcal{D}}\left[\mathcal{R}(\widehat{f})\right] - \mathcal{R}^{\star} \leq \frac{\sigma^2}{K} + 8B^2 \operatorname{diam}(\mathcal{X})^2 \left(\frac{2K}{n}\right)^{2/d}.$$
(1.21)

Note that the last term has a bad scaling with respect to the dimension because f is only assumed to be Lipschitz (and not more regular than this). It is possible at this stage to find the value of K which minimizes the right hand side of the previous inequality. This gives some theoretical guidance on the choice of K.

**Exercise 1.10.** Determine the optimal choice of K for the upper bound as a function of n and the various parameters of the model ( $\sigma$ , diam( $\mathcal{X}$ ) and B), and interpret the various scalings. Write also the associated upper bound on the expected excess risk.

**Correction.** Introduce the smooth function  $U: (0, +\infty) \to \mathbb{R}_+$  defined by

$$U(K) = \frac{\sigma^2}{K} + 8B^2 \operatorname{diam}(\mathcal{X})^2 \left(\frac{2K}{n}\right)^{2/d}$$

Note that  $U(K) \to +\infty$  as  $K \to 0$  and  $K \to +\infty$ . Moreover,

$$U'(K) = -\frac{\sigma^2}{K^2} + \frac{16}{d}B^2 \operatorname{diam}(\mathcal{X})^2 \left(\frac{2}{n}\right)^{2/d} K^{2/d-1},$$

so that U has a unique critical point

$$K^{\star} = \left(\frac{d\sigma^2}{2^{4+2/d}B^2\operatorname{diam}(\mathcal{X})^2}\right)^{d/(d+2)} n^{2/(d+2)}.$$
(1.22)

In practice, one needs of course to take the closest integer to the right hand side. In terms of scalings, note that

- $K^*$  decreases when  $f^*$  is less regular, *i.e.* B is larger (in fact, it is clear from Assumption 1.2 that the correct adimensional parameter to discuss the regularity of the function is  $B \operatorname{diam}(\mathcal{X})$ ). This is natural as less neighbors should be considered when the target function varies rapidly;
- $K^*$  needs to be increased when there is more noise, *i.e.*  $\sigma$  is larger. This is natural as averaging over more neighbors allows to reduce the statistical noise. When there is no noise ( $\sigma = 0$ ), it is in fact optimal to consider  $K^* = 1$  (the 1-nearest neighbor method is consistent);
- $K^*$  should be increased when n is increased, but slower than n, so that a data point sees a vanishing fraction of all the points.

By plugging (1.22) into the expression of the upper bound provided by the function U, we find that

$$0 \leq \mathbb{E}_{\mathcal{D}}\left[\mathcal{R}(\widehat{f})\right] - \mathcal{R}^{\star} \leq U(K_{\star}) = \frac{\sigma^2}{K_{\star}} \left(1 + \frac{d}{2}\right)$$
$$= \sigma^{4/(d+2)} \left(1 + \frac{d}{2}\right) \left(\frac{2^{4+2/d}B^2 \operatorname{diam}(\mathcal{X})^2}{d}\right)^{d/(d+2)} n^{-2/(d+2)}.$$

The previous exercise shows that

$$0 \leqslant \mathbb{E}_{\mathcal{D}}\left[\mathcal{R}(\widehat{f})\right] - \mathcal{R}^{\star} \leqslant \frac{C}{n^{2/(d+2)}}.$$

In view of the first equality in (1.17), and when d is large, this can be rewritten as

$$\sqrt{\mathbb{E}_{\mathcal{D}}\left[\left\|\widehat{f} - f^{\star}\right\|_{L^{2}(p_{\text{data}})}^{2}\right]} = O\left(n^{-1/d}\right).$$
(1.23)

The estimate on the right hand side is typical of the so-called "curse of dimensionality". To give an intuition of the scalings encountered in this context, consider the situation when we have npoints in dimension d, regularly spaced on a grid. The distance between neighboring points is of order  $n^{-1/d}$  in each direction (as there are  $n = m^d$  points overall when there are m points in each direction; the points being at a distance 1/m in each direction). When  $f^*(x')$  is approximated by the value  $f^*(x_i)$  at the closest grid point  $x_i$ , the error is of order

$$|f^{\star}(x') - f^{\star}(x_i)| \sim B ||x' - x_i|| \sim n^{-1/d},$$

which is consistent with (1.23).

# Least square regression

<b>2.1</b>	General framework of linear least square regression 26					
<b>2.2</b>	Ordinary least square regression					
	2.2.1	Derivation of the least square estimator	27			
	2.2.2	Interpretations of the least square estimator	28			
	2.2.3	Numerical solution of the ordinary least square problem	30			
	2.2.4	Statistical analysis	30			
<b>2.3</b>	2.3 Ridge regression					
	2.3.1	Capacity control	35			
	2.3.2	Ridge penalization	36			
	2.3.3	Statistical properties of ridge regression	39			
<b>2.4</b>	4 LASSO regression		<b>43</b>			
	2.4.1	Sparsity inducing regularization terms	43			
	2.4.2	Explicit expressions for one dimensional problems	45			
	2.4.3	Expressions for the Lasso least square regression problem	46			
	2.4.4	Numerical methods to find the minimizers	47			
	2.4.5	Theoretical guarantees	47			

We consider in this chapter a simple (the simplest?) model to predict real valued outputs in a regression context, namely methods based on least square regression with linear models. In contrast to other scientific fields, we focus here on the prediction capabilities of the model rather than on its statistical properties.

Various lessons and features of least square models will be encountered for more complex models, and therefore serve as basic examples to build up an intuition of statistical learning. In particular we will see how to derive estimates which allow to

- capture a bias-variance tradeoff;
- make precise how the generalization performance degrades with the dimension when no regularization is considered, and how it can remain stable when regularization is introduced;
- generalize the approach to more complex models with nonlinear features (kernel methods).

Overall, the results obtained in this chapter are simple to derive with rather basic tools of linear algebra, and no optimization algorithm is needed to find predictors since these predictors are given by closed form or analytic expressions in many cases (except for Section 2.4).

We start by presenting the general framework of linear least square regression in Section 2.1, and next introduce the ordinary least square method in Section 2.2. We then discuss two ways to prevent overfitting in least square regression, by adding either a penalization on the Euclidean norm of the parameter (ridge regression; see Section 2.3) or on its  $\ell^1$  norm (the so-called LASSO; see Section 2.4).

# 2.1 General framework of linear least square regression

We provide in this section a general framework for linear least square regression problems, in the case of one dimensional outputs  $\mathcal{Y} = \mathbb{R}$  for simplicity. We follow the presentation in [4, Section 3.2]. The aim of least square regression is to approximate the Bayes predictor  $f^*(x') = \mathbb{E}(y | x = x')$  associated with the loss function<sup>1</sup>

$$\mathcal{R}(f) = \mathbb{E}\left(|y - f(x)|^2\right).$$

Recall that the Bayes predictor is derived in Exercise 1.2. We consider to this end a parametrized family of predictors  $f_{\theta} : \mathcal{X} \to \mathcal{Y}$  for parameters  $\theta \in \Theta$ . We also replace the population loss by an empirical risk. Denoting by  $\mathcal{D} = \{(x_i, y_i)_{1 \leq i \leq n}\}$  the training set, the minimization problem to consider is therefore

$$\inf_{\theta \in \Theta} \widehat{\mathcal{R}}_n(f_\theta), \qquad \widehat{\mathcal{R}}_n(f) = \frac{1}{n} \sum_{i=1}^n \left( y_i - f_\theta(x_i) \right)^2.$$
(2.1)

We denote by  $\hat{\theta} \in \Theta$  a minimizer of the latter problem, assuming that such a minimizer indeed exists.

Until now we have described a rather general least square regression problem, which is relevant for many models including neural networks (see Chapter 8). Linear least square regression corresponds to considering functions  $f_{\theta}$  linear in  $\theta$ ; but possibly nonlinear in x. Generically, these functions are written as

$$f_{\theta}(x) = \theta^{\top} \varphi(x) \in \mathbb{R}, \qquad \theta \in \Theta = \mathbb{R}^d,$$
(2.2)

where  $\varphi : \mathcal{X} \to \mathbb{R}^d$  is a featurization function. Before reformulating the regression problem in a more abstract manner, let us first give three paradigmatic examples:

- simple linear regression corresponds to the case when  $f_{\theta}(x) = \theta_0 + \theta_1 x$  with  $x \in \mathbb{R}$  and  $\theta = (\theta_0, \theta_1) \in \mathbb{R}^2$ . Therefore,  $\varphi(x) = (1, x) \in \mathbb{R}^2$ ;
- multiple linear regression extends the previous example to the situation when  $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$ , by considering  $\varphi(x) = (1, x) \in \mathbb{R}^{d+1}$  and  $\theta \in \mathbb{R}^{d+1}$ , so that

$$f_{\theta}(x) = \theta_0 + \theta_1 x_1 + \dots + \theta_d x_d \in \mathbb{R};$$

• polynomial regression is described, for one-dimensional inputs  $x \in \mathbb{R}$ , by the featurization function  $\varphi(x) = (1, x, x^2, \dots, x^k)$ . Therefore,  $\theta \in \mathbb{R}^{k+1}$  and

$$f_{\theta}(x) = \theta_0 + \theta_1 x_1 + \dots + \theta_k x^k \in \mathbb{R};$$

• kernel regression will be mentioned later on in Section 6.4.

Let us conclude the presentation of the general framework by reformulating (2.1) for the choice (2.2) as (with some abuse of notation on  $\widehat{\mathcal{R}}_n$ )

$$\inf_{\theta \in \mathbb{R}^d} \widehat{\mathcal{R}}_n(\theta), \qquad \widehat{\mathcal{R}}_n(\theta) = \frac{1}{n} \|Y - X\theta\|_2^2, \tag{2.3}$$

where  $\|\cdot\|_2$  is the standard Euclidean norm on  $\mathbb{R}^n$ , and

$$Y = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} \in \mathbb{R}^{n \times 1}, \qquad X = \begin{pmatrix} ---\varphi(x_1) & --- \\ \vdots \\ ---\varphi(x_n) & --- \end{pmatrix} \in \mathbb{R}^{n \times d}.$$
(2.4)

Note that  $X\theta \in \mathbb{R}^n$  has the same dimension as Y.

<sup>&</sup>lt;sup>1</sup> Note that we use here the abuse of notation where the random variables x, y are not denoted by capital letters. We keep the notation X for the matrix gathering the data.

# 2.2 Ordinary least square regression

Ordinary least square regression is the simplest regression model, but it is already quite useful, and its analysis provides a lot of insight. The presentation in this section is based on [4, Sections 3.3 to 3.5] (see also Section 3.8 for more advanced topics not covered in the lectures) and [40, Section 11.2].

# 2.2.1 Derivation of the least square estimator

We make the following key assumption throughout all this section (but not for Sections 2.3 and 2.4).

**Assumption 2.1.** The dimension of the features d is smaller than the number of data points n, and the matrix X defined in (2.4) has full rank d.

For polynomial regression, this assumption is satisfied as soon as all the inputs  $(x_i)_{1 \leq i \leq n} \subset \mathbb{R}$ are distinct since X is a Vandermonde matrix.

**Exercise 2.1.** Prove that the minimization problem (2.3) is well posed (i.e. there exists a unique minimizer), and find the expression of this global minimum.

**Correction.** There are two main ways to prove that the minimization problem is well posed, and the reader should definitively have the various options in mind.

• A first approach is to note that  $\widehat{\mathcal{R}}_n$  is continuous and coercive (goes to infinity at infinity), so that it admits a global minimizer. Coercivity is proved by noting that

$$n\widehat{\mathcal{R}}_n(\theta) = \theta^\top X^\top X \theta - 2Y^\top X \theta + Y^\top Y, \qquad (2.5)$$

with  $X^{\top}X \ge \rho$  in the sense of symmetric matrices. Indeed, the matrix  $X^{\top}X$  is real, positive semidefinite and symmetric, hence diagonalizable, so that it suffices to prove that it is injective to conclude that its smallest eigenvalue  $\rho$  is positive. This follows from the fact that  $X^{\top}X\xi = 0$  for some  $\xi \in \mathbb{R}^d$  implies that  $0 = \xi^{\top}X^{\top}X\xi = ||X\xi||_2^2$ , and therefore  $\xi = 0$ since X has full rank. In view of (2.5),

$$n\widehat{\mathcal{R}}_n(\theta) \ge \rho \|\theta\|_2^2 - 2\|X^\top Y\|_2 \|\theta\|_2 + Y^\top Y$$

goes to infinity as  $\|\theta\|_2 \to +\infty$ .

Uniqueness can be established by looking at the necessary condition satisfied by a global minimizer  $\hat{\theta}$ , namely the Euler–Lagrange equation

$$0 = \nabla \widehat{\mathcal{R}}_n(\widehat{\theta}) = \frac{2}{n} \left( X^\top X \widehat{\theta} - X^\top Y \right) \in \mathbb{R}^d.$$

The latter equation has a unique solution  $\hat{\theta} = (X^{\top}X)^{-1}X^{\top}Y$ .

• A second approach is to realize that  $\widehat{\mathcal{R}}_n$  is smooth, and that its Hessian

$$\nabla^2 \widehat{\mathcal{R}}_n(\theta) = \frac{2}{n} X^\top X \tag{2.6}$$

is positive definite (by the same argument as above). Therefore,  $\widehat{\mathcal{R}}_n$  is a strongly convex function, minimized over  $\mathbb{R}^d$ , and so it admits a unique global minimizer. The characterization of the minimizer follows the same lines as above (but this characterization is not used to prove uniqueness here).

#### 28 2 Least square regression

The above exercise shows that the unique minimizer of (2.3) is

$$\widehat{\theta} = \frac{1}{n} \widehat{\Sigma}^{-1} X^{\top} Y, \qquad (2.7)$$

with

$$\widehat{\Sigma} = \frac{1}{n} X^{\top} X \in \mathbb{R}^{d \times d}$$
(2.8)

the uncentered covariance of the input data. The normalization factor 1/n ensures that all entries of the matrix  $\widehat{\Sigma}$  are of order 1 uniformly in n; and similarly for entries of  $X^{\top}Y/n$ .

#### 2.2.2 Interpretations of the least square estimator

We provide in this section various reformulations of the least square regression problem, which allow to make a connection with results seen in Statistics.

**Simple regression.** For simple regression, for which  $x_i \in \mathbb{R}$ , the expression of  $\hat{\theta}$  can be made even more explicit.

Exercise 2.2. Prove that

$$\widehat{\theta}_0 = \overline{y} - \widehat{\theta}_1 \overline{x}, \qquad \widehat{\theta}_1 = \frac{\operatorname{Cov}\left(\{(x_i)_{1 \le i \le n}\}, \{(y_i)_{1 \le i \le n}\}\right)}{\operatorname{Var}\{(x_i)_{1 \le i \le n}\}}$$

where the empirical means are

$$\overline{x} = \frac{1}{n} \sum_{i=1}^{n} x_i, \qquad \overline{y} = \frac{1}{n} \sum_{i=1}^{n} y_i,$$

while the empirical covariance and variance are defined as

$$\operatorname{Cov}\left(\{(a_k)_{1\leqslant k\leqslant K}\}, \{(b_k)_{1\leqslant k\leqslant K}\}\right) = \frac{1}{K} \sum_{k=1}^{K} a_k b_k - \left(\frac{1}{K} \sum_{k=1}^{K} a_k\right) \left(\frac{1}{K} \sum_{k=1}^{K} b_k\right),$$

and  $\operatorname{Var}\{(a_k)_{1 \leq k \leq K}\} = \operatorname{Cov}(\{(a_k)_{1 \leq k \leq K}\}, \{(a_k)_{1 \leq k \leq K}\}).$ 

Correction. We write the Euler–Lagrange equation associated with the minimization of

$$n\widehat{\mathcal{R}}_n(\theta_0,\theta_1) = \sum_{i=1}^n |y_i - \theta_0 - \theta_1 x_i|^2,$$

namely that the gradient of the former function vanishes at  $\hat{\theta}$ . By computing the partial derivatives with respect to  $\theta_0, \theta_1$  respectively, we obtain

$$\sum_{i=1}^{n} y_i - \widehat{\theta}_0 - \widehat{\theta}_1 x_i = 0, \qquad \sum_{i=1}^{n} \left( y_i - \widehat{\theta}_0 - \widehat{\theta}_1 x_i \right) x_i = 0.$$

The first equation provides the desired expression for  $\hat{\theta}_0$ , while the second one can be rewritten as follows:

$$0 = \sum_{i=1}^{n} \left[ y_i - \overline{y} - \widehat{\theta}_1 \left( x_i - \overline{x} \right) \right] x_i = \sum_{i=1}^{n} \left[ y_i - \overline{y} - \widehat{\theta}_1 \left( x_i - \overline{x} \right) \right] \left( x_i - \overline{x} \right),$$

thus leading to the claimed expression for  $\hat{\theta}_1$ .

0

Maximum likelihood estimation in the Gaussian model. In order to make a connection with Statistics, we consider a situation where it is assumed that the outputs  $y_i$  are related to the inputs  $x_i$  as

$$y_i = \theta^{\top} \varphi(x_i) + \varepsilon_i,$$

where  $\varepsilon_i \sim \mathcal{N}(0, \sigma^2)$  are independent Gaussian random variables (supposed to be one-dimensional for simplicity; but the analysis in this section can be straightforwardly extended to the case when the outputs are multidimensional). In this context, the likelihood of a realization is

$$\prod_{i=1}^{n} \sqrt{\frac{1}{2\pi\sigma^2}} \exp\left(-\frac{\varepsilon_i^2}{2\sigma^2}\right)$$

Now,  $\varepsilon_i^2 = (y_i - \theta^\top \varphi(x_i))^2$ , so that the likelihood is proportional to

$$\exp\left(-\frac{1}{2\sigma^2}\sum_{i=1}^n \left|y_i - \theta^\top \varphi(x_i)\right|^2\right) = \exp\left(-\frac{n}{2\sigma^2}\widehat{\mathcal{R}}_n(\theta)\right).$$

Finding the least square estimator is therefore equivalent in this context to finding the maximum likelihood estimator for the Gaussian model at hand.

**Orthogonal projection.** We recall here an interpretation of the vector of predictions  $\widehat{Y} = X\widehat{\theta}$  in terms of an orthogonal projection of Y.

**Proposition 2.1.** The vector of predictions  $\widehat{Y} = X(X^{\top}X)^{-1}X^{\top}Y$  is the orthogonal projection of  $Y \in \mathbb{R}^n$  onto  $\operatorname{Ran}(X) = \operatorname{Span}(X_1, \ldots, X_d) \subset \mathbb{R}^n$ , the column space of X.

Proof. Introduce the matrix  $P = X(X^{\top}X)^{-1}X^{\top}$ . Note that P is a projector since  $P^2 = P$ . Moreover, P(Xa) = Xa for any  $a \in \mathbb{R}^d$ , so that Pu = u for any  $u \in \operatorname{Ran}(X)$ . Finally,<sup>2</sup> given the expression of P, it is clear that Pu' = 0 for any  $u' \in \operatorname{Ker}(X^{\top}) = \operatorname{Ran}(X)^{\perp}$ . This shows that P is the orthogonal projector onto  $\operatorname{Ran}(X)$ .

**Prediction errors.** We characterize here the quality of the prediction on the training data set, using the coefficient of determination. More precisely, when  $\mathbf{1}_n \in \operatorname{Ran}(X)$  (which is the case for instance when the first component of  $\varphi(x)$  is 1),

$$\widehat{\mathcal{R}}_n(\widehat{\theta}) = \operatorname{Var}\{(y_i)_{1 \leq i \leq n}\}(1 - r^2), \qquad r^2 = \frac{\operatorname{Var}\{(\widehat{y}_i)_{1 \leq i \leq n}\}}{\operatorname{Var}\{(y_i)_{1 \leq i \leq n}\}} = \frac{\left\|\widehat{Y} - \overline{\widehat{y}}\mathbf{1}_n\right\|_2^2}{\|Y - \overline{y}\mathbf{1}_n\|_2^2}.$$

To prove this equality, we note first that Y and  $\widehat{Y}$  have the same means since  $Y - \widehat{Y}$  is orthogonal to  $\mathbf{1}_n \in \operatorname{Ran}(X)$ . We can then decompose  $Y - \overline{y}\mathbf{1}_n$  as the sum of  $Y - \widehat{Y}$  and  $\widehat{Y} - \overline{y}\mathbf{1}_n$  and use Pythogoras' theorem to write

$$\widehat{\mathcal{R}}_n(\widehat{\theta}) = \frac{1}{n} \left\| Y - \widehat{Y} \right\|_2^2 = \frac{1}{n} \left( \left\| Y - \overline{y} \mathbf{1}_n \right\|_2^2 - \left\| \widehat{Y} - \overline{y} \mathbf{1}_n \right\|_2^2 \right) = \operatorname{Var}\{(y_i)_{1 \le i \le n}\} - \operatorname{Var}\{(\widehat{y}_i)_{1 \le i \le n}\}.$$

The prediction error on the training data set is smaller when r is closer to 1. In fact, since the triangle  $\overline{\hat{y}}\mathbf{1}_n$ , Y,  $\widehat{Y}$  is rectangle in  $\widehat{Y}$ , the ratio of lengths defining r also corresponds to the cosine of some angle, and eventually corresponds to some normalized covariance of the data:

$$r = \frac{\left\|\widehat{Y} - \overline{\widehat{y}}\mathbf{1}_{n}\right\|_{2}}{\|Y - \overline{y}\mathbf{1}_{n}\|_{2}} = \frac{\left\langle\widehat{Y} - \overline{\widehat{y}}\mathbf{1}_{n}, Y - \overline{y}\mathbf{1}_{n}\right\rangle}{\left\|\widehat{Y} - \overline{\widehat{y}}\mathbf{1}_{n}\right\|_{2}\|Y - \overline{y}\mathbf{1}_{n}\|_{2}} = \frac{\operatorname{Cov}\left(\widehat{Y}, Y\right)}{\sqrt{\operatorname{Var}(Y)\operatorname{Var}(\widehat{Y})}},$$

with some abuse of notation for the arguments of Var and Cov.

<sup>&</sup>lt;sup>2</sup> Alternatively, one could note that  $(Y - \widehat{Y})^{\top} X a = 0$  for any  $a \in \mathbb{R}^d$ , which implies that  $(\mathrm{Id} - P)Y$  is orthogonal to  $\mathrm{Ran}(X)$ .

#### 2.2.3 Numerical solution of the ordinary least square problem

We discuss in this section how to solve (2.3) in practice. A first remark is that a good practice is to standardize the data (recall Section 1.1.4) by replacing  $x_{n,j}$  with

$$x'_{n,j} = \frac{x_{n,j} - \overline{x}_j}{\overline{\sigma}_j}, \qquad \overline{\sigma}_j = \sqrt{\operatorname{Var}\{(x_{i,j})_{1 \leq i \leq n}\}}.$$

This allows to replace X with a matrix X' where each column has mean 0 and empirical variance 1. There are then two main approaches to solve (2.3):

- rely on gradient methods to perform a direct minimization of the loss function; see Chapter 4. This may be the most efficient approach when d is large, although usually plain gradient methods are not used to find solutions, but rather conjugate gradient methods or quasi-Newton methods, typical choices considered for computational linear algebra (see [40, Section 11.2.2.3] and references therein).
- solve the linear system

$$X^{\top} X \widehat{\theta} = X^{\top} Y \tag{2.9}$$

in view of the expression of the global minimizer  $\hat{\theta}$  given by (2.7).

We discuss here the second option. The first remark is that it is not a good idea to invert the matrix  $X^{\top}X$  as this may be numerically unstable and has a large computational cost  $O(d^3)$ . Solving the linear system (2.9) is computationally much cheaper and also more stable. There are two main options to do so (see [40, Section 11.2.2.3]):

• use a singular value decomposition to write  $X = USV^{\top}$  where  $U \in \mathbb{R}^{n \times n}$  with  $U^{\top}U = \mathrm{Id}_n$ ,  $S \in \mathbb{R}^{n \times d}$  contains the  $r = \min(n, d)$  singular values  $\lambda_1, \ldots, \lambda_d$  on its main diagonal and entries 0 otherwise, and  $V \in \mathbb{R}^{d \times d}$  with  $VV^{\top} = V^{\top}V = \mathrm{Id}_d$ . The cost of performing this SVD is  $O(nd \times \min(n, d))$ .

In order to solve  $X^{\top}X\hat{\theta} = X^{\top}Y$ , we note that  $X^{\top}X = V(S^{\top}S)V^{\top}$ , where  $S^{\top}S \in \mathbb{R}^{d \times d}$  is a diagonal matrix with entries  $\lambda_1^2, \ldots, \lambda_d^2$ . The solution  $\hat{\theta}$  is found by first performing the SVD of X, then computing  $\hat{\theta} = V(S^{\top}S)^{-1}V^{\top}X^{\top}Y$ .

See also Remark 2.1 below for a discussion on how to compute predictions when  $n \leq d$  (which is not the situation we consider here).

• a second option is to perform a QR decomposition, where one writes X = QR with  $Q = [q_1| \dots |q_d] \in \mathbb{R}^{n \times d}$  and  $q_i \cdot q_j = \delta_{ij}$ , and  $R \in \mathbb{R}^{d \times d}$  an upper triangular matrix. This option is interesting when  $n \gg d$ . The QR decomposition is obtained by a pivoting strategy where one writes  $X_1 = r_{11}q_1$  (which determines  $q_1$ ) then  $X_2 = r_{12}q_1 + r_{22}q_2$  (which determines  $q_2$ ), etc. The system to solve (2.9) can be rewritten as  $R^{\top}Q^{\top}QR\hat{\theta} = R^{\top}R\hat{\theta} = X^{\top}Y = R^{\top}Q^{\top}Y$ , so that it suffices to solve the simple triangular system  $R\hat{\theta} = Q^{\top}Y$  (which has a cost  $O(d^2)$  once the QR decomposition is obtained; the cost of the QR decomposition being  $O(d^3)$ ).

In scikit-learn, the method sklean.linear\_model.fit calls the routine scipy.linalg.lstsq, which itself calls by default the LAPACK routine DGELSD, which relies on SVD.

**Remark 2.1 (Computing predictions for**  $n \leq d$ ). When  $n \leq d$  (with X of full rank n in this case), one should rewrite the problem to solve as  $XX^{\top}X\hat{\theta} = XX^{\top}Y$  and use that  $XX^{\top} = U(SS^{\top})U^{\top} \in \mathbb{R}^{n \times n}$ , where  $SS^{\top} \in \mathbb{R}^{n \times n}$  is a diagonal matrix with entries  $\lambda_1^2, \ldots, \lambda_n^2$ . This is sufficient to compute the prediction  $\widehat{Y} = X\widehat{\theta}$ .

#### 2.2.4 Statistical analysis

We discuss in this section how to prove guarantees on the performance of predictions based on ordinary least squares. Two settings can be considered:
- random design: both the inputs  $x_i$  and the outputs  $y_i$  are random. This is the classical setting of machine learning where one wants to generalize predictions to unseen data;
- fixed design: the input data  $(x_1, \ldots, x_n)$  is fixed (and therefore not random), and we consider the prediction error only at these points. In this case, the associated population risk is

$$\mathcal{R}(\theta) = \mathbb{E}_Y \left[ \frac{1}{n} \| Y - X \theta \|_2^2 \right], \qquad (2.10)$$

where the expectation is over the realizations  $y_i$  for each input  $x_i$ . This can be seen as the so-called "in-sample prediction error", and corresponds to learning the vector  $X\theta_{\star}$  of best predictions instead of a function  $\mathcal{X} \to \mathbb{R}$ . In this setting, no generalization to unseen inputs is attempted.

We consider here the fixed design setting. The main interest of this setting is that it is mathematically simpler to analyze. We will only hint at generalizations to the random design setting (see [4, Section 3.8 for further elements).

In view of the no free lunch theorem (see Chapter 10), we need assumptions on how the data is generated in order to provide guarantees. Throughout this section, we make the following assumption, which corresponds to the situation of a so-called well specified model:

$$\exists \theta_{\star} \in \mathbb{R}^d \text{ such that } y_i = \theta_{\star}^{\top} \varphi(x_i) + \varepsilon_i \text{ with } \varepsilon_i \text{ iid, } \mathbb{E}(\varepsilon_i) = 0, \ \mathbb{E}(\varepsilon_i^2) = \sigma^2.$$
(2.11)

Some conditions on the noise could weakened; in particular, it would be sufficient to assume that the  $\varepsilon_i$  are independent but not necessarily identically distributed, with  $\mathbb{E}(\varepsilon_i^2) \leq \sigma^2$ .

We show below that the minimal value of the excess risk for this model is 0 as the Bayes predictor belongs to the class of functions upon which the loss is minimized.

**Exercise 2.3.** Give the expression of the Bayes predictor.

**Correction.** Recal that the Bayes predictor is  $f^*(x) = \mathbb{E}(y | x)$  (see Exercise 1.6). In view of (2.11), it holds  $f^*(x_i) = \theta^{\top}_* \varphi(x_i)$ .

The fact that the model is well specified allows to obtain sharper bounds of order O(1/n)instead of bounds  $O(n^{-1/2})$  obtained for the general analysis of models in supervised learning (see Chapter 10.1). To write such bounds, we consider the risk averaged over the realizations of the data set, *i.e.* in expectation with respect to realizations of the outputs for fixed inputs, somewhat similarly to the situation considered for KNNs where the risk was averaged over realizations of the data set. This leads to the notion of "average excess risk", where "average" refers to realizations of the training data set, while the notion of excess risk builds upon the expected risk, which refers to realizations of test data.

The first result is a decomposition of the risk into a squared bias and a variance, as made precise in the following exercise.

Exercise 2.4 (Risk decomposition). Assume that (2.11) holds.

- (a) Prove that  $\mathcal{R}(\theta) = \sigma^2 + \frac{1}{n} (\theta \theta_\star)^\top X^\top X (\theta \theta_\star).$ (b) Deduce that  $\theta_\star$  is the unique global minimum of  $\mathcal{R}$  and compute the value of the minimal loss  $\mathcal{R}^{\star}$ .
- (c) The results of the previous questions imply that

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}\right)\right] - \mathcal{R}^{\star} = \mathbb{E}_{Y}\left[\left\|\widehat{\theta} - \theta_{\star}\right\|_{\widehat{\Sigma}}^{2}\right], \qquad (2.12)$$

where  $\|a\|_{\widehat{\Sigma}} = \sqrt{a^{\top}\widehat{\Sigma}a}$  is the Mahalanobis distance (recall the definition (2.8) of  $\widehat{\Sigma}$ ). Prove that the following bias-variance decomposition holds for the average excess risk:

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}\right)\right] - \mathcal{R}^{\star} = \left\|\mathbb{E}_{Y}\left[\widehat{\theta}\right] - \theta_{\star}\right\|_{\widehat{\Sigma}}^{2} + \mathbb{E}_{Y}\left[\left\|\widehat{\theta} - \mathbb{E}_{Y}\left(\widehat{\theta}\right)\right\|_{\widehat{\Sigma}}^{2}\right].$$
(2.13)

# Correction.

(a) In view of (2.10), and denoting by  $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_n)$ ,

$$\begin{aligned} \mathcal{R}(\theta) &= \mathbb{E}_{\varepsilon} \left[ \frac{1}{n} \left\| X(\theta_{\star} - \theta) + \varepsilon \right\|^{2} \right] \\ &= \frac{1}{n} \left\| X(\theta_{\star} - \theta) \right\|^{2} + \frac{1}{n} \mathbb{E}_{\varepsilon} \left( \|\varepsilon\|^{2} \right) \\ &= \frac{1}{n} (\theta_{\star} - \theta)^{\top} X^{\top} X(\theta_{\star} - \theta) + \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}(\varepsilon_{i}^{2}) \end{aligned}$$

which gives the claimed result.

(b) Recall that  $\widehat{\Sigma}$  is positive definite (see Exercise 2.1), so that, denoting by  $\rho > 0$  the smallest eigenvalue of this matrix,

$$\mathcal{R}(\theta) \ge \sigma^2 + \rho \left\| \theta - \theta_\star \right\|^2$$

This inequality shows that  $\theta_{\star}$  is the unique global minimizer of  $\mathcal{R}$ . Moreover,  $\mathcal{R}^{\star} = \mathcal{R}(\theta_{\star}) =$  $\sigma^2$ .

(c) As for classical bias/variance decompositions, the desired equality follows from (2.12) by introducing  $\mathbb{E}_{Y}(\hat{\theta})$  in the term on the right hand side of the previous inequality and expanding the square, making use of the fact that the cross term has a vanishing expectation:

$$\mathbb{E}_{Y}\left(\left[\widehat{\theta} - \mathbb{E}_{Y}\left(\widehat{\theta}\right)\right]^{\top}\widehat{\Sigma}\left[\theta_{\star} - \mathbb{E}_{Y}\left(\widehat{\theta}\right)\right]\right) = 0;$$

while the term  $\theta_{\star} - \mathbb{E}_{Y}(\widehat{\theta})$  is deterministic.

The next exercise summarizes some properties of  $\hat{\theta}$ .

**Exercise 2.5 (Properties of**  $\hat{\theta}$ ). Assume that (2.11) holds.

- (a) Prove that  $\hat{\theta}$  is an unbiased estimator of  $\theta_{\star}$ , i.e.  $\mathbb{E}_{Y}\left[\hat{\theta}\right] = \theta_{\star}$  (where we recall that the expectation is taken only over the realizations of the ouputs).
- (b) Show that the covariance of  $\hat{\theta}$  is

$$\operatorname{Cov}_{Y}\left(\widehat{\theta}\right) = \mathbb{E}_{Y}\left[\left(\widehat{\theta} - \theta_{\star}\right)\left(\widehat{\theta} - \theta_{\star}\right)^{\top}\right] = \frac{\sigma^{2}}{n}\widehat{\Sigma}^{-1} \in \mathbb{R}^{d \times d}.$$

# Correction.

(a) In view of (2.7), and using that  $\mathbb{E}_Y(Y) = \mathbb{E}_{\varepsilon}(X\theta_{\star} + \varepsilon) = X\theta_{\star}$  by (2.11),

$$\mathbb{E}\left[\widehat{\theta}\right] = \left(X^{\top}X\right)^{-1}X^{\top}\mathbb{E}(Y) = \left(X^{\top}X\right)^{-1}X^{\top}X\theta_{\star} = \theta$$

(b) Since  $\hat{\theta} - \mathbb{E}_Y(\hat{\theta}) = \hat{\theta} - \theta_\star = (X^\top X)^{-1} X^\top \varepsilon$  by (2.7), and using  $\mathbb{E}(\varepsilon \varepsilon^\top) = \sigma^2 \mathrm{Id}_n.$ 

it follows that

$$\operatorname{Cov}\left(\widehat{\theta}\right) = \mathbb{E}\left[\left(X^{\top}X\right)^{-1}X^{\top}\varepsilon\varepsilon^{\top}X\left(X^{\top}X\right)^{-1}\right]$$
$$= \left(X^{\top}X\right)^{-1}X^{\top}\mathbb{E}\left(\varepsilon\varepsilon^{\top}\right)X\left(X^{\top}X\right)^{-1}$$
$$= \sigma^{2}\left(X^{\top}X\right)^{-1}X^{\top}X\left(X^{\top}X\right)^{-1}$$
$$= \sigma^{2}\left(X^{\top}X\right)^{-1},$$

which gives the claimed result in view of the definition (2.8) of  $\widehat{\Sigma}$ .

We can finally use the information on the estimator  $\hat{\theta}$  to obtain guarantees on the average risk.

**Proposition 2.2 (Risk of ordinary least square estimator).** Assume that (2.11) holds. Then, the average excess risk of  $\hat{\theta}$  is

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}\right)\right] - \mathcal{R}^{\star} = \frac{\sigma^{2}d}{n}.$$

The interpretation of this result is the following:

- (i) the factor d in the numerator shows that ordinary least squares does not work well in high dimension. Some regularization is needed in order to obtain results which have a better scaling with respect to the dimension (and possibly additional assumptions on the structure of the data distribution);
- (ii) the result can be extended to the random design setting by replacing the bound of Proposition 2.2 with

$$\frac{\sigma^2}{n} \mathbb{E}\left[\operatorname{Tr}\left(\Sigma\widehat{\Sigma}^{-1}\right)\right], \qquad \Sigma = \mathbb{E}\left[\varphi(x)\varphi(x)^{\top}\right], \qquad (2.14)$$

see [4, Proposition 3.10] and Exercise 2.6 below. Obtaining estimates on the trace in the expectation however requires some non trivial results on random matrices (in particular concentration inequalities for matrices for good lower bounds on  $\widehat{\Sigma}$ ).

*Proof.* From (2.13) and Exercise 2.5(a), and since covariance matrices are symmetric,

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}\right)\right] - \mathcal{R}^{\star} = \mathbb{E}_{Y}\left[\left\|\widehat{\theta} - \theta_{\star}\right\|_{\widehat{\Sigma}}^{2}\right] = \sum_{k,k'=1}^{d} \widehat{\Sigma}_{k,k'} \mathbb{E}_{Y}\left(\left[\widehat{\theta} - \theta_{\star}\right]_{k}\left[\widehat{\theta} - \theta_{\star}\right]_{k'}\right)\right]$$
$$= \sum_{k,k'=1}^{d} \widehat{\Sigma}_{k,k'} \left[\operatorname{Cov}_{Y}(\widehat{\theta})\right]_{k,k'} = \operatorname{Tr}\left(\operatorname{Cov}_{Y}(\widehat{\theta})\widehat{\Sigma}\right).$$

We next use Exercise 2.5(b) to obtain

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}\right)\right] - \mathcal{R}^{\star} = \frac{\sigma^{2}}{n} \operatorname{Tr}\left(\operatorname{Id}_{d}\right)$$

which gives the desired equality.

Exercise 2.6 (Average excess risk in the random design setting). We consider in this exercise the random design setting, i.e. a situation in which the inputs  $x_1, \ldots, x_n$  in the data set are independently sampled according to some probability measure  $p_{\text{data}}(dx)$ , with associated labels  $y_i = \varphi(x_i)^\top \theta_\star + \varepsilon_i$ , where the random variables  $\varepsilon_i$  are independent from each other and from the inputs  $x_1, \ldots, x_n$ , with  $\mathbb{E}(\varepsilon_i) = 0$  and  $\mathbb{E}(\varepsilon_i^2) = \sigma^2$ .

(a) Prove that  $\mathcal{R}(\theta) - \mathcal{R}^{\star} = \|\theta - \theta_{\star}\|_{\Sigma}^2$  with  $\mathcal{R}^{\star} = \sigma^2$  and  $\Sigma$  defined in (2.14).

(b) Assume that  $\widehat{\Sigma}$  is invertible. Show that

$$\mathbb{E}\left[\mathcal{R}(\widehat{\theta})\right] - \mathcal{R}^{\star} = \frac{\sigma^2}{n} \mathbb{E}\left[\operatorname{Tr}\left(\varSigma\widehat{\Sigma}^{-1}\right)\right],$$

where the expectation is over the realizations of the data set.

# Correction.

(a) We write, the expectation being over realizations of  $x' \sim p_{\text{data}}$  and  $\varepsilon'$  in  $y' = \varphi(x')^{\top} \theta_{\star} + \varepsilon'$ ,

$$\mathcal{R}(\theta) = \mathbb{E}\left[\left(y' - \varphi(x')^{\top}\theta\right)^{2}\right] = \mathbb{E}\left[\left(\varphi(x')^{\top}(\theta_{\star} - \theta) + \varepsilon'\right)^{2}\right]$$
$$= (\theta - \theta_{\star})^{\top}\mathbb{E}\left[\varphi(x')\varphi(x')^{\top}\right](\theta - \theta_{\star}) + \mathbb{E}\left[(\varepsilon')^{2}\right]$$
$$= (\theta - \theta_{\star})^{\top}\Sigma(\theta - \theta_{\star}) + \sigma^{2},$$

from which the result immediately follows.

(b) In view of (2.7) and the equality  $Y = X\theta_{\star} + \varepsilon$ ,

$$\widehat{\theta} = \frac{1}{n}\widehat{\Sigma}^{-1}X^{\top}(X\theta_{\star} + \varepsilon) = \theta_{\star} + \frac{1}{n}\widehat{\Sigma}^{-1}X^{\top}\varepsilon.$$

Therefore, with an expectation both over the realizations of  $x_1, \ldots, x_n$  and  $\varepsilon_1, \ldots, \varepsilon_n$ , and using the identity (2.19) to introduce a trace,

$$\begin{split} \mathbb{E}\left[\mathcal{R}(\widehat{\theta})\right] - \mathcal{R}^{\star} &= \mathbb{E}\left[(\widehat{\theta} - \theta_{\star})^{\top} \varSigma(\widehat{\theta} - \theta_{\star})\right] = \frac{1}{n^{2}} \mathbb{E}\left[\left(\widehat{\varSigma}^{-1} X^{\top} \varepsilon\right)^{\top} \varSigma\widehat{\varSigma}^{-1} X^{\top} \varepsilon\right] \\ &= \frac{1}{n^{2}} \mathbb{E}\left[\varepsilon^{\top} X \widehat{\varSigma}^{-1} \varSigma\widehat{\varSigma}^{-1} X^{\top} \varepsilon\right] = \frac{1}{n^{2}} \mathbb{E}\left[\operatorname{Tr}\left(X^{\top} \varepsilon \varepsilon^{\top} X \widehat{\varSigma}^{-1} \varSigma\widehat{\varSigma}^{-1}\right)\right] \\ &= \frac{\sigma^{2}}{n^{2}} \mathbb{E}\left[\operatorname{Tr}\left(X^{\top} X \widehat{\varSigma}^{-1} \varSigma\widehat{\varSigma}^{-1}\right)\right], \end{split}$$

where we took an inner expectation with respect to  $\varepsilon$ . The result finally follows from the definition (2.8) of  $\widehat{\Sigma}$ .

Proposition 2.2 gives a result on the average excess risk, which is necessarily nonnegative. It is instructive to compare the estimate to the one obtained for the training risk, as provided by the following exercise.

**Exercise 2.7.** Assume that (2.11) holds with  $\varepsilon \sim \sigma \mathcal{N}(0, \mathrm{Id}_n)$ . Prove that the average excess training risk is

$$\mathbb{E}_{Y}\left[\widehat{\mathcal{R}}_{n}\left(\widehat{\theta}\right)\right] - \mathcal{R}^{\star} = -\frac{\sigma^{2}d}{n}.$$

The key comment is that the average excess training risk is negative, so that the average training risk is smaller than the Bayes risk; in fact, there is a sign change on the term of the right hand side compared to the equation in Proposition 2.2. The difference between the risk (as provided by Proposition 2.2) and the training risk gives an indication of the degree of overfitting. The larger the difference between these two quantities is, the larger the average excess risk is. This is akin to some form of generalization gap, but not strictly as we work here in the fixed design setting. The estimates provided here still allow to perform some model selection and detect overfitting.

**Correction.** In view of the expression of  $\widehat{\mathcal{R}}_n$  in (2.3),

$$\begin{aligned} \widehat{\mathcal{R}}_{n}\left(\widehat{\theta}\right) &= \frac{1}{n} \left\| Y - X\widehat{\theta} \right\|^{2} = \frac{1}{n} \left\| X\theta_{\star} + \varepsilon - X \left( X^{\top}X \right)^{-1} X^{\top}Y \right\|^{2} \\ &= \frac{1}{n} \left\| X\theta_{\star} + \varepsilon - X \left( X^{\top}X \right)^{-1} X^{\top} \left( X\theta_{\star} + \varepsilon \right) \right\|^{2} \\ &= \frac{1}{n} \left\| \left( \mathrm{Id}_{n} - X \left( X^{\top}X \right)^{-1} X^{\top} \right) \varepsilon \right\|^{2}. \end{aligned}$$

Now,  $P^{\perp} = \operatorname{Id}_n - X (X^{\top}X)^{-1} X^{\top}$  is the orthogonal projection onto  $\operatorname{Ran}(X)^{\perp}$  (see the proof of Proposition 2.1). Since  $\operatorname{Ran}(X)^{\perp}$  has dimension n - d, Cochran's theorem implies that

$$\widehat{\mathcal{R}}_n\left(\widehat{\theta}\right) = \frac{n-d}{n}\sigma^2,$$

which leads to the desired equality.

**Remark 2.2.** Note that we actually do not need the Gaussianity assumption on  $\varepsilon$  since

$$\mathbb{E}\left[\left\|P^{\perp}\varepsilon\right\|^{2}\right] = \operatorname{Tr}\left(P^{\perp}\mathbb{E}\left[\varepsilon\varepsilon^{\top}\right]P^{\perp}\right) = \sigma^{2}\operatorname{Tr}\left(P^{\perp}\right) = \sigma^{2}(n-d).$$

Exercise 2.7 allows to understand the difference between the training and the test losses. In fact, the average test loss can be reformulated as follows, by clearly distinguishing between the test set (for which labels Y' are obtained from random variables  $\varepsilon'$  independent from the noise  $\varepsilon$  used to generate the training data set Y)

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}\right)\right] = \frac{1}{n}\mathbb{E}_{\varepsilon}\left(\mathbb{E}_{\varepsilon'}\left[\left\|X(\widehat{\theta}-\theta_{\star})-\varepsilon'\right\|_{2}^{2}\right]\right) = \frac{1}{n}\mathbb{E}_{\varepsilon}\left(\mathbb{E}_{\varepsilon'}\left[\left\|X(X^{\top}X)^{-1}X^{\top}\varepsilon-\varepsilon'\right\|_{2}^{2}\right]\right),$$

where the inner expectation corresponds to  $\mathcal{R}$ , while the outer expectation corresponds to the averaging over the training sets. In contrast to this, there is a single expectation in the difference  $\mathbb{E}_Y\left[\widehat{\mathcal{R}}_n\left(\widehat{\theta}\right)\right] - \mathcal{R}^*$  between the average training loss and the Bayes risk, over  $\varepsilon$ (which amounts to setting  $\varepsilon' = \varepsilon$ ).

# 2.3 Ridge regression

This section presents a very important idea in machine learning, namely that adding a regularization (penalty) term to the empirical risk to be minimized can in fact improve the predictive performance of least square estimators. This may seem counterintuitive at first sight, but crucially relies on the fact that a small training loss may not lead to a small test loss, due to overfitting. The main motivation for regularization is therefore to avoid overfitting by reducing the effective size of the parameter space. For ridge regression, the extra penalty term is proportional to the squared Euclidean norm of the parameter. Our presentation is based on [4, Sections 3.6] and [40, Sections 4.5, 5.4 and 11.3].

Let us start by giving an example to motivate the need for regularization. Say that we aim at predicting whether the outcome of a toin cossing is heads or tails; and that out of 3 draws, we get 3 heads. Are we going to predict that all subsequent draws will give heads? Most likely not.

This example illustrates the core issue of the problem: when there are too many parameters in the model to exactly fit the data (including the inherent noise involved in the realizations of the random variables at hand), then one actually fits the empirical distribution of the data, and not the target test distribution. One should therefore prevent a perfect minimization of the training loss, which would put all the mass on the training data, in order to leave some space to cover the (unseen) test data. We next formalize this idea with the notion of capacity control, and then turn to a specific way to obtain such a capacity control, namely ridge penalization.

### 2.3.1 Capacity control

Capacity control consists in limiting the complexity, or expressivity, of the model by adding a penalization term to the empirical risk function. More precisely, the objective function to minimize becomes, for a parametrized class of functions  $\{f_{\theta}, \theta \in \Theta\}$ ,

$$\widehat{\mathcal{R}}_n(f_\theta) + \lambda \Omega(\theta), \tag{2.15}$$

for some penalization function  $\Omega: \Theta \to \mathbb{R}_+$ , and a penalization strength  $\lambda \ge 0$ . Paradigmatic examples are

- ridge penalization:  $\Omega(\theta) = \|\theta\|_2^2$ . This is also called "weight decay", for reasons that will be made precise below (see Exercise 2.9);
- LASSO (the acronym is motivated and explained in Section 2.4):  $\Omega(\theta) = \|\theta\|_1$ ;
- elastic net, which is a combination of the previous two penalizations, and therefore corresponds to adding a penalty term of the form  $\lambda_1 \|\theta\|_1 + \lambda_2 \|\theta\|_2^2$  for  $\lambda_1, \lambda_2 \ge 0$ .

Finding the penalization strengh by (cross) validation. The most appropriate value of the penalization parameter  $\lambda$  is found by cross-validation (see for instance the presentation in [40, Section 5.4.3]). We describe the procedure in the context of a large data set, as in Section 1.3.2.1, the adaptation to smaller data sets with K-fold cross validation being straightforward. More precisely, consider a training set  $\mathcal{D}_{\text{train}}$  and a validation set  $\mathcal{D}_{\text{val}}$ . One first trains the penalized empirical risk (2.15) on the training data  $\mathcal{D}_{\text{train}}$  for various values of  $\lambda$ :

$$\widehat{\theta}_{\lambda} \in \operatorname*{argmin}_{\theta \in \Theta} \left\{ \widehat{\mathcal{R}}_n(f_{\theta}) + \lambda \Omega(\theta) \right\}.$$

This defines a family of minimizers  $f_{\hat{\theta}_{\lambda}}$  indexed by  $\lambda$ . The best value of  $\lambda$  is then found as the one minimizing the *unpenalized* loss on the validation data  $\mathcal{D}_{val} = \{(x_m, y_m)_{1 \leq m \leq N_{val}}\}$ :

$$\lambda^{\star} \in \operatorname*{argmin}_{\lambda \ge 0} \left\{ \frac{1}{N_{\mathrm{val}}} \sum_{m=1}^{N_{\mathrm{val}}} \ell\left(f_{\widehat{\theta}_{\lambda}}(x_m), y_m\right) \right\}.$$

Let us emphasize here that one should not add a penalization term to the validation loss, as this loss is an approximation of the test loss, which is ultimately the quantity one wants to minimize. The whole point of the (cross) validation procedure is to determine the strength of the penalization to add to the training loss in order to prevent overfitting and obtain better predictions.

Alternative formulation of capacity control. Regularization is often introduced by adding a penalty term as in (2.15). The latter loss function can however be seen as the minimization of the Lagrangian associated with the family of constrained optimization problems (indexed by  $D \ge 0$ ):

$$\inf_{\theta \in \Theta} \left\{ \widehat{\mathcal{R}}_n(f_\theta) \mid \Omega(\theta) \leqslant D \right\}.$$
(2.16)

Note that the unconstrained minimization of  $\widehat{\mathcal{R}}_n$  is obtained in the limit  $D \to +\infty$ ; whereas capacity control becomes stronger as  $D \to 0$ . The parameter D can be seen as the dual parameter of the penalization strength  $\lambda \ge 0$ .

We will most often consider the minimization of the penalized loss (2.15), but sometimes it is useful to consider the dual perspective (2.16) (see for instance the discussion at the end of Section 2.4.1).

# 2.3.2 Ridge penalization

We discuss in this section least square regression with a ridge penalization, which corresponds to the so-called ridge regression.<sup>3</sup> One crucial interest of ridge regression is that one can cleary study the impact and benefit of regularization, as the minimizer of (2.15) has an analytical expression similar to the one obtained for ordinary least squares in (2.7).

In order to motivate once again the interest of regularization in the context of least square regression, consider the situation when  $d/n \to 1$ , or simply n = d. In this case, (2.7) simplifies as  $\hat{\theta} = X^{-1}Y$  (since  $(X^{\top}X)^{-1}X^{\top}X = \mathrm{Id}_n$  and X is invertible, so that  $X^{-1} = (X^{\top}X)^{-1}X^{\top}$ ). The prediction  $X\hat{\theta} = Y$  then gives a perfect fit to the training data. This is not a great news in terms of generalizing to unseen data points as all degrees of freedom in  $\hat{\theta}$  are used to match training data points, so that there is absolutely no flexibility to adjust for test data points. The situation is even worse for d > n, as in this case the solutions to the minimization problem (2.3) are not unique since one can add to a minimizer an arbitrary element of the kernel of X. Some of the solutions to the minimization problem can therefore have a very large norm.

 $<sup>^3</sup>$  We do not discuss here the origin of the term 'ridge', which has a long history dating back to the end of the 50s.

other motivation (to add):

update in

SGD

Analytical expression of the solution to the ridge regression problem. In this section and the end of this chapter, we no longer suppose that Assumption 2.1 holds, namely that X has full rank d. This allows in particular to consider situations where  $d \ge n + 1$ .

Definition 2.1 (Ridge least square regression). For  $\lambda > 0$ ,

$$\widehat{\theta}_{\lambda} = \operatorname*{argmin}_{\theta \in \mathbb{R}^d} \left\{ \frac{1}{n} \|Y - X\theta\|_2^2 + \lambda \|\theta\|_2^2 \right\}.$$

**Exercise 2.8.** Prove that the minimization problem in Definition 2.1 is well posed for  $\lambda > 0$  (recall that we no longer assume that X has full rank), and that the minimizer can be explicitly written as

$$\widehat{\theta}_{\lambda} = \frac{1}{n} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-1} X^{\top} Y, \qquad (2.17)$$

where  $\widehat{\Sigma}$  is still defined by (2.8).

**Correction.** Note that

$$\frac{1}{n} \|Y - X\theta\|_2^2 + \lambda \|\theta\|_2^2 \ge \lambda \|\theta\|_2^2,$$

so that the function to minimize is smooth and coercive. It therefore admits a global minimizer. Moreover, a simple computation shows that the Hessian of the function is  $2(\widehat{\Sigma} + \lambda \mathrm{Id}_d)$ , which is lower bounded (in the sense of symmetric matrices) by  $\lambda \mathrm{Id}_d$ . This implies that the function to minimize is stricly convex (in fact strongly convex), so that the minimum is unique.

The unique global minimizer is characterized by the Euler-Lagrange equation

$$\nabla_{\theta} \left( \frac{1}{n} \| Y - X \theta \|_{2}^{2} + \lambda \| \theta \|_{2}^{2} \right) \Big|_{\theta = \widehat{\theta}_{\lambda}} = \frac{2}{n} X^{\top} \left( X \widehat{\theta}_{\lambda} - Y \right) + 2\lambda \widehat{\theta}_{\lambda} = 0 \in \mathbb{R}^{d}.$$

Therefore,

$$\left(\widehat{\Sigma} + \lambda \mathrm{Id}_d\right)\widehat{\theta}_{\lambda} = \frac{1}{n}X^{\top}Y.$$
 (2.18)

Let us now remark that  $\widehat{\Sigma}$  is positive semidefinite, so that  $\widehat{\Sigma} + \lambda \operatorname{Id}_d$  is invertible. Indeed, consider  $\xi \in \mathbb{R}^d$  such that  $(\widehat{\Sigma} + \lambda \operatorname{Id}_d)\xi = 0$ . Then, by taking the scalar product with  $\xi$ ,

$$0 = \xi^{\top} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right) \xi = \frac{1}{n} \| X \xi \|_2^2 + \lambda \| \xi \|_2^2 \ge \lambda \| \xi \|_2^2,$$

which allows to conclude that  $\xi = 0$ . The conclusion is finally obtained by applying the inverse of  $\widehat{\Sigma} + \lambda \operatorname{Id}_d$  on both sides of (2.18).

**Exercise 2.9 (Motivating the terminology "weight decay").** Consider a one dimensional featurization function  $\phi : \mathcal{X} \to \mathbb{R}$ , and the associated ridge regression problem

$$\widehat{\mathcal{R}}_{n,\lambda}(\theta) = \frac{1}{n} \sum_{i=1}^{n} \left( y_i - \theta \phi(x_i) \right)^2 + \lambda \theta^2$$

for  $\theta \in \Theta = \mathbb{R}$  and  $\lambda > 0$ . Give the expression of the optimal parameter  $\widehat{\theta}_{\lambda}$  of Definition 2.1, and motivate the terminology "weight decay".

**Correction.** The function to minimize is strongly convex on  $\mathbb{R}$  for  $\lambda > 0$ , and therefore admits a unique minimizer characterized by the condition

$$\widehat{\mathcal{R}}_{n,\lambda}^{\prime}\left(\widehat{\theta}_{\lambda}\right) = \frac{2}{n} \sum_{i=1}^{n} \left(\widehat{\theta}_{\lambda}\phi(x_{i}) - y_{i}\right) \phi(x_{i}) + 2\lambda\widehat{\theta}_{\lambda} = 0,$$

so that

$$\widehat{\theta}_{\lambda} = \frac{\frac{1}{n} \sum_{i=1}^{n} y_i \phi(x_i)}{\lambda + \frac{1}{n} \sum_{i=1}^{n} \phi(x_i)^2}.$$

Alternatively, one could have obtained the same expression by writing what (2.17) is in the context of this exercise.

Compared to the parameter allowing to perform ordinary least squares (which corresponds to  $\lambda = 0$ ), there is an extra factor  $\lambda > 0$  on the denominator, added to a positive term when at least one of the  $\phi(x_i)$  is non zero. Therefore,

$$\forall 0 \leqslant \lambda \leqslant \lambda', \qquad \left| \widehat{\theta}_0 \right| \leqslant \left| \widehat{\theta}_{\lambda'} \right| \leqslant \left| \widehat{\theta}_{\lambda} \right|.$$

The magnitude (in absolute value) of the parameter therefore decreases when the regularization strength is increased, which motivates the name of the approach.

**Exercise 2.10.** Show that  $\hat{\theta}_{\lambda}$  in (2.17) can be equivalently written as

$$\widehat{\theta}_{\lambda} = X^{\top} \left( X X^{\top} + n\lambda \mathrm{Id}_n \right)^{-1} Y.$$

What could the interest of such a reformulation?

**Correction.** Let us first comment that the main benefit of this reformulation is that it provides a computationally cheaper way of obtaining the optimal parameter when  $d \ge n + 1$  (*i.e.* more features than data points), as it requires solving a linear system involving a  $n \times n$  matrix instead of a  $d \times d$  matrix.

In order to obtain the result, it suffices to prove that

$$X^{\top} \left( X X^{\top} + n\lambda \mathrm{Id}_n \right)^{-1} = \left( X^{\top} X + n\lambda \mathrm{Id}_d \right)^{-1} X^{\top}.$$

Note first that the two matrices to be inverted in the previous expression are indeed invertible, by an argument similar to the one used in the solution of Exercise 2.8. The above equality is easily seen to be true by multiplying by  $XX^{\top} + n\lambda \text{Id}_n$  on the right, and by  $X^{\top}X + n\lambda \text{Id}_d$  on the left, as one ends up with the same quantity  $X^{\top}XX^{\top} + n\lambda X^{\top}$  on both sides.

Numerical solution of the ridge regression problem. Let us first discuss how to adapt the methods described in Section 2.2.3 for ordinary least squares to ridge regression. The approach based on SVD can be used as such upon adding  $\lambda Id_d$  to the matrices  $SS^{\top}$  or  $S^{\top}S$ . For gradient methods and the approach based on QR decompositions, we first show how to reduce the minimization problem in Definition 2.1 to a minimization problem for ordinary least squares. We introduce to this end

$$\widetilde{X} = \begin{pmatrix} X \\ \sqrt{n\lambda} \mathrm{Id}_d \end{pmatrix} \in \mathbb{R}^{(n+d) \times d}, \qquad \widetilde{Y} = \begin{pmatrix} Y \\ 0_d \end{pmatrix} \in \mathbb{R}^{(n+d) \times 1}$$

Then,

$$\frac{1}{n} \|Y - X\theta\|_2^2 + \lambda \|\theta\|_2^2 = \frac{1}{n} \left\|\widetilde{Y} - \widetilde{X}\theta\right\|_2^2.$$

The right hand side of the previous equality is the typical quantity to minimize for ordinary least squares. This formulation shows that the cost of computing  $\hat{\theta}_{\lambda}$  is  $O((n+d)^3)$  for the approach

based on the QR decomposition of  $\widetilde{X}$ ; while it remains of order O(n+d) per gradient step if one prefers to resort to the (stochastic) gradient approaches described in Chapter 4.

When solutions for several (many) values of  $\lambda$  are needed, in order to perform (cross) validation in particular, it may be beneficial to consider a regularization path, where one starts by solving the problem for the largest value of  $\lambda$  at hand, and then successively solves the optimization problems for decreasing values of  $\lambda$ , using the solution found for the previous value of  $\lambda$  in order to initialize or approximate the solution for the new value of  $\lambda$  (as some form of warm start in the optimization procedure). One starts from the largest value of  $\lambda$  as the associated minimization problem is usually easier to solve that for  $\lambda$  small. In fact, in the limit  $\lambda \to +\infty$ , the solution to the minimization problem is simply  $\hat{\theta}_{\infty} = 0$ , so one expects  $\hat{\theta}_{\lambda}$  to be close to 0 for  $\lambda$  large.

# 2.3.3 Statistical properties of ridge regression

In order to study the statistical performance of ridge regression, we consider as in Section 2.2.4 the fixed design setting and study the average excess risk, namely the expectation of the excess risk

$$\mathcal{R}(\theta) - \mathcal{R}^* = \mathbb{E}_{Y'} \left[ \frac{1}{n} \|Y' - X\theta\|_2^2 \right] - \mathcal{R}^*$$

averaged over all possible training sets with inputs fixed, for the predictor based on  $\hat{\theta}_{\lambda}$  (which depends on the training set Y at hand).

**Proposition 2.3.** Suppose that (2.11) holds. Then the excess risk for ridge regression is

$$\mathbb{E}\left[\mathcal{R}\left(\widehat{\theta}_{\lambda}\right)\right] - \mathcal{R}^{\star} = \lambda^{2} \theta_{\star}^{\top} \left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-2} \widehat{\Sigma} \theta_{\star} + \frac{\sigma^{2}}{n} \mathrm{Tr}\left(\widehat{\Sigma}^{2} \left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-2}\right).$$

The decomposition on the right hand side of the above equality again allows to interpret the excess risk as the sum of a squared bias (proportional to  $\lambda^2$ ) and a variance in the predictions. This formula deserves several comments:

- the result reduces to the one of Proposition 2.2 for ordinary least squares when  $\lambda = 0$ , as there is no bias term in this case, while the trace in the variance term is  $Tr(Id_d) = d$ ;
- for large values of  $\lambda$ , the squared bias is of order 1 while the variance is of order  $1/(n\lambda^2)$ ; while for small values of  $\lambda$ , the squared bias is of order  $\lambda^2$  and the variance of order 1/n. This suggests that the value of  $\lambda$  should be chosen in order to equilibriate the two contributions, namely

$$\lambda \sim \frac{1}{\sqrt{n}}.$$

This corresponds the usual bias/variance tradeoff for minimizing the mean square error. For the above choice for the scaling of  $\lambda$ , the excess risk can be shown to scale as  $1/\sqrt{n}$  (see Proposition 2.4 below).

- the bias, of order λ, can be seen as some form of approximation error (see Section 10.1). It increases as λ increases;
- the variance term can be seen as some form of estimation error (see Section 10.1), which decreases when n and  $\lambda$  increase;
- it is useful to interpret the variance term using the concept of effective number of degrees of freedom.<sup>4</sup> More precisely, denoting by  $\hat{\sigma}_j^2$  the eigenvalues of the symmetric, positive semidefinite matrix  $\hat{\Sigma}$ , the effective number of degrees of freedom is

$$0 \leqslant \operatorname{Tr}\left(\widehat{\varSigma}^{2}\left(\widehat{\varSigma} + \lambda \operatorname{Id}_{d}\right)^{-2}\right) = \sum_{j=1}^{d} \frac{\widehat{\sigma}_{j}^{4}}{\left(\widehat{\sigma}_{j}^{2} + \lambda\right)^{2}} \leqslant d$$

<sup>&</sup>lt;sup>4</sup> Our convention here differs from the more standard one considered in [40, Section 11.3.2] or [24, Section 3.4.3] for instance, where the discussion is based on the predictions rather than on the risk.

When  $\lambda = 0$ , the effective number of degrees of freedom is d as there is no capacity control. This number converges to 0 as  $\lambda \to +\infty$ .

*Proof.* Recall the risk decomposition (2.13), applied here to  $\hat{\theta}_{\lambda}$ :

$$\mathbb{E}_{Y}\left[\mathcal{R}\left(\widehat{\theta}_{\lambda}\right)\right] - \mathcal{R}^{\star} = \left\|\mathbb{E}_{Y}\left[\widehat{\theta}_{\lambda}\right] - \theta_{\star}\right\|_{\widehat{\Sigma}}^{2} + \mathbb{E}_{Y}\left[\left\|\widehat{\theta}_{\lambda} - \mathbb{E}_{Y}\left(\widehat{\theta}_{\lambda}\right)\right\|_{\widehat{\Sigma}}^{2}\right].$$

Let us successively consider the square bias and the variance. For the bias, we use the expression (2.17) and the assumption (2.11) to write

$$\mathbb{E}_{Y}\left[\widehat{\theta}_{\lambda}\right] - \theta_{\star} = \frac{1}{n} \left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} X^{\top} X \theta_{\star} - \theta_{\star} = \left[\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \widehat{\Sigma} - \mathrm{Id}_{d}\right] \theta_{\star} = -\lambda \left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \theta_{\star},$$

so that (since matrix valued functions of  $\widehat{\Sigma}$  commute)

$$\left\|\mathbb{E}_{Y}\left[\widehat{\theta}_{\lambda}\right] - \theta_{\star}\right\|_{\widehat{\Sigma}}^{2} = \lambda^{2}\theta_{\star}^{\top}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1}\widehat{\Sigma}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1}\theta_{\star} = \lambda^{2}\theta_{\star}^{\top}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-2}\widehat{\Sigma}\theta_{\star}.$$

For the variance, we use the following equality for  $u, v \in \mathbb{R}^d$  and  $M \in \mathbb{R}^{d \times d}$  (where  $vu^{\top} \in \mathbb{R}^{d \times d}$  is the matrix with entries  $v_i u_j$  for  $1 \leq i, j \leq d$ ):

$$u^{\top} M v = \sum_{i,j=1}^{d} M_{ji} u_j v_i = \operatorname{Tr}\left[\left(v u^{\top}\right) M^{\top}\right], \qquad (2.19)$$

to write

$$\mathbb{E}_{Y}\left[\left\|\widehat{\theta}_{\lambda}-\mathbb{E}_{Y}\left(\widehat{\theta}_{\lambda}\right)\right\|_{\widehat{\Sigma}}^{2}\right] = \mathbb{E}_{\varepsilon}\left[\left\|\frac{1}{n}\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}X^{\top}\varepsilon\right\|_{\widehat{\Sigma}}^{2}\right]$$
$$=\frac{1}{n^{2}}\mathbb{E}_{\varepsilon}\left[\varepsilon^{\top}X\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\widehat{\Sigma}\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}X^{\top}\varepsilon\right]$$
$$=\frac{1}{n^{2}}\mathrm{Tr}\left[X^{\top}\mathbb{E}\left(\varepsilon\varepsilon^{\top}\right)X\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\widehat{\Sigma}\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\right]$$
$$=\frac{\sigma^{2}}{n^{2}}\mathrm{Tr}\left[X^{\top}X\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\widehat{\Sigma}\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\right]$$
$$=\frac{\sigma^{2}}{n}\mathrm{Tr}\left[\widehat{\Sigma}\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\widehat{\Sigma}\left(\widehat{\Sigma}+\lambda\mathrm{Id}_{d}\right)^{-1}\right],$$

from which the claimed result follows.

Now that Proposition 2.3 provides an explicit expression for the average excess risk, we can optimize the value of  $\lambda$  in order to obtain the tightest upper bound to the average excess risk.

**Proposition 2.4 (Upper bound on the excess risk for ridge regression).** Suppose that (2.11) holds, and consider

$$\lambda_{\star} = \frac{\sigma \sqrt{\mathrm{Tr}\left(\widehat{\Sigma}\right)}}{\|\theta_{\star}\|_{2}\sqrt{n}}.$$

Then,

$$0 \leqslant \mathbb{E}_{Y} \left[ \mathcal{R} \left( \widehat{\theta}_{\lambda_{\star}} \right) \right] - \mathcal{R}^{\star} \leqslant \frac{\sigma \sqrt{\operatorname{Tr} \left( \widehat{\Sigma} \right)} \| \theta_{\star} \|_{2}}{2\sqrt{n}}.$$
(2.20)

Let us make a few comments on this result before providing its proof:

• the upper bound (2.20) is dimension free as d does not explicitly appear (in contrast to Proposition 2.2), but only implicitly. The favorable situation is when  $\|\theta_{\star}\|_{2}$  and

$$\operatorname{Tr}\left(\widehat{\Sigma}\right) = \frac{1}{n} \sum_{i=1}^{n} \|\varphi(x_i)\|_2^2$$

both remain bounded as the dimension increases. A sufficient condition for the latter term to be bounded is that  $\|\varphi(x)\|_2 \leq \Phi < +\infty$  uniformly in  $x \in \mathcal{X}$  for some  $\Phi \in \mathbb{R}_+$  as the dimension increases;

- the convergence rate  $O(1/\sqrt{n})$  which is obtained is slower than the rate O(1/n) of Proposition 2.2 for ordinary least square regression, but it has on the other side a milder dependence on the noise ( $\sigma$  instead of  $\sigma^2$ ) and on the dimension d. The estimate may therefore be useful in the nonasymptotic regime where n is not too large, as it corresponds to a slower rate of convergence but potentially a better prefactor in the estimate;
- the expression of the optimal value  $\lambda_{\star}$  cannot be used to determine the regularization strength in practice for two reasons: (i) the values of  $\sigma$  and  $\theta_{\star}$  are usually unknown; (ii) the value of  $\lambda_{\star}$  is obtained by optimizing an upper bound, and the so-obtained value may therefore be suboptimal for the actual average (test) risk. The regularization strength should be determined in practice by (cross)validation, as discussed in Section 2.3.1. Nonetheless, the expression for  $\lambda_{\star}$ still leads to useful theoretical insights, in particular ideas on how the optimal regularization strength behaves as the number of training data points is increased.

*Proof.* Since, for any  $\lambda > 0$  and  $\mu \ge 0$ ,

$$(\mu + \lambda)^2 = \mu^2 + \lambda^2 + 2\lambda\mu \ge 4\lambda\mu,$$

and hence

$$\frac{\mu\lambda}{(\mu+\lambda)^2} \leqslant \frac{1}{4},$$

the eigenvalues of  $\lambda \widehat{\Sigma} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-2}$  are smaller than 1/4, and therefore the square bias in the equality of Proposition 2.3 can be bounded as

$$\lambda^2 \theta_{\star}^{\top} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-2} \widehat{\Sigma} \theta_{\star} \leqslant \frac{\lambda}{4} \|\theta_{\star}\|_2^2.$$

Similarly, in view of the matrix inequality (in the sense of symmetric matrices)

$$\widehat{\Sigma}^2 \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-2} = \frac{1}{\lambda} \widehat{\Sigma}^{1/2} \left[ \lambda \widehat{\Sigma} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-2} \right] \widehat{\Sigma}^{1/2} \leqslant \frac{1}{\lambda} \widehat{\Sigma}^{1/2} \left[ \frac{1}{4} \mathrm{Id}_d \right] \widehat{\Sigma}^{1/2} = \frac{1}{4\lambda} \widehat{\Sigma},$$

the variance term can be bounded as

$$\frac{\sigma^2}{n} \operatorname{Tr}\left(\widehat{\Sigma}^2\left(\widehat{\Sigma} + \lambda \operatorname{Id}_d\right)^{-2}\right) \leqslant \frac{\sigma^2}{4n\lambda} \operatorname{Tr}\left(\widehat{\Sigma}\right).$$

By gathering the two upper bounds, we obtain

$$0 \leqslant \mathbb{E}_Y \left[ \mathcal{R}\left(\widehat{\theta}_{\lambda}\right) \right] - \mathcal{R}^{\star} \leqslant \frac{\lambda}{4} \|\theta_{\star}\|_2^2 + \frac{\sigma^2}{4n\lambda} \operatorname{Tr}\left(\widehat{\varSigma}\right) := u(\lambda).$$

We can now minimize the right hand side of the previous inequality with respect to  $\lambda$ . Clearly, the function u is smooth on  $(0, +\infty)$  goes to  $+\infty$  as  $\lambda \to 0$  or  $\lambda \to +\infty$ . There exists therefore a global minimizer. Global minimizers are characterized by the Euler–Lagrange equation

$$u'(\lambda_{\star}) = 0 = \frac{1}{4} \|\theta_{\star}\|_{2}^{2} - \frac{\sigma^{2}}{4n\lambda_{\star}^{2}} \operatorname{Tr}\left(\widehat{\Sigma}\right).$$

This leads to the expression of  $\lambda_{\star}$ . The upper bound for the average excess risk is easily seen to be equal to  $\lambda_{\star} \|\theta_{\star}\|_{2}^{2}/2$ , which is (2.20).

**Exercise 2.11.** The aim of this exercise is to characterize how  $\operatorname{Tr}(\widehat{\Sigma})$  scales with respect to the dimension d for the particular example of polynomial regression on  $\mathcal{X} = [-1, 1]$ , in the limit  $n \to +\infty$ . Recall that in this case  $\varphi(x) = (1, x, x^2, \dots, x^{d-1})$  when  $\theta \in \mathbb{R}^d$ .

(1) Compute the entries of the limiting matrix  $\Sigma = \lim_{n \to +\infty} \widehat{\Sigma}$  when the inputs  $x_i$  are independently and identically distributed according to the uniform distribution on  $\mathcal{X}$ .

(2) Determine the asymptotic behavior of  $\operatorname{Tr}(\Sigma)$  as  $d \to +\infty$ , and interpret the result.

# Correction.

(1) For  $1 \leq k, k' \leq d$ , the Law of Large Numbers implies that

$$\begin{split} \Sigma_{k,k'} &= \lim_{n \to +\infty} \frac{1}{n} \sum_{i=1}^{n} \varphi_k(x_i) \varphi_{k'}(x_i) = \int_{-1}^{1} \varphi_k(x) \varphi_{k'}(x) \frac{dx}{2} = \frac{1}{2} \int_{-1}^{1} x^{k+k'-2} dx \\ &= \frac{1}{2(k+k'-1)} \left[ 1 - (-1)^{k+k'-1} \right]. \end{split}$$

(2) With the result of the previous question,

$$\operatorname{Tr}(\varSigma) = \sum_{k=1}^{d} \varSigma_{k,k} = \sum_{k=1}^{d} \frac{1}{2k-1} \sim \frac{1}{2} \log d,$$

since

$$\frac{1}{2}\log(2d-1) = \frac{1}{2}\int_{1}^{d} \frac{1}{y-1/2} \, dy \leqslant \sum_{k=1}^{d} \frac{1}{2k-1}$$
$$= 1 + \sum_{k=1}^{d-1} \frac{1}{2k+1} \leqslant 1 + \frac{1}{2}\int_{0}^{d-2} \frac{1}{y+1/2} \, dy = 1 + \frac{1}{2}\log(2d-3).$$

If  $\|\theta_{\star}\|_2$  is bounded as  $d \to +\infty$ , the upper bound obtained on the average excess risk therefore scales as  $\sqrt{\log(d)/n}$  for ridge regression, instead of the bound scaling as d/n for ordinary least squares. The bound is therefore much better when

$$\frac{\log(d)}{d^2} \ll \frac{1}{n}.$$

Moreover, the bound can be used for  $d \ge n$ .

**Exercise 2.12 (Ridge regression in the random design setting).** We consider the same setting as in Exercise 2.6 in the context of ridge regression. Prove that

$$\mathbb{E}\left[\mathcal{R}(\widehat{\theta}_{\lambda})\right] - \mathcal{R}^{\star} = \lambda^{2} \mathbb{E}\left[\theta_{\star}^{\top}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \mathcal{E}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \theta_{\star}\right] + \frac{\sigma^{2}}{n} \mathbb{E}\left[\mathrm{Tr}\left(\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-2} \widehat{\Sigma} \mathcal{E}\right)\right].$$

where the expectation is over the realizations of the data set.

**Correction.** We follow the computations of Exercise 2.12 to write (the expectation being both over the realizations of  $x_1, \ldots, x_n$  and  $\varepsilon_1, \ldots, \varepsilon_n$ )

$$\mathbb{E}\left[\mathcal{R}(\widehat{\theta}_{\lambda})\right] - \mathcal{R}^{\star} = \mathbb{E}\left[(\widehat{\theta}_{\lambda} - \theta_{\star})^{\top} \Sigma(\widehat{\theta}_{\lambda} - \theta_{\star})\right],$$

with

#### 2.4 LASSO regression 43

$$\widehat{\theta}_{\lambda} = \frac{1}{n} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-1} X^{\top} (X \theta_{\star} + \varepsilon) = \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-1} \widehat{\Sigma} \theta_{\star} + \frac{1}{n} \left( \widehat{\Sigma} + \lambda \mathrm{Id}_d \right)^{-1} X^{\top} \varepsilon,$$

so that

$$\widehat{\theta}_{\lambda} - \theta_{\star} = -\lambda \left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \theta_{\star} + \frac{1}{n} \left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} X^{\top} \varepsilon.$$

Therefore,

$$\mathbb{E}\left[\mathcal{R}(\widehat{\theta}_{\lambda})\right] - \mathcal{R}^{\star} = \lambda^{2} \mathbb{E}\left[\theta_{\star}^{\top}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \mathcal{L}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} \theta_{\star}\right] \\ + \frac{1}{n^{2}} \mathbb{E}\left[\left(\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} X^{\top} \varepsilon\right)^{\top} \mathcal{L}\left(\widehat{\Sigma} + \lambda \mathrm{Id}_{d}\right)^{-1} X^{\top} \varepsilon\right].$$

The second term on the right hand side of the above equality can be rewritten as

$$\frac{1}{n^{2}} \mathbb{E} \left[ \left( \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} X^{\top} \varepsilon \right)^{\top} \Sigma \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} X^{\top} \varepsilon \right] \\ = \frac{1}{n^{2}} \mathbb{E} \left[ \varepsilon^{\top} X \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} \Sigma \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} X^{\top} \varepsilon \right] \\ = \frac{1}{n^{2}} \mathbb{E} \left[ \operatorname{Tr} \left( X^{\top} \varepsilon \varepsilon^{\top} X \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} \Sigma \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} \right) \right] \\ = \frac{\sigma^{2}}{n} \mathbb{E} \left[ \operatorname{Tr} \left( \widehat{\Sigma} \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} \Sigma \left( \widehat{\Sigma} + \lambda \operatorname{Id}_{d} \right)^{-1} \right) \right],$$

from which the claimed result follows.

# 2.4 LASSO regression

LASSO was introduced by Tibshirani in 1996 (see [53]). The acronym stands for "Least Absolute Shrinkage and Selection Operator". As the name indicates, an essential aim of Lasso regularization is to perform feature selection, and build sparse predictors that depend only a small number of the components of the feature vector  $\varphi(x)$ . This is useful for two reasons:

- (i) it makes models more interpretable (as it allows to decide which features are important for predictions), and/or also allows to add extra features, whose relevance is unclear, as irrelevant features will automatically be discarded;
- (ii) as for ridge regression, it also allows to reduce the dimension dependence in the guarantee bounds on the excess risk (recall for instance the bounds  $O(n^{-1/d})$  for K-nearest neighbors obtained in (1.23), and the bound  $\sigma^2 d/n$  of Proposition 2.2 for ordinary least square regression).

There are however two difficulties with this approach:

- (i) the identity of the relevant variables is not known beforehand. They need therefore to be identified;
- (ii) Lasso works only if the Bayes predictor itself involves only a small number of features (at least approximatively).

The presentation in this section follows [40, Section 11.4] and [4, Chapter 8].

# 2.4.1 Sparsity inducing regularization terms

As in the remainder of this chapter, we focus here on linear methods, and therefore aim at solving

$$\min_{\theta \in \mathbb{R}^d} \mathbb{E}\left[ \ell\left(y, \varphi(x)^\top \theta\right) \right]$$

In fact, we will often specify our analysis to the regression setting for which  $\ell(y, z) = ||y - z||_2^2$  is the squared Euclidean norm, but the framework discussed here is general enough to cover other situations.

As motivated in Section 2.3, a regularization term  $\Omega(\theta)$  should be added to the training loss in order to prevent overfitting. We discuss two options here, which both lead to sparse solutions:

• a first option to consider  $\Omega(\theta) = \|\theta\|_0$ , which is the number of non zero entries of  $\theta$ . This naturally favors sparse solutions where many entries of the optimal parameter  $\theta$  are 0. However, the associated optimization problem

$$\min_{\theta \in \mathbb{R}^d} \left\{ \widehat{\mathcal{R}}_n(\theta) + \lambda \|\theta\|_0 \right\}$$

is difficult to solve as it essentially requires some form of combinatorial approach in order to identify the set of relevant  $k = \|\theta\|_0$  variables, or rely on dedicated algorithms such as greedy methods where features are introduced one by one and chosen to minimize the loss ("orthogonal matching pursuit"; see for instance [49, Section 25.1.2]).

• another option is to rely on a  $\ell^1$  penalty term

$$\Omega(\theta) = \|\theta\|_1 = \sum_{k=1}^d |\theta_k|.$$
(2.21)

The main interest of this approach is that the function appearing in the minimization problem

$$\min_{\theta \in \mathbb{R}^d} \left\{ \widehat{\mathcal{R}}_n(\theta) + \lambda \|\theta\|_1 \right\}$$
(2.22)

is convex when  $\widehat{\mathcal{R}}_n$  is convex (see Exercise 2.13 below). Gradient methods (as described in Chapter 4) can then be used to approximate the minimizer. It is however not clear at first sight why the  $\ell^1$  penalty (2.21) leads to a sparse minimizer. We next motivate this. The combination of the sparsification property, and the fact that the regularized loss to be minimized has good properties (convexity) makes the method particularly attractive.

**Exercise 2.13.** Prove that the function  $\theta \mapsto \widehat{\mathcal{R}}_n(\theta) + \lambda \|\theta\|_1$  is convex on  $\mathbb{R}^d$  for the empirical risk function  $\widehat{\mathcal{R}}_n$  defined in (2.3) and any  $\lambda \ge 0$ .

**Correction.** The function  $\widehat{\mathcal{R}}_n$  is convex, as can be seen for instance from the fact that its Hessian is positive semidefinite (recall (2.6); here X is not necessarily of full rank). The function  $\theta \mapsto \|\theta\|_1$  is also convex (coming back to the definition of convex functions and relying on the triangle inequality satisfied by any norm). The result then follows from the fact that the sum of two convex functions is convex.

Draw picture Motivating the sparsity inducing effect of  $\ell^1$  regularization. Balls for the  $\ell^1$  norm have extremal points which "stick out more". These extremal points correspond to sparse elements  $\theta \in \mathbb{R}^d$  where one component at least is 0. The tangency or contact between isolines of  $\widehat{\mathcal{R}}_n$  and the unit ball in the  $\ell^1$  norm tends to happen at the corners or vertices of the unit ball. For ridge regression on the other hand, the relevant unit ball to consider in the one in the  $\ell^2$  norm, whose boundary is a hypersphere, which is fully symmetric and does not have sparse extremal points.

In order to make contact between the above considerations and the minimization problem (2.22), we recall that (2.22) can be seen as some dual formulation of the constrained minimization problem (2.16). The latter setting is exactly the one considered above. This motivates why sparse solutions are expected for  $\ell^1$  regularization, but not for  $\ell^2$  regularization.

#### 2.4.2 Explicit expressions for one dimensional problems

In order to gain some intuition on the behavior of solutions to the Lasso minimization problem (2.22), we consider the simplest possible setting, namely the situation when  $\theta \in \mathbb{R}$  and there is a single one dimensional data point x = 1 with associated output  $y \in \mathbb{R}$ . In this case, the function to minimize is

$$F_{\lambda}(\theta) = \frac{1}{2}(y-\theta)^2 + \lambda|\theta|, \qquad (2.23)$$

where we introduced a factor 1/2 in front of the first term for convenience. The next exercise shows how to analytically obtain the unique global minimizer of  $F_{\lambda}$ .

#### **Exercise 2.14.** Consider the function $F_{\lambda}$ defined in (2.23).

- (1) Prove that  $F_{\lambda}$  admits a unique minimizer.
- (2) Compute the left and right derivatives of  $F_{\lambda}$  at any  $\theta \in \mathbb{R}$ .
- (3) Prove that the minimizer  $\theta_{\lambda}^{\star}$  of  $F_{\lambda}$  is

$$\theta_{\lambda}^{\star} = \begin{vmatrix} y + \lambda & \text{for } y \leqslant -\lambda, \\ 0 & \text{for } -\lambda \leqslant y \leqslant \lambda, \\ y - \lambda & \text{for } y \geqslant \lambda, \end{vmatrix}$$
(2.24)

# Correction.

(1) The function  $F_{\lambda}$  is the sum of the convex function  $\theta \mapsto \lambda |\theta|$  and the strongly convex function  $\theta \mapsto (y - \theta)^2/2$  (as can be seen for instance by computing the second derivative of this function, which is equal to 1); so that  $F_{\lambda}$  is strongly convex on  $\mathbb{R}$  and therefore admits a unique minimizer.

Alternatively, one could have noted that  $F_{\lambda}$  is smooth and coercive, and hence admits a minimizer. The uniqueness comes from the fact that  $\theta \mapsto (y - \theta)^2/2$  is strictly convex and  $\theta \mapsto \lambda |\theta|$  is convex, so overall  $F_{\lambda}$  is strictly convex.

(2) For  $\theta < 0$ , the function  $\theta \mapsto F_{\lambda}(\theta) = \frac{1}{2}(y-\theta)^2 - \lambda\theta$  is smooth, with derivative  $F'_{\lambda}(\theta) = \theta - y - \lambda$ . For  $\theta > 0$ , the function  $\theta \mapsto F_{\lambda}(\theta) = \frac{1}{2}(y-\theta)^2 + \lambda\theta$  is smooth, with derivative  $F'_{\lambda}(\theta) = \theta - y + \lambda$ . The only point where  $F_{\lambda}$  is not continuously differentiable is  $\theta = 0$ , for which the left and right derivatives are respectively

$$F_{\lambda}'(0^{-}) = \lim_{\substack{\eta \to 0 \\ \eta > 0}} F_{\lambda}(-\eta) = -y - \lambda, \qquad F_{\lambda}'(0^{+}) = \lim_{\substack{\eta \to 0 \\ \eta > 0}} F_{\lambda}(\eta) = -y + \lambda.$$

- (3) We distinguish three situations, depending on the values of y:
  - when  $y > \lambda$ , then  $F'_{\lambda}(y \lambda) = 0$  for the value  $y \lambda > 0$ , and  $F'_{\lambda}(\theta) \neq 0$  for  $\theta \neq y \lambda$ . This shows that the unique minimizer  $\theta^*_{\lambda}$  is equal to  $y - \lambda > 0$ ;
  - when  $y < -\lambda$ , then  $F'_{\lambda}(y + \lambda) = 0$  for the value  $y + \lambda < 0$ , and  $F'_{\lambda}(\theta) \neq 0$  for  $\theta \neq y + \lambda$ . This shows that the unique minimizer  $\theta^{\star}_{\lambda}$  is equal to  $y + \lambda < 0$ ;
  - when  $|y| \leq \lambda$ , then  $F'_{\lambda} < 0$  on  $(-\infty, 0)$  and  $F'_{\lambda} > 0$  on  $(0, +\infty)$ . This shows that the global minimizer is  $\theta^*_{\lambda} = 0$ .
  - The collection of these three scenarios indeed leads to (2.24).

The expression (2.24) for  $\theta^{\star}_{\lambda}$  can be written in a more condensed way as

$$\theta_{\lambda}^{\star} = \operatorname{sign}(y) \max\left(|y| - \lambda, 0\right).$$

This corresponds to a soft thresholding function, where the value of y is shifted by an amount  $\lambda > 0$  towards 0 (*i.e.* one add  $\lambda$  if y < 0 and subtracts  $\lambda$  if y > 0) unless y is too small, in which case it is set to 0.

#### 2.4.3 Expressions for the Lasso least square regression problem

We consider in this section the general Lasso least square regression problem, where one seeks to minimize

$$\widehat{\mathcal{R}}_{n,\lambda}(\theta) = \frac{1}{2n} \|Y - X\theta\|_2^2 + \lambda \|\theta\|_1 = \frac{1}{2n} \sum_{i=1}^n \left(y_i - \theta^\top \varphi(x_i)\right)^2 + \lambda \sum_{k=1}^d |\theta_k|.$$

As in (2.23), we introduce a factor 1/2 in front of the empirical risk in order to simplify the algebra. This is not a restriction of generality as it amounts to rescaling the values of  $\lambda$  by a factor 2. The function  $\widehat{\mathcal{R}}_{n,\lambda}$  is convex as the sum of two convex functions, and strongly convex when X has full rank.

The optimality condition satisfied by a global minimizer reads

$$0_d \in \partial \widehat{\mathcal{R}}_{n,\lambda}(\theta^\star_\lambda),$$

more precise reference needed where  $\partial \widehat{\mathcal{R}}_{n,\lambda}$  is the subgradient (see Chapter 4 for background material on subgradients). Since  $\widehat{\mathcal{R}}_{n,\lambda}$  is smooth on  $(\mathbb{R} \setminus \{0\})^d$ , the components of the subgradient are simply the partial derivatives for the components of  $\theta$  which are non zero.

In order to derive necessary conditions of optimality, let us next consider variations with respect to the component  $\theta_k$  when  $\theta_k = 0$ . We fix to this end the components  $\theta_1, \ldots, \theta_{k-1}, \theta_{k+1}, \ldots, \theta_d$ , and introduce the function

$$\mathscr{R}_{k}(\theta_{k}) = \widehat{\mathcal{R}}_{n,\lambda} \left(\theta_{1}, \dots, \theta_{k-1}, \theta_{k}, \theta_{k+1}, \dots, \theta_{d}\right)$$
$$= \frac{1}{2n} \sum_{i=1}^{n} \left(a_{ik} - \theta_{k} \varphi_{k}(x_{i})\right)^{2} + \lambda \left|\theta_{k}\right| + L_{k}(\theta_{1}, \dots, \theta_{k-1}, \theta_{k+1}, \dots, \theta_{d}),$$

where  $L_k$  does not depend on  $\theta_k$ , and

$$a_{ik} = y_i - \sum_{\ell \neq k} \theta_\ell \varphi_\ell(x_i).$$
(2.25)

We next rewrite  $\mathscr{R}_k(\theta_k)$  in a manner similar to (2.23) as follows:

$$\begin{aligned} \mathscr{R}_k(\theta_k) &= \left(\frac{1}{2n}\sum_{i=1}^n \varphi_k(x_i)^2\right) \theta_k^2 - \frac{1}{n} \left(\sum_{i=1}^n a_{ik}\varphi_k(x_i)\right) \theta_k + \lambda \left|\theta_k\right| + \widetilde{L}_k(\theta_1, \dots, \theta_{k-1}, \theta_{k+1}, \dots, \theta_d) \\ &= \left(\frac{1}{n}\sum_{i=1}^n \varphi_k(x_i)^2\right) \left[\frac{1}{2} \left(\theta_k - \frac{\sum_{i=1}^n a_{ik}\varphi_k(x_i)}{\sum_{i=1}^n \varphi_k(x_i)^2}\right)^2 + \frac{n\lambda}{\sum_{i=1}^n \varphi_k(x_i)^2} \left|\theta_k\right|\right] + \widehat{L}_k(\theta_1, \dots, \theta_{k-1}, \theta_{k+1}, \dots, \theta_d). \end{aligned}$$

In view of (2.24), we therefore find that the optimal parameter  $\theta_{\lambda}^{\star}$  necessarily satisfies

$$\theta_{k,\lambda}^{\star} = \begin{vmatrix} \sum_{i=1}^{n} a_{ik}^{\star} \varphi_k(x_i) + n\lambda \\ \sum_{i=1}^{n} \varphi_k(x_i)^2 & \text{for } \frac{1}{n} \sum_{i=1}^{n} a_{ik}^{\star} \varphi_k(x_i) \leqslant -\lambda, \\ 0 & \text{for } -\lambda \leqslant \frac{1}{n} \sum_{i=1}^{n} a_{ik}^{\star} \varphi_k(x_i) \leqslant \lambda, \\ \frac{\sum_{i=1}^{n} a_{ik}^{\star} \varphi_k(x_i) - n\lambda}{\sum_{i=1}^{n} \varphi_k(x_i)^2} & \text{for } \frac{1}{n} \sum_{i=1}^{n} a_{ik}^{\star} \varphi_k(x_i) \geqslant \lambda. \end{aligned}$$
(2.26)

where  $a_{ik}^{\star}$  is (2.25) evaluated at  $\theta = \theta_{\lambda}^{\star}$ .

The latter expression motivates the term "shrinkage" in the expression "least absolute shrinkage" for Lasso, as large coefficients (in absolute value) are shrunk towards 0. This also suggests that the Lasso estimator is biased. This may not be an issue for two reasons:

- (1) we want a good prediction, and do not care so much about estimating the true parameter  $\theta_{\star}$  or recovering the true model (which is the aim of high dimensional statistics);
- (2) introducing some bias may be helpful in view of the bias/variance tradeoff, especially if the magnitude of the bias is determined using cross validation.

#### 2.4.4 Numerical methods to find the minimizers

The optimality conditions (2.26) are cumbersome since the components  $\theta_{k',\lambda}^{\star}$  for  $k' \neq k$  are hidden in the coefficients  $a_{ik}^{\star}$ . The numerical methods to find solutions to (2.22) are therefore based on other ideas than solving the nonlinear equation (2.26) (see [40, Section 11.4.9] for further precisions):

- a popular approach, implemented for instance in the default method for Lasso regression in scikit-learn, is *coordinate descent*: at each iteration, a coordinate to update is selected (for instance by cycling over all components or by choosing it at random), and a one dimensional minimization, with other coordinates fixed, is performed. The interest of this approach is that the one dimensional minimization can be performed analytically, following the computations in Section 2.4.3. This leads to expressions close to (2.26) with coefficients  $a_{ik}$  evaluated at the current value of the parameter  $\theta$ ;
- another approach is based on a projected gradient method, for which the parameter  $\theta \in \mathbb{R}^d$  is decomposed as  $\theta = \theta^+ \theta^-$ , where  $\theta^+, \theta^- \in \mathbb{R}^d_+$  have only nonnegative components. The optimization problem (2.22) can then be reformulated as the minimization of a smooth function under positivity constraints:

$$\min_{\theta^+,\theta^- \in \mathbb{R}^d_+} \left\{ \widehat{\mathcal{R}}_n \left( \theta^+ - \theta^- \right) + \lambda \sum_{k=1}^d \left( \theta^+_k + \theta^-_k \right) \right\}.$$

This amounts to doubling the size of the optimization problem and adding 2*d* positivity constraints  $\theta_k^+, \theta_k^- \ge 0$  for  $1 \le k \le d$ . These constraints are satisfied by projecting the updates of gradient methods onto the set of admissible values for  $\theta^+, \theta^-$ ;

- *iterative soft thresholding* is based on proximal gradient strategies, which builds upon a framework to regularize non differentiable terms in a function to optimize;
- one can also rely on *continuation/homotopy methods*, based on the remark that  $\theta_{\lambda}^{\star}$  is piecewise affine in  $\lambda$ . The idea is then to start from large values of  $\lambda$ , and build the path of solutions by computing break points one after the other.

#### 2.4.5 Theoretical guarantees

In order to provide theoretical upper bounds on the excess risk for Lasso regression, we consider once again the fixed design setting, as in Sections 2.2.4 and 2.3.3. A key point of the analysis performed in this section (as for ridge regression in Section 2.3.3) is to carefully keep track of the dimensionality dependence, possibly upon going from "fast rates" O(1/n) to "slow rates"  $O(1/\sqrt{n})$ .

We still assume for the analysis that (2.11) holds, for some  $\theta_{\star}$  which is sparse – so that  $\|\theta_{\star}\|_1$  remains bounded as the dimension d increases. We start by considering a general norm  $\Omega(\theta)$  for the regularization term (recall that  $\Omega(\theta) = \|\theta\|_2$  for ridge regression, and  $\Omega(\theta) = \|\theta\|_1$  for Lasso; see Section 2.3.1). The dual norm is denoted by  $\Omega^*$  and defined as

$$\Omega^*(z) = \sup \left\{ z^\top \theta \mid \Omega(\theta) \leqslant 1 \right\}$$

**Exercise 2.15.** Fix  $1 \leq p \leq +\infty$  and consider

$$\Omega(\theta) = \|\theta\|_p = \left(\sum_{k=1}^d |\theta_k|^p\right)^{1/p}$$

for  $1 \leq p < +\infty$  and, for  $p = +\infty$ ,

$$\|\theta\|_{\infty} = \max_{1 \leqslant k \leqslant d} |\theta_k|.$$

Show that  $\Omega^*(z) = ||z||_{p'}$  where  $1 \leq p' \leq +\infty$  is the exponent conjugate to p:

$$\frac{1}{p} + \frac{1}{p'} = 1.$$

**Correction.** Fix  $z \in \mathbb{R}^d$ . Hölder's inequality implies that  $z^{\top}\theta \leq ||z||_{p'}||\theta||_p$ , so that  $0 \leq \Omega^*(z) \leq ||z||_{p'}$ . It therefore suffices to find  $\theta_z \in \mathbb{R}^d$  such that  $z^{\top}\theta_z = ||z||_{p'}$  and  $\Omega(\theta_z) = ||\theta_z||_p = 1$ . We distinguish three situations:

• for  $p = +\infty$  (so that p' = 1), one can choose  $\theta_{z,k} = \operatorname{sign}(z_k)$ , in which case  $\|\theta_z\|_{\infty} = 1$  and

$$z^{\top} \theta_z = \sum_{k=1}^d |z_k| = \|z\|_1$$

• for p = 1 (so that  $p' = +\infty$ ), one can choose  $\theta_{z,k} = \operatorname{sign}(z_{k_z^*}) \mathbf{1}_{k=k_z^*}$  where the integer  $k_z^* \in \{1, \ldots, d\}$  is such that  $\|z\|_{\infty} = |z_{k_z^*}|$ . In this case,  $\|\theta_z\|_1 = 1$  and

$$z^{\top}\theta_z = \left|z_{k_z^*}\right| = \|z\|_{\infty};$$

• for  $p \in (1, +\infty)$ , one can choose

$$\theta_{z,k} = \frac{|z_k|^{p'-1} \operatorname{sign}(z_k)}{\|z\|_{p'}^{p'/p}}$$

Note that, since p = p'/(p'-1) so that p(p'-1) = p',

$$\|\theta_z\|_p^p = \sum_{k=1}^d \frac{|z_k|^{p(p'-1)}}{\|z\|_{p'}^{p'}} = 1,$$

and

$$z^{\top}\theta_{z} = \frac{1}{\|z\|_{p'}^{p'/p}} \sum_{k=1}^{d} |z_{k}|^{p'-1} \operatorname{sign}(z_{k}) z_{k} = \|z\|_{p'}^{p'(1-1/p)} = \|z\|_{p'}.$$

This finally gives the desired result.

**Exercise 2.16.** Prove that, for any  $u, v \in \mathbb{R}^d$ , it holds  $|u^{\top}v| \leq \Omega(u)\Omega^*(v)$ .

**Correction.** It suffices to prove the result for  $v \neq 0$  (the result being trivial for v = 0). Note that  $\Omega(v) > 0$  as  $\Omega$  is a norm, so that, upon introducing

$$\theta_v = \frac{v}{\Omega(v)},$$

one has  $\Omega(\theta_v) \leq 1$ , and hence, by definition of  $\Omega^*$ ,

$$|u^{\top}v| = \Omega(v) \left|u^{\top}\theta_{v}\right| \leqslant \Omega(v)\Omega^{*}(u),$$

which is the desired inequality.

We need some preliminary results on the excess risk

$$\mathcal{R}\left(\widehat{\theta}_{\lambda}\right) - \mathcal{R}^{\star} = \frac{1}{n} \left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2}^{2} = \left(\widehat{\theta}_{\lambda} - \theta_{\star}\right)^{\top} \widehat{\Sigma}\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right)$$

for Lasso regression (recall Exercise 2.4), where  $\hat{\theta}_{\lambda}$  is a minimizer<sup>5</sup> of  $\theta \mapsto \frac{1}{2n} \|Y - X\theta\|_2^2 + \lambda \Omega(\theta)$ .

Lemma 2.1. Suppose that (2.11) holds.

(a) If  $\Omega^*(X^{\top}\varepsilon) \leqslant n\lambda/2$ , then

$$\Omega\left(\widehat{\theta}_{\lambda}\right) \leqslant 3\Omega(\theta_{\star}), \qquad \frac{1}{n} \left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2}^{2} \leqslant 3\lambda\Omega(\theta_{\star}).$$

(b) In all cases,

$$\frac{1}{n} \left\| X \left( \widehat{\theta}_{\lambda} - \theta_{\star} \right) \right\|_{2}^{2} \leqslant \frac{4}{n} \| \varepsilon \|_{2}^{2} + 4\lambda \Omega(\theta_{\star}).$$

*Proof.* By definition of  $\hat{\theta}_{\lambda}$ ,

$$\left\|Y - X\widehat{\theta}_{\lambda}\right\|_{2}^{2} \leq \left\|Y - X\theta_{\star}\right\|_{2}^{2} + 2n\lambda\Omega\left(\theta_{\star}\right) - 2n\lambda\Omega\left(\widehat{\theta}_{\lambda}\right).$$

In view of (2.11), the latter inequality can be rewritten as

$$\left\| X\left(\theta_{\star} - \widehat{\theta}_{\lambda}\right) + \varepsilon \right\|_{2}^{2} \leq \|\varepsilon\|_{2}^{2} + 2n\lambda\Omega\left(\theta_{\star}\right) - 2n\lambda\Omega\left(\widehat{\theta}_{\lambda}\right).$$

By expanding the square on the left hand side,

$$\left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2}^{2} \leq 2\varepsilon^{\top} X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) + 2n\lambda \Omega\left(\theta_{\star}\right) - 2n\lambda \Omega\left(\widehat{\theta}_{\lambda}\right).$$

We distinguish two cases at this stage:

• When  $\Omega^*(X^{\top}\varepsilon) \leq n\lambda/2$ , the latter right hand side can be bounded in view of Exercise 2.16 by

$$\begin{split} 2\Omega^{*}\left(\boldsymbol{X}^{\top}\boldsymbol{\varepsilon}\right)\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}-\boldsymbol{\theta}_{\star}\right)+2n\lambda\boldsymbol{\Omega}\left(\boldsymbol{\theta}_{\star}\right)-2n\lambda\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}\right)&\leqslant n\lambda\left[\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}-\boldsymbol{\theta}_{\star}\right)+2\boldsymbol{\Omega}\left(\boldsymbol{\theta}_{\star}\right)-2\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}\right)\right]\\ &\leqslant n\lambda\left[\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}\right)+\boldsymbol{\Omega}\left(\boldsymbol{\theta}_{\star}\right)+2\boldsymbol{\Omega}\left(\boldsymbol{\theta}_{\star}\right)-2\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}\right)\right]\\ &=n\lambda\left[3\boldsymbol{\Omega}\left(\boldsymbol{\theta}_{\star}\right)-\boldsymbol{\Omega}\left(\widehat{\boldsymbol{\theta}}_{\lambda}\right)\right], \end{split}$$

from which the first claimed inequalities follows.

• In the general case, we use first a discrete Cauchy–Schwarz inequality then a discrete Young inequality to write

$$\left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2}^{2} \leq 2 \left\| \varepsilon \right\|_{2} \left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2} + 2n\lambda\Omega\left(\theta_{\star}\right) \leq 2 \left\| \varepsilon \right\|_{2}^{2} + \frac{1}{2} \left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2}^{2} + 2n\lambda\Omega\left(\theta_{\star}\right),$$

from which the second claimed inequality follow.

This allows to conclude the proof.

<sup>5</sup> Mind the factor 1/2 here. In contrast, the excess risk relies on the standard quadratic error. Overall this amounts to a rescaling of  $\lambda$ .

We can now apply the previous result to Lasso regression by choosing  $\Omega(\theta) = \|\theta\|_1$ . Then,  $\Omega^* (X^{\top}\varepsilon) = \|X^{\top}\varepsilon\|_{\infty}$  is a maximum over the absolute values of 2d (possibly dependent) random variables  $Z_k^{\pm} = \pm (X^{\top}\varepsilon)_k$  for  $1 \leq k \leq d$ . Classical results in probability theory (see for instance [4, Section 1.2.4]) show that this quantity scales as  $\sqrt{\log(2d)}$  multiplied by the magnitude of the random variables at hand. Since  $Z_k^{\pm}$  is the sum of *n* independent random variables, the Central Limit Theorem suggests that each random variable  $Z_k^{\pm}$  has values of order  $\sqrt{n}$ . Overall, this suggests that  $\Omega^* (X^{\top}\varepsilon)$  is of order  $\sqrt{n\log(2d)}$  so that, in order to apply Lemma 2.1(a), one should choose

$$\lambda \sim \sqrt{\frac{\log(d)}{n}}.$$

In order to state a precise result, we denote by  $\|\widehat{\Sigma}\|_{\infty}$  the largest element of  $\widehat{\Sigma}$  in absolute value.

**Proposition 2.5.** Suppose that (2.11) holds true, with  $\varepsilon$  a Gaussian vector with covariance  $\sigma^2 \mathrm{Id}_n$ . Then, for

$$\lambda = \frac{2\sigma}{\sqrt{n}} \sqrt{2\|\widehat{\Sigma}\|_{\infty}} \sqrt{\log(d) + \log\left(\frac{1}{\delta}\right)},\tag{2.27}$$

and

$$\widehat{\theta}_{\lambda} \in \operatorname*{argmin}_{\theta \in \mathbb{R}^d} \left\{ \frac{1}{2n} \|Y - X\theta\|_2^2 + \lambda \varOmega(\theta) \right\},$$

it holds, with probability larger or equal than  $1 - \delta$ ,

$$\mathcal{R}\left(\widehat{\theta}_{\lambda}\right) - \mathcal{R}^{\star} = \frac{1}{n} \left\| X\left(\widehat{\theta}_{\lambda} - \theta_{\star}\right) \right\|_{2}^{2} \leqslant 6 \|\theta_{\star}\|_{1} \frac{\sigma}{\sqrt{n}} \sqrt{2\|\widehat{\Sigma}\|_{\infty}} \sqrt{\log(d) + \log\left(\frac{1}{\delta}\right)}.$$

This result deserves various comments. The first point is that the bound provided by Proposition 2.5 is a so-called PAC bound (where PAC stands for "probably approximately correct"; see Section 10.3), *i.e.* it is a bound which holds with a high probability for the realizations under consideration. It can be considered as more useful than bounds in average as provided by Propositions 2.2 and 2.4 for ordinary least square regression and ridge regression, respectively, and [4, Exercise 8.8] for Lasso regression.

The dimension explicitly appears only through the a factor  $\sqrt{\log(d)}$ , which is a quite mild dependence, in particular if  $n \gg \log(d)$ . Note also that  $\theta_{\star}$  needs to be sparse for the estimate to be useful (*i.e.* in order for  $\|\theta_{\star}\|_1$  not to change as the dimension d of  $\varphi(x)$  is increased).

Various extensions are possible, in particular obtaining bounds with "fast rates" O(1/n) for the average excess risk (see [4, Sections 8.3.3-8.3.4]), and extending the analysis to the random design setting (see [4, Section 8.3.5]).

*Proof.* For any  $1 \leq k \leq d$ , the random variable  $(X^{\top}\varepsilon)_k$  is a Gaussian random variable with mean 0 and variance  $n\sigma^2 \widehat{\Sigma}_{kk}$ . Recall that, for  $Z \sim \mathcal{N}(0,1)$  and  $t \geq 0$ , it holds  $\mathbb{P}(|Z| \geq t) \leq e^{-t^2/2}$  (see Exercise 2.17 below). Using the latter inequality, and the union bound to control  $\|X^{\top}\varepsilon\|_{\infty}$ ,

$$\mathbb{P}\left(\Omega^*\left(X^{\top}\varepsilon\right) > \frac{n\lambda}{2}\right) = \mathbb{P}\left(\left\|X^{\top}\varepsilon\right\|_{\infty} > \frac{n\lambda}{2}\right)$$
$$\leqslant \sum_{k=1}^{d} \mathbb{P}\left(\left|\left(X^{\top}\varepsilon\right)_{k}\right| > \frac{n\lambda}{2}\right) = \sum_{k=1}^{d} \mathbb{P}\left(\frac{\left|\left(X^{\top}\varepsilon\right)_{k}\right|}{\sqrt{n\sigma^{2}\widehat{\Sigma}_{kk}}} > \frac{\lambda\sqrt{n}}{2\sigma\sqrt{\widehat{\Sigma}_{kk}}}\right)$$
$$\leqslant \sum_{k=1}^{d} \exp\left(-\frac{n\lambda^{2}}{8\sigma^{2}\widehat{\Sigma}_{kk}}\right) \leqslant d\exp\left(-\frac{n\lambda^{2}}{8\sigma^{2}\|\widehat{\Sigma}\|_{\infty}}\right).$$

The right hand side of the last inequality is equal to  $\delta$  for

$$\frac{n\lambda^2}{8\sigma^2 \|\widehat{\Sigma}\|_{\infty}} = \log\left(\frac{d}{\delta}\right),\,$$

which leads to (2.27). Now, when  $\Omega^*(X^{\top}\varepsilon) \leq n\lambda/2$  (which happens with a probability larger or equal than  $1-\delta$ ), we can apply item (a) of Lemma 2.1 to obtain the claimed bound on the excess risk. 

**Exercise 2.17.** Prove that, for  $Z \sim \mathcal{N}(0,1)$  and  $t \ge 0$ , it holds  $\mathbb{P}(|Z| \ge t) \le e^{-t^2/2}$ .

**Correction.** Note first that  $\mathbb{P}(|Z| \ge t) \le 2\mathbb{P}(Z \ge t)$ . Now, using that  $e^{-st} \le 1$  for  $s, t \ge 0$ , **Correction.** Note first that  $\mathbb{P}(|Z| \ge t) \le 2\mathbb{P}(Z \ge t)$ . Now, using that  $e^{-st} \le 1$  for  $s, t \ge 0$ ,  $\mathbb{P}(Z \ge t) = \frac{1}{\sqrt{2\pi}} \int_{t}^{+\infty} e^{-s^{2}/2} ds = \frac{1}{\sqrt{2\pi}} \int_{0}^{+\infty} e^{-(s+t)^{2}/2} ds = e^{-t^{2}/2} \int_{0}^{+\infty} e^{-st} e^{-s^{2}/2} \frac{ds}{\sqrt{2\pi}}$   $\le e^{-t^{2}/2} \int_{0}^{+\infty} e^{-s^{2}/2} \frac{ds}{\sqrt{2\pi}} = \frac{1}{2} e^{-t^{2}/2},$ from which the desired upper bound follows.

# Classification with logistic regression

3.1	<b>Conve</b> 3.1.1	exification of the lossConvex surrogates and $\Phi$ -risk3.1.1.1Classification with real valued functions3.1.1.2Convex surrogates	<b>53</b> 53 54 54
	319	3.1.1.3 $\Phi$ -risk	$55 \\ 56$
3.2	5.1.2 Logist 3.2.1 3.2.2	Binary logistic regression as a discriminative probabilistic model Binary logistic regression through a classification-calibrated convex	<b>61</b> 61
	$3.2.3 \\ 3.2.4$	$\begin{array}{l} \text{surrogate} \\ \text{Estimating } \theta \text{ for prediction} \\ \text{Multiclass logistic regression} \\ \end{array}$	62 63 64

We discuss in this chapter show how to perform classification with a method called logistic regression. The presentation mainly focuses on binary classification, but we also mention how to perform muticlass classification. We start by providing in Section 3.1 a mathematical framework which allows to conviently minimize loss functions to perform classification; and then describe in Section 3.2 a method based on a linear model to actually perform classification.

# 3.1 Convexification of the loss

For simplicity of exposition, we restrict ourselves to binary classification on  $\mathcal{Y} = \{-1, 1\}$  in this section. The loss function which naturally apppears in classification problems (see Section 1.2.1.1) is the 0-1 loss. Minimizing the associated training loss is however a difficult problem as it is combinatorial in nature (one needs to consider the  $2^n$  possibilities for the labels for binary classification for instance). We show in this section how the 0-1 loss for classification problems can be replaced by a surrogate convex loss function which upper bounds the 0-1 loss, and admits the same Bayes predictor as for the 0-1 loss. The interest of convex losses is that their optimization is easier to perform (see Chapter 4). Our presentation is based on [4, Section 4.1] and [39, Section 4.7].

Two main questions can be asked at this stage:

- which convex surrogate functions are admissible? This is discussed in Section 3.1.1;
- how does the risk associated with the convex surrogate inform us on the original risk? This is made precise in Section 3.1.2.

# 3.1.1 Convex surrogates and $\Phi$ -risk

We discuss in this section how to perform classification with real valued functions (see Section 3.1.1.1), in order to replace the 0-1 loss by a convex upper bound  $\Phi$  (see Section 3.1.1.2). Admissible convex upper bounds should have the same Bayes predictor as the original problem, which motivates the concept of "classification calibrated" convex surrogates  $\Phi$  (see Section 3.1.1.3).

# 3.1.1.1 Classification with real valued functions

Consider a real valued function  $g: \mathcal{X} \to \mathbb{R}$ . Denoting by  $\mathcal{Y} = \{-1, 1\}$ , binary classification can be performed using a function  $f: \mathcal{X} \to \mathcal{Y}$  defined from g as f = sign(g), *i.e.* 

$$f(x) = \begin{vmatrix} 1 & \text{if } g(x) > 0, \\ -1 & \text{if } g(x) < 0, \\ 0 & \text{if } g(x) = 0. \end{vmatrix}$$

Note that there is some arbitrariness in the choice of values for f(0). We consider here f(0) = 0 so that the prediction is always wrong for g(x) = 0 (note that, strictly speaking, one would need to extend the output space to  $\{-1, 0, 1\}$  instead of  $\{-1, 1\}$ ). Another option would be to choose at random whether f(x) = 1 or -1 when g(x) = 0.

With some abuse of notation, we denote by  $\mathcal{R}(g)$  the risk associated with g:

$$\mathcal{R}(g) = \mathbb{E}\left[\mathbf{1}_{f(X)\neq Y}\right], \qquad f = \operatorname{sign}(g). \tag{3.1}$$

In order to make contact with generalizations of the latter risk, we rewrite it as

$$\mathcal{R}(g) = \mathbb{P}\left[\operatorname{sign}(g(X)) \neq Y\right] = \mathbb{E}\left[\mathbf{1}_{Yg(X)} \leqslant 0\right] = \mathbb{E}\left[\Phi_{0-1}(Yg(X))\right],\tag{3.2}$$

where

$$\Phi_{0-1}(u) = \mathbf{1}_{u \leqslant 0}.$$

There are infinitely many Bayes predictors associated with the original risk (3.1). To derive one, we introduce, as in Section 1.2.2.1,

$$\eta(x') = \mathbb{P}(Y = 1 | X = x').$$
(3.3)

Then,

$$\mathcal{R}(g) = \mathbb{E}\left[\eta(X)\mathbf{1}_{g(X)\leqslant 0} + (1-\eta(X))\mathbf{1}_{g(X)\geqslant 0}\right]$$

A Bayes predictor  $g^*$  is therefore characterized by the conditions  $g^*(x) \leq 0$  if  $\eta(x) \leq 1/2$ and  $g^*(x) \geq 0$  if  $\eta(x) \geq 1/2$ . One possible Bayes predictor is

$$g^{\star}(x) = \eta(x) - \frac{1}{2}.$$

**Remark 3.1.** Note that the Bayes predictor  $f^*$  for the 0-1 loss is  $sign(g^*)$ , in accordance with the result of Exercise 1.6.

#### 3.1.1.2 Convex surrogates

The key idea behind considering convex surrogates is to replace the function  $\Phi_{0-1}$  appearing in the reformulation (3.2) of the original loss by a convex function  $\Phi : \mathbb{R} \to \mathbb{R}_+$  and to minimize the  $\Phi$ -risk

$$\mathcal{R}_{\Phi}(g) = \mathbb{E}\left[\Phi(Yg(X))\right]. \tag{3.4}$$

The function  $\Phi$  should be an upper bound of the 0-1 loss:

$$0 \leqslant \Phi_{0-1} \leqslant \Phi. \tag{3.5}$$

Let us give three important examples:

• the quadratic loss  $\Phi(u) = (u-1)^2$ : in this case  $\Phi(yg(x)) = (y-g(x))^2$  since  $y^2 = 1$ . This situation corresponds to the standard regression paradigm, where the prediction is based on the sign of the function  $g : \mathcal{X} \to \mathcal{Y}$  appearing in the regression problem. One possible issue of this approach is the overpenalization of large values of |yg(x)|;

• the logistic loss  $\Phi(u) = \log(1 + e^{-u})/\log(2)$  (see Exercise 3.1 below): in this case,

$$\Phi(yg(x)) = -\frac{1}{\log(2)} \log\left(\frac{1}{1 + e^{-yg(x)}}\right) = -\frac{1}{\log(2)} \log\sigma(yg(x)), \tag{3.6}$$

where  $\sigma$  is the sigmoid function

$$\sigma(v) = \frac{1}{1 + e^{-v}}.$$
(3.7)

This is the convex surrogate used in Section 3.2 for logistic regression, without the normalization factor log(2);

• the hinge loss  $\Phi(u) = \max(1 - u, 0)$  used in Chapter 6 for classification with support vector machines. The squared hinge loss  $\Phi(u) = \max(1 - u, 0)^2$  is sometimes used to have a smoother model.

**Exercise 3.1.** Make precise the asymptotic behavior of the function  $\Psi(u) = \log(1 + e^{-u})$ .

Correction. Straightforward computations give

$$\Psi(u) \underset{u \to -\infty}{\sim} -u, \qquad \Psi(u) \underset{u \to +\infty}{\sim} e^{-u}.$$

#### 3.1.1.3 *P*-risk

The  $\Phi$ -risk (3.4) can be rewritten as (by taking conditional expectations on Y)

$$\mathcal{R}_{\Phi}(g) = \mathbb{E}\left[\Phi(Yg(X))\right] = \mathbb{E}\left[\eta(X)\Phi(g(X)) + (1 - \eta(X))\Phi(-g(X))\right]$$

Therefore,

$$\mathcal{R}_{\Phi}(g) = \mathbb{E}\left[C_{\eta(X)}(g(X))\right],\tag{3.8}$$

where we recall the definition (3.3) of  $\eta$ , and

$$C_{\zeta}(\alpha) = \zeta \Phi(\alpha) + (1 - \zeta) \Phi(-\alpha).$$

A first requirement on the convex surrogate  $\Phi$  is that the Bayes predictor for the  $\Phi$ -risk is the same as for the original risk  $\mathcal{R}$ . The Bayes predictor for the  $\Phi$ -risk is characterized as follows:

$$\forall x \in \mathcal{X}, \qquad g_{\Phi}^{\star}(x) \in \operatorname*{argmin}_{\alpha \in \mathbb{R}} \Big\{ \eta(x)\Phi(\alpha) + (1-\eta(x))\Phi(-\alpha) \Big\}.$$
(3.9)

The fact that this Bayes predictor coincides with the Bayes predictor for the original risk translates into the following requirements:

$$\zeta > \frac{1}{2} \iff \underset{\alpha \in \mathbb{R}}{\operatorname{argmin}} C_{\zeta}(\alpha) \subset (0, +\infty),$$
  

$$\zeta < \frac{1}{2} \iff \underset{\alpha \in \mathbb{R}}{\operatorname{argmin}} C_{\zeta}(\alpha) \subset (-\infty, 0),$$
(3.10)

so that the prediction, which is given by the sign of the argmin, is 1 when  $\zeta > 1/2$  and -1 when  $\zeta < 1/2$ . A function  $\Phi$  which satisfies these requirements is said to be *classification-calibrated*. The next proposition gives simple sufficient conditions to this end when  $\Phi$  is convex. Note that the convexity of  $\Phi$  implies that  $C_{\zeta}$  is convex for any  $0 \leq \zeta \leq 1$ .

**Proposition 3.1.** Consider a convex function  $\Phi : \mathbb{R} \to \mathbb{R}$ . The function  $\Phi$  is classificationcalibrated (i.e. (3.10) is satisfied) if and only if  $\Phi$  is differentiable at 0 and  $\Phi'(0) < 0$ .

#### 56 3 Classification with logistic regression

*Proof.* Since  $C_{\zeta}$  is a convex function from  $\mathbb{R}$  to  $\mathbb{R}$ , it is continuous and admits left/right derivatives at any point (see Exercise 3.2). Therefore, the minimum of  $C_{\zeta}$  is attained on  $(0, +\infty)$  if and only if  $C'_{\zeta}(0^+) < 0$ ; and attained on  $(-\infty, 0)$  if and only if  $C'_{\zeta}(0^-) > 0$ . This shows that (3.10) is equivalent to

$$\zeta > \frac{1}{2} \iff C'_{\zeta}(0^{+}) = \zeta \Phi'(0^{+}) - (1 - \zeta) \Phi'(0^{-}) < 0,$$
  

$$\zeta < \frac{1}{2} \iff C'_{\zeta}(0^{-}) = \zeta \Phi'(0^{-}) - (1 - \zeta) \Phi'(0^{+}) > 0.$$
(3.11)

When  $\Phi$  is differentiable at 0 and  $\Phi'(0) < 0$ , then  $C_{\zeta}$  is also differentiable at 0, with  $C'_{\zeta}(0) = (2\zeta - 1)\Phi'(0)$ . The equivalences in (3.11) are then easily seen to hold.

Conversely, assume that  $\Phi$  is classification-calibrated, and hence that (3.11) holds. By considering a sequence  $(\zeta_n)_{n\geq 0} \subset (1/2, 1]$  converging to 1/2, it follows, by passing to the limit in the first line of (3.11) that  $\Phi'(0^+) - \Phi'(0^-) \leq 0$ . Now, since  $\Phi$  is convex, it also holds that  $\Phi'(0^+) \geq \Phi'(0^-)$ , so that one can conclude that  $\Phi'(0^+) = \Phi'(0^-)$ . This prove that  $\Phi$  is differentiable at 0, and also that  $C'_{\zeta}(0) = (2\zeta - 1)\Phi'(0)$ . Moreover,  $\Phi'(0) < 0$  in order to meet the conditions in (3.11). This allows to conclude the proof of the claimed equivalence.

**Exercise 3.2.** Consider a convex function  $\Phi : \mathbb{R} \to \mathbb{R}$ . We prove in this exercise that the right derivative is well defined. The fact that the left derivative is well defined can be obtained by similar estimates.

(i) Fix  $x_0, x_1, x_2 \in \mathbb{R}$  with  $x_0 < x_1 < x_2$ . By considering  $x_1$  as a convex combination of  $x_0, x_2$ , prove that

$$\frac{\Phi(x_1) - \Phi(x_0)}{x_1 - x_0} \leqslant \frac{\Phi(x_2) - \Phi(x_0)}{x_2 - x_0}.$$

- (ii) Prove that the left hand side of the previous inequality is bounded from below uniformly in  $x_1 \in (0, x_2]$ .
- (iii) Deduce that the right derivative of  $\Phi$  at  $x_0$  is well defined.

**Correction.** (i) Note that

$$x_1 = (1 - \alpha)x_0 + \alpha x_2, \qquad \alpha = \frac{x_1 - x_0}{x_2 - x_0}, \qquad 1 - \alpha = \frac{x_2 - x_1}{x_2 - x_0}.$$

Now, the inequality to be proved is equivalent to

$$(x_2 - x_0)\Phi(x_1) \leqslant (x_1 - x_0) \left[\Phi(x_2) - \Phi(x_0)\right] + \Phi(x_0)(x_2 - x_0),$$

which corresponds to the inequality  $\Phi(x_1) \leq (1-\alpha)\Phi(x_0) + \alpha\Phi(x_2)$  by dividing both sides by  $x_2 - x_0$ .

(ii) Fix  $x_{\star} < x_0$ . Then, using the result of the previous question,

$$\frac{\Phi(x_1) - \Phi(x_0)}{x_1 - x_0} \ge \frac{\Phi(x_0) - \Phi(x_\star)}{x_0 - x_\star} > -\infty,$$

which provides the desired uniform bound.

(iii) The function  $x_1 \mapsto (\Phi(x_1) - \Phi(x_0))/(x_1 - x_0)$  is nondecreasing and bounded from below. It therefore admits a limit when  $x_1 \to x_0$  with  $x_1 > x_0$ . This limit is, by definition, the right derivative of  $\Phi$  at  $x_0$ .

# 3.1.2 Relationship between original risk and $\Phi$ -risk

We have seen in Section 3.1.1.3 that the Bayes predictors for the  $\Phi$ -risk and the original risk are the same. In practice, one minimizes the  $\Phi$ -risk associated with the (classification calibrated) convex

surrogate  $\Phi$ , but is ultimately interested in performance guarantees on the original risk, based on the 0-1 loss. This motivates exploring the relationship between these two risks. Ideally, one would want an inequality of the form

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant H\left(\mathcal{R}_{\Phi}(g) - \mathcal{R}_{\Phi}^{\star}\right),$$

where  $\mathcal{R}_{\Phi}^{\star}$  is the risk associated with any Bayes predictor for  $\mathcal{R}_{\Phi}$ , and H is some increasing function such that H(0) = 0, called *calibration function*. In order to establish such as a result, we reformulate the excess risk using the Bayes predictor, and then upper bound this quantity with the  $\Phi$ -risk.

**Lemma 3.1.** For any Bayes predictor  $g^* : \mathcal{X} \to \mathbb{R}$  and any function  $g : \mathcal{X} \to \mathbb{R}$  such that  $\mathbb{P}(g(X) = 0) = 0$ , it holds

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant \mathbb{E}\left[|2\eta(X) - 1| \mathbf{1}_{g(X)g^{\star}(X)\leqslant 0}\right].$$
(3.12)

The condition  $\mathbb{P}(g(X) = 0) = 0$  is a technical condition related to our choice for sign(0). Other conventions can be considered, in which case the result can be slightly different (see for instance [39, Lemma 4.5]).

*Proof.* We rewrite the excess risk using conditional expectations as

$$\mathcal{R}(g) - \mathcal{R}^{\star} = \mathbb{E}\left[\mathbb{E}\left(\mathbf{1}_{Y \neq \operatorname{sign}(g(X))} - \mathbf{1}_{Y \neq \operatorname{sign}(g^{\star}(X))} \mid X\right)\right],$$

and then consider the situations where the predictions are different, namely  $\eta(X) > 1/2$  (*i.e.*  $g^{\star}(X) > 0$ ) and  $g(X) \leq 0$ , or  $\eta(X) < 1/2$  (*i.e.*  $g^{\star}(X) < 0$ ) and  $g(X) \geq 0$ . We discard the case when  $\eta(X) = 1/2$  since  $g^{\star}(X) = 0$  in this case and therefore  $\mathbf{1}_{Y \neq \operatorname{sign}(g(X))} - \mathbf{1}_{Y \neq \operatorname{sign}(g^{\star}(X))} = \mathbf{1}_{Y \neq \operatorname{sign}(g(X))} - \mathbf{1} \leq 0$ .

Since  $\mathbb{P}(g(X) = 0) = 0$ , it suffices to treat the case when  $g(X) \neq 0$ . In the first case  $(\eta(X) > 1/2)$  and g(X) < 0, the conditional expectation is

$$\mathbb{E}\left(\mathbf{1}_{Y\neq\operatorname{sign}(g(X))} - \mathbf{1}_{Y\neq\operatorname{sign}(g^{\star}(X))} \middle| X\right) = \mathbb{P}(Y\neq-1 \,|\, X) - \mathbb{P}(Y\neq1 \,|\, X) = 2\eta(X) - 1.$$

Similarly, in the second case  $(\eta(X) < 1/2 \text{ and } g(X) > 0)$ , the conditional expectation is equal to  $1 - 2\eta(X)$ . Overall, this can be summarized as

$$\mathbb{E}\left(\mathbf{1}_{Y\neq\operatorname{sign}(g(X))} - \mathbf{1}_{Y\neq\operatorname{sign}(g^{\star}(X))} \mid X\right) \mathbf{1}_{g(X)\neq 0} \mathbf{1}_{g^{\star}(X)\neq 0} = |2\eta(X) - 1| \mathbf{1}_{g(X)g^{\star}(X)<0} \mathbf{1}_{g(X)\neq 0} \mathbf{1}_{g^{\star}(X)\neq 0},$$

which leads to the claimed estimate by taking the expectation over X.

**Exercise 3.3.** Under the same conditions as in Lemma 3.1, prove that  $0 \leq \mathcal{R}(g) - \mathcal{R}^* \leq \mathbb{E}[|2\eta(X) - 1 - g(X)|].$ 

**Correction.** We rely on (3.12). It suffices to consider the case when g and  $g^*$  do not have the same signs, *i.e.*  $\eta(X) > 1/2$  and g(X) < 0, or  $\eta(X) < 1/2$  and g(X) > 0. In the former case,

$$|2\eta(X) - 1| = 2\eta(X) - 1 \leq 2\eta(X) - 1 - g(X) = |2\eta(X) - 1 - g(X)|;$$

while, in the latter case,

$$|2\eta(X) - 1| = 1 - 2\eta(X) \leq 1 - 2\eta(X) + g(X) = |2\eta(X) - 1 - g(X)|.$$

The upper bound of Lemma 3.1 can now be leveraged to obtain bounds on the excess risk in terms of the excess  $\Phi$ -risk, by adapting [39, Theorem 4.7].

#### 583 Classification with logistic regression

**Theorem 3.1.** Consider a convex function  $\Phi : \mathbb{R} \to \mathbb{R}$ , which is nonincreasing on  $\mathbb{R}_{-}$ , and assume that there exist  $s \ge 1$  and  $c \in \mathbb{R}_+$  such that

$$\forall x \in \mathcal{X}, \qquad \left| \eta(x) - \frac{1}{2} \right|^s \leqslant c^s \left[ \Phi(0) - C_{\eta(x)} \left( g_{\Phi}^{\star}(x) \right) \right], \tag{3.13}$$

for some Bayes predictor  $g_{\Phi}^{\star}$  associated with  $\mathcal{R}_{\Phi}$  (recall (3.9)). Then, for any  $g: \mathcal{X} \to \mathbb{R}$  such that  $\mathbb{P}(g(X) = 0) = 0$ ,

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant 2c \left( \mathcal{R}_{\Phi}(g) - \mathcal{R}_{\Phi}^{\star} \right)^{1/s}.$$

We refer to Exercise 3.5 below for examples of functions  $\Phi$  satisfying the assumptions of this theorem. See also Exercise 3.4 for checking that the right hand side of (3.13) is indeed nonnegative.

*Proof.* We start from (3.12) and then use Jensen's inequality to write

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant \mathbb{E}\left[|2\eta(X) - 1|\mathbf{1}_{g(X)g^{\star}(X)\leqslant 0}\right] \leqslant \left(\mathbb{E}\left[|2\eta(X) - 1|^{s}\mathbf{1}_{g(X)g^{\star}(X)\leqslant 0}\right]\right)^{1/s}.$$

We next upper bound the latter right hand side using (3.13) to obtain

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant 2c \left( \mathbb{E} \left[ \left( \Phi(0) - C_{\eta(x)} \left( g_{\Phi}^{\star}(x) \right) \right) \mathbf{1}_{g(X)g^{\star}(X) \leqslant 0} \right] \right)^{1/s}.$$

We claim that

$$\Phi(0)\mathbf{1}_{g(X)g^{\star}(X)\leqslant 0} \leqslant C_{\eta(x)}(g(x))\mathbf{1}_{g(X)g^{\star}(X)\leqslant 0}.$$
(3.14)

If this is indeed the case, then, since  $C_{\eta(x)}(g(x)) - C_{\eta(x)}(g_{\Phi}^{\star}(x)) \ge 0$  by definition of  $g_{\Phi}^{\star}$ ,

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant 2c \left( \mathbb{E} \left[ \left( C_{\eta(x)}(g(x)) - C_{\eta(x)}\left(g_{\Phi}^{\star}(x)\right) \right) \mathbf{1}_{g(X)g^{\star}(X)\leqslant 0} \right] \right)^{1/s} \\ \leqslant 2c \left( \mathbb{E} \left[ C_{\eta(x)}(g(x)) - C_{\eta(x)}\left(g_{\Phi}^{\star}(x)\right) \right] \right)^{1/s} = 2c \left( \mathcal{R}_{\Phi}(g) - \mathcal{R}_{\Phi}^{\star} \right)^{1/s},$$

where the last equality follows from (3.8).

In order to conclude the proof, it remains to show (3.14). Choosing the Bayes predictor  $g^{\star}(x) =$  $\eta(x) - 1/2$ , and since  $\Phi$  is nonincreasing on  $\mathbb{R}_-$ , we obtain, for any  $x \in \mathcal{X}$  such that  $g(x)g^{\star}(x) \leq 0$ ,

$$\begin{split} \varPhi(0) &\leqslant \varPhi(2g(x)g^{\star}(x)) = \varPhi\left([2\eta(x) - 1]g(x)\right) = \varPhi\left(\eta(x)g(x) + (1 - \eta(x))(-g(x))\right) \\ &\leqslant \eta(x)\varPhi(g(x)) + (1 - \eta(x))\varPhi(-g(x)) = C_{\eta(x)}(g(x)), \end{split}$$

where we used the convexity of  $\Phi$  to pass to the second line. This gives the claimed bound (3.14) and allows to conclude the proof. Π

**Exercise 3.4.** Prove that  $\Phi(0) - C_{\eta(x)}(g^{\star}_{\Phi}(x)) \ge 0$  for any  $x \in \mathcal{X}$ .

**Correction.** In view of the definition of  $g_{\Phi}^{\star}(x)$  as the minimizer of  $\alpha \mapsto C_{\eta(x)}(\alpha)$ , it holds  $C_{\eta(x)}(g_{\Phi}^{\star}(x)) \leq C_{\eta(x)}(0) = \Phi(0)$ , from which the result directly follows.

$$C_{\eta(x)}\left(g_{\Phi}^{\star}(x)\right) \leqslant C_{\eta(x)}(0) = \Phi(0),$$

**Exercise 3.5.** Prove that the following functions are classification-calibrated and satisfy the assumptions of Theorem 3.1:

(1)  $\Phi(u) = (1-u)^2;$ (2)  $\Phi(u) = \max(1 - u, 0);$ (3)  $\Phi(u) = \log(1 + e^{-u})$  (note that we are not normalizing  $\Phi$  here in order for (3.5) to hold). For the latter case, introduce the function  $\psi: [0,1] \to \mathbb{R}$  defined as

$$\psi(t) = \log(2) + t\log(t) + (1-t)\log(1-t) - \frac{1}{2}(2t-1)^2$$

and note that  $\psi(t) \ge 0$  for all  $t \in [0, 1]$ .

**Correction.** The three functions  $\Phi$  which are considered are obviously nonincreasing on  $\mathbb{R}_-$ . The first and last ones are clearly convex (for instance because they are smooth and their respective derivatives  $u \mapsto 2(u-1)$  and  $u \mapsto -1 + 1/(1 + e^{-u})$  are increasing). The second one is convex as the maximum of the convex functions  $u \mapsto 1-u$  and  $u \mapsto 0$ . Indeed, consider  $\Phi(u) = \max(\Phi_1(u), \Phi_2(u))$  with  $\Phi_1, \Phi_2 : \mathbb{R} \to \mathbb{R}$  convex. Fix  $u, v \in \mathbb{R}$  and  $t \in [0, 1]$ . Then, by convexity of  $\Phi_1$ ,

$$\Phi_1((1-t)u + tv) \leqslant (1-t)\Phi_1(u) + t\Phi_1(v) \leqslant (1-t)\Phi(u) + t\Phi(v).$$

A similar inequality holds for  $\Phi_2$ , so that

$$\begin{aligned} \varPhi((1-t)u + tv) &= \max \left\{ \varPhi_1((1-t)u + tv), \varPhi_2((1-t)u + tv) \right\} \\ &\leq (1-t)\varPhi_1(u) + t\varPhi_1(v) \leqslant (1-t)\varPhi(u) + t\varPhi(v), \end{aligned}$$

which allows to conclude that  $\Phi$  is convex.

The three functions are also easily seen to be classification-calibrated in application of Proposition 3.1 as they are all smooth around 0 with  $\Phi'(0) < 0$ . It therefore suffice to check that (3.13) is satisfied. This requires to first determine some Bayes predictor, based on the characterization (3.9).

(1) For  $\Phi(u) = (1-u)^2$ , a Bayes predictor is

$$g_{\varPhi}^{\star}(x) \in \operatorname*{argmin}_{\alpha \in \mathbb{R}} \Big\{ \eta(x)(1-\alpha)^2 + (1-\eta(x))(1+\alpha)^2 \Big\} = \operatorname*{argmin}_{\alpha \in \mathbb{R}} \Big\{ \alpha^2 - 2\alpha(2\eta(x)-1) \Big\},$$

so that  $g_{\Phi}^{\star}(x) = 2\eta(x) - 1$ . Therefore,

$$\Phi(0) - C_{\eta(x)} \left(g_{\Phi}^{\star}(x)\right) = 1 - 4\eta(x)(1 - \eta(x))^2 - 4(1 - \eta(x))\eta(x)^2$$
  
= 1 - 4\eta(x)(1 - \eta(x))[1 - \eta(x) + \eta(x)] = (2\eta(x) - 1)^2,

which shows that (3.13) holds with s = 2 and c = 1/2. This implies that

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant \sqrt{\mathcal{R}_{\Phi}(g) - \mathcal{R}_{\Phi}^{\star}}.$$

(2) For  $\Phi(u) = \max(1 - u, 0)$ , a Bayes predictor is

$$g_{\varPhi}^{\star}(x) \in \operatorname*{argmin}_{\alpha \in \mathbb{R}} \Big\{ \eta(x) \max(1-\alpha, 0) + (1-\eta(x)) \max(1+\alpha, 0) \Big\}.$$

The argument of the minimum is a continuous piecewise affine function, equal to  $\eta(x)(1-\alpha)$  for  $\alpha \leq -1$  and  $(1 - \eta(x))(1 + \alpha)$  for  $\alpha \geq 1$ . The minimum of this function is therefore attained either at  $\alpha = -1$  or  $\alpha = 1$ . In fact,

$$g_{\Phi}^{\star}(x) = \begin{vmatrix} -1 & \text{if } \eta(x) < 1/2 \\ 1 & \text{if } \eta(x) > 1/2 \end{vmatrix}$$

while  $g_{\Phi}^{\star}(x)$  can be any number in [-1,1] if  $\eta(x) = 1/2$ . Therefore,

$$\begin{split} \varPhi(0) - C_{\eta(x)} \left( g_{\varPhi}^{\star}(x) \right) &= 1 - \eta(x) \max \left( 1 - g_{\varPhi}^{\star}(x), 0 \right) - \left( 1 - \eta(x) \right) \max \left( 1 + g_{\varPhi}^{\star}(x), 0 \right) \\ &= \begin{vmatrix} 1 - 2\eta(x) & \text{if } g_{\varPhi}^{\star}(x) = -1 \text{ } i.e. \ \eta(x) < 1/2, \\ 2\eta(x) - 1 & \text{if } g_{\varPhi}^{\star}(x) = 1 \text{ } i.e. \ \eta(x) > 1/2, \\ 0 & \text{if } \eta(x) = 1/2. \end{vmatrix}$$

# 60 3 Classification with logistic regression

This means that  $\Phi(0) - C_{\eta(x)}(g_{\Phi}^{\star}(x)) = |2\eta(x) - 1|$ , which shows that (3.13) holds with s = 1 and c = 1/2. This implies that

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant \mathcal{R}_{\Phi}(g) - \mathcal{R}_{\Phi}^{\star}$$

(3) For  $\Phi(u) = \log(1 + e^{-u})$ , a Bayes predictor is

$$g_{\Phi}^{\star}(x) \in \operatorname*{argmin}_{\alpha \in \mathbb{R}} \Big\{ \eta(x) \log \left( 1 + \mathrm{e}^{-\alpha} \right) + (1 - \eta(x)) \log \left( 1 + \mathrm{e}^{\alpha} \right) \Big\}.$$

We consider  $\eta(x = \notin \{0, 1\}$ . The function  $\alpha \mapsto C_{\eta(x)}(\alpha)$  to be minimized is smooth and coercive (as it behaves as  $-\eta(x)\alpha$  for  $\alpha \to -\infty$  and  $(1 - \eta(x))\alpha$  for  $\alpha \to +\infty$ ). There exists therefore a global minimizer. The optimality condition for a minimizer reads

$$0 = C'_{\eta(x)}(\alpha) = -\eta(x)\frac{e^{-\alpha}}{1 + e^{-\alpha}} + (1 - \eta(x))\frac{e^{\alpha}}{1 + e^{\alpha}} = \frac{-\eta(x)e^{-\alpha} + 1 - \eta(x)}{1 + e^{-\alpha}},$$

so that

$$\mathrm{e}^{-\alpha} = \frac{1 - \eta(x)}{\eta(x)},$$

and finally

$$g_{\varPhi}^{\star}(x) = \log\left(\frac{\eta(x)}{1-\eta(x)}\right)$$

Therefore, using that the function  $\psi$  has nonnegative values,

$$\begin{split} \varPhi(0) - C_{\eta(x)}\left(g_{\varPhi}^{\star}(x)\right) &= \log(2) - \eta(x)\log\left(1 + \frac{1 - \eta(x)}{\eta(x)}\right) - (1 - \eta(x))\log\left(1 + \frac{\eta(x)}{1 - \eta(x)}\right) \\ &= \log(2) + \eta(x)\log\eta(x) + (1 - \eta(x))\log(1 - \eta(x)) \\ &\geqslant \frac{1}{2}(2\eta(x) - 1)^2 = 2\left(\eta(x) - \frac{1}{2}\right)^2, \end{split}$$

which shows that (3.13) holds with s = 2 and  $c = 1/\sqrt{2}$ . This implies that

$$0 \leqslant \mathcal{R}(g) - \mathcal{R}^{\star} \leqslant \sqrt{2\left(\mathcal{R}_{\varPhi}(g) - \mathcal{R}_{\varPhi}^{\star}\right)}.$$

To prove that  $\psi(t) \ge 0$  for  $t \in [0, 1]$ , we note that  $\psi(1/2) = 0$ , and next that

$$\psi'(t) = \log(t) - \log(1-t) - 4t + 2,$$

so  $\psi'(1/2) = 0$ . Moreover,

$$\psi''(t) = \frac{1}{t} + \frac{1}{1-t} - 4 = \frac{1}{t(1-t)} - 4 \ge 0.$$

This shows that  $\psi$  is convex, hence it is above its tangent at t = 1/2, the latter tangent being the horizontal line. This allows to obtained the claimed nonnegativity of the function.

**Impact on approximation and estimation errors.** Let us conclude this section by discussing two points and issues raised by the estimate of Theorem 3.1:

• For the same classification problem, several convex surrogates can be used... The predictor  $g_{\Phi}^{\star}$  and the associated classification function  $f_{\Phi}^{\star} = \text{sign}(g_{\Phi}^{\star})$  will however be different. This may have an impact on the approximation error via the class of functions which are considered (see Section 10.1).

• Note that s = 2 for smooth losses (a rather generic fact, see [4, Section 4.1.4]); while s = 1 for a non smooth surrogate such as the hinge loss. This means that the possibly nice convergence rate for  $\mathcal{R}_{\Phi}(g) - \mathcal{R}_{\Phi}^{\star}$  will be degraded by the application of the square-root when considering the prediction function  $f = \operatorname{sign}(g)$ . On the other hand, smooth losses are easier to optimize, so that there is a trade-off to be found here.

# 3.2 Logistic regression

The presentation of logistic regression is based mostly on [40, Chapter 10] and [8, Section 4.3]. There are two ways to introduce the method for binary classification, by motivating it first as a discriminative probabilistic model (see Section 3.2.1), and then connecting this heuristic approach to the framework of Section 3.1 (see Section 3.2.2). In both case, predictions are based on a nonlinear function combined with an affine one, whose parameter need to be estimated, as we discuss in Section 3.2.3. We conclude the presentation in Section 3.2.4 by showing how to perform multiclass classification.

**Remark 3.2.** Various extensions of the basic logistic regression can be found in the literature. See for instance [40, Chapter 10] for discussions on robust logistic regression (when outliers are present in the data set), Bayesian logistic regression (where the a posteriori distribution of  $\theta$  is computed in order to quantify the uncertainty in the predictions), hierarchical classification, ...

## 3.2.1 Binary logistic regression as a discriminative probabilistic model

The aim of discriminative probabilistic models<sup>1</sup> is to learn the conditional distributions of the data  $\mathbb{P}(Y = c \mid X = x)$  for  $c \in \mathcal{Y}$ . We consider here binary classification with  $\mathcal{Y} = \{0, 1\}$ .

The idea behind logistic regression is to approximate the conditional distribution as

$$p_{\theta}(y=1 \mid x) = \sigma \left( w^{\top} x + b \right), \qquad \theta = (w,b) \in \mathbb{R}^{d} \times \mathbb{R}, \tag{3.15}$$

where we recall that  $\sigma$  is the sigmoid function (3.7). Note that  $p_{\theta}(y = 1 | x) \in [0, 1]$  thanks to the properties of the sigmoid function. A linear classifier is then obtained by predicting the label associated with the largest estimated conditional probability:

$$f_{\theta}(x) = \begin{vmatrix} 1 & \text{if } p_{\theta}(y=1 \mid x) > p_{\theta}(y=0 \mid x), \\ 0 & \text{if } p_{\theta}(y=1 \mid x) < p_{\theta}(y=0 \mid x), \end{vmatrix}$$

with some rule to break ties when  $p_{\theta}(y=1 \mid x) = p_{\theta}(y=0 \mid x)$ . This can be summarized as

$$f_{\theta}(x) = \mathbf{1}_{\left\{ \log \left( \frac{p_{\theta}(y=1 \mid x)}{p_{\theta}(y=0 \mid x)} \right) > 0 \right\}} = \mathbf{1}_{\left\{ w^{\top} x + b > 0 \right\}},$$

which corresponds to a historic model of classification known as the perceptron. The hyperplane  $w^{\top}x + b = 0$  is called the decision boundary. The vector  $w \in \mathbb{R}^d$  is normal to this hyperplane, while b is the offset from the origin.

**Remark 3.3.** Note that we present here a model with a symmetric rule for classification, i.e. the threshold value of  $p_{\theta}(y = 1 | x)$  for prediction is 1/2. In some applications, asymmetric thresholds  $p_{\theta}(y = 1 | x) > \delta$  should be considered; e.g. for spam classification it makes sense to consider  $\delta > 1/2$  if y = 1 corresponds to spam emails, as one would prefer to not classify an email as spam rather than classifying as spam a genuine email. In this case the decision boundary is  $w^{\top}x + b = \sigma^{-1}(\delta)$ .

<sup>&</sup>lt;sup>1</sup> There are two classes of probabilistic models: discriminative and generative (see [40, Section 9.4] for a discussion on advantages and disadvantages of both approaches). We will see generative models later on. In essence, they model probability distributions  $p_{\text{data}}(x, y)$  as  $p_{\text{data}}(y)p_{\text{data}}(x|y)$ , with  $p_{\text{data}}(x|y)$  the generative part (obtaining new inputs from labels y). In contrast, discrimative models rely on the conditional probability  $p_{\text{data}}(y|x)$  of the label y given the input x.

### 62 3 Classification with logistic regression

It should be clear from the presentation that logistic regression allows to classify data which is linearly separable. For more general datasets, one needs to resort to some featurization function  $\varphi(x)$  with values in  $\mathbb{R}^D$ , and consider the model

$$p_{\theta}(y=1 \mid x) = \sigma \left( w^{\top} \varphi(x) + b \right), \qquad \theta = (w,b) \in \mathbb{R}^{D} \times \mathbb{R}.$$

A simple example would be two-dimensional data points, the ones with labels 1 being enclosed in the unit disk, while the ones with labels 0 lie outside the disk. In this case, the featurization function  $\varphi(x) = (x_1^2, x_2^2)$  would allow to linearly separate the dataset. In fact, in the general case, one should learn the optimal feature function on top of the parameters  $\theta$ ; for instance by using a neural network (see Chapter 8). The idea is that, if the (nonlinear) featurization function is sufficiently good, then a simple classification method based on a linear model, such as logistic regression, can be sufficient to obtain good prediction performances.

The steepness of the sigmoid is controlled by the magnitude of w (but does not depend on b). The larger  $||w||_2$  is, the sharper the transition, which increases the risk of overfitting. The relative values of  $p_{\theta}(y = 1 | x)$  and  $p_{\theta}(y = 0 | x)$  can be used to give some indication on the confidence of the prediction.

The parameter  $\theta$  to use for predictions is determined by maximizing the log-likelihood of the data for the Bernoulli model under consideration; see Section 3.2.3.

# 3.2.2 Binary logistic regression through a classification-calibrated convex surrogate

We reinterpret in this section the method as introduced in Section 3.2.1 in the light of the mathematical framework of Section 3.2.2. We consider here the situation when  $\mathcal{Y} = \{-1, 1\}$  (which is not the setting considered in Section 3.2.1).

Let us first reformulate the model of Section 3.2.1 for  $\mathcal{Y} = \{-1, 1\}$ . Equation 3.2.1 remains unchanged, while (see Exercise 3.6 below)

$$p_{\theta}(y = -1 \mid x) = 1 - \sigma \left( w^{\top} \varphi(x) + b \right) = \sigma \left( - \left[ w^{\top} \varphi(x) + b \right] \right).$$

This can be summarized as

$$p_{\theta}(y \mid x) = \sigma \left( y g_{\theta}(x) \right), \qquad g_{\theta}(x) = \theta^{\top} \phi(x), \tag{3.16}$$

where

$$\theta = (w, b) \in \mathbb{R}^d, \qquad \phi(x) = (\varphi(x), 1) \in \mathbb{R}^d$$

Let us next introduce the convex surrogate

$$\Phi(u) = \log(1 + e^{-u}), \qquad (3.17)$$

which is a classification-calibrated convex surrogate in view of Exercise 3.5. Moreover, recalling (3.6) (up to an unimportant multiplicative factor log 2), we obtain, for some decision function g (not necessarily affine at this stage)

$$\Phi(yg(x)) = -\log\sigma(yg(x)).$$

The associated risk is

$$\mathcal{R}_{\Phi}(g) = \mathbb{E}\left[-\log \sigma(yg(x))\right],$$

so that the corresponding empirical risk for the affine decision function  $g_{\theta}$  is

$$\widehat{\mathcal{R}}_{\Phi,n}(\theta) = -\frac{1}{n} \sum_{i=1}^{n} \log \sigma \left( y_i g_{\theta}(x_i) \right).$$

Minimizing the empirical  $\Phi$ -risk can, in view of (3.16), be understood as maximizing the loglikelihood of the data in the framework of the discriminative model; see also (3.18) below.

#### 3.2.3 Estimating $\theta$ for prediction

We come back to the setting of Section 3.2.1. In practice, one finds an estimation  $\hat{\theta} = (\hat{w}, \hat{b})$  of  $\theta$ and performs predictions on unseen data points  $x' \in \mathcal{X}$  based on  $\mathbf{1}_{\{\hat{\theta}^{\top}\phi(x')>0\}}$ . The value of  $\hat{\theta}$  is found by minimizing some empirical risk, possibly with the addition of some regularization term.

Logistic regression, as described in Section 3.2.1, relies on a Bernoulli model. Recall that for a Bernoulli random variable  $Z \sim \mathcal{B}(p)$ , the likelihood of observing the realization z for Z is  $p^z(1-p)^{1-z}$ . The negative log-likelihood of the training data set for the discriminative probabilistic model of Section 3.2.1 is therefore

$$\widehat{\mathcal{R}}_{n}(\theta) = -\frac{1}{n} \sum_{i=1}^{n} \log\left(\sigma\left(g_{\theta}(x_{i})\right)^{y_{i}} \left[1 - \sigma\left(g_{\theta}(x_{i})\right)\right]^{1-y_{i}}\right) = \frac{1}{n} \sum_{i=1}^{n} H\left(y_{i}, \sigma\left(g_{\theta}(x_{i})\right)\right), \quad (3.18)$$

with H the cross-entropy

$$H(y,\hat{y}) = -y\log\hat{y} - (1-y)\log(1-\hat{y}).$$
(3.19)

The quantity  $\widehat{\mathcal{R}}_n(\theta)$  should be minimized in order to find  $\widehat{\theta}$ . Let us however emphasize that, as discussed in Section 3.2.2, this can be equivalently seen as minimizing some empirical  $\Phi$ -risk, which allows to understand logistic regression in the usual setting of supervised learning.

**Exercise 3.6.** *Prove the following statements:* 

- $\sigma(-t) = 1 \sigma(t)$  for any  $t \in \mathbb{R}$ ;
- $\sigma' = \sigma(1 \sigma);$
- for  $z \in \{0,1\}$  fixed, the function  $\psi_z : t \mapsto H(z,\sigma(t))$  is convex on  $\mathbb{R}$ .

Deduce that the function  $h_i: \theta \mapsto H(y_i, \sigma(g_\theta(x_i)))$  is convex.

**Correction.** The first equality follows from

$$\sigma(-t) = \frac{1}{1 + e^t} = \frac{e^{-t}}{1 + e^{-t}} = 1 - \frac{1}{1 + e^{-t}} = 1 - \sigma(t).$$

For the second one, we use the above equality to write

$$\sigma'(t) = \frac{e^{-t}}{(1+e^{-t})^2} = \sigma(t)\frac{e^{-t}}{1+e^{-t}} = (1-\sigma(t))\sigma(t).$$

Finally,

$$\sigma(t)^{z}(1-\sigma(t))^{1-z} = \left(\frac{e^{t}}{1+e^{t}}\right)^{z} \left(\frac{1}{1+e^{t}}\right)^{1-z} = \frac{e^{zt}}{1+e^{t}}$$

so that, recalling the definition (3.17),

$$\psi_z(t) = -\log\left[\sigma(t)^z (1 - \sigma(t))^{1-z}\right] = -zt + \Phi(-t).$$
(3.20)

This shows that  $\psi_z$  is the sum of a linear function and a convex function, which implies that  $\psi_z$  is convex. Finally, the function  $h_i$  is convex as the composition of the affine function  $\theta \mapsto g_\theta(x_i)$  and the convex function  $\psi_{y_i}$ .

In view of Exercise 3.6, the empirical risk (3.18) is convex. This implies that this function should be rather easy to minimize, and that gradient methods such as those presented in Chapter 4 should work fine. Newton or quasi-Newton methods could also be used. These two options are considered in the method implemented in scikit-learn.<sup>2</sup> The computation of the gradient and of the Hessian is made precise in the following exercise.

 $^2$  See Section 1.1.11.3 of https://scikit-learn.org/stable/modules/linear\_model.html#logistic-regression

# 64 3 Classification with logistic regression

**Exercise 3.7.** Compute the gradient and the Hessian of the function  $h_i : \theta \mapsto H(y_i, \sigma(g_\theta(x_i)))$ . Deduce the expression of the gradient and Hessian of  $\widehat{\mathcal{R}}_n$ .

**Correction.** Recall that  $h_i(\theta) = \psi_{y_i}(g_{\theta}(x_i))$ , with the notation of Exercise 3.6. First, using (3.20),

$$\psi'_{y_i}(t) = -y_i - \Phi'(-t) = \frac{e^t}{1 + e^t} - y_i = \sigma(t) - y_i.$$

Since  $\nabla_{\theta} g_{\theta}(x_i) = \phi(x_i)$ , we then obtain

$$\nabla_{\theta} h_i(\theta) = \psi'_{y_i}(g_{\theta}(x_i)) \nabla_{\theta} g_{\theta}(x_i) = \left[\sigma(g_{\theta}(x_i)) - y_i\right] \phi(x_i).$$

Finally,

$$\nabla_{\theta}\widehat{\mathcal{R}}_{n}(\theta) = \frac{1}{n} \sum_{i=1}^{n} \left[ \sigma(g_{\theta}(x_{i})) - y_{i} \right] \phi(x_{i}) \in \mathbb{R}^{d}.$$

Similarly, since  $\nabla_{\theta}\sigma(g_{\theta}(x_i)) = \sigma(g_{\theta}(x_i))[1 - \sigma(g_{\theta}(x_i))]\phi(x_i)$ , one finds

$$\nabla_{\theta}^{2}\widehat{\mathcal{R}}_{n}(\theta) = \frac{1}{n} \sum_{i=1}^{n} \sigma(g_{\theta}(x_{i})) [1 - \sigma(g_{\theta}(x_{i}))] \phi(x_{i}) \phi(x_{i})^{\top} \in \mathbb{R}^{d \times d}.$$

The latter matrix is the sum of positive semidefinite matrices, so that the Hessian is positive semidefinite (as expected, since the function under consideration is convex).

In practice, it is usually beneficial to normalize the data in some way, for instance by standardizing it, or by using a min-max scaler, as explained in Section 1.1.4. For the same reasons as for linear least square regression (recall Sections 2.3 and 2.4), one should consider adding regularization terms whose magnitude is determined by cross validation. These regularization terms are typically  $\|\theta\|_2^2$ , or  $\|\theta\|_1$  if the solution is expected to be sparse. When  $\theta = (w, b)$  and the last component of the feature function  $\phi$  is 1, in which case b is simply an offset that allows to recenter the data, it is customary to consider a regularization only on the part w, which determines the steepness of the decision function.

# 3.2.4 Multiclass logistic regression

We finally discuss how to use logistic regression to predict labels in situations where there are 3 labels or more, building upon the discussion in Section 1.2.1.1. More precisely, we assume that there are  $K \ge 3$  classes, with  $\mathcal{Y} = \{1, \ldots, K\}$ , and recall the expression of the softmax function, defined for  $(a_1, \ldots, a_K) \in \mathbb{R}^K$  as

$$S_K(a_1, \dots, a_K) = \left(\frac{e^{a_1}}{\sum_{k=1}^K e^{a_k}}, \frac{e^{a_2}}{\sum_{k=1}^K e^{a_k}}, \dots, \frac{e^{a_K}}{\sum_{k=1}^K e^{a_k}}\right) \in \mathbb{R}^K.$$

Predictions are based on the values of the decision function

$$G_{\theta}(x) = W\phi(x) + b \in \mathbb{R}^{K}, \qquad W \in \mathbb{R}^{K \times d}, \quad b \in \mathbb{R}^{K},$$

by computing the k-th component of the vector  $S_K$ , namely

$$p_{\theta}(y = k \,|\, x) = S_{K,k}(G_{\theta}(x)).$$

The label y' predicted for a new input x' is

$$y' \in \operatorname*{argmax}_{1 \leqslant k \leqslant K} p_{\theta} \left( y = k \,|\, x' \right)$$

**Remark 3.4.** The discussion in Remark 1.1 on the case K = 2 justifies a posteriori the choice of the sigmoid function in (3.15) for logistic regression in the binary case. In essence, the choice of this function can be traced back to the choice of the exponential function in the softmax to transform real numbers into positive ones.

The loss function to consider for the training of the model is still the negative log-likelihood of the data, which generalizes (3.18). More precisely, the likelihood of the outputs  $(y_1, \ldots, y_n)$  conditionally on the inputs  $(x_1, \ldots, x_n)$  is

$$\prod_{i=1}^{n} \prod_{k=1}^{K} S_{K,k} (G_{\theta}(x))^{\underline{y}_{i,k}},$$

where  $\underline{y}_i \in \{0, 1\}^K$  is the one-hot encoding of  $y_i$ , *i.e.* the vector whose components are all equal to 0 except the  $y_i$ -th component which is equal to 1 (see the discussion in Section 1.1.4). Therefore, the empirical risk to consider reads

$$\widehat{\mathcal{R}}_n(\theta) = \frac{1}{n} \sum_{i=1}^n H_K\Big(\underline{y}_i, S_K\left(G_\theta(x_i)\right)\Big),$$

with  $H_K$  the cross-entropy

$$H_K\left(\underline{y}, \underline{\hat{y}}\right) = -\sum_{k=1}^{K} \underline{y}_k \log \underline{\hat{y}}_k.$$
(3.21)

Discussions similar to the ones written for binary logistic regression apply here concerning the training of the model (see [40, Section 10.3.2] for expressions of the gradient and the Hessian, the model being still convex) and its regularization.
# (Stochastic) Gradients methods

4.1	Minir	nization problems to solve and general strategy	
	4.1.1	Formulation of the problem	
	4.1.2	Types of problems to solve	
	4.1.3	Stochastic vs. deterministic methods	
4.2	Simple gradient descent		
	4.2.1	Analysis for a simple case: convex Lipschitz losses	
	4.2.2	Other situations	
4.3	Deter	ministic methods beyond simple gradient descent	
4.4	Stoch	astic gradient descent and its extensions	
	4.4.1	Choice of stepsizes/learning rates	
	4.4.2	Variance reduction	
	4.4.3	Theoretical analysis	
	4.4.4	Momentum versions	

We discuss in this chapter how to minimize (regularized) training losses to find the parameters  $\hat{\theta}$  allowing to perform predictions. This is done by using gradient methods, often stochastic gradient dynamics, and their extensions – including momentum methods. The reason why gradient methods are preferred is that the optimization problems under consideration are often set in very high dimensional spaces, which makes (quasi-)Newton methods impractical. Our presentation is based on [40, Chapter 8], [4, Chapter 5] and [49, Chapter 14].

We start by presenting the general framework of minimization problems in machine learning in Section 4.1, listing in particular properties of optimization problems encountered when training models, and discussing stochastic approaches where gradients are estimated with minibatches of the data points. We next analyze gradient descent in Section 4.2, in a manner amenable to adaptations to stochastic versions; and discuss extensions in Section 4.3, in particular based on the introduction of momentum variables. We finally present the main results of this chapter in Section 4.4, namely stochastic gradient dynamics and its momentum extensions such as the celebrated Adam method [30].

# 4.1 Minimization problems to solve and general strategy

We generically formulate the minimization problem to solve as

$$\widehat{\theta} \in \operatorname*{argmin}_{\theta \in \Theta} F(\theta), \tag{4.1}$$

for some function  $F : \Theta \to \mathbb{R}$ , which depends on the training set. Typically,  $\Theta = \mathbb{R}^d$ , so we are in the context of continuous optimization, even for classification problems (thanks to the use of surrogate functions as discussed in Section 3.1).

#### 68 4 (Stochastic) Gradients methods

The form of this optimization problem is motivated in Section 4.1.1. We next discuss the various types of minimization problems considered in machine learning in Section 4.1.2. We finally distinguish in Section 4.1.3 between deterministic gradient methods, where the full training set is used, and stochastic methods, where the gradient is approximated over a randomly chosen fraction of the data set.

#### 4.1.1 Formulation of the problem

We consider in this chapter the minimization of (regularized) training losses to find parameters used for predictions. More precisely,

$$\widehat{\theta}_{\lambda} \in \operatorname*{argmin}_{\theta \in \Theta} \left\{ \widehat{\mathcal{R}}_{n}(\theta) + \lambda \Omega(\theta) \right\}, \qquad \widehat{\mathcal{R}}_{n}(\theta) = \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i}, f_{\theta}(x_{i})\right), \tag{4.2}$$

where  $\Omega: \mathcal{X} \to \mathbb{R}_+$  determines the regularization and  $\lambda \ge 0$  is the strength of the regularization term. Predictions are then performed based on  $f_{\hat{\theta}_{\lambda}}$ , for a value of  $\lambda$  chosen by (cross) validation, as discussed in Section 2.3.1. The problem (4.2) can be written as (4.1) upon setting  $F = \widehat{\mathcal{R}}_n + \lambda \Omega$ .

**Remark 4.1.** As discussed in Section 2.3.1 as well, an alternative formulation of the minimization problem is

$$\widehat{\theta}_D \in \operatorname*{argmin}_{\theta \in \Theta} \left\{ \widehat{\mathcal{R}}_n(\theta) \, \middle| \, \Omega(\theta) \leqslant D \right\},\tag{4.3}$$

but we will restrict ourselves in this chapter to unconstrained problems of the form (4.2).

A very important remark at this stage is that the aim is not to obtain the best minimizer to (4.1) (*i.e.* the best training error) but the smallest test error, in order to avoid overfitting. To elaborate on this statement, denote by

$$\mathcal{R}(f) = \mathbb{E}\left[\ell(y, f(x))\right]$$

the risk associated with a predictor f, and assume that the minimization of the risk over predictors  $f_{\theta}$  admits a global minimizer, namely

$$\mathcal{R}(f_{\theta_{\star}}) = \inf_{\theta \in \Theta} \mathcal{R}(f_{\theta}).$$

Consider  $\lambda = 0$  for simplicity. Then, the risk associated with a minimizer  $\hat{\theta}$  of the optimization problem (4.2) can be decomposed as

$$\mathcal{R}\left(f_{\widehat{\theta}}\right) - \inf_{\theta \in \Theta} \mathcal{R}(f_{\theta}) = \mathcal{R}\left(f_{\widehat{\theta}}\right) - \mathcal{R}\left(f_{\theta_{\star}}\right) \\ = \underbrace{\mathcal{R}\left(f_{\widehat{\theta}}\right) - \widehat{\mathcal{R}}_{n}\left(f_{\widehat{\theta}}\right)}_{\text{estimation error}} + \underbrace{\widehat{\mathcal{R}}_{n}\left(f_{\widehat{\theta}}\right) - \widehat{\mathcal{R}}_{n}\left(f_{\theta_{\star}}\right)}_{\text{optimization error}} + \underbrace{\widehat{\mathcal{R}}_{n}\left(f_{\theta_{\star}}\right) - \mathcal{R}\left(f_{\theta_{\star}}\right)}_{\text{estimation error}}.$$
(4.4)

This equality deserves several comments:

- first, the characterization of the quality of the optimization relies on function values (and not on the distance between approximate minimizers, magnitude of gradients, etc);
- it suffices to aim for an accuracy in the optimization procedure of the order of the estimation error, which is typically of order  $1/\sqrt{n}$  (from an argument based on the Central Limit Theorem);
- concerning the optimization procedure, the key point is to ensure that the minimizer  $\hat{\theta}$  found by the numerical method is close enough to the best predictor  $\theta_{\star}$  in the class, measured in terms of function values. This motivates looking at the validation loss during the minimization, as a proxy for the (test) loss. In practice, one should stop the optimization based on values of the validation loss, which corresponds to the procedure known as *early stopping*; see below.

**Early stopping.** We describe here the procedure known as early stopping, referring to [40, Section 13.5.1] for further references. The first step is to split the data into an actual training set and a validation set. Optimization is performed on the training set only. Typically, the training loss decreases over the iterations of the algorithm; while the validation loss initially decreases but then increases at some point, or at least stops decreasing. This is a sign that the minimization algorithm enters a regime where the training parameters are adjusted to fit the noise in the training data rather than the structure of the data points.

Early stopping consists in ending the optimization procedure when the validation loss no longer decreases, in order to avoid overfitting. More precisely, one defines a certain number of epochs, called patience, after which the optimization is stopped if the validation loss has not further decreased (see Section 4.1.3 for the definition of the notion of an epoch). The final value of  $\theta$  which is retained is the one for which the validation loss is minimal. This means that, in practical implementations of the method, one keeps track of the minimal value of the validation loss has not decreased.

#### 4.1.2 Types of problems to solve

Let us list here various characteristics of the optimization problems encountered in machine learning:

• convex vs. non-convex functionals: the loss functions for some problems are convex, hence easier to minimize – for instance linear regression with Lasso (see Section 2.4) or logistic regression (see Section 3.2.3). Other problems are inherently non convex. A prominent example is the training of neural networks (see Chapter 8). For non convex problems, optimization methods typically converge to local minima.

The theoretical analysis of optimization methods is often performed for (strongly) convex functionals, in order to write proofs and obtain clean statements. The associated numerical algorithms can be used for non convex problems, although there are no convergence guarantees in this case.

- constrained vs. unconstrained problems: most of the time we will consider unconstrained problems, where the functional to minimize includes some regularization term, as in (4.2). We will only briefly mention constrained problems such as (4.3) in very specific situations.
- smooth vs. non-smooth optimization: in certain situations, the function F to minimize in (4.1) is not continously differentiable, so that the gradient is not well defined everywhere. This is the case for instance for (4.2) with a regularization given by  $\Omega(\theta) = \|\theta\|_1$ , or if the loss function  $\ell$  relies on the hinge loss (see Section 3.1.1.2). The function to minimize is however often the sum of a regular part and a non smooth one. One option to fall back on a smooth optimization problem is to rely on proximal methods (see for instance [40, Section 8.6]). Another option, which we consider here, is to replace gradients in the numerical methods at hand by subgradients (see below).

**Sub-gradients.** We recall here the notion of subgradients, and compute the subgradients of some functions appearing in machine learning problems. For a function  $F : \mathbb{R}^d \to \mathbb{R}$ , an element  $G \in \mathbb{R}^d$  is a subgradient of F at  $\theta \in \mathbb{R}^d$  if

$$\forall z \in \mathbb{R}^d, \qquad F(z) \ge F(\theta) + G^{\perp}(z - \theta).$$

The set of subgradients at  $\theta \in \mathbb{R}^d$  is called the subdifferential at  $\theta$ , and is denoted by  $\partial F(\theta)$ . The subgradient is non-empty for convex functions F. If F is moreover differentiable at  $\theta \in \mathbb{R}^d$ , then  $\partial F(\theta) = \{\nabla F(\theta)\}$ .

**Exercise 4.1.** Compute the subdifferential of  $F : \mathbb{R} \to \mathbb{R}_+$  defined by  $F(\theta) = |\theta|$ .

**Correction.** The function F is  $C^1$  on  $\mathbb{R} \setminus \{0\}$ , so that  $\partial F(\theta) = \{F'(\theta)\}$  for  $\theta \neq 0$ , with  $F'(\theta) =$  $\mathbf{1}_{\theta>0} - \mathbf{1}_{\theta<0}$ . For  $\theta = 0$ , we consider  $F(z) - F(0) - Gz = |z| - Gz \ge 0$  for all  $z \in \mathbb{R}$  if and only if  $G \in [-1, 1]$ . This shows that  $\partial F(0) = [-1, 1]$ .

**Exercise 4.2 (Hinge loss).** Fix  $y \in \mathbb{R}$  and  $x \in \mathbb{R}^d$ . Compute the subdifferential with respect to  $\theta \in \mathbb{R}^d$  of the function  $F(\theta) = \max(0, 1 - y\theta^{\top}x)$  based on the hinge loss.

**Correction.** The function is  $C^1$  for  $\theta \in \mathbb{R}^d$  such that  $1 - y\theta^\top x \neq 0$ , so that  $\partial F(\theta) = \{\nabla F(\theta)\}$ . When  $1 - y\theta^{\top}x > 0$ ,

$$\nabla F(\theta) = -yx,$$

while  $\nabla F(\theta) = 0$  for  $1 - y\theta^{\top}x < 0$ . Consider next  $\theta \in \mathbb{R}^d$  such that  $\theta^{\top}x = y$  (so that  $1 - y\theta^{\top}x = 0$ ). The subgradient of the hinge loss  $f: t \mapsto \min(1-t, 0)$  at t = 1 is  $\partial f(0) = [-1, 0]$  since, for  $\alpha \in [-1, 0]$ , it holds

$$\forall t < 1, \qquad f(t) - f(1) - \alpha(t - 1) = (1 - t)(1 + \alpha) \ge 0,$$

while, for t > 1, one has  $f(t) - f(1) - \alpha(t-1) = -\alpha(t-1) \ge 0$ . On the other hand, one can check that one of the two previous inequalities does not hold if  $\alpha \notin [-1,0]$ . Since F is the composition of f with the smooth function  $\theta \mapsto yx^{\top}\theta$ , one finally obtains  $\partial F(\theta) = \{\alpha yx, -1 \leq \alpha \leq 0\}$ . Indeed, for  $G = \alpha y x \in \partial F(\theta)$ ,

$$F(\eta) - F(\theta) - G^{\top}(\eta - \theta) = f\left(yx^{\top}\eta\right) - f\left(yx^{\top}\theta\right) - \alpha yx^{\top}(\eta - \theta) = f(t) - f(1) - \alpha(t - 1)$$
  
with  $t = yx^{\top}\eta$ .

#### 4.1.3 Stochastic vs. deterministic methods

The computation of the full gradient  $\nabla F(\theta)$  for functions such as those appearing in (4.2) has a computational cost proportional to n because one needs to sum over all data points. Since approximate minimizers are sufficient (in view of the discussion after (4.4)), a precise gradient is not needed, in particular at the early stages of the minimization procedure. An unbiased estimate of the gradient may be sufficient.

Along some sequence of parameters  $(\theta_t)_{t\geq 0}$  obtained with the optimization method at hand, we look for a stochastic approximation  $\hat{g}_t(\theta)$  of the gradient satisfying

$$\mathbb{E}\left[\widehat{g}_t(\theta_{t-1}) \mid \theta_{t-1}\right] = \nabla F(\theta_{t-1}). \tag{4.5}$$

Let us next describe three typical strategies to this end, for the specific case when F is the empirical training loss  $\mathcal{R}_n$  appearing in (4.2):

(i) stochastic approximation is a historical strategy, dating back to Robbins–Monro, which corresponds to setting

$$\widehat{g}_t(\theta) = \nabla_{\theta} \ell \left( y_{I_t}, f_{\theta}(x_{I_t}) \right),$$

where  $I_t$  are independent and identically distributed random variable with uniform distribution in  $\{1, \ldots, n\}$ . It is clear in this context that (4.5) holds. The main interest of this approach is that the computation of the estimator of the gradient is unexpensive as a single data point is needed. On the other hand, the estimator may suffer from a large variance and therefore not be a good enough approximation of the gradient in order to drive the parameter to local minima of the loss function.

(ii) minibatching is an extension of stochastic averaging where several data points are chosen. More precisely, the method relies on sampling a random set  $\mathcal{I}_t$  of m indices sampled at random from  $\{1, \ldots, n\}$ , with or without replacement. The estimator of the gradient is

$$\widehat{g}_t(\theta) = \frac{1}{m} \sum_{i \in \mathcal{I}_t} \nabla_{\theta} \ell\left(y_i, f_{\theta}(x_i)\right)$$

The computational cost of this estimator is proportional to m. Note also that minibatching reduces to stochastic approximation when m = 1, and to a full gradient computation when m = n and the sampling of the indices is performed without replacement. Further stastistical properties of minibatching are studied in Exercise 4.3.

(iii) shuffling is often used in practice, although it does not allow to write (4.5). More precisely, the set of n data points is first randomly permutated, then decomposed in n/m batches of sizes m (assuming that the former number is an integer). One then computes the gradient over each of the n/m batches, and updates the parameter  $\theta$  using some gradient method. This leads to the very important notion of *epoch*, which corresponds to the associated n/m elementary gradient steps. This means that one passes per construction once (and only once) over all data points in one epoch of this procedure; in contrast to estimators based on sampling with or without replacement with a minibatch of size m, for which a data point can be seen several times or none over n/m elementary steps.

Exercise 4.3 (Statistical properties of minibatching). Denote the full gradient by

$$g(\theta) = \frac{1}{n} \sum_{i=1}^{n} G_i(\theta), \qquad G_i(\theta) = \nabla_{\theta} \ell\left(y_i, f_{\theta}(x_i)\right) \in \mathbb{R}^d,$$

and consider the minibatched estimator the gradient for the random set  $\mathcal{I}^m = \{i_1, \ldots, i_m\} \subset \{1, \ldots, n\}$ , namely

$$\widehat{g}^m(\theta) = \frac{1}{m} \sum_{j=1}^m G_{i_j}(\theta).$$

Introduce also the empirical covariance

$$\Sigma(\theta) = \frac{1}{n-1} \sum_{i=1}^{n} \left( G_i(\theta) - g(\theta) \right) \left( G_i(\theta) - g(\theta) \right)^\top \in \mathbb{R}^{d \times d}.$$

(a) Show that  $\mathbb{E}\left[\left(G_{i_j}(\theta) - g(\theta)\right)\left(G_{i_k}(\theta) - g(\theta)\right)^{\top}\right] = 0$  when  $j \neq k$  and sampling is performed uniformly in  $\{1, \ldots, n\}$  with replacement.

(b) Deduce that  $\operatorname{Cov}\left(\widehat{g}^{m}(\theta)\right) = \frac{1}{m}\left(1 - \frac{1}{n}\right)\Sigma(\theta)$  when sampling is performed with replacement. (c) We next consider sampling without replacement for the end of this exercise. Show that

$$\operatorname{Var}\left(\mathbf{1}_{k\in\mathcal{I}^{m}}\right) = \frac{m}{n}\left(1-\frac{m}{n}\right), \qquad \operatorname{Cov}\left(\mathbf{1}_{j\in\mathcal{I}^{m}},\mathbf{1}_{k\in\mathcal{I}^{m}}\right) = -\frac{m(n-m)}{n^{2}(n-1)}.$$

(d) Prove that

$$\operatorname{Cov}\left(\widehat{g}^{m}(\theta)\right) = \frac{1}{m^{2}} \sum_{i=1}^{n} \widetilde{G}_{i}(\theta) \widetilde{G}_{i}(\theta)^{\top} \operatorname{Var}\left(\mathbf{1}_{i \in \mathcal{I}^{m}}\right) + \frac{1}{m^{2}} \sum_{i \neq j} \widetilde{G}_{i}(\theta) \widetilde{G}_{j}(\theta)^{\top} \operatorname{Cov}\left(\mathbf{1}_{i \in \mathcal{I}^{m}}, \mathbf{1}_{j \in \mathcal{I}^{m}}\right),$$

where  $\widetilde{G}_i(\theta) = G_i(\theta) - g(\theta)$ . (e) Conclude that  $\operatorname{Cov}\left(\widehat{g}^m(\theta)\right) = \frac{1}{m}\left(1 - \frac{m}{n}\right)\Sigma(\theta)$  for sampling without replacement. 724 (Stochastic) Gradients methods

#### Correction.

(a) The random variables  $i_j$  and  $i_k$  are independent when  $j \neq k$ , so that

$$\mathbb{E}\left[\left(G_{i_j}(\theta) - g(\theta)\right)\left(G_{i_k}(\theta) - g(\theta)\right)^{\top}\right] = \mathbb{E}\left[G_{i_j}(\theta) - g(\theta)\right] \mathbb{E}\left[G_{i_k}(\theta) - g(\theta)\right]^{\top} = 0,$$

since  $\mathbb{E}\left[G_{i_j}(\theta)\right] = \mathbb{E}\left[G_{i_j}(\theta)\right] = g(\theta)$ . (b) It holds, using first the result of the first question, then the fact that the random indices  $i_j$ are identically distributed, and finally that  $i_1$  is uniformly distributed on  $\{1, \ldots, n\}$ ,

$$\operatorname{Cov}\left(\widehat{g}^{m}(\theta)\right) = \frac{1}{m^{2}} \sum_{j,k=1}^{m} \mathbb{E}\left[\left(G_{i_{j}}(\theta) - g(\theta)\right)\left(G_{i_{k}}(\theta) - g(\theta)\right)^{\top}\right]$$
$$= \frac{1}{m^{2}} \sum_{j=1}^{m} \mathbb{E}\left[\left(G_{i_{j}}(\theta) - g(\theta)\right)\left(G_{i_{j}}(\theta) - g(\theta)\right)^{\top}\right]$$
$$= \frac{1}{m} \mathbb{E}\left[\left(G_{i_{1}}(\theta) - g(\theta)\right)\left(G_{i_{1}}(\theta) - g(\theta)\right)^{\top}\right]$$
$$= \frac{1}{mn} \sum_{i=1}^{n} \left(G_{i}(\theta) - g(\theta)\right)\left(G_{i}(\theta) - g(\theta)\right)^{\top}$$
$$= \frac{1}{m} \frac{n-1}{n} \Sigma(\theta),$$

from which the claimed equality follows.

(c) The first equality is obtained by noting that the probability that a given index k belongs to  $\mathcal{I}^m$  is the ratio of the number of sets of size m containing the index k under consideration (choose m-1 indices among n-1 remaining indices) over the total number of subsets of size m (choose m indices among n):

$$\mathbb{E}\left(\mathbf{1}_{k\in\mathcal{I}^{m}}\right) = \frac{\binom{n-1}{m-1}}{\binom{n}{m}} = \frac{m}{n}.$$

Therefore,

$$\operatorname{Var}\left(\mathbf{1}_{k\in\mathcal{I}^{m}}\right) = \mathbb{E}\left(\mathbf{1}_{k\in\mathcal{I}^{m}}\right) - \mathbb{E}\left(\mathbf{1}_{k\in\mathcal{I}^{m}}\right)^{2} = \frac{m}{n} - \left(\frac{m}{n}\right)^{2} = \frac{m}{n}\left(1 - \frac{m}{n}\right).$$

For the second equality, we compute the probability that a given couple of indices (j, k)belongs to  $\mathcal{I}^m$ . This is the ratio of the number of sets of size m containing the indices j, kunder consideration (choose m-2 indices among n-2 remaining indices) over the total number of subsets of size m (choose m indices among n):

$$\mathbb{E}\left(\mathbf{1}_{(j,k)\in\mathcal{I}^m\times\mathcal{I}^m}\right) = \frac{\binom{n-2}{m-2}}{\binom{n}{m}} = \frac{m(m-1)}{n(n-1)}.$$

Therefore,

$$\operatorname{Cov}\left(\mathbf{1}_{j\in\mathcal{I}^{m}},\mathbf{1}_{k\in\mathcal{I}^{m}}\right) = \mathbb{E}\left(\mathbf{1}_{(j,k)\in\mathcal{I}^{m}\times\mathcal{I}^{m}}\right) - \left(\frac{m}{n}\right)^{2} = \frac{m(m-1)}{n(n-1)} - \left(\frac{m}{n}\right)^{2}$$
$$= \frac{n(m^{2}-m) - (n-1)m^{2}}{(n-1)n^{2}} = \frac{m^{2}-nm}{(n-1)n^{2}},$$

which leads to the desired result.

(d) The formula is directly obtained by noting that

$$\widehat{g}^m(\theta) - g(\theta) = \frac{1}{m} \sum_{k=1}^n G_k(\theta) \mathbf{1}_{k \in \mathcal{I}^m} - g(\theta) = \frac{1}{m} \sum_{k=1}^n \widetilde{G}_k(\theta) \mathbf{1}_{k \in \mathcal{I}^m},$$

and expanding the product appearning in the covariance.

(e) We plug the expressions found in the third question into the equality from the fourth question. This gives

$$\operatorname{Cov}\left(\widehat{g}^{m}(\theta)\right) = \frac{1}{m^{2}} \frac{m}{n} \left(1 - \frac{m}{n}\right) \sum_{i=1}^{n} \widetilde{G}_{i}(\theta) \widetilde{G}_{i}(\theta)^{\top} - \frac{1}{m^{2}} \frac{m(n-m)}{n^{2}(n-1)} \sum_{i \neq j} \widetilde{G}_{i}(\theta) \widetilde{G}_{j}(\theta)^{\top} = \frac{1}{m} \left(1 - \frac{m}{n}\right) \left[\frac{1}{n} \sum_{i=1}^{n} \widetilde{G}_{i}(\theta) \widetilde{G}_{i}(\theta)^{\top} - \frac{1}{n(n-1)} \sum_{i \neq j} \widetilde{G}_{i}(\theta) \widetilde{G}_{j}(\theta)^{\top}\right].$$

Now,

$$\sum_{i=1}^{n} \widetilde{G}_i(\theta) = 0,$$

so that

$$0 = \left(\sum_{i=1}^{n} \widetilde{G}_{i}(\theta)\right) \left(\sum_{i=1}^{n} \widetilde{G}_{i}(\theta)\right)^{\top} = \sum_{i \neq j} \widetilde{G}_{i}(\theta) \widetilde{G}_{j}(\theta)^{\top} + \sum_{i=1}^{n} \widetilde{G}_{i}(\theta) \widetilde{G}_{i}(\theta)^{\top}.$$

Therefore,

$$\operatorname{Cov}\left(\widehat{g}^{m}(\theta)\right) = \frac{1}{m}\left(1 - \frac{m}{n}\right)\left[\frac{1}{n} + \frac{1}{n(n-1)}\right]\sum_{i=1}^{n}\widetilde{G}_{i}(\theta)\widetilde{G}_{i}(\theta)^{\top},$$

which leads to the desired conclusion.

# 4.2 Simple gradient descent

We study in this section gradient descent, which consists in iterating

$$\theta_t = \theta_{t-1} - \gamma_t \nabla F(\theta_{t-1}), \tag{4.6}$$

starting from some given value of  $\theta_0$ , and using a sequence of learning rates or stepsize sequences  $(\gamma_t)_{t \ge 0}$ . These learning rates can be either kept constant, decay according to a given schedule, or found at each step by some line search procedure.

**Remark 4.2.** When the function F to optimize is not smooth, the descent direction should be replaced by any element of the subgradient, namely

$$\theta_t = \theta_{t-1} - \gamma_t v_{t-1}, \qquad v_{t-1} \in \partial F(\theta_{t-1}).$$

The motivation for gradient dynamics is that it can be considered as a time discretization of the continuous dynamics  $\dot{\theta}(t) = -\nabla F(\theta(t))$ , which is such that

$$\frac{d}{dt} \Big[ F(\theta(t)) \Big] = -|\nabla F(\theta(t))|^2 \leqslant 0.$$

#### 74 4 (Stochastic) Gradients methods

The values of F are therefore nonincreasing along the gradient dynamics. There are of course many generalizations of gradient descent, such as conjugate gradient, but our aim here is to understand the algorithms used in machine learning such as stochastic gradient descent and its refinements and extensions.

We therefore provide in this section various convergence results for function values, in scenarios typical of machine learning applications; starting with the simplest case in Section 4.2.1, namely convex Lipschitz losses; and then discuss other settings in Section 4.2.2.

#### 4.2.1 Analysis for a simple case: convex Lipschitz losses

The presentation in this section is taken from [49, Section 14.1]. We consider a convex and  $\rho$ -Lipschitz target F, and perform gradient descent with a fixed learning rate  $\gamma > 0$ . The aim is to obtain bounds on  $F(\bar{\theta}_t) - F(\eta_\star)$ , where  $\eta_\star$  is a minimizer<sup>1</sup> of F, and

$$\overline{\theta}_t = \frac{1}{t} \sum_{s=0}^{t-1} \theta_s \tag{4.7}$$

is obtained by iterate averaging (this makes sense for convex functions, but should not be employed for non-convex ones; see also Section 4.2.2 for convergence results not relying on iterate averaging). Note that, by convexity,

$$F(\overline{\theta}_t) - F(\eta_\star) \leqslant \frac{1}{t} \sum_{s=0}^{t-1} F(\theta_s) - F(\eta_\star).$$
(4.8)

Moreover, still by the convexity of F, it holds  $F(\theta) - F(\eta_{\star}) \leq (\theta - \eta_{\star})^{\top} \nabla F(\theta)$ , so that

$$F(\overline{\theta}_t) - F(\eta_\star) \leqslant \frac{1}{t} \sum_{s=0}^{t-1} (\theta_s - \eta_\star)^\top \nabla F(\theta_s).$$

The next lemma allows to bound the right hand side of the previous equality. We state it in a general manner, with descent directions  $v_s$  that need not be  $\nabla F(\theta_s)$ , as this result will also be used later on in Section 4.4.3 for stochastic gradient descent.

**Lemma 4.1.** Consider an arbitrary sequence of vectors  $v_0, \ldots, v_{t-1} \in \mathbb{R}^d$ , and a sequence of parameters iteratively defined from a given  $\theta_0 \in \mathbb{R}^d$  as  $\theta_t = \theta_{t-1} - \gamma v_{t-1}$ . Then,

$$\sum_{s=0}^{t-1} (\theta_s - \eta_\star)^\top v_s \leqslant \frac{\|\theta_0 - \eta_\star\|^2}{2\gamma} + \frac{\gamma}{2} \sum_{s=0}^{t-1} \|v_s\|^2.$$

In particular, if  $||v_s|| \leq \rho$  and  $\gamma = \frac{||\theta_0 - \eta_\star||}{\rho\sqrt{t}}$ , then

$$\frac{1}{t} \sum_{s=0}^{t-1} (\theta_s - \eta_\star)^\top v_s \leqslant \frac{\rho \|\theta_0 - \eta_\star\|}{\sqrt{t}}.$$

Note that the above result is not an "anytime" result as the learning rate depend on the horizon t and the bound is valid only for this time t.

*Proof.* We rewrite  $(\theta_s - \eta_\star)^\top v_s$  as the difference of two terms at subsequent times in order to introduce a telescopic sum. More precisely,

<sup>&</sup>lt;sup>1</sup> We do not denote the minimizer of F by  $\theta_{\star}$  as this notation is kept for the argmin of the risk  $\theta \mapsto \mathcal{R}(f_{\theta})$ , and not the empirical risk.

4.2 Simple gradient descent

$$(\theta_s - \eta_\star)^\top v_s = \frac{1}{\gamma} (\theta_s - \eta_\star)^\top \gamma v_s = \frac{1}{2\gamma} \left( -\|\theta_s - \eta_\star - \gamma v_s\|^2 + \|\theta_s - \eta_\star\|^2 + \gamma^2 \|v_s\|^2 \right)$$
  
=  $\frac{1}{2\gamma} \left( -\|\theta_{s+1} - \eta_\star\|^2 + \|\theta_s - \eta_\star\|^2 \right) + \frac{\gamma}{2} \|v_s\|^2.$ 

Therefore,

$$\sum_{s=0}^{t-1} (\theta_s - \eta_\star)^\top v_s = \frac{1}{2\gamma} \left( \|\theta_0 - \eta_\star\|^2 - \|\theta_t - \eta_\star\|^2 \right) + \frac{\gamma}{2} \sum_{s=0}^{t-1} \|v_s\|^2,$$

which leads to the first inequality to prove. The second inequality is proved by first using the upper bound on  $||v_s||$  to write

$$\frac{1}{t}\sum_{s=0}^{t-1} (\theta_s - \eta_\star)^\top v_s \leqslant \frac{\|\theta_0 - \eta_\star\|^2}{2\gamma t} + \frac{\gamma \rho^2}{2},$$

and then minimizing the right hand side with respect to  $\gamma$ . The optimal value is characterized by

$$\frac{\rho^2}{2} = \frac{\|\theta_0 - \eta_\star\|^2}{2\gamma^2 t},$$

which indeed leads to the claimed expression for  $\gamma$  and allows to conclude.

We can conclude this section by applying the result of Lemma 4.1 for  $v_s = \nabla F(\theta_s)$ . Note that  $\|\nabla F(\theta_s)\| \leq \rho$  when F is  $\rho$ -Lipschitz (this is easy to prove when  $F \in C^1$ , and can be extended to subgradients as discussed in [49, Lemma 14.7]).

**Corollary 4.1.** Consider a convex and  $\rho$ -Lipschitz function F, and a learning rate  $\gamma = \frac{\|\theta_0 - \eta_\star\|}{\alpha\sqrt{t}}$ . Then, for the simple gradient dynamics (4.6),

$$F(\overline{\theta}_t) - F(\eta_\star) \leqslant \frac{\rho \|\theta_0 - \eta_\star\|}{\sqrt{t}}.$$

This means that, to achieve an error of at most  $\varepsilon > 0$  in function values, it suffices to run gradient descent for  $t \ge \rho^2 \|\theta_0 - \eta_\star\|^2 / \varepsilon^2$  steps.

#### 4.2.2 Other situations

We review in this section some of the results presented in [4, Section 5.2], to which we refer for a more complete presentation. The results we highligh are anytime results, which do not require iterate averaging:

• Consider  $F \in C^1$  which is L-smooth and  $\mu$ -strongly convex, namely

$$\forall \eta \in \mathbb{R}^d, \qquad \left| F(\eta) - F(\theta) + \nabla F(\theta)^\top (\eta - \theta) \right| \leq \frac{L}{2} \|\eta - \theta\|_2^2.$$
 (LS).

and

$$\forall \eta \in \mathbb{R}^d, \qquad F(\eta) \ge F(\theta) + \nabla F(\theta)^\top (\eta - \theta) + \frac{\mu}{2} \|\eta - \theta\|_2^2.$$
 (SC).

One can prove that L is the largest eigenvalue of  $\nabla^2 F$  (in modulus for non-convex functions) when  $F \in C^2$ , while  $\mu$  is the smallest eigenvalues of  $\nabla^2 F$ . In this context, gradient descent with the learning rate  $\gamma = 1/L$  leads to an exponential convergence in function values:

$$F(\theta_t) - F(\eta_\star) \leq e^{-t/\kappa} \left( F(\theta_0) - F(\eta_\star) \right),$$

where  $\kappa = L/\mu \ge 1$  is a conditioning number. Typically,  $\mu$  is small and of the order of the regularization strength for ridge regression for instance.

75

- 76 4 (Stochastic) Gradients methods
- When F is L-smooth, convex and  $\gamma = 1/L$ , then

$$F(\theta_t) - F(\eta_\star) \leqslant \frac{L}{2t} \|\theta_0 - \eta_\star\|^2.$$

• The convergence rate is smaller for less smooth functions F. When F is convex and  $\rho$ -Lipschitz continuous, and upon choosing

$$\gamma_t = \frac{\|\theta_0 - \eta_\star\|}{\rho\sqrt{t}}$$

one can show that

$$\min_{0 \leqslant s \leqslant t-1} F(\theta_s) - F(\eta_\star) \leqslant \frac{\rho \|\theta_0 - \eta_\star\|}{2\sqrt{t}} (2 + \log t).$$

This result is the counterpart of the estimates of Corollary 4.1 when no averaging of the parameter is performed. Note that iterate averaging may be a better option for convex targets if the function F itself is expensive to evaluate (so that one would not want to compute its values at all timesteps; this is particularly true for stochastic methods where minibatching is used in order to avoid going over all data points).

# 4.3 Deterministic methods beyond simple gradient descent

Gradient descent methods are often not used as such in practice, because their convergence can be quite slow. There are two principal ways to accelerate them in applications of machine learning:

- (i) rely on second order methods such as Newton and quasi-Newton methods, although this may require to compute the Hessian (often computationally too expensive) or store the gradients over several iterations (which can require a large storage capacity when the parameters are of high dimension as for deep neural networks);
- (ii) introduce some momentum in the optimization method.

The current practice in machine learning is more towards the second approach, and we will therefore present only methods related to this option. The key idea is to introduce a new variable  $\omega$ , thought of as a momentum (mass times velocity in physics), and to go from the non-inertial dynamics  $\dot{\theta}(t) = -\nabla F(\theta(t))$  to the inertial dynamics

$$\begin{cases} \theta(t) = \omega(t), \\ \dot{\omega}(t) = -\nabla F(\theta(t)) - \xi \omega(t), \end{cases}$$
(4.9)

where  $\xi > 0$  is some friction coefficient. The inertia inherent in the dynamics motivates the terminology "heavy ball method" which is sometimes used to refer to this class of evolution equations. In fact, (4.9) corresponds to Polyak averaging (which dates back to 1964), and can also be seen as some Hamiltonian dynamics with added dissipation through viscous friction. In particular, denoting by H the Hamiltonian,

$$\frac{dH(\theta(t),\omega(t))}{dt} = -\xi\omega(t)^2 \leqslant 0, \qquad H(\theta,\omega) = F(\theta) + \frac{1}{2}\omega^2.$$

Let us first further elaborate on the notion of inertia, and motivate that (4.9) can be seen as some gradient dynamics with memory. We integrate the second equation in (4.9) by noting that

$$\frac{d}{dt}\left(\mathrm{e}^{\xi t}\omega(t)\right) = -\mathrm{e}^{\xi t}\nabla F(\theta(t)),$$

so that

$$\omega(t) = e^{-\xi t} \omega(0) - \int_0^t e^{-\xi(t-s)} \nabla F(\theta(s)) \, ds$$

Therefore, for  $\omega(0) = 0$  (or for t large, since the first term disappers in this limit),

$$\dot{\theta}(t) = -\int_0^t e^{-\xi(t-s)} \nabla F(\theta(s)) \, ds.$$

This should be compared to gradient dynamics where  $\nabla F(\theta(t))$  appears on the right hand side. It is clear from the latter equation that (4.9) can be considered as some gradient dynamics where the drift is averaged in time.

In order to relate the dissipative Hamiltonian dynamics (4.9) to the gradient dynamics, note that one can rewrite the second equation as

$$\ddot{\theta}(t) + \xi \dot{\theta}(t) = -\nabla F(\theta(t)),$$

and introduce the time-rescaled solution  $\theta_{\xi}(t) = \theta(\xi t)$  to obtain (see Exercise 4.4)

$$\theta_{\xi}(t) = \theta_{\xi}(0) - \int_{0}^{t} \nabla F\left(\theta_{\xi}(s)\right) \, ds + \mathcal{O}\left(\frac{1}{\xi}\right). \tag{4.10}$$

This motivates that the limiting dynamics  $\theta_{\infty}(t)$  is a solution to the non inertial dynamics. Exercise 4.4. Prove the equality (4.10).

Correction. Note that the dynamics can be integrated as

$$\xi(\theta(t) - \theta(0)) = \dot{\theta}(0) - \dot{\theta}(t) - \int_0^t \nabla F(\theta(s)) \, ds$$

so that

$$\theta(\xi t) = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \int_0^t \nabla F(\theta(\xi s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(\xi s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(\xi s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{\omega(0) - \omega(\xi t)}{\xi} - \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s)) \, ds = \theta(0) + \frac{1}{\xi} \int_0^{\xi t} \nabla F(\theta(s))$$

from which the result easily follows.

Nesterov accelerated gradient. Appropriate discretizations of (4.9) lead to numerical methods with decay rates  $1/t^2$  for smooth (non strongly) convex functions, instead of the slower rate 1/t obtained with gradient methods; see [4, Exercise 5.15]. This motivates the terminology "accelerated", in particular given that the rate  $1/t^2$  is known to be optimal. Nesterov's method is one such instance of a discretization of (4.9). It is can be formulated as (see for instance [51, Appendix A.1])

$$\begin{cases} \omega_{t+1} = \beta \omega_t - \gamma \nabla F(\theta_t + \beta \omega_t), \\ \theta_{t+1} = \theta_t + \omega_{t+1}. \end{cases}$$
(4.11)

This discrete time evolution can be rewritten in order to make the relationship with (4.9) more apparent. One introduces to this end  $\omega_t = \sqrt{\gamma} \widehat{\omega}_t$  and writes  $\beta = e^{-\xi \sqrt{\gamma}}$ , so that

$$\begin{cases} \widehat{\omega}_{t+1} = \widehat{\omega}_t + \frac{\beta - 1}{\sqrt{\gamma}} \sqrt{\gamma} \widehat{\omega}_t - \sqrt{\gamma} \nabla F \left( \theta_t + \beta \sqrt{\gamma} \widehat{\omega}_t \right), \\ \theta_{t+1} = \theta_t + \sqrt{\gamma} \widehat{\omega}_{t+1}, \end{cases}$$

or equivalently

$$\begin{cases} \frac{\widehat{\omega}_{t+1} - \widehat{\omega}_t}{\sqrt{\gamma}} = -\nabla F \left( \theta_t + \beta \sqrt{\gamma} \widehat{\omega}_t \right) - \frac{1 - \beta}{\sqrt{\gamma}} \widehat{\omega}_t, \\ \frac{\theta_{t+1} - \theta_t}{\sqrt{\gamma}} = \widehat{\omega}_{t+1}, \end{cases}$$

which corresponds to some discretization of (4.9) with a time step  $\sqrt{\gamma}$  since  $(1-\beta)/\sqrt{\gamma} = \xi + O(\sqrt{\gamma})$ . In practice,  $\gamma$  and  $\beta$  have to be carefully chosen so that the dynamics has a stable behavior. The above derivation, which suggests that  $\beta = e^{-\xi\sqrt{\gamma}}$  for some friction coefficient  $\xi > 0$ , provides some guidelines to this end.

77

## 4.4 Stochastic gradient descent and its extensions

We turn in this section to stochastic gradient dynamics (SGD) and their variations. These methods are currently the most commonly used to minimize empirical risk functions, and result from replacing  $\nabla F(\theta)$  in gradient dynamics by some stochastic approximation, typically obtained via minibatching, as described in Section 4.1.3. More precisely, we consider the dynamics

$$\theta_t = \theta_{t-1} - \gamma_t \widehat{g}_t(\theta_{t-1}), \qquad \mathbb{E}\left[\widehat{g}_t(\theta_{t-1}) \,|\, \theta_{t-1}\right] = \nabla F(\theta_{t-1}). \tag{4.12}$$

One can rewrite the estimator of the gradient as

$$\widehat{g}_t(\theta) = \nabla F(\theta) + \Sigma(\theta)^{1/2} Z_t, \qquad \Sigma(\theta) = \mathbb{E}\left[ \left( \widehat{g}_t(\theta) - \nabla F(\theta) \right) \left( \widehat{g}_t(\theta) - \nabla F(\theta) \right)^\top \right] \in \mathbb{R}^{d \times d},$$

where  $Z \in \mathbb{R}^d$  is a random vector with mean 0 and unit covariance. Typically,  $\Sigma(\theta)$  is of order 1/mwhen consider minibatching with batches of size m. When the stepsize is fixed to some value  $\gamma > 0$ , one can show that the discrete dynamics (4.12) is similar in behavior to solutions of the following stochastic differential equation observed at times multiples of  $\gamma$  (see for instance [48]):

$$d\theta_s = -\nabla F(\theta_s) \, ds + \sqrt{\gamma} \Sigma(\theta_s)^{1/2} dW_s.$$

This corresponds to gradient dynamics perturbed by some noise, whose magnitude diminishes as the stepsize is reduced or the minibatch size m is increased.

#### 4.4.1 Choice of stepsizes/learning rates

The timestep/learning rate in (4.12) should be chosen as a compromise between not being too small, otherwise the dynamics takes too much time to converge, but also not being too large, otherwise there is too much gradient noise and the dynamics may be unstable (similarly to what happens for standard gradient dynamics). A practical approach to setting the timestep consists in starting with a small learning rate, performing the minimization for a (small) subset<sup>2</sup> of the data for a fixed (rather small) number of steps, and then increasing the timestep until the best value achieved for each learning rate starts to increase. The timestep will then be chosen to be of the order of magnitude of the timestep leading to the smallest value of the target function for the chosen subset of data, possibly reduced a bit for cautiousness, in order to take into account that the prediction for the optimal timestep may not be fully reliable as only a subset of the data was used.

Some practitioners advocate using learning schedules. A first option is to resort to decreasing timesteps, inspired by proofs of convergence in the context of stochastic approximation algorithms à la Robbins–Monro [46]. Typical conditions on the learning rates are

$$\sum_{t \geqslant 1} \gamma_t^2 < +\infty, \qquad \sum_{t \geqslant 1} \gamma_t = +\infty,$$

which can be achieved for instance with  $\gamma_t = t^{-\alpha}$  with  $1/2 < \alpha \leq 1$ . Another option is to consider optimization in various stages where the learning rate is kept constant, then decreased after a certain number of iterations (either fixed or depending on the behavior of the values of the function to optimize). For deep learning, a practical recommendation can be to start with initially small learning rates as one may be far away from a local minimum, so that the gradient of the loss function can be large; then to increase the learning rate when convergence starts (*i.e.* as the gradient of the function to minimize starts to decrease in magnitude) in order to favor a better exploration of the loss landscape; and finally reduce the learning rate at the end to fine tune the convergence to a clearly defined local minimum.

<sup>&</sup>lt;sup>2</sup> More generally, cross-validation can be done with a fraction of the data only to coarsely identify interesting regions of parameter space, before fine tuning with more (all) data points.

#### 4.4.2 Variance reduction

It can be beneficial to reduce the minibatching noise in order to accelerate the convergence (from a theoretical perspective, go from convergence rates of order 1/t, typical of stochastic gradient dynamics, to exponential convergence rates for gradient dynamics on strongly convex targets). An idea to this end is to rely on control variate techniques where the estimator for the gradient is replaced by

$$\widetilde{g}_t(\theta) = \widehat{g}_t(\theta) + \nabla F(\theta) - \widehat{g}_t(\theta),$$

with  $\tilde{\theta}$  fixed and  $\nabla F(\tilde{\theta})$  computed exactly. Note that the same minibatch is used for the two estimators on the right hand side of the previous equality. By construction,  $\tilde{g}_t$  is an unbiased estimator of  $\nabla F$ . Moreover,

$$\widetilde{g}_t(\theta) = \nabla F(\widetilde{\theta}) + \frac{1}{m} \sum_{i \in \mathcal{I}_t} \left[ \nabla_{\theta} \ell\left( y_i, f_{\theta}(x_i) \right) - \nabla_{\theta} \ell\left( y_i, f_{\widetilde{\theta}}(x_i) \right) \right],$$

so that the random part of the estimator comes only from the second term, which is small when  $\theta$  is close to  $\tilde{\theta}$  (typically with a variance of order  $\|\theta - \tilde{\theta}\|^2$ ). In practice, this is done by computing the exact gradient for certain iterates  $\theta_{kt_0}$  (with  $t_0$  some fixed period) and using the control variate for the next gradient steps  $kt_0 \leq t \leq (k+1)t_0$ . This method is known as "stochastic variance reduced gradient" (SVRG). A popular extension of this approach is provided by SAGA ("stochastic averaged gradient accelerated"), where  $\nabla F(\tilde{\theta})$  is not computed except for the initial condition, and the control variate is obtained by combining this single exact gradient and an average over the estimators of the gradient in subsequent steps.

#### 4.4.3 Theoretical analysis

The convergence analysis of the SGD method (4.12) is similar to the derivation performed in Section 4.2.1 for deterministic gradient dynamics, with expectations at well chosen places. These expectations are with respect to realizations of the sampling noise in the estimator of the gradient, typically through the choice of the minibatch (which is conditionally independent of  $\theta_{t-1}$ ). We rely in particular on Lemma 4.1, where the vectors  $v_t$  are replaced by the estimators  $\hat{g}_{t+1}(\theta_t)$  of the gradient.

**Theorem 4.1.** Assume that F is convex, and that, for all  $s \ge 1$ ,

$$\mathbb{E}\left[\left\|\widehat{g}_{s}(\theta_{s-1})\right\|^{2}\right] \leqslant \rho^{2}, \qquad \mathbb{E}\left[\widehat{g}_{s}(\theta_{s-1}) \mid \theta_{s-1}\right] = \nabla F(\theta_{s-1}).$$
(4.13)

Fix  $t \ge 1$  and consider the learning rate

$$\gamma = \frac{\|\theta_0 - \eta_\star\|}{\rho\sqrt{t}}.$$

Then, the following convergence result holds for the averaged parameter  $\overline{\theta}_t$  defined in (4.7):

$$0 \leqslant \mathbb{E}\left[F(\overline{\theta}_t)\right] - F(\eta_\star) \leqslant \frac{\rho \|\theta_0 - \eta_\star\|}{\sqrt{t}}.$$

*Proof.* By convexity, it still holds, by taking expectations in (4.8),

$$0 \leqslant \mathbb{E}\left[F(\overline{\theta}_t)\right] - F(\eta_\star) \leqslant \frac{1}{t} \sum_{s=0}^{t-1} \mathbb{E}\left[F(\theta_s)\right] - F(\eta_\star).$$

Now,

$$F(\theta_s) - F(\eta_\star) \leqslant (\theta_s - \eta_\star)^\top \nabla F(\theta_s) = (\theta_s - \eta_\star)^\top \mathbb{E}\left[\widehat{g}_{s+1}(\theta_s) \,|\, \theta_s\right],$$

#### 80 4 (Stochastic) Gradients methods

so that, by the law of total expectations,

$$\mathbb{E}\left[F(\theta_s)\right] - F(\eta_\star) \leqslant \mathbb{E}\left[\left(\theta_s - \eta_\star\right)^\top \widehat{g}_{s+1}(\theta_s)\right]$$

Therefore, using Lemma 4.1 with  $v_s = \hat{g}_{s+1}(\theta_s)$ ,

$$\mathbb{E}\left[F(\overline{\theta}_t)\right] - F(\eta_\star) \leqslant \frac{\|\theta_0 - \eta_\star\|^2}{2\gamma t} + \frac{\gamma}{2t} \sum_{s=0}^{t-1} \mathbb{E}\left[\|\widehat{g}_{s+1}(\theta_s)\|^2\right] \leqslant \frac{\|\theta_0 - \eta_\star\|^2}{2\gamma t} + \frac{\gamma\rho^2}{2}.$$

The conclusion then follows by same arguments as those leading to Corollary 4.1.

Let us conclude this section with some remarks. First, it is a good practice to perform parameter averaging only after some burn-in period, when the current values of the parameters are rather close from each other. This avoids having some bias related to initially irrelevant values of the parameter. Let us next note that the convergence estimate which is obtained for SGD is similar to the one obtained for gradient dynamics (with some adaptation in the assumptions, in the sense that the inequality in (4.13) replaces the condition that F is  $\rho$ -Lipschitz), but for a numerical method which is much cheaper as only (possibly crude) unbiased estimators of the full gradient are needed. This motivates very strong gains in performance when using SGD.

**Exercise 4.5.** We provide in this exercise a sufficient condition for the first condition in (4.13) to hold. More precisely, consider minibatching for the loss function  $\widehat{\mathcal{R}}_n$  in (4.2), and assume that

$$\forall i \in \{1, \dots, n\}, \qquad \|\nabla_{\theta} \ell(y_i, f_{\theta}(x_i))\| \leq \rho.$$

Prove that the first condition in (4.13) holds and that  $F = \widehat{\mathcal{R}}_n$  is  $\rho$ -Lipschitz.

Correction. Note that, by a Cauchy–Schwarz inequality,

$$\left\|\widehat{g}_{t}(\theta_{t-1})\right\|^{2} = \left\|\frac{1}{m}\sum_{k\in\mathcal{I}_{t}}\nabla_{\theta}\ell\left(y_{j},f_{\theta}(x_{j})\right)\right\|^{2} \leqslant \frac{1}{m}\sum_{k\in\mathcal{I}_{t}}\left\|\nabla_{\theta}\ell\left(y_{j},f_{\theta}(x_{j})\right)\right\|^{2} \leqslant \rho^{2},$$

so that the same bound holds in expectation. The fact that  $\widehat{\mathcal{R}}_n$  is  $\rho$ -Lipschitz follows directy from the fact that it is the average of the functions  $\theta \mapsto \ell(y_i, f_\theta(x_i))$ , which are  $\rho$ -Lipschitz since their gradients are uniformly bounded by  $\rho$ .

#### 4.4.4 Momentum versions

Let us conclude this section by presenting momentum extensions of SGD, which also include some form of preconditioning in order to tackle the anisotropy of the loss landscape. For gradient dynamics, a preconditioning consists in considering the dynamics

$$\dot{\theta}(t) = -M^{-1}\nabla F(\theta(t)).$$

When  $F(\theta) = \theta^{\top} S\theta/2$ , a good preconditioner is M = S as the dynamics then reduces to  $\dot{\theta}(t) = -\theta(t)$ , for which all components evolve on the same timescale. In general, a convenient choice of preconditioning for  $\theta$  close to a local minimum  $\eta_{\star}$  is  $M = \nabla^2 F(\eta_{\star})$  since the gradient of the loss function can be locally approximated as  $\nabla F(\theta) \approx \nabla^2 F(\eta_{\star})(\theta - \eta_{\star})$ . In practice, inverting the Hessian or an approximation of it can be quite expensive, so that diagonal preconditioners are typically considered.

Adam [30], introduced in 2014, is currently the most common optimization algorithm in machine learning, in particular for large scale neural network models. The name is derived from "adaptive moment estimation", as it is a momentum extension of SGD, with an adaptive procedure which provides some form of diagonal preconditioning. More precisely, the update of the parameter is performed in a componentwise manner as

$$\theta_{t,k} = \theta_{t-1,k} - \gamma \frac{v_{t-1,k}}{\varepsilon + \sqrt{s_{t-1,k}}},$$

with

$$v_t = \beta_1 v_{t-1} + (1 - \beta_1) \widehat{g}_t(\theta_{t-1}), \qquad s_{t,k} = \beta_2 s_{t-1,k} + (1 - \beta_2) \widehat{g}_{t,k}(\theta_{t-1})^2,$$

where  $\varepsilon > 0$  is some (small) regularization parameter, and  $\beta_1, \beta_2 \in [0, 1]$ . The vector v represents some averaged sum of the estimated gradients, which allows to introduce some inertia in the dynamics. The vector  $s_t$  performs some time averaging of the second moment of the components of the gradients, which allows to renormalize the magnitude of the gradient in order to have updates in the parameters of order  $\gamma$  somewhat independently of whether the dynamics is in a region of small or large gradients. In fact, as the variables v and s are initialized as  $v_0 = 0$  and  $s_0 = 0$ , it is customary to correct for the bias incurred by this initialization by considering

$$\widehat{v}_t = \frac{v_t}{1 - \beta_1^t}, \qquad \widehat{s}_t = \frac{s_t}{1 - \beta_2^t},$$
(4.14)

and update the parameter as

$$\theta_{t,k} = \theta_{t-1,k} - \gamma \frac{v_{t-1,k}}{\varepsilon + \sqrt{s_{t-1,k}}}.$$

See Exercise 4.6 for further motivation.

The default values for  $\gamma$ ,  $\varepsilon$ ,  $\beta_1$  and  $\overline{\beta_2}$  in implementations of Adam are  $\overline{\gamma} = 10^{-3}$ ,  $\varepsilon = 10^{-6}$ ,  $\overline{\beta}_1 = 0.9$  and  $\overline{\beta}_2 = 0.999$ . In fact, similarly to the discussion after (4.11), the latter two coefficients should be thought of as  $\beta_i = e^{-\gamma/\alpha_i}$ , and should therefore be modified when changing the learning rate to  $\gamma \neq \overline{\gamma}$  as  $\beta_i = \overline{\beta}_i^{\gamma/\overline{\gamma}}$  (see the discussion in [13, Section 5.2]).

**Exercise 4.6.** Give the expression of  $v_t$  in terms of  $v_0$  and  $\hat{g}_s(\theta_{s-1})$  for  $1 \leq s \leq t$ , and motivate the renormalization procedure (4.14).

Correction. A simple induction shows that

$$v_t = (1 - \beta_1) \left[ \widehat{g}_t(\theta_{t-1}) + \beta_1 \widehat{g}_{t-1}(\theta_{t-2}) + \dots + \beta_1^{t-2} \widehat{g}_2(\theta_1) + \beta_1^{t-1} \widehat{g}_1(\theta_0) \right],$$

so that

$$\widehat{v}_{t} = \frac{\widehat{g}_{t}(\theta_{t-1}) + \beta_{1}\widehat{g}_{t-1}(\theta_{t-2}) + \dots + \beta_{1}^{t-1}\widehat{g}_{1}(\theta_{0})}{1 + \beta_{1} + \dots + \beta_{1}^{t-1}}$$

can be understood as a weighted average of the previous gradient estimators, which actually corresponds to some convex combination.

Add also the continuous version of the dynamics: Belloto da Silva/Gazeau 2020

$5.1 \\ 5.2$	Motivating the need for dimensionality reduction			
5.3	PCA in practice			
<b>5.4</b>	Interpretation of PCA 8			
5.5	Extensions of PCA			
	5.5.1	High dimensional data	91	
	5.5.2	Kernel PCA	91	
	5.5.3	Probabilistic PCA	93	

We study in this chapter a first instance of unsupervised learning (*i.e.* "learning without a teacher"). The aim is to infer properties of the distribution  $p_{data}(dx)$  from the (unlabelled) data points  $x_1, \ldots, x_n$ . For low dimensional situations, one could approximate  $p_{data}(dx)$  using kernel density estimators. Such approaches are however not an option for high dimensional situations typical of machine learning applications. In order to make progress, one therefore has to settle for rather crude global models (*e.g.* represent  $p_{data}$  as a mixture of Gaussian distributions), or to give up on identifying the distribution itself but rather find its important modes by some cluster analysis (see Chapter 9). Another option, which we consider here, is to identify a low dimensional manifold around which data points are concentrated. This is done mostly in a linear setting in this chapter, although we hint at some nonlinear extensions.

More precisely, after motivating why it is interesting to reduce the dimension of the data points in Section 5.1, we derive Principal Component Analysis (PCA) through a perspective based on minimizing a reconstruction error in Section 5.2. We then make precise the usual formulation of PCA in Section 5.3, and elaborate on its interpretation in Section 5.4. We finally discuss in Section 5.5 various extensions and generalizations of this method, including kernel PCA, probabilistic PCA and factor analysis, and nonlinear versions of PCA. The presentation here is based on [40, Sections 20.1 and 20.2], [6, Chapters 15 and 21] and [8, Chapter 12].

# 5.1 Motivating the need for dimensionality reduction

One may want to reduce the dimensionality of the data points  $x \in \mathcal{X}$  for various reasons:

- *computational*: the data can be compressed in a preprocessing step to perform faster subsequent operations and/or apply other machine learning algorithms for regression or classification;
- *featurization* is related to the previous item: in certain applications, such as finding so-called reaction coordinates or collective variables in molecular dynamics, it is of interest to find (non)linear functions of the data points which provide relevant inputs for other methods.
- *data exploration/visualization*: high dimensional data cannot be easily represented in a way amenable for humans to understand its organization. It is on the other hand much easier to

plot 2 or 3 dimensional data points, and develop some intuition on how the data points are organized. This makes sense only if the chosen dimensions are particularly representative.

Beyond these motivations, one actually expects that the data points lie around some low dimensional manifold, and that the probability measure from which these data points are distributed is concentrated around this manifold. Indeed, if this was not the case, learning algorithms would not be able to make successful predictions – think for instance of image recognition, even binarized images of low resolution, as the binarized version of the MNIST data set in which images are composed of  $28 \times 28$  pixels. There are  $2^{28 \times 28} \approx 10^{236}$  possible images, which is an enormous number in comparison to the size of the training data set (about 60,000 instances). We however note that many pixels are irrelevant (pixels around the boundary are mostly white), and also expect that only some pixels have a crucial role in digit recognition (a very low resolution image can be enough to distinguish between digits).

Given these elements of context, we think in this chapter of data points  $x \in \mathcal{X} = \mathbb{R}^{1 \times d}$  as obtained through some latent model (recall that x is considered as a line vector). More precisely, we start by recentering the data by the empirical average of the sample at hand. The idea is next to recover x from a lower dimensional variable  $z = xU^{\top} + \varepsilon \in \mathbb{R}^{1 \times \ell}$  with  $U \in \mathbb{R}^{\ell \times d}$  and  $\varepsilon$  some random variable with values in  $\mathbb{R}^{1 \times \ell}$  (if no centering was considered, one would have to add a bias to the above linear model; see Lemma 5.1). The notation  $\ell$  is chosen to emphasize that this number is the dimension of the latent space. We consider  $\ell \leq d - 1$ , and in fact we want  $\ell$  to be much smaller than d. The latent model corresponds to considering an encoding function  $\mathbb{R}^{1 \times \ell} \to \mathbb{R}^{1 \times \ell}$ acting as  $x \mapsto xU^{\top}$ , and a decoding function  $\mathbb{R}^{1 \times \ell} \to \mathbb{R}^{1 \times d}$  acting as  $z \mapsto zU$  (it is not clear at this stage why the decoder should use U and not a generic matrix  $\mathbb{R}^{\ell \times d}$ ; this is justified below in Lemma 5.1). If no noise is present, one recovers PCA. Noise can however be injected into the model in order to have a generative viewpoint, and construct a distribution of data points from a distribution of latent variables z, see Section 5.5.3.

### 5.2 Deriving PCA from reconstruction error

We justify in this section the latent model described at the end of Section 5.1, and show how PCA can be understood from the minimization of some reconstruction error. We consider to this end a general linear dimensionality reduction scheme, where a data point  $x \in \mathbb{R}^{1 \times d}$  is first embedded in  $\mathbb{R}^{1 \times \ell}$  as z = xE (encoding), then lifted back to  $\mathbb{R}^{1 \times d}$  as zD (decoding), and finally shifted by  $b \in \mathbb{R}^{1 \times d}$ . The matrices  $E \in \mathbb{R}^{d \times \ell}$  and  $D \in \mathbb{R}^{\ell \times d}$ , and the bias  $b \in \mathbb{R}^{1 \times d}$  are found by minimizing the reconstruction loss, which corresponds to considering the following training loss:

$$\widehat{\mathcal{R}}_n(E, D, b) = \frac{1}{n} \sum_{i=1}^n \|x_i - (x_i E D + b)\|_2^2.$$
(5.1)

In fact, the following lemma shows that bias b is known and related to the empirical average of the data

$$\overline{x}_n = \frac{1}{n} \sum_{i=1}^n x_i.$$

Moreover, the matrix E for optimal solutions is related to the matrix D of optimal solutions; and furthermore D has specific orthogonality properties (see [49, Lemma 23.1]).

Lemma 5.1. The optimal reconstruction loss can be obtained as

$$\inf_{\substack{E \in \mathbb{R}^{d \times \ell} \\ D \in \mathbb{R}^{\ell \times d} \\ b \in \mathbb{R}^{1 \times d}}} \widehat{\mathcal{R}}_n \left( E, D, b \right) = \min_{\substack{V \in \mathbb{R}^{\ell \times d} \\ VV^\top = \mathrm{Id}_\ell}} \left\{ \frac{1}{n} \sum_{i=1}^n \left\| x_i - \overline{x}_n - \left( x_i - \overline{x}_n \right) V^\top V \right\|_2^2 \right\}.$$

*Proof.* Fix  $(E, D) \in \mathbb{R}^{d \times \ell} \times \mathbb{R}^{\ell \times d}$ . Let us first perform the minimization over  $b \in \mathbb{R}^d$ . It is easy to see that the function  $b \mapsto \widehat{\mathcal{R}}_n(E, D, b)$  is smooth and strongly convex, hence admits a unique minimizer characterized by the Euler–Lagrange equation

$$0 = \nabla_b \widehat{\mathcal{R}}_n \left( E, D, b^\star \right) = -\frac{2}{n} \sum_{i=1}^n \left[ x_i - \left( x_i E D + b^\star \right) \right].$$

This shows that

$$b^{\star} = \overline{x}_n \left( \mathrm{Id}_d - ED \right).$$

Therefore, upon introducing  $\tilde{x}_i = x_i - \overline{x}_n$ ,

$$\inf_{\substack{E \in \mathbb{R}^{d \times \ell} \\ D \in \mathbb{R}^{\ell \times d} \\ b \in \mathbb{R}^{1 \times d}}} \widehat{\mathcal{R}}_n(E, D, b) = \inf_{\substack{E \in \mathbb{R}^{d \times \ell} \\ D \in \mathbb{R}^{\ell \times d}}} \frac{1}{n} \sum_{i=1}^n \|\widetilde{x}_i - \widetilde{x}_i ED\|_2^2.$$

Denote next by  $\mathcal{V}$  the range of the mapping  $\mathbb{R}^d \ni x \mapsto xED$ . Note that  $\mathcal{V}$  is a vector subspace of  $\mathbb{R}^d$  of dimension  $L_{\mathcal{V}} \leqslant \ell$ . Denoting by  $\Pi_{\mathcal{V}}$  the orthogonal projection onto  $\mathcal{V}$ , it holds, upon decomposing  $\tilde{x}_i - \tilde{x}_i ED$  as the sum of  $\tilde{x}_i - \Pi_{\mathcal{V}} \tilde{x}_i \in \mathcal{V}^{\perp}$  and  $\Pi_{\mathcal{V}} \tilde{x}_i - \tilde{x}_i ED \in \mathcal{V}$  and using Pythagore's theorem,

$$\frac{1}{n}\sum_{i=1}^{n}\|\widetilde{x}_{i}-\widetilde{x}_{i}ED\|_{2}^{2} \ge \frac{1}{n}\sum_{i=1}^{n}\|\widetilde{x}_{i}-\Pi_{\mathcal{V}}\widetilde{x}_{i}\|_{2}^{2}.$$

This shows that

$$\inf_{\substack{E \in \mathbb{R}^{d \times \ell} \\ D \in \mathbb{R}^{\ell \times d}}} \frac{1}{n} \sum_{i=1}^{n} \|\widetilde{x}_{i} - \widetilde{x}_{i} ED\|_{2}^{2} \ge \inf_{\substack{\mathcal{V}_{\ell} \subset \mathbb{R}^{1 \times d} \\ \dim(\mathcal{V}_{\ell}) \leqslant \ell}} \frac{1}{n} \sum_{i=1}^{n} \|\widetilde{x}_{i} - \Pi_{\mathcal{V}_{\ell}} \widetilde{x}_{i}\|_{2}^{2}.$$

Since a vector subspace of dimension smaller than  $\ell$  is always included in a subspace of dimension exactly  $\ell$ , and since  $||x - \Pi_{\mathcal{V}} x||_2 \leq ||x - \Pi_{\mathcal{V}'} x||_2$  when  $\mathcal{V}' \subset \mathcal{V}$  for two vector subspaces  $\mathcal{V}, \mathcal{V}' \subset \mathbb{R}^{1 \times d}$ , we can in fact restrict the minimization on the right hand side of the previous inequality to vector subspaces of dimension exactly equal to  $\ell$ :

$$\inf_{\substack{E \in \mathbb{R}^{d \times \ell} \\ D \in \mathbb{R}^{\ell \times d}}} \frac{1}{n} \sum_{i=1}^{n} \|\widetilde{x}_{i} - \widetilde{x}_{i} ED\|_{2}^{2} \ge \inf_{\substack{\mathcal{V}_{\ell} \subset \mathbb{R}^{1 \times d} \\ \dim(\mathcal{V}_{\ell}) = \ell}} \frac{1}{n} \sum_{i=1}^{n} \|\widetilde{x}_{i} - \Pi_{\mathcal{V}_{\ell}} \widetilde{x}_{i}\|_{2}^{2}.$$

Consider next a vector subspace  $\mathcal{V}_{\ell}$  of dimension  $\ell$ . Introduce an orthonormal basis  $(v_1, \ldots, v_{\ell})$  of  $\mathcal{V}_{\ell}$ , and the matrix  $V \in \mathbb{R}^{\ell \times d}$  whose lines are the vectors  $v_k \in \mathbb{R}^{1 \times d}$ . Then,

$$\Pi_{\mathcal{V}_{\ell}} x = \sum_{k=1}^{\ell} \left( x v_k^{\top} \right) v_k = x V^{\top} V.$$

Since the matrix V is characterized by the condition  $VV^{\top} = \mathrm{Id}_{\ell}$  (which simply encodes the fact that the family  $(v_1, \ldots, v_{\ell})$  is orthonormal), we find that

$$\inf_{\substack{E \in \mathbb{R}^{d \times \ell} \\ D \in \mathbb{R}^{\ell \times d}}} \frac{1}{n} \sum_{i=1}^{n} \|\widetilde{x}_{i} - \widetilde{x}_{i} ED\|_{2}^{2} \ge \inf_{\substack{V \in \mathbb{R}^{\ell \times d} \\ VV^{\top} = \mathrm{Id}_{\ell}}} \frac{1}{n} \sum_{i=1}^{n} \|\widetilde{x}_{i} - \widetilde{x}_{i} V^{\top} V\|_{2}^{2}.$$

Note at this stage that the function to minimize on the right hand side of the previous inequality is continuous, and that the minimization is performed over a compact set since

$$\forall k \in \{1, \dots, \ell\}, \qquad \sum_{j=1}^{d} V_{k,j^2} = 1,$$

so that all coefficients of V are bounded by 1 in absolute value. There exists therefore a minimizer  $V^*$  to the right hand side. The left hand side then also admits a minimizer  $E^* = V^*$ and  $D^* = (V^*)^{\top}$ . This finally shows that the initial minimization problem admits a minimizer  $(E^*, D^*, b^*)$ , and allows to conclude.

In the remainder of this section, in order to alleviate the notation, we will assume that the data has been recentered in a pre-processing step, so that  $\overline{x}_n = 0$ . Lemma 5.1 then shows that the minimization of the reconstruction loss  $\widehat{\mathcal{R}}_n$  defined in (5.1) can be equivalently rewritten as

$$\min\left\{\frac{1}{n}\sum_{i=1}^{n}\left\|x_{i}-x_{i}V^{\top}V\right\|_{2}^{2}, \ V \in \mathbb{R}^{\ell \times d}, \ VV^{\top} = \mathrm{Id}_{\ell}\right\}.$$
(5.2)

Upon writing

$$V = \begin{pmatrix} & & v_1 & & \\ & \vdots & \\ & & v_\ell & & \end{pmatrix},$$

with  $v_i v_j^{\top} = \delta_{ij}$  from the condition  $VV^{\top} = \mathrm{Id}_{\ell}$ , we note that

$$x_i V^\top V = \sum_{k=1}^{\ell} \left( x_i v_k^\top \right) v_k$$

is the orthogonal projection of  $x_i$  onto  $\text{Span}(v_1, v_2, \ldots, v_\ell)$ . The properties of orthogonal projections imply that

$$\|x_{i} - x_{i}V^{\top}V\|_{2}^{2} = \|x_{i}\|_{2}^{2} - \|x_{i}V^{\top}V\|_{2}^{2}, \qquad (5.3)$$

which can also be seen from a direct computation. This shows that (5.2) is equivalent to

$$\max\left\{\frac{1}{n}\sum_{i=1}^{n}\left\|x_{i}V^{\top}V\right\|_{2}^{2}, \ V \in \mathbb{R}^{\ell \times d}, \ VV^{\top} = \mathrm{Id}_{\ell}\right\}.$$
(5.4)

The interpretation of the latter maximization problem is that one wants to maximize the variance of the data points projected onto a vector subspace of dimension  $\ell$  (this interpretation in terms of variance being valid when the data has been recentered, see Section 5.3).

The formulations (5.4) or (5.2) allow us to solve the problem in an explicit manner. For pedagogical reasons, we start by presenting the solution for one dimensional latent spaces, and then generalize the procedure to latent spaces of arbitrary dimension  $\ell \leq d-1$ .

**One dimensional latent spaces.** When the latent space is one dimensional, *i.e.*  $\ell = 1$ , the problem (5.4) can be reformulated as

$$\max\left\{\frac{1}{n}\sum_{i=1}^{n}\|(x_{i}\cdot v)v\|^{2} \ \middle| \ v\in\mathbb{R}^{1\times d}, \ \|v\|^{2}=1\right\}.$$

The functional to maximize can be rewritten as

$$\frac{1}{n}\sum_{i=1}^{n} \|(x_i \cdot v)v\|^2 = \frac{1}{n}\sum_{i=1}^{n} (x_i \cdot v)^2 = v\widehat{\Sigma}v^{\top}, \qquad \widehat{\Sigma} = \frac{1}{n}\sum_{i=1}^{n} x_i^{\top}x_i \in \mathbb{R}^{d \times d}.$$
 (5.5)

The Euler–Lagrange equation associated with the maximization problem therefore reads

$$u\widehat{\Sigma} = \lambda u$$

where  $\lambda \in \mathbb{R}$  is the Lagrange multiplier associated with the normalization constraint. This shows that u is an eigenvector associated with the symmetric positive matrix  $\widehat{\Sigma}$ . In order to maximize  $u\widehat{\Sigma}u^{\top} = \lambda ||u||^2 = \lambda$ , one therefore chooses a normalized eigenvector associated with the largest eigenvalue of  $\widehat{\Sigma}$ . Generalization to a latent space of arbitrary dimension. The derivation in the case  $\ell = 1$  suggests to consider the eigenvalues  $\lambda_1 \ge \lambda_2 \ge \ldots \ge \lambda_d$  and associated normalized eigenvectors  $(u_j)_{1 \le j \le d}$  of the symmetric positive matrix  $\widehat{\Sigma}$ :

$$\lambda_j u_j = u_j \widehat{\Sigma}, \qquad \|u_j\|_2^2 = 1.$$

**Theorem 5.1.** The reconstruction error

$$\min\left\{\frac{1}{n}\sum_{i=1}^{n}\left\|x_{i}-x_{i}V^{\top}V\right\|_{2}^{2}, \ V \in \mathbb{R}^{\ell \times d}, \ VV^{\top} = \mathrm{Id}_{\ell}\right\}$$
(5.6)

is minimized for

$$U = \begin{pmatrix} --- u_1 & --- \\ \vdots \\ --- u_\ell & --- \end{pmatrix},$$

with  $(u_1, \ldots, u_\ell)$  an orthonormal family of eigenvectors associated with the  $\ell$  largest eigenvalues of  $\widehat{\Sigma}$ . Moreover, the minimal reconstruction error is

$$\frac{1}{n}\sum_{i=1}^{n} \left\|x_{i} - x_{i}U^{\top}U\right\|_{2}^{2} = \sum_{k=\ell+1}^{d} \lambda_{k}.$$
(5.7)

**Remark 5.1.** There is of course no uniqueness result on the minimizer matrix  $U \in \mathbb{R}^{\ell \times d}$ , since, for example,  $U^{\top}U$  remains the same by changing U to RU for any orthogonal matrix  $R \in \mathbb{R}^{\ell \times \ell}$ .

*Proof.* The equality (5.7) is obtained by evaluating the terms on the right hand side of (5.3) for V = U. Indeed, from the orthonormality of the basis  $(u_1, \ldots, u_\ell)$  and the equality

$$x_i U^\top U = \sum_{k=1}^{\ell} \left( u_k x_i^\top \right) u_k,$$

it follows that

$$\frac{1}{n}\sum_{i=1}^{n} \left\| x_{i}U^{\top}U \right\|^{2} = \frac{1}{n}\sum_{i=1}^{n}\sum_{k=1}^{\ell} \left( u_{k}x_{i}^{\top} \right)^{2} = \sum_{k=1}^{\ell} u_{k} \left( \frac{1}{n}\sum_{i=1}^{n}x_{i}^{\top}x_{i} \right) u_{k}^{\top}$$
$$= \sum_{k=1}^{\ell} u_{k}\widehat{\Sigma}u_{k}^{\top} = \sum_{k=1}^{\ell} \lambda_{k} \|u_{k}\|_{2}^{2} = \sum_{k=1}^{\ell} \lambda_{k}.$$

Similarly,

$$\frac{1}{n}\sum_{i=1}^{n}\|x_i\|^2 = \sum_{k=1}^{d}\lambda_k,$$

from which the claimed equality follows.

To prove that U is an optimal solution of (5.6), we consider an arbitrary matrix  $V \in \mathbb{R}^{\ell \times d}$  such that  $VV^{\top} = \mathrm{Id}_{\ell}$ . Denoting by  $v_1, \ldots, v_{\ell}$  the lines of V,

$$\frac{1}{n}\sum_{i=1}^{n} \left\|x_{i} - x_{i}V^{\top}V\right\|_{2}^{2} = \frac{1}{n}\sum_{i=1}^{n} \left\|x_{i}\right\|_{2}^{2} - \frac{1}{n}\sum_{i=1}^{n} \left\|x_{i}V^{\top}V\right\|_{2}^{2} = \frac{1}{n}\sum_{i=1}^{n} \left\|x_{i}\right\|_{2}^{2} - \frac{1}{n}\sum_{i=1}^{n}\sum_{k=1}^{\ell} \left(x_{i}v_{k}^{\top}\right)^{2}.$$

Now, upon expanding  $v_k$  on the orthonormal basis  $(u_1, \ldots, u_d)$  of eigenvectors of  $\widehat{\Sigma}$  associated with the eigenvalues  $\lambda_1 \ge \ldots \ge \lambda_d$  sorted in decreasing order:

$$x_i v_k^{\top} = \sum_{k'=1}^{a} \left( v_k u_{k'}^{\top} \right) \left( x_i u_{k'}^{\top} \right),$$

,

we can rewrite (using that  $u_{k'}\widehat{\Sigma}u_{k''}^{\top} = \lambda_{k'}\delta_{k',k''}$ )

$$\frac{1}{n}\sum_{i=1}^{n}\sum_{k=1}^{\ell} (x_{i}v_{k}^{\top})^{2} = \frac{1}{n}\sum_{i=1}^{n}\sum_{k=1}^{\ell}\sum_{k',k''=1}^{d} (v_{k}u_{k'}^{\top}) (x_{i}u_{k'}^{\top}) (v_{k}u_{k''}^{\top}) (x_{i}u_{k''}^{\top})$$
$$= \sum_{k=1}^{\ell}\sum_{k',k''=1}^{d} (v_{k}u_{k'}^{\top}) (v_{k}u_{k''}^{\top}) u_{k'}\widehat{\Sigma}u_{k''}^{\top}$$
$$= \sum_{k'=1}^{d}\lambda_{k'}\sum_{k=1}^{\ell} (v_{k}u_{k'}^{\top})^{2} = \sum_{k'=1}^{d}\lambda_{k'} ||u_{k'}V^{\top}V||_{2}^{2}.$$

We claim that the latter quantity is smaller than  $\lambda_1 + \cdots + \lambda_\ell$ , which will allow us to conclude in view of (5.7). Denote by

$$\alpha_k = \left\| u_k V^\top V \right\|_2^2,$$

which is the squared norm of the projection of  $u_k$  onto  $\text{Span}(v_1, \ldots, v_\ell)$ . In particular,  $0 \leq \alpha_k \leq 1$  for any  $1 \leq k \leq d$ . Moreover,

$$\sum_{k'=1}^{d} \alpha_{k'} = \sum_{k'=1}^{d} \sum_{k=1}^{\ell} \left( u_{k'} v_k^{\top} \right)^2 = \sum_{k=1}^{\ell} \sum_{k'=1}^{d} \left( u_{k'} v_k^{\top} \right)^2 = \sum_{k=1}^{\ell} \|v_k\|_2^2 = \ell.$$

Therefore,

$$\max_{\substack{V \in \mathbb{R}^{\ell \times d} \\ VV^{\top} = \mathrm{Id}_{\ell}}} \frac{1}{n} \sum_{i=1}^{n} \left\| x_i V^{\top} V \right\|_2^2 \leqslant \max_{\substack{\alpha \in [0,1]^d \\ \alpha_1 + \dots + \alpha_d = \ell}} \sum_{k=1}^{d} \lambda_k \alpha_k.$$

The maximum on the right hand side is clearly attained by choosing  $\alpha_1 = \cdots = \alpha_\ell = 1$  and  $\alpha_{\ell+1} = \cdots = \alpha_d = 0$ , the value of the maximum being then  $\lambda_1 + \cdots + \lambda_\ell$ . This allows to conclude the proof.

# 5.3 PCA in practice

We summarize in this section how PCA is performed in practice, given a data set  $\{x_1, \ldots, x_n\} \subset \mathbb{R}^d$ and a fixed dimension  $\ell$  for the latent space.

(1) A first remark is that the data needs to be centered. When the features are heterogeneous in scale (for instance because they correspond to measurements of quantities in different units), it is also a good practice to standardize it; see for instance the discussion in [40, Section 20.1.3.1]. This amounts to replacing the covariance matrix by the correlation matrix, namely replacing the attributes  $x_i^j$  by

$$\widetilde{x}_i^j = \frac{x_i^j - \overline{x}_n^j}{\sigma_{n,j}}, \qquad \overline{x}_n = \frac{1}{n} \sum_{i=1}^n x_i, \qquad \sigma_{n,j}^2 = \frac{1}{n} \sum_{i=1}^n \left( x_i^j - \overline{x}_n^j \right)^2.$$

The dataset  $\{\tilde{x}_1, \ldots, \tilde{x}_n\}$  has by construction a vanishing empirical mean and unit empirical variances for all attributes.

(2) The second step is to construct the empirical covariance matrix associated with the renormalized data, namely

$$\widehat{\boldsymbol{\varSigma}} = \frac{1}{n} \sum_{i=1}^{n} \widetilde{\boldsymbol{x}}_i^\top \widetilde{\boldsymbol{x}}_i \in \mathbb{R}^{d \times d},$$

and to diagonalize this symmetric positive matrix:

$$u_k \widehat{\Sigma} = \lambda_k u_k, \qquad ||u_k||^2 = 1, \qquad \lambda_1 \ge \lambda_2 \ge \ldots \ge \lambda_d \ge 0.$$
 (5.8)

- (3) We next compute the embeddings of the data points in the latent space:  $z_i = x_i U_{\ell}^{\top} \in \mathbb{R}^{1 \times \ell}$ , where  $U_{\ell} \in \mathbb{R}^{\ell \times d}$  is the matrix whoses lines are the first  $\ell$  eigenvectors  $u_1, \ldots, u_{\ell}$ .
- (4) Approximate reconstructions of the data points are given by

$$\widehat{x}_i^j = \overline{x}_n^j + \sigma_{n,j} \left( z_i U_\ell \right)_i \in \mathbb{R}^{1 \times d}.$$

Predictions  $\widehat{x'}$  for test points  $x' \in \mathcal{X}$  are performed as

$$\widehat{x'}^{j} = \overline{x}_{n}^{j} + \sigma_{n,j} \left[ (x' - \overline{x}_{n}) U_{\ell}^{\top} U_{\ell} \right]_{j} \in \mathbb{R}^{1 \times d}.$$

#### 5.4 Interpretation of PCA

For notational simplicity, we denote by  $\{x_1, \ldots, x_n\} \subset \mathbb{R}^d$  the recentered data set instead, which can be renormalized or not. The principal components are the vectors

$$c_k = \begin{pmatrix} u_k x_1^\top \\ \vdots \\ u_k x_n^\top \end{pmatrix}$$

associated with the eignvectors of  $\widehat{\Sigma}$  and the normalized data set. The various components of  $c_k$  give the score of each data point on the *k*th loading  $u_k$ . These loadings are linear combination of features, and can therefore be thought of as "implicit features".

Some properties of the score vector are given in the next proposition. In order to state it, we introduce the empirical correlation for two sequences  $a = (a_i)_{1 \leq i \leq n}$  and  $b = (b_i)_{1 \leq i \leq n}$  with vanishing empirical averages:

$$\operatorname{Corr}(a,b) = \frac{\displaystyle\sum_{i=1}^{n} a_i b_i}{\displaystyle\sqrt{\displaystyle\sum_{i=1}^{n} a_i^2} \sqrt{\displaystyle\sum_{i=1}^{n} b_i^2}}.$$

**Proposition 5.1.** The principal components have a vanishing empirical correlation:

$$\forall 1 \leq k \neq k' \leq d, \qquad \operatorname{Corr}\left(c_k, c_{k'}\right) = 0. \tag{5.9}$$

Moreover, denoting by  $X^j$  the *j*th column of the data matrix X (i.e.  $X^j$  is the vector of the *j*-th features),

$$\forall j \in \{1, \dots, d\}, \qquad \sum_{k=1}^{d} \operatorname{Corr} \left(c_k, X^j\right)^2 = 1.$$
 (5.10)

The proof of this result is the content of Exercise 5.1. The second property corresponds to the so-called *correlation sphere*. In particular, the vector  $(\operatorname{Corr}(c_k, X^j))_{1 \leq k \leq \ell}$  of features which are retained by PCA is necessarily inside the unit ball. For  $\ell = 2$ , the vector  $(\operatorname{Corr}(c_1, X^j), \operatorname{Corr}(c_2, X^j))$  lies in a disk in the so-called first factorial plane. Depending on the position of this vector in the unit ball, various situations can be distinguished:

- if Corr  $(c_1, X^j)$  is close to 1 in absolute value and Corr  $(c_2, X^j)$  is small in absolute value, then the *j*-th feature correlates well with the score on the first loading;
- if Corr  $(c_2, X^j)$  is close to 1 in absolute value and Corr  $(c_1, X^j)$  is small in absolute value, then the *j*-th feature correlates well with the score on the second loading;
- if both Corr  $(c_1, X^j)$  and Corr  $(c_2, X^j)$  are small in absolute value, then the *j*-th feature is not correlated with the scores on the first two loadings.

In essence, this representation allows to check whether a feature correlates more to  $c_1$  or  $c_2$ , or maybe none of them. Once the features are interpreted this way, one can also look at the projection of data points themselves, *i.e.* consider the scatter plot of the vectors  $(c_{1,i}, c_{2,i})$  for  $1 \leq i \leq n$ . This allows to group data points depending on whether they are better explained by  $c_1$  (in which case they would be close to  $(\pm 1, 0)$ ), by  $c_2$  (in which case they would be close to  $(0, \pm 1)$ ), or by none of these quantities. This allows to interpret the data in terms of the relevant features identified by the correlation sphere. It is good to focus on an example to gain some intuition here; for instance the decathlon data, or MNIST digits (see the discussions around [24, Figure 14.23] and [40, Figure 20.2] for instance).

**Exercise 5.1.** The aim of this exercise is to prove Proposition 5.1.

(a) Prove (5.9) by computing  $c_k \cdot c_{k'}$ . (b) Show that Corr  $(c_k, X^j) = \frac{\sqrt{\lambda_k} u_k^j}{\sigma_{n,j}}$ . (c) Prove (5.10).

#### Correction.

(a) Note that

$$\frac{1}{n}c_k \cdot c_{k'} = \frac{1}{n}\sum_{i=1}^n (u_k \cdot x_i)(x_{k'} \cdot x_i) = u_k \left(\frac{1}{n}\sum_{i=1}^n x_i^\top x_i\right) u_{k'}^\top = u_k \widehat{\Sigma} u_{k'}^\top = \lambda_k \delta_{k,k'}.$$

(b) We compute

$$\frac{1}{n}c_k \cdot X^j = \frac{1}{n}\sum_{i=1}^n \left(u_k x_i^{\top}\right) x_i^j = \sum_{k'=1}^d \left(\frac{1}{n}\sum_{i=1}^n x_i^{k'} x_i^j\right) u_k^{k'} = \sum_{k'=1}^d u_k^{k'} \widehat{\Sigma}_{k',j} = \left[u_k \widehat{\Sigma}\right]_j = \lambda_k u_k^j.$$

The conclusion follows by noting that

$$\frac{1}{n} \|c_k\|_2^2 = \lambda_k$$

by the first question, and

$$\frac{1}{n}\left\|X^{j}\right\|^{2} = \sigma_{n,j}^{2}$$

by definition of the empirical variance (recall that the data has been recentered here). (c) The result of the previous question implies that

$$\sum_{k=1}^{d} \operatorname{Corr}\left(c_{k}, X^{j}\right)^{2} = \sum_{k=1}^{d} \frac{\lambda_{k}\left(u_{k}^{j}\right)^{2}}{\sigma_{n,j}^{2}} = \frac{1}{\sigma_{n,j}^{2}} \sum_{k=1}^{d} \left[u_{k}\widehat{\Sigma}\right]_{j} u_{k}^{j} = \frac{1}{\sigma_{n,j}^{2}} \sum_{k,j'=1}^{d} \left[\widehat{\Sigma}\right]_{j',j} u_{k}^{j'} u_{k'}^{j}$$

Now, since  $(u_k)_{1 \leq k \leq d}$  is an orthonormal basis, it holds that

$$\sum_{k=1}^{d} u_k^{\top} u_k = \mathrm{Id}_d$$

as can be checked by applying a line vector x to the left of the previous equality and noting that the left hand side becomes the decomposition of x onto the basis  $(u_k)_{1 \leq k \leq d}$ . It follows that

$$\sum_{k=1}^{d} u_k^{j'} u_k^{j} = \left[ \sum_{k=1}^{d} u_k^{\top} u_k \right]_{j',j} = [\mathrm{Id}_d]_{j',j} = \delta_{j',j}.$$

Therefore,

$$\sum_{k=1}^{d} \operatorname{Corr} \left( c_k, X^j \right)^2 = \frac{1}{\sigma_{n,j}^2} \left[ \widehat{\Sigma} \right]_{j,j} = 1.$$

**Choosing the number**  $\ell$  of latent dimensions. The reconstruction error is always better when more latent dimensions are considered, even for the test data. There is no issue of overfitting. The common practice is to identify a sufficiently small value of  $\ell$  by looking at the scree plot, where eigenvalues  $\lambda_k$  are plotted as a function of k. One identifies a knee or elbow in this curve, where the decrease of the eigenvalues slows down so that the marginal increase in the sum  $\lambda_1 + \cdots + \lambda_k$ becomes smaller with k. Alternatively, this corresponds to looking at the fraction of explained variance

$$f_k = \frac{\sum_{k'=1}^{\kappa} \lambda_{k'}}{\sum_{k'=1}^{d} \lambda_{k'}}$$

as a function of k, and identifying the index  $\ell$  after which the gains become smaller.

# 5.5 Extensions of PCA

We present in this section various extensions of PCA. We start by discussing in Section 5.5.1 how to adapt the method to high dimensional data (*i.e.* situations in which the dimension d of the features is much larger than the number n of data points). We next show how to include some nonlinearity by using kernel functions, as made precise in Section 5.5.2. This method can be thought of nonlinearly pre-processing the data and then performing some linear dimensionality reduction. A genuinely nonlinear extension of PCA, based on autoencoder neural networks, is discussed later on in Section 8.3. Finally, we present in Section 5.5.3 a generative method based on PCA, known as probabilitic PCA.

#### 5.5.1 High dimensional data

PCA requires diagonalizing the matrix  $\widehat{\Sigma} \in \mathbb{R}^{d \times d}$  defined in (5.5), which has a computational cost  $O(d^3)$ . When there are many features d, even more than the number of data points n, the diagonalization of the matrix can become too expensive, and it would be better to reformulate PCA with a digonalization of a  $n \times n$  matrix, similarly to what is done for linear least square regression in Exercise 2.10. The key point here is to note that the eigenvalue problem (5.8) can be rewritten as follows, upon introducing the design matrix  $X \in \mathbb{R}^{n \times d}$  whose *i*th row is  $x_i$  and multiplying (5.8) by  $X^{\top}$  on the right:

$$uX^{\top}\Gamma = \lambda uX^{\top}, \qquad \Gamma = \frac{1}{n}XX^{\top} \in \mathbb{R}^{n \times n}.$$

More explicitly,  $\Gamma_{ij} = x_i x_j^{\top}$ . This suggests to solve eigenvalue problems for  $\Gamma$ , namely find  $v \in \mathbb{R}^{1 \times n} \setminus \{0\}$  and  $\lambda > 0$  such that  $v\Gamma = \lambda v$ . Then,

$$vX\widehat{\Sigma} = v\Gamma X = \lambda vX,$$

so that vX is an eigenvector of  $\widehat{\Sigma}$ . Note that  $vX \neq 0$  otherwise  $v\Gamma = 0 = \lambda v$ , so v = 0, which is not possible. This shows that normalized eigenvectors u of  $\widehat{\Sigma} \in \mathbb{R}^{d \times d}$  are obtained from eigenvectors v of  $\Gamma \in \mathbb{R}^{n \times n}$  by multiplying them by X on the right, and renormalizing the so-obtained vector as  $vX/||vX||_2 \in \mathbb{R}^{1 \times d}$ .

#### 5.5.2 Kernel PCA

Kernel methods allow to go beyond linear methods by introducing some nonlinearity through an implicit featurization procedure. We here only briefly discuss how this works, as a more complete presentation is provided in Section 6.4.2.

Before presenting the implicit featurization, let us start by considering an explicit featurization based on some feature function  $\phi : \mathcal{X} \to \mathbb{R}^{1 \times D}$ . In this context, the design matrix  $\widehat{\Sigma}$  in (5.5) is replaced by

$$\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^{n} \phi(x_i)^{\top} \phi(x_i) \in \mathbb{R}^{D \times D}.$$

In fact, some recentering has to be performed in order for this matrix to be interpreted as some covariance matrix, see the discussion in [40, Section 20.4.6].

The value of D can be large, in fact infinite, but, as in Section 5.5.1, this will not pose an issue as one can formule PCA based on the diagonalization of a  $n \times n$  matrix. The eigenvalue equation to be solved in this context is

$$\lambda v = v\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^{n} \left( v\phi(x_i)^{\top} \right) \phi(x_i) \in \mathbb{R}^{1 \times D}.$$

This shows that v is a linear combination of  $\phi(x_1), \ldots, \phi(x_n)$ , and can therefore be written as

$$v = \sum_{i=1}^{n} a_i \phi(x_i).$$

By plugging this equality in the eigenproblem to solve, one finds

$$\frac{1}{n}\sum_{i,j=1}^{n}a_{j}\left(\phi(x_{j})\phi(x_{i})^{\top}\right)\phi(x_{i}) = \lambda\sum_{j=1}^{n}a_{j}\phi(x_{j}).$$

We multiply this equality by  $\phi(x_k)^{\top}$  on the right, and introduce the symmetric matrix  $K \in \mathbb{R}^{n \times n}$  with entries  $K_{ij} = \mathcal{K}(x_i, x_j)$ , where

$$\mathcal{K}(x, x') = \phi(x)\phi(x')^{\top} \in \mathbb{R}.$$

This leads to

$$\frac{1}{n}\sum_{i,j=1}^{n}a_{j}K_{ji}K_{ik} = \lambda\sum_{j=1}^{n}a_{j}K_{jk}$$

Since  $1 \leq k \leq n$  is arbitrary, the eigenvalue problem can finally be reformulated as

$$aK^2 = n\lambda aK, \qquad a \in \mathbb{R}^{1 \times n}, \qquad K = K^\top \in \mathbb{R}^{n \times n},$$

which can be solved by diagonalizing the matrix K. Once the matrix is diagonalized and the  $\ell$  leading eigenvectors  $a_1, \ldots, a_\ell$  are found, reconstruction can be performed with the orthonormal basis of vectors

$$u_{k} = \frac{\sum_{i=1}^{n} a_{k,i}\phi(x_{i})}{\left\|\sum_{i=1}^{n} a_{k,i}\phi(x_{i})\right\|_{2}}.$$

More precisely, the reconstruction is based on the scores

$$u_{k}\phi(x')^{\top} = \frac{\sum_{i=1}^{n} a_{k,i}\phi(x_{i})\phi(x')^{\top}}{\left\|\sum_{i=1}^{n} a_{k,i}\phi(x_{i})\right\|_{2}}.$$

Since

$$\left\|\sum_{i=1}^{n} a_{k,i} \phi(x_i)\right\|_{2}^{2} = \sum_{i,j=1}^{n} a_{k,i} a_{k,j} \phi(x_i) \phi(x_j)^{\top} = a_k K a_k^{\top},$$

we finally obtain

$$u_k \phi\left(x'\right)^{\top} = \frac{\sum_{i=1}^n a_{k,i} \mathcal{K}\left(x', x_i\right)}{\sqrt{a_k K a_k^{\top}}}.$$

Note that, crucially, the explicit featurization  $\phi$  is not needed, as the only quantity which appears is the kernel function  $\mathcal{K}(x, x')$ . It therefore suffices to make precise this kernel function. As mentioned above, we elaborate on this in Section 6.4.2.

In the context of kernel PCA, the test data point x' is replaced by the knowledge of the  $\ell$  scores  $u_k \phi(x')^{\top}$  for  $1 \leq k \leq \ell$ . The reconstruction of the test point x' itself is not straightforward as it is based on some form of regression problem, see for instance the discussion in [6, Section 15.7]. Applications of kernel PCA are given in the work [47] which introduced the method. One major aim is to featurize the data in a nonlinear way as postprocessing step, for example for denoising images.

#### 5.5.3 Probabilistic PCA

PCA allows to embed training points  $x_i$  into a lower dimensional space, but does not allow to generate new points. Probabilistic PCA (PPCA) is a generative method which provides a simple model for  $p_{\text{data}}(x)$ , known as "predictive distribution". More precisely, considering in this section realizations under the predictive distribution as column vectors, the model is obtained as

$$X = WZ + \varepsilon + b \in \mathbb{R}^d, \qquad Z \sim \mathcal{N}(0, \mathrm{Id}_\ell), \qquad \varepsilon \sim \mathcal{N}(0, \sigma^2 \mathrm{Id}_d),$$

for  $\sigma > 0$  and some matrix  $W \in \mathbb{R}^{d \times \ell}$ , and with  $\varepsilon$  independent of Z. Note that X is a Gaussian random vector, with mean b and covariance matrix  $WW^{\top} + \sigma^2 \mathrm{Id}_d$  since

$$\mathbb{E}[X] = \mathbb{E}\left[WZ + \varepsilon + b\right] = W\mathbb{E}[Z] + b = b,$$

and

$$\operatorname{Var}(X) = \mathbb{E}\left[ \left( WZ + \varepsilon \right) \left( WZ + \varepsilon \right)^{\top} \right] = \mathbb{E}\left[ WZZ^{\top}W^{\top} \right] + \mathbb{E}\left[ \varepsilon \varepsilon^{\top} \right] = W\mathbb{E}\left[ ZZ^{\top} \right] W^{\top} + \sigma^{2} \operatorname{Id}_{d}.$$

Probabilistic PCA is therefore fully determined by the parameters W, b and  $\sigma^2$ . These parameters can be found by a maximum likelihood approach (either by reducing the question to an eigenvalue problem, or by using an expectation/maximization approach for high dimensional problems; see [8, Section 12.2]).

Factor analysis is a generative extension of PCA, which goes beyond probabilistic PCA by working with a diagonal covariance matrix  $\mathcal{V}$  instead of an isotropic covariance  $\sigma^2 \mathrm{Id}_d$ . The generative model under consideration is obtained by decoding a latent variable distributed according to an isotropic Gaussian random variable, then adding a bias to recenter the data, as well as some Gaussian noise whose intensity depends on the coordinate at hand; see Exercise 5.2 below.

Exercise 5.2 (Factor analysis). We consider the generative model

$$X = WZ + \varepsilon + b, \qquad Z \sim \mathcal{N}(0, \mathrm{Id}_{\ell}), \qquad \varepsilon \sim \mathcal{N}(0, \mathcal{V}), \tag{5.11}$$

with  $Z, \varepsilon$  independent random variables,  $X, b \in \mathbb{R}^d$ ,  $W \in \mathbb{R}^{d \times \ell}$ ,  $\mathcal{V} \in \mathbb{R}^{d \times d}$ ,  $Z \in \mathbb{R}^{\ell}$ , with  $\ell \leq d-1$  the dimension of the latent space. The covariance matrix for the noise variable  $\varepsilon$  is assumed to be diagonal:

$$\mathcal{V} = \begin{pmatrix} v_1 \ 0 \ \dots \\ 0 \ v_2 \ 0 \ \dots \\ & \ddots \\ 0 \ \dots \ 0 \ v_d \end{pmatrix},$$

with entries  $v_i \ge 0$ . The idea is to determine  $W, b, \mathcal{V}$  from the dataset  $\{x_i\}_{1 \le i \le n} \subset \mathbb{R}^d$  by some maximum likelihood procedure (as for PCA, we work here in an unsupervised setting, so the data set does not come with labels).

- (a) Prove that X is a Gaussian random variable with mean b and covariance  $WW^{\top} + \mathcal{V}$ . How many free parameters enter the definition of the covariance matrix? How does this number compare to the number of parameters used to model the covariance of X without any structural assumption?
- (b) Show that the generative model is unchanged if W is replaced by WR, with  $R \in \mathbb{R}^{\ell \times \ell}$  an orthogonal matrix (i.e. a matrix R such that  $R^{\top}R = RR^{\top} = \mathrm{Id}_{\ell}$ ).

The previous question suggests that it is not relevant to look for W alone. In view of the previous question, a better quantity to consider is the part  $WW^{\top}$  of the covariance  $\Sigma_W = WW^{\top} + \mathcal{V}$ . We assume in the remainder of this exercise that  $v_k > 0$  for any  $1 \leq k \leq d$ .

- (c) Prove that  $\Sigma_W$  is invertible.
- (d) Write the negative log-likelihood  $L(W, \mathcal{V}, b)$  corresponding to the dataset  $\{x_i\}_{1 \leq i \leq n}$  for the model (5.11) (assuming that the data points are independent).
- (e) Find the expression of b by minimizing  $L(W, \mathcal{V}, b)$  for  $W, \mathcal{V}$  fixed.
- (f) Show that finding W, V amounts to minimizing

$$\mathscr{L}(W, \mathcal{V}) = \operatorname{Tr}\left(\Sigma_W^{-1}\widehat{C}\right) + \ln\left(\det \Sigma_W\right)$$

where  $\widehat{C} \in \mathbb{R}^{d \times d}$  is the empirical covariance of the data set  $\{x_i\}_{1 \leq i \leq n}$ .

The algorithm to find  $WW^{\top}$  and  $\mathcal{V}$  proceeds in an iterative manner, by alternately minimizing first over  $\mathcal{V}$  for W fixed, then over W for  $\mathcal{V}$  fixed, until convergence is reached. We make precise how to perform these two steps in the next questions.

The following three questions are more difficult and can be skipped. The last questions should however be treated.

(g) We start by minimizing  $\mathscr{L}(W, \mathcal{V})$  for W fixed, assuming that this minimization problem is well posed. Using the following identities for symmetric matrices X, M (see for instance [40, Section 7.8.7.3])

$$\nabla_X \left[ \operatorname{Tr} \left( X^{-1} M \right) \right] = -X^{-1} M X^{-1}, \qquad \nabla_X \left[ \ln(\det X) \right] = X^{-1}, \tag{5.12}$$

show that the minimizer  $\mathcal{V}$  satisfies

$$\operatorname{diag}\left(\Sigma_W^{-1}\widehat{C}\Sigma_W^{-1}\right) = \operatorname{diag}\left(\Sigma_W^{-1}\right),\tag{5.13}$$

where diag(M) for  $M \in \mathbb{R}^{d \times d}$  is the diagonal matrix with entries  $M_{kk}$  on the diagonal.

(h) We fix  $\mathcal{V}$ , and consider the change of unknown matrix  $\mathscr{W} = U\mathcal{V}^{-1/2}W \in \mathbb{R}^{d \times \ell}$ , where U is an orthogonal matrix obtained from the diagonalization of the symmetric positive matrix  $\mathcal{V}^{-1/2}\widehat{C}\mathcal{V}^{-1/2}$ , i.e.

$$\mathcal{V}^{-1/2}\widehat{C}\mathcal{V}^{-1/2} = U^{\top}\Lambda U, \tag{5.14}$$

where  $U \in \mathbb{R}^{d \times d}$  is an orthogonal matrix, and  $\Lambda \in \mathbb{R}^{d \times d}$  is a diagonal matrix with entries  $\lambda_i \ge 0$ , ranked in order of decreasing magnitudes. Show that the minimization of  $\mathscr{L}(W, \mathcal{V})$  over  $W \in \mathbb{R}^{d \times \ell}$  is equivalent to the minimization over  $\mathscr{W} \in \mathbb{R}^{d \times \ell}$  of

$$\mathfrak{L}_{\mathcal{V}}(\mathscr{W}) = \operatorname{Tr}\left[\left(\mathscr{W}\mathscr{W}^{\top} + \operatorname{Id}_{d}\right)^{-1}\Lambda\right] + \ln\left[\det\left(\mathscr{W}\mathscr{W}^{\top} + \operatorname{Id}_{d}\right)\right].$$

2

Ş

(i) Using (5.12), it can be shown that minimizers of  $\mathfrak{L}_{\mathcal{V}}(\mathcal{W})$  satisfy

$$\left(\mathscr{W}\mathscr{W}^{\top} + \mathrm{Id}_{d}\right)^{-1} \Lambda \left(\mathscr{W}\mathscr{W}^{\top} + \mathrm{Id}_{d}\right)^{-1} \mathscr{W} = \left(\mathscr{W}\mathscr{W}^{\top} + \mathrm{Id}_{d}\right)^{-1} \mathscr{W}.$$
 (5.15)

This equality can be reformulated in terms of the variable W as  $W = \widehat{C}\Sigma_W^{-1}W$ . Use the latter identity and (5.13) to establish that diag  $(\Sigma_W^{-1}\widehat{C} - \mathrm{Id}_d) = 0$ . By multiplying the previous equation by  $\mathcal{V}$  and using the identity  $\mathcal{V}\Sigma_W^{-1} = \mathrm{Id}_d - WW^{\top}\Sigma_W^{-1}$ , prove finally that

$$\mathcal{V} = \operatorname{diag}\left(\widehat{C} - WW^{\top}\right). \tag{5.16}$$

Upon multiplying (5.15) by  $\mathscr{W}\mathscr{W}^{\top} + \mathrm{Id}_d$  on the left, and by  $\mathscr{W}^{\top} (\mathscr{W}\mathscr{W}^{\top} + \mathrm{Id}_d)$  on the right, it follows that

$$\Lambda \mathscr{W} \mathscr{W}^{\top} = \left( \mathscr{W} \mathscr{W}^{\top} + \mathrm{Id}_d \right) \mathscr{W} \mathscr{W}^{\top}.$$
(5.17)

The symmetric positive matrix  $\mathscr{W}\mathscr{W}^{\top} \in \mathbb{R}^{d \times d}$  can be diagonalized as  $\mathscr{W}\mathscr{W}^{\top} = V^{\top}DV$ , for some diagonal matrix D with nonnegative entries  $d_k$  (ranked in decreasing order), and an orthogonal matrix V.

- (j) Consider an eigenvector  $\xi_k$  of  $\mathscr{W}\mathscr{W}^{\top}$  associated with the eigenvalue  $d_k$ . Write the equation on  $\xi_k, d_k$  obtained from (5.17).
- (k) How should the eigenvalues  $d_k$  be chosen in order to minimize  $\mathfrak{L}_{\mathcal{V}}(\mathscr{W})$ ?

In order to find W for  $\mathcal{V}$  fixed, one can choose

$$W = \mathcal{V}^{1/2} U^{\top} \begin{pmatrix} | & | & \cdots & | \\ \xi_1 & \xi_2 & \cdots & \xi_\ell \\ | & | & \cdots & | \end{pmatrix} \operatorname{diag} \left( \sqrt{\lambda_1 - 1}, \dots, \sqrt{\lambda_\ell - 1} \right),$$
(5.18)

with  $\{\xi_k\}_{1 \leq k \leq \ell}$  the eigenvectors of  $\Lambda$  associated with the largest eigenvalues. All the quantities in the previous equation depend only on  $\mathcal{V}$  and  $\widehat{C}$ . The equation (5.18), together with (5.16), forms the basis for the iterative method discussed in [58, Section 2.2] to find the parameters in factor analysis. One then alternates between (5.16) (which says how to update  $\mathcal{V}$  for W given) and (5.18) (which says how to update W for  $\mathcal{V}$  given).

#### Correction.

(a) The random variable X is a linear combination of independent Gaussian random variables, so it is also a Gaussian random variable. To fully characterize its law, it suffices to compute its mean and covariance. Note first that  $\mathbb{E}(X) = b$  since  $\mathbb{E}(\varepsilon_j) = 0$  for any  $1 \leq j \leq d$ and  $\mathbb{E}(Z_k) = 0$  for any  $1 \leq k \leq \ell$ . Moreover, since the random variables  $Z, \varepsilon$  are independent and of mean 0,

$$Cov(X) = \mathbb{E}\left[ (X - b)(X - b)^{\top} \right] = \mathbb{E}\left[ (WZ + \varepsilon)(WZ + \varepsilon)^{\top} \right]$$
$$= W\mathbb{E}\left( ZZ^{\top} \right) W^{\top} + \mathbb{E}\left( \varepsilon\varepsilon^{\top} \right) = WW^{\top} + \mathcal{V}.$$

The free parameters correspond to the *d* diagonal entries of  $\mathcal{V}$ , and the  $d \times \ell$  entries of *W*. There are therefore  $d(\ell + 1)$  free parameters. When  $\ell \ll d$ , this is much smaller than the d(d+1)/2 parameters that would be used to model a full covariance matrix for *X*.

- (b) When W is replaced by WR, the random variable X is still Gaussian with mean 0. Its covariance is, in view of the previous question,  $WRR^{\top}W^{\top} + \mathcal{V} = WW^{\top} + \mathcal{V}$  because R is orthogonal. This shows that the law of X is unchanged, *i.e.* the generative model is unchanged.
- (c) Consider  $\xi \in \mathbb{R}^d$  such that  $\Sigma_W \xi = 0$ . Then,

$$0 = \xi^{\top} \Sigma_W \xi = \left\| W^{\top} \xi \right\|^2 + \xi^{\top} \mathcal{V} \xi \ge \sum_{k=1}^d v_k \xi_k^2.$$

Since  $v_k > 0$  for all  $1 \leq k \leq d$ , this shows that  $\xi_k = 0$  for all  $1 \leq k \leq d$ . In conclusion,  $\xi = 0$ , so that the matrix  $\Sigma_W$  has null kernel and is hence invertible.

(d) The likelihood for a given data point  $x \in \mathbb{R}^d$  for the distribution  $\mathcal{N}(b, \Sigma_W)$  is

$$p_{b,\Sigma_W}(x) = (2\pi)^{-d/2} (\det \Sigma_W)^{-1/2} \exp\left(-\frac{1}{2}(x-b)^\top \Sigma_W^{-1}(x-b)\right)$$

The negative log-likelihood for the data set is therefore

$$L(W, \mathcal{V}, b) = -\sum_{i=1}^{n} \log p_{b, \mathcal{L}_{W}}(x_{i}) = \sum_{i=1}^{n} \frac{1}{2} (x_{i} - b)^{\top} \mathcal{L}_{W}^{-1}(x_{i} - b) + \frac{n}{2} \log(\det \mathcal{L}_{W}) + \frac{nd}{2} \log(2\pi).$$

(e) The function  $b \mapsto L(W, \mathcal{V}, b)$  is a quadratic function, with positive definite Hessian  $n\Sigma_W^{-1}$ (in view of Question (c)). It is therefore strongly convex, hence the minimization of  $b \mapsto L(W, \mathcal{V}, b)$  is well posed (alternatively, one could see that the function is stricly convex, continuous, and coercive). The minimizer  $b_\star$  for  $W, \mathcal{V}$  fixed is characterized by the condition

$$0 = \nabla_b L(W, \mathcal{V}, b_\star) = \sum_{i=1}^n \Sigma_W^{-1}(x_i - b_\star).$$

Since  $\Sigma_W$  is invertible, this shows that

$$b_{\star} = \frac{1}{n} \sum_{i=1}^{n} x_i$$

is the empirical mean of the data points. Note that  $b_{\star}$  does not depend on  $W, \mathcal{V}$ .

(f) By plugging the expression of  $b_{\star}$  in  $L(W, \mathcal{V}, b)$ , we see that  $W, \mathcal{V}$  are found by minimizing the following function (where we removed a constant term and multiplied by a factor 2/n):

$$\frac{1}{n}\sum_{i=1}^{n}(x_{i}-b_{\star})^{\top}\boldsymbol{\Sigma}_{W}^{-1}(x_{i}-b_{\star})+\log(\det\boldsymbol{\Sigma}_{W})=\operatorname{Tr}\left(\boldsymbol{\Sigma}_{W}^{-1}\widehat{C}\right)+\log(\det\boldsymbol{\Sigma}_{W}),$$

where

$$\widehat{C} = \frac{1}{n} \sum_{i=1}^{n} (x_i - b_\star) (x_i - b_\star)^\top \in \mathbb{R}^{d \times d}$$

is the empirical covariance of the data set.

(g) The derivatives (5.13) imply that, for any diagonal matrix  $\mathscr{D} \in \mathbb{R}^{d \times d}$ ,

$$\langle \nabla_{\mathcal{V}} \mathscr{L}(W, \mathcal{V}), \mathscr{D} \rangle = \operatorname{Tr} \left( \nabla_{\mathcal{V}} \mathscr{L}(W, \mathcal{V})^{\top} \mathscr{D} \right) = -\operatorname{Tr} \left( \Sigma_{W}^{-1} \widehat{C} \Sigma_{W}^{-1} \mathscr{D} \right) + \operatorname{Tr} \left( \Sigma_{W}^{-1} \mathscr{D} \right).$$

- This immediately gives the desired equality since the gradient vanishes at critical points. (h) The expression of  $\mathfrak{L}_{\mathcal{V}}(\mathscr{W})$  is obtained by noting that  $W = \mathcal{V}^{1/2} U^{\top} \mathscr{W}$  and recalling  $U^{-1} = U^{\top}$ , so that
  - $$\begin{split} \boldsymbol{\Sigma}_{W}^{-1}\widehat{\boldsymbol{C}} &= \left(\boldsymbol{\mathcal{V}}^{1/2}\boldsymbol{U}^{\top}\boldsymbol{\mathscr{W}}\boldsymbol{\mathscr{W}}^{\top}\boldsymbol{U}\boldsymbol{\mathcal{V}}^{1/2} + \boldsymbol{\mathcal{V}}\right)^{-1}\widehat{\boldsymbol{C}} = \boldsymbol{\mathcal{V}}^{-1/2}\left(\boldsymbol{U}^{\top}\boldsymbol{\mathscr{W}}\boldsymbol{\mathscr{W}}^{\top}\boldsymbol{U} + \mathrm{Id}_{d}\right)^{-1}\boldsymbol{\mathcal{V}}^{-1/2}\widehat{\boldsymbol{C}} \\ &= \boldsymbol{\mathcal{V}}^{-1/2}\boldsymbol{U}^{\top}\left(\boldsymbol{\mathscr{W}}\boldsymbol{\mathscr{W}}^{\top} + \mathrm{Id}_{d}\right)^{-1}\boldsymbol{U}\boldsymbol{\mathcal{V}}^{-1/2}\widehat{\boldsymbol{C}} = \boldsymbol{\mathcal{V}}^{-1/2}\boldsymbol{U}^{\top}\left(\boldsymbol{\mathscr{W}}\boldsymbol{\mathscr{W}}^{\top} + \mathrm{Id}_{d}\right)^{-1}\boldsymbol{\Lambda}\boldsymbol{U}\boldsymbol{\mathcal{V}}^{1/2}, \end{split}$$

where we used that  $U\mathcal{V}^{-1/2}\widehat{C} = \Lambda U\mathcal{V}^{1/2}$  in the last line, an equality obtained from (5.14). Therefore, by the cyclicity of the trace,

$$\operatorname{Tr}\left(\Sigma_{W}^{-1}\widehat{C}\right) = \operatorname{Tr}\left[\left(\mathscr{W}\mathscr{W}^{\top} + \operatorname{Id}_{d}\right)^{-1}\Lambda\right]$$

Similar computations show that

$$\begin{split} \boldsymbol{\varSigma}_{W} &= \boldsymbol{\mathcal{V}}^{1/2} \left( \boldsymbol{U}^{\top} \boldsymbol{\mathscr{W}} \boldsymbol{\mathscr{W}}^{\top} \boldsymbol{U} + \mathrm{Id}_{d} \right) \boldsymbol{\mathcal{V}}^{1/2} = \boldsymbol{\mathcal{V}}^{1/2} \left( \boldsymbol{U}^{\top} \boldsymbol{\mathscr{W}} \boldsymbol{\mathscr{W}}^{\top} \boldsymbol{U} + \boldsymbol{U}^{\top} \boldsymbol{U} \right) \boldsymbol{\mathcal{V}}^{1/2} \\ &= \boldsymbol{\mathcal{V}}^{1/2} \boldsymbol{U}^{\top} \left( \boldsymbol{\mathscr{W}} \boldsymbol{\mathscr{W}}^{\top} + \mathrm{Id}_{d} \right) \boldsymbol{U} \boldsymbol{\mathcal{V}}^{1/2}, \end{split}$$

so that

$$\det \Sigma_W = \det(\mathcal{V}) \det \left[ U^\top \left( \mathscr{W} \mathscr{W}^\top + \mathrm{Id}_d \right) U \right] = \det(\mathcal{V}) \det \left( \mathscr{W} \mathscr{W}^\top + \mathrm{Id}_d \right)$$

Since the term det  $\mathcal{V}$  is an unimportant multiplicative factor (which becomes a constant additive term when taking the logarithm), we obtain the desired conclusion.

(i) We multiply (5.13) by the diagonal matrix  $\mathcal{V}$  (noting that  $\operatorname{diag}(M)\mathcal{V} = \operatorname{diag}(M\mathcal{V})$ ) and then use the identity  $\Sigma_W^{-1}\mathcal{V} = \operatorname{Id}_d - \Sigma_W^{-1}WW^{\top}$  to write

$$0 = \operatorname{diag}\left(\Sigma_W^{-1}\widehat{C}\Sigma_W^{-1}\mathcal{V} - \Sigma_W^{-1}\mathcal{V}\right) = \operatorname{diag}\left(\Sigma_W^{-1}\widehat{C} - \Sigma_W^{-1}\widehat{C}\Sigma_W^{-1}WW^{\top} - \operatorname{Id}_d + \Sigma_W^{-1}WW^{\top}\right).$$

The second and fourth terms on the right hand side of the previous equality cancel since  $W = \widehat{C} \Sigma_W^{-1} W$ . Therefore,

$$\operatorname{diag}\left(\Sigma_W^{-1}\widehat{C} - \operatorname{Id}_d\right) = 0$$

Remultiplying this equation by  $\mathcal{V}$  on the left and using the identity  $\mathcal{V}\Sigma_W^{-1} = \mathrm{Id}_d - WW^{\top}\Sigma_W^{-1}$ , we obtain

$$\mathcal{V} = \operatorname{diag}\left(\mathcal{V}\mathcal{\Sigma}_{W}^{-1}\widehat{C}\right) = \operatorname{diag}\left[\left(\operatorname{Id}_{d} - WW^{\top}\mathcal{\Sigma}_{W}^{-1}\right)\widehat{C}\right] = \operatorname{diag}\left(\widehat{C} - WW^{\top}\right),$$

where the last equality relies again on  $W = \widehat{C} \Sigma_W^{-1} W$ .

(j) By applying both sides of the equality (5.17) to an eigenvector  $\xi_k \neq 0$ , it follows that

$$d_k\Lambda\xi_k = (d_k+1)d_k\xi_k.$$

This shows that either  $d_k = 0$ , or  $\xi_k$  is an eigenvector of the diagonal matrix  $\Lambda$  and  $d_k = \lambda_k - 1 \ge 0$ .

(k) By computing the trace using an orthonormal eigenbasis of  $\mathscr{W}\mathscr{W}^{\top}$ ,

$$\mathfrak{L}_{\mathcal{V}}(\mathscr{W}) = \sum_{k=1}^{d} \frac{\lambda_k}{1+d_k} + \log(1+d_k).$$

Now,

$$\frac{\lambda_k}{1+d_k} + \log(1+d_k) = \begin{cases} 1+\log(\lambda_k) & \text{if } d_k = \lambda_k - 1, \\ \lambda_k & \text{if } d_k = 0. \end{cases}$$

Since  $\log \lambda < \lambda$  for any  $\lambda > 0$ , it follows that, in order to minimize  $\mathfrak{L}_{\mathcal{V}}(\mathscr{W})$ , one should set  $d_k = \lambda_k - 1$  for  $1 \leq k \leq \ell$ , and  $d_k = 0$  for  $\ell + 1 \leq k \leq d$ .

# Support vector machines

6.1	Linea	r classification for separable data sets	
	6.1.1	Hard margin SVM	
	6.1.2	Primal optimization problem 100	
	6.1.3	Dual formulation	
	6.1.4	Optimization algorithms 103	
	6.1.5	Learning guarantees by leave-one-out analysis	
6.2	Linea	r classification for non separable data sets 105	
	6.2.1	Soft margin SVM 105	
	6.2.2	Characterization of minimizers 106	
	6.2.3	Dual formulation 107	
6.3	Multi	Multiclass SVM 108	
6.4	Kerne	el SVM 109	
	6.4.1	Principle of the method 109	
	6.4.2	Kernel methods 110	

We present in this chapter a method to perform classification by linearly separating data sets, using the method called support vector machines. The name of this method comes from the fact that a small or sparse subset of the training points is retained to construct decision functions whose signs allow to make predictions for the labels of new data points.

We start in Section 6.1 by describing the method for data sets which can be linearly separated in an exact manner; and then make precise in Section 6.2 how to extend the approach to situations where there are some outliers in the data set which prevent an exact linear separation. We also discuss options to perform multiclass classification in Section 6.3. Our presentation for these sections is based on [40, Section 17.3], [39, Chapter 5] and [8, Section 7.1].

In fact, most data sets are not linearly separable as such, even in an approximate manner. Some form of featurization is needed to transform the data. We discuss in Section 6.4 how to perform SVM with implicit features based on kernel methods (see [40, Section 17.1], [39, Chapter 6], [49, Chapter 16] and [4, Chapter 7] for further references on kernel methods).

## 6.1 Linear classification for separable data sets

We work here in the context of supervised learning, and aim at performing classification for binary data, with label space  $\mathcal{Y} = \{-1, 1\}$  (see Section 6.3 for an extension to multiclass classification). We consider a training data set  $\mathcal{D}_{\text{train}} = \{(x_i, y_i), 1 \leq i \leq n\} \subset \mathcal{X} \times \mathcal{Y}$ , with  $\mathcal{X} \subset \mathbb{R}^d$ . As for various other learning methods, it is useful to normalize/standardize the inputs before using the methods described in this chapter.

The aim is to find a prediction function minimizing the risk  $\mathcal{R}(f) = \mathbb{E}[\mathbf{1}_{y \neq f(x)}]$ . The method described here is based on linear classifiers, *i.e.* functions  $x \mapsto \operatorname{sign}(w^{\top}x+b)$  with  $w \in \mathbb{R}^d$  and  $b \in \mathbb{R}$ .

100 6 Support vector machines

We first make precise the geometric interpretation of the problem in Section 6.1.1, and then formulate the associated minimization of the loss in primal and dual forms, respectively in Sections 6.1.2 and 6.1.3. We then discuss how to train the model in Section 6.1.4, and conclude by providing theoretical guarantees on the learning problem in Section 6.1.5.

#### 6.1.1 Hard margin SVM

We assume in all this section that the data set  $\mathcal{D}_{\text{train}}$  can be linearly separated, *i.e.* that there exists  $\theta = (w, b) \in \mathbb{R}^d \times \mathbb{R}$  such that

$$\forall 1 \leq i \leq n, \qquad y_i \left( w^\top x_i + b \right) > 0. \tag{6.1}$$

Data points with labels -1 are therefore separated from data points with labels 1 by the hyperplane of equation  $w^{\top}x + b = 0$  (also called "decision boundary"). However, there may be infinitely many such separating hyperplanes. The idea is to choose the one with the largest margin, *i.e.* the distance of the closest point of the data set to the separating plane should be as large as possible in order to obtain the solution to the classification problem which is the most robust to perturbations in the inputs  $(x_i)_{1 \le i \le n}$  of the data set.

**Definition 6.1 (Geometric margin).** Consider  $w \in \mathbb{R}^d \setminus \{0\}$  and  $b \in \mathbb{R}$ . The geometric margin  $\rho_g$  of a linear classifier  $g(x) = w^{\top}x + b$  at a point  $x \in \mathbb{R}^d$  is its Euclidean distance to the hyperplane  $g^{-1}\{0\} = \{\tilde{x} \in \mathbb{R}^d \mid w^{\top}\tilde{x} + b = 0\}$ :

$$\rho_g(x) = \frac{|w^\top x + b|}{\|w\|_2}.$$
(6.2)

The geometric margin for a data set  $\mathcal{D}_{train}$  is

$$\rho_g(\mathcal{D}_{\text{train}}) = \min_{1 \le i \le n} \rho_g(x_i).$$

The equality (6.2) is not immediately clear. The reader is asked to prove it in the next exercise.

**Exercise 6.1.** *Prove* (6.2).

**Correction.** Decompose x as  $x_{\perp} + rw/||w||_2$ , with  $x_{\perp}$  the orthogonal projection of x onto the affine subspace  $g^{-1}\{0\}$ . A simple computation shows that  $g(x) = g(x_{\perp}) + r||w||_2 = r||w||_2$ , and  $\rho_g(x) = ||x - x_{\perp}||_2 = |r|$ . Therefore,

$$\rho_g(x) = \frac{|g(x)|}{\|w\|_2},$$

which is the claimed result.

#### 6.1.2 Primal optimization problem

SVM is based on finding the maximum margin hyperplane. Since  $|w^{\top}x_i + b| = y_i(w^{\top}x_i + b)$  in view of (6.1), the maximization of the margin for the training data set can be formulated as

$$\max_{(w,b)\in\mathbb{R}^d\times\mathbb{R}}\min_{1\leqslant i\leqslant n}\left\{\frac{y_i(w^\top x_i+b)}{\|w\|_2} \mid \forall 1\leqslant i\leqslant n, \ y_i(w^\top x_i+b)>0\right\}.$$

This optimization problem can be seen as a maximization over (w, b) with *n* inequality constraints. Since  $y_i(w^{\top}x_i + b)/||w||_2$  is invariant by rescaling of w, b (*i.e.* the function value is the same if w, b are replaced by  $\alpha w, \alpha b$  with  $\alpha > 0$ ), it is possible to normalize the learning problem by requiring that

$$\min_{1 \le i \le n} y_i(w^\top x_i + b) = 1$$

This leads to the reformulation

$$\max_{(w,b)\in\mathbb{R}^d\times\mathbb{R}}\left\{\frac{1}{\|w\|_2} \mid \forall 1\leqslant i\leqslant n, \ y_i(w^\top x_i+b)\geqslant 1\right\}.$$

In fact, as maximizing  $1/||w||_2$  is equivalent to minimizing  $||w||_2^2$ , the optimization problem can finally be stated in primal form as

$$\min_{(w,b)\in\mathbb{R}^d\times\mathbb{R}}\left\{\frac{1}{2}\|w\|_2^2 \mid \forall 1\leqslant i\leqslant n, \ y_i\left(w^\top x_i+b\right)\geqslant 1\right\}.$$
(6.3)

Well posedness. The optimization problem (6.3) is the minimization of the convex function  $(w,b) \mapsto ||w||_2^2$  under affine constraints  $c_i(w,b) = y_i(w^{\top}x_i + b) - 1 \ge 0$ . The constraints are therefore qualified. Showing the existence of a solution however requires some care as the functional to minimize is not coercive in the *b* variable. We prove to this end that the minimization can be restricted to a bounded set, in which case the existence of a minimizer is clear since the function to minimize is continuous.

To prove that the minimization can be restricted to a bounded set, we consider an element  $(\tilde{w}, \tilde{b}) \in \mathbb{R}^d \times \mathbb{R}$  which satisfies the constraints  $c_i(\tilde{w}, \tilde{b}) \ge 0$  for all  $1 \le i \le n$ . It is then sufficient to restrict the minimization over  $w \in \mathbb{R}^d$  to the minimization over  $w \in \overline{B(0, R)}$  with R = $\|\tilde{w}\|_2 < +\infty$ . Next, when  $y_i = 1$ , the condition  $c_i(w, b) \ge 0$  is equivalent to  $-b \le w^\top x_i - 1$ ; while, when  $y_i = -1$ , it is equivalent to  $b \le -w^\top x_i - 1$ . Finally,

$$\begin{aligned} |b| &= \max(b, -b) \leqslant \max_{1 \leqslant i \leqslant n} \max \left\{ w^{\top} x_i - 1, -w^{\top} x_i - 1 \right\} \\ &\leqslant \max_{1 \leqslant i \leqslant n} \|w\|_2 \|x_i\|_2 \leqslant \left( \max_{1 \leqslant i \leqslant n} \|x_i\|_2 \right) \|\widetilde{w}\| := B < +\infty. \end{aligned}$$

This allows to conclude that the minimization problem (6.3) can be restricted to  $(w, b) \in \overline{B(0, R)} \times [-B, B]$ .

Uniqueness cannot be asserted at this stage, although one minimizes a continuous, convex function over a closed, bounded, convex set in finite dimension. In fact, the global minimizer may not be unique, although there cannot be local minima. Any global minimizer satisfies the following necessary conditions of optimality:

$$\begin{cases} w - \sum_{i=1}^{n} \alpha_i \nabla_w c_i(w, b) = 0, \\ \sum_{i=1}^{n} \alpha_i \partial_b c_i(w, b) = 0, \\ \forall 1 \leq i \leq n, \qquad \alpha_i \geq 0, \quad \alpha_i \left[ y_i \left( w^\top x_i + b \right) - 1 \right] = 0, \end{cases}$$

which can be written more explicitly as

$$\begin{cases} w = \sum_{i=1}^{n} \alpha_{i} y_{i} x_{i}, \\ \sum_{i=1}^{n} \alpha_{i} y_{i} = 0, \\ \forall 1 \leq i \leq n, \qquad \alpha_{i} \geq 0, \quad \alpha_{i} \left[ y_{i} \left( w^{\top} x_{i} + b \right) - 1 \right] = 0. \end{cases}$$

$$(6.4)$$

#### 102 6 Support vector machines

**Support vectors.** The last condition in (6.4) asserts that  $\alpha_i = 0$  when  $c_i(w, b) > 0$ , *i.e.* the constraint is not active. A value  $\alpha_i \neq 0$  can be obtained only for points  $x_i$  on the marginal hyperplanes  $\{x \in \mathbb{R}^d | w^{\top}x + b \pm 1\}$ . Data points which do not lie on one of the two marginal hyperplanes do not affect the definition of w; and in fact not the definition of b either, as shown in the next exercise. It is useful for this to introduce the set of active constraints

$$\mathcal{S} = \{ 1 \leqslant i \leqslant n : \alpha_i \neq 0 \}.$$
(6.5)

Then,

$$w = \sum_{i \in \mathcal{S}} \alpha_i y_i x_i.$$

**Exercise 6.2.** Show that  $b = \frac{1}{|S|} \sum_{i \in S} \left( y_i - \sum_{j \in S} \alpha_j y_j x_j^\top x_i \right).$ 

**Correction.** We start from the equality  $y_i (w^{\top} x_i + b) = 1$  for any  $i \in S$ , which we multiply by  $y_i$ , using that  $y_i^2 = 1$ :  $b = y_i - w^{\top} x_i = y_i - \sum \alpha_j y_j x_j^{\top} x_i$ ,

$$b = y_i - w^\top x_i = y_i - \sum_{j \in \mathcal{S}} \alpha_j y_j x_j^\top x_i,$$

from the first equality in (6.4) and the definition of  $\mathcal{S}$ . The desired result is then obtained by summing these equalities over indices  $i \in S$  and dividing by |S|.

#### 6.1.3 Dual formulation

The motivation for turning to a dual formulation of the primal problem (6.3) is twofold:

- it is useful to formulate the problem an infinite dimensional setting based on kernel functions (see Section 6.4);
- on the algorithmic side, some reference numerical methods solve the dual problem rather than the primal one (see Section 6.1.4).

To write the dual problem, we introduce the Lagrangian

$$L(w, b, \alpha) = \frac{1}{2} \|w\|_2^2 - \sum_{i=1}^n \alpha_i \left[ y_i \left( w^\top x_i + b \right) - 1 \right].$$

The primal minimization problem (6.1.4) corresponds to

$$\min_{(w,b)\in \mathbb{R}^d\times \mathbb{R}} \max_{\alpha\in \mathbb{R}^n_+} L(w,b,\alpha).$$

It always holds that

$$\max_{\alpha \in \mathbb{R}^{n}_{+}} \min_{(w,b) \in \mathbb{R}^{d} \times \mathbb{R}} L(w,b,\alpha) \leqslant \min_{(w,b) \in \mathbb{R}^{d} \times \mathbb{R}} \max_{\alpha \in \mathbb{R}^{n}_{+}} L(w,b,\alpha),$$
(6.6)

as can be seen from the inequality

$$L(w, b, \alpha) \leqslant \max_{\alpha' \in \mathbb{R}^n_+} L(w, b, \alpha'),$$

then taking the minimum over  $(w, b) \in \mathbb{R}^d \times \mathbb{R}$  first on the left hand side, next on the right hand side, and finally considering the maximum over  $\alpha \in \mathbb{R}^n_+$  on the left hand side. In fact, there is no duality gap in (6.6), *i.e.* the inequality in (6.6) is in fact an equality. This comes from the fact
that the constraints are qualified and convex, and the objective function is convex. The primal optimization problem (6.3) can finally be reformulated in dual form as

$$\max_{\alpha \in \mathbb{R}^{n}_{+}} F_{\alpha}(w, b), \qquad F_{\alpha}(w, b) = \min_{(w, b) \in \mathbb{R}^{d} \times \mathbb{R}} \left\{ \frac{1}{2} \|w\|_{2}^{2} - \sum_{i=1}^{n} \alpha_{i} \left[ y_{i} \left( w^{\top} x_{i} + b \right) - 1 \right] \right\}.$$
(6.7)

In fact, the inner minimization can be performed explicitly, as shown in the following exercise.

**Exercise 6.3.** Prove that the dual problem rewrites

$$\max_{\alpha \in \mathbb{R}^n_+} \left\{ \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j \left( x_i^\top x_j \right) y_i y_j \ \left| \ \sum_{i=1}^n \alpha_i y_i = 0 \right\} \right\}.$$
(6.8)

In its formulation (6.8), the dual optimization problem is the maximization of a quadratic concave function under affine constraints. The function is indeed concave since the Hessian, which is the matrix with entries  $-(x_i^{\top}x_j)y_iy_j$ , is negative:

$$\forall \xi \in \mathbb{R}^n, \qquad -\sum_{i,j}^n (x_i^\top x_j) y_i y_j \xi_i \xi_j = - \left\| \sum_{i=1}^n \xi_i y_i x_i \right\|^2 \leqslant 0.$$

**Correction.** Note that  $F_{\alpha}(w,b) = F_{\alpha}^{1}(w) + F_{\alpha}^{2}(b) + F_{\alpha}^{3}$  with

$$F_{\alpha}^{1}(w) = \frac{1}{2} \|w\|_{2}^{2} - \sum_{i=1}^{n} \alpha_{i} y_{i} w^{\top} x_{i}, \qquad F_{\alpha}^{2}(b) = -\sum_{i=1}^{n} \alpha_{i} y_{i} b, \qquad F_{\alpha}^{3} = \sum_{i=1}^{n} \alpha_{i} x_{i} b,$$

The minimization over b of  $F_{\alpha}$  is equivalent to the minimization of  $F_{\alpha}^2$ . Now, the infimum of the latter function is  $-\infty$  unless

$$\sum_{i=1}^{n} \alpha_i y_i = 0.$$

The maximization over  $\alpha \in \mathbb{R}^n_+$  in (6.7) can thus be restricted to elements satisfying the above constraint, in which case  $F^2_{\alpha}(b) = 0$ . Therefore, (6.7) is equivalent to

$$\max_{\alpha \in \mathbb{R}^{n}_{+}} \left\{ \min_{w \in \mathbb{R}^{d}} \left\{ \frac{1}{2} \|w\|_{2}^{2} - \sum_{i=1}^{n} \alpha_{i} \left( y_{i} w^{\top} x_{i} - 1 \right) \right\} \left| \sum_{i=1}^{n} \alpha_{i} y_{i} = 0 \right\}.$$

The inner minimization over  $w \in \mathbb{R}^d$  is the unconstrained minimization of a quadratic functional in w, which admits a unique global minimizer

$$w_{\alpha} = \sum_{i=1}^{n} \alpha_i y_i x_i.$$

Since

$$\frac{1}{2} \|w_{\alpha}\|_{2}^{2} - \sum_{i=1}^{n} \alpha_{i} \left( y_{i} w_{\alpha}^{\top} x_{i} - 1 \right) = \sum_{i=1}^{n} \alpha_{i} - \frac{1}{2} \|w_{\alpha}\|^{2},$$

we obtained the claimed expression.

#### 6.1.4 Optimization algorithms

Both the primal and dual problems are quadratic programming problems (*i.e.* they correspond to minimizing quadratic functions under affine constraints). There is a huge literature on solvers

for these problems, but the challenge is to avoid relying on generic quadratic programming methods, whose computational cost scales as  $O(n^3)$ . The implementation in scikit-learn relies on the C library libsvm (see [10]). It relies on a version of the algorithm known as Sequential Minimal Optimization, where the dual problem (6.8) is analytically solved when the minimization is performed only over two coefficients  $\alpha_k$ ,  $\alpha_\ell$ , the other ones being fixed; going over all possible couples sequentially or randomly.

### 6.1.5 Learning guarantees by leave-one-out analysis

We provide bounds on the expected risk in this section. The material here is more advanced and can be skipped at first reading. We consider the framework of K-fold cross validation for K = n, which corresponds to the so-called leave-one-out cross validation (recall Section 1.3.2.2). It consists in taking out one point of the training data set, performing training on the remaining n-1 points, computing the prediction error on a validation set composed of the point that was taken out; and averaging this over the n possible choices for the single point validation set. More precisely,

$$\widehat{\mathcal{R}}_{\text{LOO}}(\mathcal{D}_{\text{train}}) = \frac{1}{n} \sum_{i=1}^{n} \mathbf{1}_{f_{\mathcal{D}_{\text{train}} \setminus \{(x_i, y_i)\}}(x_i) \neq y_i},$$
(6.9)

where  $f_{\mathcal{D}_{\text{train}} \setminus \{(x_i, y_i)\}}$  is the prediction function obtained from training on the data set  $\mathcal{D}_{\text{train}} \setminus \{(x_i, y_i)\}$  (*i.e.* the predictor  $x \mapsto \text{sign}(w^\top x + b)$  with w, b found by solving (6.3) or (6.8) with  $\mathcal{D}_{\text{train}}$  replaced by  $\mathcal{D}_{\text{train}} \setminus \{(x_i, y_i)\}$ ).

In order to characterize the validation error, we note that removing a data point which is not a support vector does not change the solution. Bounds on the validation error are therefore related to the fraction of support vectors. We write here bounds in expectation over realizations of the training set  $\mathcal{D}_{\text{train}}$ , distributed according to  $p_{\text{data}}^{\otimes n}$ . We start by providing a general result on the average leave-one-out error (see [39, Lemma 5.3]).

**Lemma 6.1.** The average leave-one-out validation error for a data set  $\mathcal{D}_{\text{train}}^n$  of size  $n \ge 2$  is an unbiased estimator of the average expected risk of a data set  $\mathcal{D}_{\text{train}}^{n-1}$  of size n-1:

$$\mathbb{E}_{\mathcal{D}_{\text{train}}^{n} \sim p_{\text{data}}^{\otimes n}} \left[ \widehat{\mathcal{R}}_{\text{LOO}} \left( \mathcal{D}_{\text{train}}^{n} \right) \right] = \mathbb{E}_{\mathcal{D}_{\text{train}}^{n-1} \sim p_{\text{data}}^{\otimes n-1}} \left[ \mathcal{R} \left( f_{\mathcal{D}_{\text{train}}^{n-1}} \right) \right],$$

where  $f_{\mathcal{D}_{\text{train}}^{n-1}}$  is the prediction function obtained by training over  $\mathcal{D}_{\text{train}}^{n-1}$ .

*Proof.* By using first the definition of the leave-one-out error, then the independence of the data points,

$$\mathbb{E}_{\mathcal{D}_{\text{train}}^{n} \sim p_{\text{data}}^{\otimes n}} \left[ \widehat{\mathcal{R}}_{\text{LOO}} \left( \mathcal{D}_{\text{train}}^{n} \right) \right] = \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}_{\mathcal{D}_{\text{train}}^{n} \sim p_{\text{data}}^{\otimes n}} \left[ \mathbf{1}_{f_{\mathcal{D}_{\text{train}}^{n} \setminus \{(x_{i}, y_{i})\}}(x_{i}) \neq y_{i}} \right] \\ = \mathbb{E}_{\mathcal{D}_{\text{train}}^{n} \sim p_{\text{data}}^{\otimes n}} \left[ \mathbf{1}_{f_{\mathcal{D}_{\text{train}}^{n} \setminus \{(x_{1}, y_{1})\}}(x_{1}) \neq y_{1}} \right] \\ = \mathbb{E}_{\mathcal{D}_{\text{train}}^{n-1} \sim p_{\text{data}}^{\otimes n-1}} \left( \mathbb{E}_{(x_{1}, y_{1}) \sim p_{\text{data}}} \left[ \mathbf{1}_{f_{\mathcal{D}_{\text{train}}^{n-1}}(x_{1}) \neq y_{1}} \right] \right)$$

This leads to the claimed result since the expectation between brackets on the right hand side of the previous equality is  $\mathcal{R}(f_{\mathcal{D}_{\text{train}}^{n-1}})$ .

The interest of the leave-one-out estimator for SVM (which is very expensive to compute in general unless there is some algebraic simplification which makes it unnecessary to solve noptimization problems) is that it can be easily related to the fraction of support points in the data. **Theorem 6.1.** Assume that the training sets  $\mathcal{D}_{\text{train}}^n$  are linearly separable for any  $n \ge 2$ . Recall the definition (6.5) of the set of active constraints, and denote by  $N_{\text{SV}}(\mathcal{D}_{\text{train}}) = |\mathcal{S}|$  the number of support vector defining the prediction function

$$f_{\mathcal{D}_{\text{train}}}(x) = \operatorname{sign}\left(\sum_{i\in\mathcal{S}} \alpha_i y_i x_i^{\top} x + b\right).$$

Then the expected risk can be bounded as

$$\mathbb{E}_{\mathcal{D}_{\text{train}}^{n} \sim p_{\text{data}}^{\otimes n}} \left[ \mathcal{R}\left( f_{\mathcal{D}_{\text{train}}^{n}} \right) \right] \leqslant \mathbb{E}_{\mathcal{D}_{\text{train}}^{n+1} \sim p_{\text{data}}^{\otimes n+1}} \left\lfloor \frac{N_{\text{SV}}(\mathcal{D}_{\text{train}}^{n+1})}{n+1} \right\rfloor$$

*Proof.* By Lemma 6.1, the left hand side of the inequality to prove is equal to

$$\mathbb{E}_{\mathcal{D}_{\text{train}}^{n+1} \sim p_{\text{data}}^{\otimes n+1}} \left[ \widehat{\mathcal{R}}_{\text{LOO}} \left( \mathcal{D}_{\text{train}}^{n+1} \right) \right]$$

Now, in the leave-one-out estimator (6.9), the predictors which appear in the sum are unchanged when the data point which is removed is not a support vector, and the prediction at this point is therefore correct as the data sets are assumed to be linearly separable. To misclassify, the removed data point needs to be a support vector, so that

$$\widehat{\mathcal{R}}_{\text{LOO}}\left(\mathcal{D}_{\text{train}}^{n+1}\right) \leqslant \frac{N_{\text{SV}}(\mathcal{D}_{\text{train}}^{n+1})}{n+1},$$

from which the result follows by taking expectations over realizations of the data set on both sides.  $\hfill\square$ 

# 6.2 Linear classification for non separable data sets

Data sets are often not linearly separable. The constraints then need to be relaxed in order to allow for some misclassification of the data. The deviation to the constraint  $y_i(w^{\top}x_i + b) \ge 1$  is controlled using slack variables  $\xi_i \ge 0$ , namely

$$y_i\left(w^{\top}x_i+b\right) \ge 1-\xi_i.$$

Outliers correspond to  $\xi_i > 0$ . The decision function then equilibriates between large margins (which may lead to a higher number of outliers) and the limitation of the slack. This approach is known as soft margin SVM.

#### 6.2.1 Soft margin SVM

The primal optimization problem generalizes (6.3) as

$$\min_{(w,b,\xi)\in\mathbb{R}^d\times\mathbb{R}\times\mathbb{R}^n_+}\left\{\frac{1}{2}\|w\|_2^2 + C\sum_{i=1}^n\xi_i^p \ \middle|\ \forall 1\leqslant i\leqslant n, \ y_i\left(w^\top x_i + b\right)\geqslant 1-\xi_i\right\},\tag{6.10}$$

where  $p \ge 1$  is some exponent and C > 0 allows to control the amount of slack. The parameter C, which is akin to some regularization parameter, is typically determined by (cross) validation. Hard SVM is recoverd in the limit  $C \to +\infty$ , where no slack is allowed. Large values of C allow to focus attention on points close to the decision boundaries, as outliers with values of the slack not sufficiently small are strongly penalized.

The most common choices for the exponent p are p = 1, which is the value we consider in the sequel, and p = 2.

106 6 Support vector machines

**Remark 6.1 (Interpretation as convex surrogate).** The minimization problem (6.10) can be rewritten as

$$\min_{(w,b)\in\mathbb{R}^d\times\mathbb{R}}\left\{\frac{1}{n}\sum_{i=1}^n \max\left\{0, 1-y_i\left(w^{\top}x_i+b\right)\right\}^p + \frac{1}{2nC}\|w\|_2^2\right\}$$

This formulation makes it apparent that this amounts to a classification problem using convex surrogate functions

$$\Phi_p(u) = \max\{0, 1-u\}^p,$$

and a ridge penalization with strength  $\lambda = 1/(2nC)$ . The cases p = 1 and p = 2 respectively correspond to the hinge loss and the squared hinge loss.

#### 6.2.2 Characterization of minimizers

We write here the characterization of the minimizer of (6.10) for p = 1. The computations are similar to those leading to (6.4). More precisely, the minimization problem under consideration corresponds to minimizing the function

$$f(w, b, \xi) = \frac{1}{2} \|w\|^2 + C \sum_{i=1}^n \xi_i,$$

under the constraints

$$c_i(w,b,\xi) = y_i\left(w^\top x_i + b\right) - 1 + \xi_i \ge 0, \qquad \widetilde{c}_i(w,b,\xi) = \xi_i \ge 0.$$

The associated necessary conditions to be satisfied by a global minimizer read

$$\begin{cases} \nabla_w f(w,b,\xi) - \sum_{i=1}^n \alpha_i \nabla_w c_i(w,b,\xi) = 0, \\ \partial_b f(w,b,\xi) - \sum_{i=1}^n \alpha_i \partial_b c_i(w,b,\xi) = 0, \\ \forall 1 \leqslant i \leqslant n, \qquad \partial_{\xi_i} f(w,b,\xi) - \alpha_i \partial_{\xi_i} c_i(w,b,\xi) - \beta_i \partial_{\xi_i} \widetilde{c}_i(w,b,\xi) = 0, \\ \alpha_i c_i(w,b,\xi) = \beta_i \widetilde{c}_i(w,b,\xi) = 0, \\ \alpha_i \geqslant 0, \quad \beta_i \geqslant 0, \end{cases}$$

so that

$$w = \sum_{i=1}^{n} \alpha_i y_i x_i,$$

with

$$\sum_{i=1}^{n} \alpha_i y_i = 0,$$

and

$$\forall 1 \leqslant i \leqslant n, \qquad \alpha_i + \beta_i = C, \qquad \alpha_i c_i(w, b, \xi) = \beta_i \widetilde{c}_i(w, b, \xi) = 0$$

Recalling the definition (6.5) of the set of active constraints, we find that

$$w = \sum_{i \in \mathcal{S}} \alpha_i y_i x_i.$$

An index  $1 \leq i \leq n$  belongs to S only if  $y_i(w^{\top}x_i+b)-1+\xi_i=0$ , *i.e.* the constraint  $c_i(w,b,\xi) \geq 0$  is active. There are therefore two types of support vectors:

•  $\xi_i = 0$  and  $y_i(w^{\top}x_i + b) = 1$ , so that  $x_i$  belongs to one of the two marginal hyperplanes, as for separable data sets;

•  $\xi_i > 0$ , in which case  $x_i$  is an outlier. One necessarily has  $\beta_i = 0$  in this case, so that  $\alpha_i = C$ .

An alternative interpretation of the result is obtained by looking at the values of  $\alpha_i \in [0, C]$ :

- if  $\alpha_i = 0$ , the data point is ignored for the prediction;
- if  $\alpha_i \in (0, C)$ , then  $\xi_i = 0$ . The data point is used for the prediction, and lies on one of the two marginal hyperplanes. As in Exercise 6.2, it can be shown that the value of b is obtained from data points on the marginal hyperplanes.
- if  $\alpha_i = C$ , then the data point can lie (but need not to) inside the margin region, either correctly classified if  $\xi_i < 1$ , or misclassified if  $\xi_i > 1$ .

#### 6.2.3 Dual formulation

To obtain the dual formulation of the primal problem (6.10), we introduce the Lagrangian

$$L(w, b, \xi, \alpha, \beta) = \frac{1}{2} \|w\|_2^2 + C \sum_{i=1}^n \xi_i - \sum_{i=1}^n \alpha_i \left[ y_i \left( w^\top x_i + b \right) - 1 + \xi_i \right] - \sum_{i=1}^n \beta_i \xi_i.$$

The dual problem reads

$$\max_{\alpha,\beta\in\mathbb{R}^n_+} \left\{ \min_{(w,b,\xi)\in\mathbb{R}^d\times\mathbb{R}\times\mathbb{R}^n_+\}} L(w,b,\xi,\alpha,\beta) \right\}.$$

We can then follow the same derivation as in Section 6.1.3, and obtain the following result, similar to the one of Exercise 6.3.

Exercise 6.4. Prove that the dual version of the primal problem (6.10) writes

$$\max_{\alpha \in [0,C]^n} \left\{ \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j \left( x_i^\top x_j \right) y_i y_j \ \left| \ \sum_{i=1}^n \alpha_i y_i = 0 \right\}.$$
(6.11)

Note that (6.11) coincides with (6.8) except that the components of  $\alpha$  are restricted to lie in the interval [0, C] instead of  $\mathbb{R}_+$ .

**Correction.** Note that 
$$L(w, b, \xi, \alpha, \beta) = F_{\alpha,\beta}^{1}(w) + F_{\alpha,\beta}^{2}(b) + F_{\alpha,\beta}^{3}(\xi) + F_{\alpha,\beta}^{4}$$
 with  
 $F_{\alpha,\beta}^{1}(w) = \frac{1}{2} \|w\|_{2}^{2} - \sum_{i=1}^{n} \alpha_{i} y_{i} w^{\top} x_{i}, \qquad F_{\alpha,\beta}^{2}(b) = -\sum_{i=1}^{n} \alpha_{i} y_{i} b, \qquad F_{\alpha,\beta}^{3}(\xi) = \sum_{i=1}^{n} (C - \alpha_{i} - \beta_{i}) \xi_{i},$ 

and

$$F_{\alpha,\beta}^4 = \sum_{i=1}^n \alpha_i.$$

The minimization of  $L(w, b, \xi, \alpha, \beta)$  over  $(w, b, \xi) \in \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^n_+$  for  $\alpha, \beta$  fixed therefore amounts to minimizing the various functions  $F^k_{\alpha,\beta}$  for k = 1, 2, 3 with respect to their argument. A simple computation shows that the optimality condition for  $F^1_{\alpha,\beta}$  reads

$$w_{\alpha,\beta}^{\star} = \sum_{i=1}^{n} \alpha_i y_i x_i,$$

and

$$\min_{w \in \mathbb{R}^d} F^1_{\alpha,\beta}(w) = F^1_{\alpha,\beta}\left(w^{\star}_{\alpha,\beta}\right) = -\frac{1}{2} \left\| \sum_{i=1}^n \alpha_i y_i x_i \right\|^2.$$

Moreover, the minimum of  $F_{\alpha,\beta}^2$  is  $-\infty$  unless

108 6 Support vector machines

$$\sum_{i=1}^{n} \alpha_i y_i = 0$$

The vectors  $\alpha \in \mathbb{R}^n_+$  which do not satisfy the latter condition can therefore be discarded when maximizing over  $(\alpha, \beta) \in \mathbb{R}^n_+ \times \mathbb{R}^n_+$ . Finally, the minimum of  $F^3_{\alpha,\beta}$  over  $\xi \in \mathbb{R}^n_+$  is  $-\infty$ if  $\alpha_i + \beta_i > C$  for some index  $1 \leq i \leq n$ . The maximization over  $(\alpha, \beta) \in \mathbb{R}^n_+ \times \mathbb{R}^n_+$  can therefore be restricted to vectors for which  $0 \leq \alpha_i + \beta_i \leq C$ . In this case, the minimum of  $F^3_{\alpha,\beta}$  is obtained for  $\xi = 0$  and it is equal to 0. The dual problem therefore reads

$$\max_{\alpha,\beta\in\mathbb{R}^n_+} \left\{ -\frac{1}{2} \left\| \sum_{i=1}^n \alpha_i y_i x_i \right\|^2 + F_{\alpha,\beta}^4 \left\| \alpha_i + \beta_i \leqslant C \right\}.$$

Note that  $\beta_i$  only appears through the constraint  $\alpha_i + \beta_i \leq C$ . The maximization over  $\alpha, \beta \in \mathbb{R}^n_+$  satisfying  $\alpha_i + \beta_i \leq C$  for all  $1 \leq i \leq n$  is therefore equivalent to the maximization over  $\alpha \in [0, C]^n$ , which allows to conclude.

# 6.3 Multiclass SVM

SVM is intrinsically a method for binary classification. It can however be transformed for multiclass classification. There are two main strategies to this end: combine pairwise decision functions for all possible couples of labels ("one-vs-one" strategy), or base predictions on decision functions where one label is separated from all other ones ("one-vs-all" approach). We successively describe both strategies, for a situation where the labels of the data set belong to  $\mathcal{Y} = \{1, \ldots, K\}$ .

**One-vs-one.** For any  $1 \leq k, k' \leq K$  with  $k \neq k'$ , we consider the prediction functions  $f_{k,k'}$  obtained by training SVM on the data set containing only points for which the label is either k or k'. For a test input x', the prediction  $y' \in \mathcal{Y}$  is the index which has obtained the largest number of assignments for the K(K-1)/2 classifiers under consideration. More precisely,

$$y' \in \operatorname*{argmax}_{y \in \mathcal{Y}} \left\{ \sum_{1 \leqslant k < k' \leqslant K} \mathbf{1}_{f_{k,k'}(x') = y} \right\}.$$

This approach is however expensive when K is large as one needs to train K(K-1)/2 models.

**One-vs-all.** We select an index  $k \in \{1, ..., K\}$ , and transform the initial training set into a training set where there are only two classes, one corresponding to data points with labels k, and the other one gathering all data points whose labels differ from k. This amounts to considering the inputs  $x_i$  for  $1 \le i \le n$  with new labels  $\tilde{y}_i \in \{-1, 1\}$ , with

$$\widetilde{y}_i = \mathbf{1}_{y_i = k} - \mathbf{1}_{y_i \neq k}.$$

We then perform SVM as described in the binary case. The corresponding decision function is denoted by  $g_k$ . For a new test input x', one then selects the label based on some form of maximal separation, for instance

$$y' \in \operatorname*{argmax}_{1 \leqslant k \leqslant K} g_k(x'),$$

i.e. the class label which lead to the largest margin; or

$$y' \in \underset{1 \leq k \leq K}{\operatorname{argmax}} \log \frac{p_k(y=1 \mid x')}{p_k(y=-1 \mid x')}$$

where  $p_k$  is some probabilistic model (whose hyperparameters should be calibrated with some maximum likelihood procedure on a separate validation set in order to avoid overfitting; this is known as Platt scaling, see [40, Sections 17.3.5 and 17.3.7]). There are however two difficulties in this approach. The first one is class imbalance, as there will typically be many more points with labels in  $\{1, \ldots, K\} \setminus \{k\}$  than points with labels k. The second one is that the various the values of the decision functions  $g_k$  are difficult to compare as the corresponding classification problems may have quite different scales; similarly, the probabilities  $p_k(y = \pm 1 | x')$  may be hard to compare.

# 6.4 Kernel SVM

There are two difficulties with classification using SVM, and in general with other linear methods, either for classification (think of logistic regression), regression (think of least square regression) or dimensionality reduction (think of PCA):

- the data may not be well represented with a linear model. In the context of classification, this means that the data set may not be linearly separable. On the other hand, it could become linearly separable once a nonlinear featurization has been performed with some feature function  $\phi$ . This function may have values in very high dimensional spaces, as it may be necessary to go to quite high dimensions in order to linearly separate the data. In any case, this approach allows to obtain non affine decision boundaries;
- the computational cost of the method can become important as the dimension of the number of features is increased. The kernel trick, as described below (and already seen in a particular case in Section 5.5.2), allows to alleviate this issue by ensuring that the cost of the numerical methods depends on the number of data points n and not on the dimension of the feature function.

We start by presenting the principle of Kernel SVM in Section 6.4.1, and then give a mathematically more precise introduction to kernel methods in Section 6.4.2.

#### 6.4.1 Principle of the method

The principle of kernel methods is to rewrite all predictions using expressions such as

$$\mathcal{K}(x, x') = \phi(x) \cdot \phi(x'), \tag{6.12}$$

which should be understood as the scalar product between two elements of a (possibly infinite dimensional) Hilbert space. For SVMs, predictions using kernel methods rely on  $w^{\top}\phi(x) + b$ , where the expressions of w, b are the ones obtained in Sections 6.1 and 6.2 but with  $x_i$  replaced by  $\phi(x_i)$ :

$$w = \sum_{i \in \mathcal{S}} \alpha_i y_i \phi(x_i), \qquad b = \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \left( y_i - \sum_{j \in \mathcal{S}} \alpha_j y_j \phi(x_j)^\top \phi(x_i) \right).$$

Therefore,

$$w^{\top}\phi(x) + b = \sum_{i \in \mathcal{S}} \alpha_i y_i \mathcal{K}(x, x_i) + \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \left( y_i - \sum_{j \in \mathcal{S}} \alpha_j y_j \mathcal{K}(x_j, x_i) \right)$$

can be computed with the knowledge of the kernel  $\mathcal{K}$  only. Moreover, the values of the coefficients  $\alpha \in \mathbb{R}^n_+$  can be found by considering the dual problem derived similarly to (6.11):

$$\max_{\alpha \in [0,C]^n} \left\{ \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j y_i y_j \mathcal{K}(x_i, x_j) \ \left| \ \sum_{i=1}^n \alpha_i y_i = 0 \right\} \right\}.$$

add plot of decision function in this NL context 110 6 Support vector machines

This makes it possible to substantially enlarge the dimension of the features, and in fact consider infinite dimensional features, as long as the associated kernel  $\mathcal{K}$  is easy to evaluate. In fact, in practice, one chooses the kernel function  $\mathcal{K}$ , in which case the associated features are implicit. Indeed, the feature function  $\phi$  needs not be known explicitly as it suffices to rely on a result saying that there exist a Hilbert space  $\mathcal{H}$  and a feature function  $\phi : \mathcal{X} \to \mathcal{H}$  associated with a kernel  $\mathcal{K} : \mathcal{X}^2 \to \mathbb{R}$  according to (6.12). This is discussed in the next section.

#### 6.4.2 Kernel methods

We provide in this section the rigorous definition of kernel functions, and give conditions characterizing admissible kernels.

**Definition 6.2.** A kernel  $\mathcal{K} : \mathcal{X}^2 \to \mathbb{R}$  is said to be positive definite symmetric (PDS) if, for any  $(x_1, \ldots, x_n) \in \mathcal{X}^n$ , the matrix  $K \in \mathbb{R}^{n \times n}$  with entries  $K_{i,j} = \mathcal{K}(x_i, x_j)$  is symmetric positive semidefinite (i.e. for any  $t \in \mathbb{R}^n$ , it holds  $t^{\top}Kt \ge 0$ , or equivalently K is real symmetric with nonnegative eigenvalues).

The following important result, due to Aronszajn in 1950, characterizes PDS kernels.

**Theorem 6.2.** The function  $\mathcal{K} : \mathcal{X}^2 \to \mathbb{R}$  is a PDS kernel if and only if there exist a Hilbert space  $\mathcal{H}$  and a function  $\phi : \mathcal{X} \to \mathcal{H}$  such that

$$\mathcal{K}(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}.$$
(6.13)

Proof. Assume that there exist a Hilbert space  $\mathcal{H}$  and a function  $\phi : \mathcal{X} \to \mathcal{H}$  such that  $\mathcal{K}(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$ . Fix  $(x_1, \ldots, x_n) \in \mathcal{X}^n$ , and consider the matrix  $K \in \mathbb{R}^{n \times n}$  with entries  $K_{i,j} = \mathcal{K}(x_i, x_j)$ . Note first that K is symmetric by the symmetry of the scalar product. Next, for  $t \in \mathbb{R}^n$ ,

$$t^{\top}Kt = \sum_{i,j=1}^{n} \mathcal{K}(x_i, x_j) t_i t_j = \left\langle \sum_{i=1}^{n} t_i \phi(x_i), \sum_{j=1}^{n} t_j \phi(x_j) \right\rangle_{\mathcal{H}} = \left\| \sum_{i=1}^{n} t_i \phi(x_i) \right\|_{\mathcal{H}}^2 \ge 0,$$

so that K is positive semidefinite. This allows to conclude that  $\mathcal{K}$  is a PDS kernel.

Assume conversely that  $\mathcal{K}$  is a PDS kernel. We first construct a Hilbert space  $\mathcal{H}$  of functions with argument in  $\mathcal{X}$  by completion of finite linear combinations of elementary functions  $x \mapsto \mathcal{K}(x, x_i)$ :

$$\mathcal{H}_0 = \left\{ \sum_{i \in I} t_i \mathcal{K}(\cdot, x_i), \ t_i \in \mathbb{R}, x_i \in \mathcal{X}, |I| < +\infty \right\}.$$

This vector space is endowed with the inner product

$$\langle f,g\rangle_{\mathcal{H}} = \sum_{i=1}^{n} \sum_{j=1}^{m} t_i t'_j \mathcal{K}(x_i,x'_j), \qquad f = \sum_{i=1}^{n} t_i \mathcal{K}(\cdot,x_i), \quad g = \sum_{j=1}^{m} t'_j \mathcal{K}(\cdot,x'_j).$$

Note that  $\langle f, g \rangle_{\mathcal{H}}$  does not depend on the representative of f as

$$\langle f,g \rangle_{\mathcal{H}} = \sum_{j=1}^m t_j f(x'_j);$$

similarly it depends only on g and not on its chosen representative. Moreover,  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  is a symmetric bilinear form since  $\mathcal{K}(x_i, x'_j) = \mathcal{K}(x'_j, x_i)$  (this can be seen by noting that  $\mathcal{K}(x_1, x_2) = K_{12} = K_{21} = \mathcal{K}(x_2, x_1)$  since the 2 × 2 matrix build from  $\mathcal{K}$  with only two points is symmetric; the equality holds for any two points). In addition,  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  is nonnegative (*i.e.*  $\langle f, f \rangle_{\mathcal{H}} \ge 0$  for any  $f \in \mathcal{H}_0$ ). Also,

$$\langle \mathcal{K}(\cdot, x), f \rangle_{\mathcal{H}} = f(x)$$

and

$$\langle \mathcal{K}(\cdot, x), \mathcal{K}(\cdot, x') \rangle_{\mathcal{H}} = \mathcal{K}(x, x'),$$

so that one can choose  $\phi(x) = \mathcal{K}(\cdot, x)$  in order to satisfy (6.13).

To show that the inner product is a scalar product, it remains to prove that it is definite. Fix to this end  $f \in \mathcal{H}_0$  such that  $\langle f, f \rangle_{\mathcal{H}} = 0$ . By a Cauchy–Schwarz inequality,<sup>1</sup> for a given  $x \in \mathcal{X}$ ,

$$\langle f, \mathcal{K}(\cdot, x) \rangle_{\mathcal{H}}^2 \leqslant \langle f, f \rangle_{\mathcal{H}} \langle \mathcal{K}(\cdot, x), \mathcal{K}(\cdot, x) \rangle_{\mathcal{H}} = \langle f, f \rangle_{\mathcal{H}} \mathcal{K}(x, x), \tag{6.14}$$

so that  $f(x) = \langle f, \mathcal{K}(\cdot, x) \rangle_{\mathcal{H}} = 0$ . Since x was arbitrary, we obtain f = 0.

The proof is concluded by defining  $\mathcal{H}$  as the completion of  $\mathcal{H}_0$  for the norm induced by the scalar product  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ .

Note that (6.14) ensures that  $f \mapsto \langle f, \mathcal{K}(\cdot, x) \rangle_{\mathcal{H}}$  is a Lipschitz mapping from  $\mathcal{H}_0$  to  $\mathbb{R}$ , and hence can be extended to a Lipschitz mapping on  $\mathcal{H}$ . In particular,

$$\forall f \in \mathcal{H}, \qquad f(x) = \langle \mathcal{K}(\cdot, x), f \rangle_{\mathcal{H}}. \tag{6.15}$$

This property motivates the terminology *Reproducing Kernel Hilbert Space* (RKHS) for the Hilbert space  $\mathcal{H}$ . The latter Hilbert space is also called feature space, with associated feature map  $\phi : \mathcal{X} \to \mathcal{H}$  given by  $\phi(x) = \mathcal{K}(\cdot, x) \in \mathcal{H}$ . Note that  $\mathcal{H}$  contains only continuous functions as one needs to make sense of pointwise values through (6.15). It is typically a space of rather smooth functions, such as a Sobolev space of sufficiently high order.

We next give two examples of PDS kernels, and make precise the associated feature maps.

#### **Example 6.1 (Polynomial kernels).** For $p \in \mathbb{N}$ , define

$$\mathcal{K}(x, x') = (1 + x \cdot x')^{p}.$$
(6.16)

Let us construct, for pedagogical reasons, the feature map when d = 2 and p = 2. More precisely, for  $x = (x_1, x_2) \in \mathbb{R}^2$ , introduce

$$\phi(x) = \left(x_1^2, x_2^2, \sqrt{2}x_1x_2, \sqrt{2}x_1, \sqrt{2}x_2, 1\right)^\top \in \mathbb{R}^6.$$

Then, for  $x, x' \in \mathbb{R}^2$ ,

$$\begin{split} \phi(x) \cdot \phi(x') &= x_1^2 (x_1')^2 + x_2^2 (x_2')^2 + 2x_1 x_1' x_2 x_2' + 2x_1 x_1' + 2x_2 x_2' + 1 = (1 + x_1 x_1' + x_2 x_2')^2 \\ &= (1 + x \cdot x')^2 \,, \end{split}$$

which is indeed (6.16) for p = 2.

**Example 6.2 (Gaussian kernels).** For  $\sigma > 0$ , define

$$\mathcal{K}(x, x') = e^{-\|x - x'\|^2 / (2\sigma^2)}$$

This kernel can be obtained from

$$\widetilde{\mathcal{K}}(x, x') = \mathrm{e}^{-x \cdot x' / \sigma^2}$$

by a renormalization procedure:

$$\mathcal{K}(x,x') = \frac{\widetilde{\mathcal{K}}(x,x')}{\sqrt{\widetilde{\mathcal{K}}(x,x)}\sqrt{\widetilde{\mathcal{K}}(x',x')}}$$

<sup>&</sup>lt;sup>1</sup> Note indeed that this inequality holds for any nonnegative symmetric bilinear form (which may not be a scalar product) on a vector space  $\mathcal{Z}$ , as can be seen by considering  $0 \leq \langle z_1 + tz_2, z_1 + tz_2 \rangle_{\mathcal{Z}} = \langle z_1, z_1 \rangle_{\mathcal{Z}} + 2t \langle z_1, z_2 \rangle_{\mathcal{Z}} + t^2 \langle z_2, z_2 \rangle_{\mathcal{Z}}$ , so that the determinant of the polynomial in t has to be nonpositive, which leads to the claimed inequality  $\langle z_1, z_2 \rangle_{\mathcal{Z}}^2 \leq \langle z_1, z_1 \rangle_{\mathcal{Z}} \langle z_2, z_2 \rangle_{\mathcal{Z}}$ .

To prove that  $\mathcal{K}$  is PDS, it therefore suffices in view of [39, Lemma 6.9] to prove that  $\widetilde{\mathcal{K}}$  is PDS. The latter statement can be obtained either by relying on Fourier analysis using Bochner's theorem (see [39, Theorem 6.24] or [4, Theorem 7.3]), or by explicitly constructing  $\phi$  by noting that

$$\widetilde{\mathcal{K}}(x,x') = \sum_{p=0}^{+\infty} \frac{(x \cdot x')^p}{p! \sigma^{2p}}.$$

There are of course many other examples of PDS kernels. Given PDS kernels, one can also construct new PDS kernels by various rules (sums and products of PDS kernels, composition of PDS kernels with power series with nonnegative coefficients, etc).

**Exercise 6.5.** Consider a PDS kernel  $\mathcal{K}$ . Prove that

$$\forall x, x' \in \mathcal{X}, \qquad \mathcal{K}(x, x')^2 \leqslant \mathcal{K}(x, x) \mathcal{K}(x', x').$$

**Correction.** The  $2 \times 2$  matrix

$$K = \begin{pmatrix} \mathcal{K}(x,x) & \mathcal{K}(x,x') \\ \mathcal{K}(x,x') & \mathcal{K}(x',x') \end{pmatrix}$$

is symmetric positive since  $\mathcal{K}$  is a PDS kernel. In particular,  $\det(K) \ge 0$ . The result then follows from the equality  $\det(K) = \mathcal{K}(x, x)\mathcal{K}(x', x') - \mathcal{K}(x, x')^2$ .

We conclude this section by a very important result, the representer theorem, which states that predictors can be written as a linear combination of  $\{\mathcal{K}(\cdot, x_i)\}_{1 \leq i \leq n}$ . The following result allows to generalize manipulations performed in Section 6.4.1 for SVM.

**Theorem 6.3 (Representer theorem).** Consider a Hilbert space  $\mathcal{H}$  and a feature map  $\phi : \mathcal{X} \to \mathcal{H}$ . Fix  $(x_1, \ldots, x_n) \in \mathcal{X}^n$  and consider  $\Psi : \mathbb{R}^{n+1} \to \mathbb{R}$  nondecreasing with respect to the last variable. Then,

$$f^{\star} \in \operatorname*{argmin}_{f \in \mathcal{H}} \Psi\left(\langle f, \phi(x_1) \rangle_{\mathcal{H}}, \dots, \langle f, \phi(x_n) \rangle_{\mathcal{H}}, \|f\|_{\mathcal{H}}^2\right)$$

is of the form

$$f^{\star} = \sum_{i=1}^{n} t_i \phi(x_i) \in \mathcal{H}.$$

This result can be applied to PDS kernel  $\mathcal{K} : \mathcal{X}^2 \to \mathbb{R}$  by considering the associated RHKS  $\mathcal{H}$ and  $\phi(x_i) = \mathcal{K}(\cdot, x_i)$ . In this case, any minimizer

$$f^{\star} \in \operatorname*{argmin}_{f \in \mathcal{H}} \Psi\left(f(x_1), \dots, f(x_n), \|f\|_{\mathcal{H}}^2\right)$$

can be written as

$$f^{\star} = \sum_{i=1}^{n} t_i \mathcal{K}(\cdot, x_i).$$

*Proof.* Introduce the following finite dimensional vector subspace of  $\mathcal{H}$ :

$$\mathcal{H}_{x_1,\dots,x_n} = \left\{ \sum_{i=1}^n t_i \phi(x_i), \ t \in \mathbb{R}^n \right\} \subset \mathcal{H},$$

which is the linear span of the feature vectors associated with the training data set. Any  $f \in \mathcal{H}$  can be orthogonally decomposed as

$$f = f_{x_1,\dots,x_n} + f^{\perp}, \qquad f_{x_1,\dots,x_n} \in \mathcal{H}_{x_1,\dots,x_n}, \qquad \left\langle f_{x_1,\dots,x_n}, f^{\perp} \right\rangle_{\mathcal{H}} = 0.$$

Then,

$$\langle f, \phi(x_i) \rangle_{\mathcal{H}} = \langle f_{x_1, \dots, x_n}, \phi(x_i) \rangle_{\mathcal{H}},$$

and

$$\|f\|_{\mathcal{H}}^{2} = \|f_{x_{1},\dots,x_{n}}\|_{\mathcal{H}}^{2} + \|f^{\perp}\|_{\mathcal{H}}^{2} \ge \|f_{x_{1},\dots,x_{n}}\|_{\mathcal{H}}^{2}$$

Therefore, since  $\Psi$  is nondecreasing with respect to its last argument,

$$\Psi\left(\langle f, \phi(x_1) \rangle_{\mathcal{H}}, \dots, \langle f, \phi(x_n) \rangle_{\mathcal{H}}, \|f\|_{\mathcal{H}}^2\right)$$
  
$$\geq \Psi\left(\langle f_{x_1, \dots, x_n}, \phi(x_1) \rangle_{\mathcal{H}}, \dots, \langle f_{x_1, \dots, x_n}, \phi(x_n) \rangle_{\mathcal{H}}, \|f_{x_1, \dots, x_n}\|_{\mathcal{H}}^2\right)$$

This inequality shows that the minimization over  $f \in \mathcal{H}$  can be restricted to a minimization over the subspace  $\mathcal{H}_{x_1,\ldots,x_n}$ , which allows to conclude the proof.

Let us apply Theorem 6.3 to typical training problems in supervised learning, which read (with some squared penalty term)

$$\min_{f \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_i, f(x_i)\right) + \lambda \|f\|_{\mathcal{H}}^2 \right\}.$$

The minimization can be restricted to functions of the form

$$f = \sum_{i=1}^{n} t_i \mathcal{K}(\cdot, x_i),$$

which therefore amounts to minimizing over  $t = (t_1, \ldots, t_n) \in \mathbb{R}^n$ . Note that the penalty term then reads

$$||f||_{\mathcal{H}}^2 = \sum_{i,j=1}^n t_i t_j \mathcal{K}(x_i, x_j) = t^\top K t,$$

where  $K \in \mathbb{R}^{n \times n}$  is the matrix with entries  $K_{ij} = \mathcal{K}(x_i, x_j)$ . The minimization problem can then be seen as some generalized ridge penalization:

$$t^{\star} \in \operatorname*{argmin}_{t \in \mathbb{R}^{n}} \left\{ \frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i}, (Kt)_{i}\right) + \lambda t^{\top} Kt \right\},\$$

where we used the equality

$$(Kt)_i = \sum_{j=1}^n t_j \mathcal{K}(x_i, x_j).$$

Predictions are finally performed with the optimal set of coefficients  $t^{\star}$  as

$$f^{\star}(x') = \sum_{i=1}^{n} t_i^{\star} \mathcal{K}\left(x', x_i\right).$$

# Trees and ensemble methods

7.1	Regression and classifications trees 115		
	7.1.1	Principle of the method 115	
	7.1.2	Regression trees	
	7.1.3	Classification trees	
	7.1.4	Some practical points 120	
7.2	Bagging and averaging techniques 121		
	7.2.1	Motivation: the ideal setting 122	
	7.2.2	Bagging	
	7.2.3	Random forests	
7.3	Boosting 123		
	7.3.1	General philosophy	
	7.3.2	Gradient boosting	

We discuss in this chapter how to perform predictions (either classification or regression) with trees. We start by presenting regression and classification trees in Section 7.1, and next discuss how to improve the performance of possibly weak learners such as trees by combining prediction functions, either by some form of averaging (see Section 7.2) or by some greedy reweighting procedure known as boosting (see Section 7.3).

The presentation of trees is based on [40, Chapter 18], [24, Chapter 9] (in particular Section 9.2), [8, Section 14.4] and [49, Chapter 18]; while the elements on ensemble methods and boosting are taken from [40, Chapter 18], [4, Chapter 10], [24, Chapter 11], [8, Sections 14.2 and 14.3] and [39, Chapter 7].

# 7.1 Regression and classifications trees

We show in this section how to perform predictions with decision trees. We start by making precise the principle of the method in Section 7.1.1, and then specify the strategy to regression trees in Section 7.1.2 and next to classification trees in Section 7.1.3. We conclude by discussing a few practical points of the method in Section 7.1.4.

We think of inputs having ordered values (real variables, or, in the case of discrete inputs, values which can be compared such as shoe sizes, ages, etc.); see Section 7.1.4 for a discussion on genuine categorical data.

# 7.1.1 Principle of the method

The idea behind decision trees is to recursively partition the input space  $\mathcal{X}$  and define a local model in each resulting region. Each region is associated with a leaf of the tree. The partitioning

#### 116 7 Trees and ensemble methods

is done with a series of questions to hierarchically decompose the input space, and group samples with the same labels (for classification) or with similar target values (for regression).

We consider recursive binary trees to simplify the presentation, and also to avoid data fragmentation (if there are too many possible choices at each step, the end nodes/leaves can be associated with too small a number of instances). The construction is performed using axis parallel splits, *i.e.* separations are done based on values of certain features only, and not linear combinations of features (which can make it difficult to learn certain separations: think of diagonally separated data in a feature space of dimension 2...). Restricting to axis parallel splits is better for the computational scalability of the method, allows for an easy randomization of the algorithm, and also makes the interpretation easier.

The construction can be formalized by considering a set of nested decision rules. At each node j of the decision tree, the feature  $x_{k_j}$  of an input x is compared to a treshold value  $t_j$ , *i.e.* one checks whether  $x_{k_j} \leq t_j$  or  $x_{k_j} > t_j$ . Outputs are specified at leaves. For the example of Figure 7.1, the ouputs are specified in the regions

$$R_1 = \{x_1 \leqslant t_1, \ x_2 \leqslant t_2\}, \qquad R_2 = \{x_1 > t_1, \ x_2 \leqslant t_4\}, \qquad \dots$$

The right part of Figure 7.1 illustrates a key advantage of the method, namely its interpretability.



Fig. 7.1: Left: Binary tree. Right: Associated regions.

## 7.1.2 Regression trees

We consider regression for inputs  $\mathcal{X} = \mathbb{R}^d$  and ouputs  $\mathcal{Y} = \mathbb{R}$  for simplicity. Regression trees implement prediction functions of the form

$$f_{\theta}(x') = \sum_{j=1}^{J} w_j \mathbf{1}_{x' \in R_j},$$

when there are J regions  $R_1, \ldots, R_J$  (corresponding to J leaves in the decision tree). The parameter  $\theta$  gathers the regions and the values  $w_j \in \mathbb{R}$  predicted in these regions:

$$\theta = \{R_j, w_j\}_{1 \le j \le J}.$$

For (unregularized) least square regression, the values  $(w_j)_{1 \leq j \leq J}$  are in fact easy to obtain once the regions  $R_1, \ldots, R_J$  are determined. This is made precise in the next exercise.

Exercise 7.1. For least square regression, prove that

$$w_j = \frac{\sum_{i=1}^n y_i \mathbf{1}_{x_i \in R_j}}{\sum_{i=1}^n \mathbf{1}_{x_i \in R_j}}$$

is the empirical average of the outputs associated with inputs belonging to  $R_i$ .

Correction. The loss function to consider for regression problems is

$$\frac{1}{n}\sum_{i=1}^{n}(y_i - f_{\theta}(x_i))^2 = \frac{1}{n}\sum_{j=1}^{J}\sum_{i \in I_j}(y_i - w_j)^2,$$

where  $I_j$  is the set of indices for which the data points are in the region  $R_j$ :

$$I_j = \{i \in \{1, \ldots, n\} : x_i \in R_j\}.$$

Note that the minimization over  $(w_1, \ldots, w_J)$  can in fact be rewritten as a minimization of  $w_j$  for each region  $R_j$  since

$$\min_{(w_1,\dots,w_J)\in\mathbb{R}^J}\left\{\frac{1}{n}\sum_{i=1}^n (y_i - f_\theta(x_i))^2\right\} = \frac{1}{n}\sum_{j=1}^J \min_{w_j\in\mathbb{R}}\left\{\sum_{i\in I_j} (y_i - w_j)^2\right\}.$$

Now, the minimization problem

$$\min_{w_j \in \mathbb{R}} \left\{ \sum_{i \in I_j} (y_i - w_j)^2 \right\}$$

is well posed (the function to minimize is strongly convex in finite dimension, in particular continuous and coercive) and the Euler–Lagrange condition characterizing the minimizer  $w_j^{\star}$  reads

$$\sum_{i \in I_j} \left( y_i - w_j^\star \right) = 0$$

This leads to the claimed result.

The previous exercise shows that the key point to construct regression trees is to find the partition  $R_1, \ldots, R_J$ . In general, finding the optimal partition is infeasible. In practice, one resorts to a greedy strategy where one node is grown at the time. Popular implementations of this idea are CART, ID3 and C4.5 for instance. The tree is then pruned in order to avoid overfitting.

**Growing the tree.** Let us explain how to grow a tree when being at node j. We introduce to this end the set  $\mathcal{D}_j$  of data points  $(x_i, y_i)$  which are still considered at this node, and denote by  $I_j$ 

#### 118 7 Trees and ensemble methods

the associated indices. We then consider the following possible partitions of this set, where we separate points based on the value of the k-th feature of the input:

$$\mathcal{D}_{j}^{-}(k,t) = \{(x_{i}, y_{i}) \in \mathcal{D}_{j} : x_{i,k} \leq t\}, \qquad \mathcal{D}_{j}^{+}(k,t) = \{(x_{i}, y_{i}) \in \mathcal{D}_{j} : x_{i,k} > t\}.$$

We choose k and t so that the following loss function is maximally decreased when partitioning  $\mathcal{D}_{i}$ :

$$(k_{j}, t_{j}) \in \underset{\substack{1 \leq k \leq d \\ t \in \mathbb{R}}}{\operatorname{argmin}} \left\{ \min_{\substack{w_{-} \in \mathbb{R} \\ (x_{i}, y_{i}) \in \mathcal{D}_{j}^{-}(k, t)}} \sum_{\substack{(y_{i} - w_{-})^{2} + \min_{w_{+} \in \mathbb{R}}} \sum_{\substack{(x_{i}, y_{i}) \in \mathcal{D}_{j}^{+}(k, t)}} (y_{i} - w_{+})^{2} \right\}$$
$$= \underset{\substack{1 \leq k \leq d \\ t \in \mathbb{R}}}{\operatorname{argmin}} \left\{ \sum_{\substack{(x_{i}, y_{i}) \in \mathcal{D}_{j}^{-}(k, t)}} \left(y_{i} - w_{j, k, t}^{-}\right)^{2} + \sum_{\substack{(x_{i}, y_{i}) \in \mathcal{D}_{j}^{+}(k, t)}} \left(y_{i} - w_{j, k, t}^{+}\right)^{2} \right\},$$

where  $w_{j,k,t}^{\pm}$  are the empirical averages of the ouputs on the sets  $\mathcal{D}_{j}^{\pm}(k,t)$ :

$$w_{j,k,t}^{\pm} = \frac{1}{|\mathcal{D}_{j}^{\pm}(k,t)|} \sum_{(x_{i},y_{i})\in\mathcal{D}_{j}^{\pm}(k,t)} y_{i}$$

The optimal values for k and t can be determined with a computational cost  $O(dn_j \log n_j)$ when  $\mathcal{X} = \mathbb{R}^d$  and there are  $n_j = |I_j|$  data points to consider (see [49, Section 18.2.3]). A naive implementation can be performed with a cost  $O(dn_j^2)$ . The factor d comes from the fact that one needs to consider all the possible values  $1 \leq k \leq d$ . For a given value of k, the optimal treshold could be found by sorting the values  $x_{1,k}, \ldots, x_{n,k}$  for the  $n_j$  data points in  $\mathcal{D}_j$  (cost  $O(n_j \log n_j)$ ) and then looping over the  $n_j$  possibles values  $(x_{i,k})_{i \in I_j}$  for t. Note indeed that it suffices to consider values of t coinciding with the value of a component of a data point. For each value of t, the loss function to evaluate can be evaluated with a computational cost  $O(n_j)$ .

**Pruning the tree.** Trees are grown until some minimal node size is reached (*e.g.* 5 data points) or the maximal depth is attained. It is good practice then to prune the tree back in order to avoid overfitting. Note indeed that large enough trees can always perfectly fit the data, but they are then too tightly adjusted to the data at hand and suffer from poor generalization capabilities.



Fig. 7.2: Left: Reference tree  $T_0$ . Right: Pruned tree where node 1 has been collapsed.

Complexity pruning is one way of proceeding (see [24, Section 9.2.2]). To present the method, denote by  $T_0$  the tree to be pruned. A subtree T of  $T_0$  is any tree that can be obtained by collapsing

a certain number of internal (*i.e.* non-terminal) nodes; see Figure 7.2 for an illustration. The metric to find the right tree size is

$$C_{\alpha}(T) = \sum_{t=1}^{|T|} N_t(T)Q_t(T) + \alpha |T|$$

where  $\alpha \ge 0$  is some regularization parameter, |T| is the number of terminal nodes of the tree T,  $N_t(T)$  is the number of data inputs at node t, and  $Q_t(T)$  is the empirical variance at this node:

$$Q_t(T) = \frac{1}{N_t(T)} \sum_{i:x_i \in R_t} (y_i - w_t)^2, \qquad w_t = \frac{1}{N_t(T)} \sum_{i:x_i \in R_t} y_i.$$

For a given value of  $\alpha \ge 0$  one then looks for the subtree  $T_{\alpha}$  of  $T_0$  which minimizes  $C_{\alpha}(T)$ . Large values of  $\alpha$  favor smaller trees, while smaller values of  $\alpha$  allow for larger trees. The initial tree is recovered for  $\alpha = 0$ . To find  $T_{\alpha}$  in practice, one resorts to weakest-link pruning, which consists in collapsing the node that produces the smallest per-node increase in  $C_0(T)$ ; and continues until the single node (root) tree is produced. This produces a finite sequence of subtrees which contains  $T_{\alpha}$ (see [1, Section 1.10.9] for further details).

Let us finally conclude by saying that pruning is not always performed in practice. It makes sense only if one wants a single nice decision tree. In practice, as discussed in Sections 7.2 and 7.3 in particular, one often relies on ensembles of rather small trees, which are individually not providing great predictions (and therefore need not be pruned).

#### 7.1.3 Classification trees

We discuss here how to perform classification with trees, for a label set  $\mathcal{Y} = \{1, \ldots, K\}$ . We rely on the empirical probabilities of  $k \in \{1, \ldots, K\}$  in  $\mathcal{D}$ :

$$\widehat{p}_k(\mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{(x_i, y_i) \in \mathcal{D}} \mathbf{1}_{y_i = k}.$$

Using the same notation as in Section 7.1.2, the prediction at the *j*th terminal node is  $\underset{1 \leq k \leq K}{\operatorname{argmax}} \widehat{p}_k(\mathcal{D}_j)$ 

in the region  $R_i$ .

As for regression trees, splitting is based on the values of the features of data points; but one needs another measure in the greedy approach to select the feature and treshold. The general form of this selection procedure is the following, at node j:

$$(k_j, t_j) \in \operatorname*{argmin}_{\substack{1 \leq k \leq d \\ t \in \mathbb{R}}} \left\{ \frac{|\mathcal{D}_j^-(k, t)|}{|\mathcal{D}_j|} \mathscr{C}\left(\mathcal{D}_j^-(k, t)\right) + \frac{|\mathcal{D}_j^+(k, t)|}{|\mathcal{D}_j|} \mathscr{C}\left(\mathcal{D}_j^+(k, t)\right) \right\},$$

where  $\mathscr{C}(\mathcal{D})$  is some cost function relation to a set of data points  $\mathcal{D}$ . Various choices can be considered, which all measure the purity of the labels. More precisely, one can consider

- the misclassification error  $1 \max_{1 \leq k \leq K} \widehat{p}_k(\mathcal{D})$ . This expression is motivated by the fact that, for a terminal node, one would predict the argument of the max; the Gini index  $\sum_{k=1}^{K} \widehat{p}_k(\mathcal{D}) (1 \widehat{p}_k(\mathcal{D})) = 1 \sum_{k=1}^{K} \widehat{p}_k(\mathcal{D})^2$ . The interpretation of this cost function is that it gives the expected error rate. Indeed, the probability to make a mistake on the label k is the product of oberving the label, namely  $\hat{p}_k(\mathcal{D})$ , multiplied by the probability to misclassify it, namely  $1 - \hat{p}_k(\mathcal{D})$ ; the total error rate is then the sum over k of the latter quantity;
- the cross entropy or deviance  $-\sum_{k=1}^{K} \widehat{p}_k(\mathcal{D}) \log \widehat{p}_k(\mathcal{D}).$

#### 120 7 Trees and ensemble methods

Pruning is traditionally done with the misclassification error; while the splitting is done with the Gini index or cross entropy. All these cost functions favor pure nodes with a single population. This can be easily understood when K = 2, as the cost functions are respectively  $1 - \max(p, 1-p)$ , 2p(1-p) and  $-p \log p - (1-p) \log(1-p)$ .

Exercise 7.2. Consider the set

$$\mathscr{P} = \left\{ (p_1, \dots, p_K) \in [0, 1]^K \mid \sum_{k=1}^K p_k = 1 \right\}.$$

Show that the mappings

$$\mathscr{P} \ni (p_1, \dots, p_K) \mapsto \sum_{k=1}^K p_k (1-p_k), \qquad \mathscr{P} \ni (p_1, \dots, p_K) \mapsto -\sum_{k=1}^K p_k \log p_k,$$

have nonnegative values, and that  $p^{\star,\ell} = (\mathbf{1}_{k=\ell})_{1 \leq k \leq K}$  are the only minimizers of these functions.

**Correction.** The first mapping obviously has nonnegative values since  $z(1-z) \ge 0$  for  $z \in [0, 1]$ . Moreover,

$$\sum_{k=1}^{K} p_k (1-p_k) = 0 \iff \forall k \in \{1, \dots, K\}, \ p_k (1-p_k) = 0 \iff \forall k \in \{1, \dots, K\}, \ p_k \in \{0, 1\}.$$

Since  $p_1 + \cdots + p_K = 1$ , there is one and only one index  $\ell \in \{1, \ldots, K\}$  such that  $p_\ell = 1$ , while  $p_k = 0$  for  $k \neq \ell$ .

For the second mapping, we note that  $-z \log z \ge 0$  for  $z \in [0, 1]$  with  $-z \log z = 0$  if and only if  $z \in \{0, 1\}$  (since the function vanishes for  $z \in \{0, 1\}$  and is strictly concave on (0, 1) as can be seen from computing the second derivative, which is  $z \mapsto -1/z$ ; the nonnegativity itself is already a consequence of  $\log z \le 0$  for  $z \in (0, 1]$ ). The conclusion then follows from the same reasoning as for the first mapping.

## 7.1.4 Some practical points

We discuss in this section various practical points.

**Categorical values.** In some cases, predictions (either regression or classification) are performed with inputs which cannot be ordered or compared, for instance colors of flowers. There are  $2^{K-1}-1$  possible partitions of the K labels in two groups (see Exercise 7.3), which is too large a number to be tackled unless K is very small. For regression, an option is to order categories by increasing mean of the output and split them as if they were an ordered predictor. This can also be done for binary classification, but is more complicated for multiclass classification, see the discussion in [24, Section 9.2.4].

**Exercise 7.3.** Show by induction that there are  $2^{K-1} - 1$  possible partitions of the K labels in two groups.

**Correction.** Denote by  $S_K$  the number of partitions of  $\{1, \ldots, K\}$  composed of two sets. Note that  $S_2 = 1$ . Moreover,

$$S_{K+1} = 2S_K + 1,$$

since the new element K + 1 can be put in any of the subsets for  $\{1, \ldots, K\}$ , with two choices per partition; or can be put in a subset by itself (*i.e.* which corresponds to partitioning the set into  $\{1, \ldots, K\}$  and  $\{K + 1\}$ ). Therefore,

$$S_{K+1} + 1 = 2(S_K + 1),$$

so that  $S_K + 1 = 2^{K-2}(S_2 + 1)$ , which leads to the claimed formula.

**Regularization.** As mentioned above, the training error of trees can be brought down to 0 by considering trees deep enough, in the absence of label noise. A heuristic approach to prevent this is to fix the maximal depth of the tree, and/or the minimal number of data points at the leaves. Another option is to perform backpruning, as discussed in Section 7.1.2.

Variance of the predictors. Trees can have a very low bias as they are flexible enough to accomodate any training set and loss function provided the tree is deep enough. This suggests that the associated predictors will have a large variance, i.e. will be unstable with respect to changes in the training data points. This is due in particular to their hierarchical structure, since a modification at a node in the tree during the learning process impacts all the leaves attached to this node (think for example of an error or a difference at the root node, which will lead to two completely different trees a priori). The high variance of predictors based on decision trees motivates considering ensemble methods; see Section 7.2.

**Choice of loss function.** As for other models, it is possible to consider asymmetric loss functions or splitting criteria, in order to put more emphasis on the correct classification of certain labels, or regression of certain values (think of cancer prediction or spam classification for instance).

Final discussion: advantages and disadvantages of decision trees. We conclude this section by listing the main advantages and disadvantages of trees for classification and regression. Advantages include

- the method is easily interpretable, as the data is broken up directly based on the features. This • is probably the main interest of the approach:
- it can handle data of mixed types, with discrete, continuous and categorical inputs;
- there is no need to standardize the data as the method is scale invariant by construction;
- it automatically performs feature selection; •
- it is relatively robust to outliers, as the tresholds can be considered as some forms of medians • or quantiles (although the predicted values are still based on averages, for regression);
- the method is fast to fit and scales well to large data sets.

All these characteristics make trees an interesting option for data mining. Trees however has some disadvantages, in particular:

- the predictions are usually not very accurate compared to other models, due to the greedy • approach used in the fit;
- the variance of the predictions is high due to some instability with respect to changes in the data inputs;
- the predictions are not smooth for regression, as the predictors are piecewise constant over regions of parameter space;
- there is some limitation in expressivity due to the axis splits (think of 2-dimensional data that would need to be separated along the diagonal: this requires rather deep trees).

# 7.2 Bagging and averaging techniques

The idea behind ensemble methods is to combine the results of predictors trained on different datasets, "as independent as possible", in order to obtain a better predictor – or at least a predictor of the same quality as with other techniques, but in reduced wall clock time (as computations on many simple predictors can be parallelized). The prediction obtained from ensemble methods is either an actual average for regression, or a majority vote for classification. Let us emphasize that,

#### 122 7 Trees and ensemble methods

although we present this approach in this chapter devoted to trees, it is in fact a general philosophy which can be used in many contexts.

Ensemble methods make sense for local averaging methods such as K nearest neighbors, and for nonlinear models obtained from empirical risk minimization. They do not make sense on the other hand for affine models obtained from empirical risk minimization as the average of affine models is still an affine model. Our presentation is based on regression in a supervised context, but can be extended to classification when considering the  $\Phi$ -risk associated with convex surrogates.

#### 7.2.1 Motivation: the ideal setting

If M independent and identically distributed data sets  $\mathcal{D}_1, \ldots, \mathcal{D}_M$  are available (which is of course often too strong an assumption...), one can construct a predictor  $\hat{f}$  by averaging the predictors  $\hat{f}_{\theta_m}$  trained on each data set  $\mathcal{D}_m$ :

$$\widehat{f}(x') = \frac{1}{M} \sum_{m=1}^{M} \widehat{f}_{\theta_m}(x').$$

Note that the parameters  $\theta_m$  are independent and identically distributed. Therefore,  $\hat{f}$  has the same bias as the underlying base predictors, but a variance of order O(1/M) (more precisely, the variance of the base predictors divided by M).

For example, for K nearest neighbors, the upper bound (1.21) obtained for a single predictor should be modified as

$$0 \leqslant \mathbb{E}_{\mathcal{D}}\left[\mathcal{R}(\widehat{f})\right] - \mathcal{R}^{\star} \leqslant \frac{\sigma^2}{KM} + 8B^2 \operatorname{diam}(\mathcal{X})^2 \left(\frac{2K}{n}\right)^{2/d}$$

*i.e.* the variance term has been divided by M. The optimal value of K then scales as  $M^{-d/(d+2)}n^{2/(d+2)}$ , and the optimal excess risk as  $(Mn)^{-2/(d+2)}$ . This bound is similar to the one obtained for single data set of size Mn. The main interest of averaging in this context is that the evaluation of the averaged predictor can be less expensive, or can be parallelized.

## 7.2.2 Bagging

The situation considered in Section 7.2.1 is an idealized one as one usually does not have independent data sets. This motivates artificially creating data sets with some randomness by replicating them from a baseline data set. This is precisely the idea behing bootstraping, which motivates the terminology "bagging" (bootstrap aggregating). More precisely, given a data set  $\mathcal{D} = \{(x_i, y_i)_{1 \leq i \leq n}\}$ , one constructs M data sets by sampling with replacement from  $\mathcal{D}$ . Of course, there is a non zero probability that two data points are identical in the data set. The average proportion of different data points can in fact be made precise, as the following exercise shows.

**Exercise 7.4.** What is the number of different data points seen in average, in the limit  $n \to +\infty$ ?

**Correction.** The probability not to select a given data point among the *n* ones is (n-1)/n = 1 - 1/n. The probability not to select it over *n* draws is then

$$p_{\rm ns} = \left(1 - \frac{1}{n}\right)^n \xrightarrow[n \to +\infty]{} {\rm e}^{-1}.$$

Therefore, the probability to select a given data point is

$$p_{\rm s} = 1 - p_{\rm ns} = 1 - \left(1 - \frac{1}{n}\right)^n \xrightarrow[n \to +\infty]{} 1 - {\rm e}^{-1}.$$

The number of different data points seen in average is finally  $np_s$ . The fraction of different data points converges to  $1 - e^{-1} \approx 0.63$ . This means that about 37% of the data points are not present in the data set.

Note that "out-of-bag" instances (*i.e.* data points which do not show up in the random data set which is sampled) can be used for validation. One interest of bagging is that it prevents algorithms to focus too much on some data points. This gives more robustness to the learning procedure and ensures a better generalization. The method works better for unstable estimators, with high variance; see for instance the mathematical analysis for 1-nearest neighbors presented in [4, Section 10.1.2].

### 7.2.3 Random forests

Random forests mix bagging, as an ensemble of decision trees is learned on bootstrapped samples of data, and random projections, since at each node of the trees splits are based only on a random subset of all features. The motivation for the latter extra source of randomness is to decorrelate the predictors as much as possible in order to benefit from the decrease in variance offered by ensemble methods. We refer to [7, 37] for a more detailed presentation and elements on the mathematical analysis of the approach.

# 7.3 Boosting

We present the approach for real-valued outputs. This is natural for regression; for classification it means that one works with convex surrogates to work out classification rules from real-valued outputs (see Section 3.1). The bottom line of the method is to sequentially fit an additive model as

$$f_T(x) = \sum_{t=1}^T \alpha_t F_t(x),$$
 (7.1)

with possibly some reweighting of the data in some cases. We first present the general philosophy of the method in Section 7.3.1, in particular the celebrated AdaBoost method and some of its variations; and then discuss gradient boosting in Section 7.3.2.

#### 7.3.1 General philosophy

Boosting is a greedy procedure where the predictor (7.1) is updated by looking for an extra term which maximally minimizes the empirical risk. More precisely, at iteration  $t \ge 1$ , one first solves the minimization problem

$$(\alpha_t, \theta_t) = \underset{\substack{\alpha \in \mathbb{R} \\ \theta \in \Theta}}{\operatorname{argmin}} \left\{ \frac{1}{n} \sum_{i=1}^n \ell \Big( y_i, f_{t-1}(x_i) + \alpha F(x_i; \theta) \Big) \right\},$$
(7.2)

and then sets

$$f_t = f_{t-1} + \alpha_t F(\cdot, \theta_t).$$

The procedure is initialized in some way, typically by setting  $f_0 = 0$ . A regularization term could be added to the empirical risk to minimize, but we do not introduce such a component here for simplicity of exposition. The details on how the minimization in (7.2) is exactly done depends on the choice of loss function, which is why we next consider two particular cases. Gradient boosting, presented in Section 7.3.2, allows for a more general approach which does not rely on the form of the loss function.

**Least-square boosting.** We consider the loss function  $\ell(y, z) = |y - z|^2$ . Introduce the residual  $r_{i,t}$  at step t for the prediction associated with  $x_i$ , namely

$$r_{i,t} = y_i - f_{t-1}(x_i).$$

In this context, the minimization problem (7.2) can be reformulated as

124 7 Trees and ensemble methods

$$(\alpha_t, \theta_t) = \operatorname*{argmin}_{\substack{\alpha \in \mathbb{R} \\ \theta \in \Theta}} \left\{ \frac{1}{n} \sum_{i=1}^n |r_{i,t} - \alpha F(x_i; \theta)|^2 \right\},\$$

which is the usual least-square fit of  $\{(r_{i,t})_{1 \leq i \leq n}\}$  by  $F(x_i, \theta)$  (when integrating the constant  $\alpha = 1$  into the definition of F). See for instance [40, Figure 18.6] for an illustration.

AdaBoost. This the historical example of boosting, used for binary classification on  $\mathcal{Y} = \{-1, 1\}$ . The base classifier functions  $F(\cdot; \theta)$  have values in  $\mathcal{Y}$ . One works here with the convex surrogate loss function  $\Phi(u) = e^{-u}$ , the associated elementary loss being  $\ell(y, z) = e^{-yz}$ . Note in particular that  $\mathbf{1}_{y\neq z} = \mathbf{1}_{yz\leqslant 0} \leqslant \ell(y, z)$ , in accordance with the discussion in Section 3.1.1.2. In this context, the minimization problem (7.2) can then be reformulated as

$$(\alpha_t, \theta_t) = \operatorname*{argmin}_{\substack{\alpha \in \mathbb{R} \\ \theta \in \Theta}} \left\{ \frac{1}{n} \sum_{i=1}^n \omega_{i,t} \mathrm{e}^{-\alpha y_i F(x_i;\theta)} \right\}, \qquad \omega_{i,t} = \mathrm{e}^{-y_i f_{t-1}(x_i)},$$

the associated classifier being

$$\operatorname{sign}(f_T) = \operatorname{sign}\left(\sum_{t=1}^T \alpha_t F(\cdot, \theta_t)\right).$$

In fact, the function to minimize can be rewritten as

$$\frac{1}{n}\sum_{i=1}^{n}\omega_{i,t}e^{-\alpha y_{i}F(x_{i};\theta)} = \left(\sum_{i:F(x_{i};\theta)=y_{i}}\omega_{i,t}\right)e^{-\alpha} + \left(\sum_{i:F(x_{i};\theta)\neq y_{i}}\omega_{i,t}\right)e^{\alpha}$$
$$= \left(e^{\alpha} - e^{-\alpha}\right)\sum_{i=1}^{n}\omega_{i,t}\mathbf{1}_{F(x_{i};\theta)\neq y_{i}} + e^{-\alpha}\sum_{i=1}^{n}\omega_{i,t}.$$

The latter formulation makes it clear that the optimization over  $\theta$  can be performed independently of  $\alpha$ , and corresponds to minimizing some training error for the 0-1 loss and a data set with weights. In fact, we will see below that  $\alpha_t$  is positive, so that, ideally, one would like to consider

$$\theta_t \in \operatorname*{argmin}_{\theta \in \Theta} \left\{ \sum_{i=1}^n \omega_{i,t} \mathbf{1}_{F(x_i;\theta) \neq y_i} \right\}.$$
(7.3)

As explained in Section 3.1, the latter optimization problem is difficult to solve. It is also not necessary to solve it too precisely in the context of boosting, which relies on improving weak classifiers such as trees of limited depth (assuming that such weak classifiers can indeed be obtained for the problem at hand).

Once  $\theta_t$  has been found, one can minimize over  $\alpha$ , as the following exercise shows.

Exercise 7.5. Find the minimizer of

$$\alpha_t \in \operatorname*{argmin}_{\alpha \in \mathbb{R}} \left\{ \frac{1}{n} \sum_{i=1}^n \omega_{i,t} \mathrm{e}^{-\alpha y_i F(x_i;\theta_t)} \right\}.$$

It is useful to introduce the classification error on the weighted data set:

$$E_{t} = \frac{\sum_{i=1}^{n} \omega_{i,t} \mathbf{1}_{F(x_{i};\theta_{t})\neq y_{i}}}{\sum_{i=1}^{n} \omega_{i,t}} \in [0,1].$$
(7.4)

Check that  $\alpha_t > 0$  when  $E_t \in (0, 1/2)$  (which corresponds to a better prediction than guessing at random).

Correction. The minimization problem is equivalent to

$$\min_{\alpha \in \mathbb{R}} \left\{ \left( e^{\alpha} - e^{-\alpha} \right) E_t + e^{-\alpha} \right\}.$$

The function to minimize is continuous and can be rewritten as  $E_t e^{\alpha} - (1 - E_t)e^{-\alpha}$ . It therefore goes to  $+\infty$  as  $\alpha \to \pm\infty$  when  $E_t \in (0, 1)$ . There exists in this case a global minimizer  $\alpha_t$ , which necessarily satisfies the Euler-Lagrange condition

$$E_t \mathrm{e}^{\alpha_t} = (1 - E_t) \mathrm{e}^{-\alpha_t},$$

so that

$$\alpha_t = \frac{1}{2} \log\left(\frac{1 - E_t}{E_t}\right). \tag{7.5}$$

One can check that the latter expression remains valid in the limits  $E_t \to 0$  or  $E_t \to 1$ . It is also apparent that  $\alpha \in (0, +\infty)$  when  $E_t \in (0, 1/2)$ .

Overall, starting from  $\omega_{i,1} = 1$  and  $f_0 = 0$ , the procedure therefore amounts to iterating the following steps for  $1 \leq t \leq T$ :

- find an approximate solution  $\theta_t$  to (7.3) (such that  $0 \leq E_t < 1/2$  with  $E_t$  defined in (7.4));
- choose the value  $\alpha_t$  given in Exercise 7.5;
- update the classifier as  $f_t = f_{t-1} + \alpha_t F(\cdot; \theta_t);$
- update the weights as

$$\omega_{i,t+1} = \begin{vmatrix} \omega_{i,t} e^{2\alpha_t} & \text{if } y_i \neq F(x_i; \theta_t), \\ \omega_{i,t} & \text{if } y_i = F(x_i; \theta_t). \end{vmatrix}$$

Note that the latter update is equivalent (up to an unimportant multiplicative constant) to setting  $\omega_{i,t} e^{\alpha_t}$  if  $y_i \neq F(x_i; \theta_t)$  and  $\omega_{i,t} e^{-\alpha_t}$  if  $y_i = F(x_i; \theta_t)$ .

In this procedure, misclassified data points get larger weights, and will therefore count more in the next iteration. However,  $(1 - E_t)e^{-\alpha_t} = E_te^{\alpha_t}$ , so the mass of the well classified and misclassified points are the same, but there are less and less misclassified points as their weight increases.

Let us conclude this section by a result showing that the training error decreases exponentially fast with the number of rounds of boosting.

Theorem 7.1. The training error is upper bounded as

$$\widehat{\mathcal{R}}(f_T) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{y_i f_T(x_i) \leqslant 0} \leqslant \exp\left(-2 \sum_{t=1}^T \left[\frac{1}{2} - E_t\right]^2\right).$$

In particular, if there exists  $\gamma \in [0, 1/2]$  such that

$$\forall t \in \{1, \dots, T\}, \qquad E_t \leqslant \frac{1}{2} - \gamma, \tag{7.6}$$

then  $\widehat{\mathcal{R}}(f_T) \leq e^{-2\gamma^2 T}$ .

The condition (7.6) means that the approximate solution to (7.3) is a weak classifier, slightly better than guessing at random. The result supports the claim that the optimization over  $\theta$  in (7.3) need not be performed too precisely.

*Proof.* Since the method is unchanged when multiplying the weights by a constant factor, we can consider the following weight update:

126 7 Trees and ensemble methods

$$\omega_{i,t+1} = \omega_{i,t} \frac{\mathrm{e}^{-\alpha_t y_i F(x_i;\theta_t)}}{Z_t}, \qquad Z_t = \sum_{j=1}^n \omega_{j,t} \,\mathrm{e}^{-\alpha_t y_j F(x_j;\theta_t)}.$$

Moreover, we start from

$$\omega_{i,1} = \frac{1}{n}.$$

These normalizations ensure that

$$\forall t \in \mathbb{N}, \qquad \sum_{i=1}^{n} \omega_{i,t} = 1.$$
(7.7)

Note next that

$$\frac{1}{n} e^{-y_i f_T(x_i)} = e^{-\alpha_T y_i F(x_i;\theta_T)} \dots e^{-\alpha_1 y_i F(x_i;\theta_1)} \omega_{i,1}$$
$$= e^{-\alpha_T y_i F(x_i;\theta_T)} \dots e^{-\alpha_2 y_i F(x_i;\theta_2)} \omega_{i,2} Z_1$$
$$= e^{-\alpha_T y_i F(x_i;\theta_T)} \dots e^{-\alpha_3 y_i F(x_i;\theta_3)} \omega_{i,3} Z_2 Z_1 = \dots = Z_T \dots Z_1 \omega_{i,T+1}.$$

Then, since  $\mathbf{1}_{u \leq 0} \leq e^{-u}$  for all  $u \in \mathbb{R}$ ,

$$\widehat{\mathcal{R}}(f_T) \leqslant \frac{1}{n} \sum_{i=1}^n e^{-y_i f_T(x_i)} = \sum_{i=1}^n Z_1 \dots Z_T \, \omega_{i,T+1} = Z_1 \dots Z_T.$$

Now, using first (7.4) and (7.7), and then (7.5),

$$Z_t = \sum_{i=1}^n \omega_{i,t} e^{-\alpha_t y_i F(x_i;\theta_t)} = e^{-\alpha_t} (1 - E_t) + e^{\alpha_t} E_t = 2\sqrt{E_t (1 - E_t)}.$$

We next use that, for  $u \in [0, 1]$ ,

$$2\sqrt{u(1-u)} = \sqrt{1 - 4\left(\frac{1}{2} - u\right)^2} \leqslant \exp\left(-2\left(\frac{1}{2} - u\right)^2\right)$$

since  $1 - v \leq e^{-v}$  so that  $\sqrt{1 - v} \leq e^{-v/2}$  for any  $v \in [0, 1]$ . The combination of the latter three (in)equalities then gives the claimed result.

#### 7.3.2 Gradient boosting

One limitation of AdaBoost and its variants is that the problem (7.2) has to be solved, with a procedure depending on the elementary loss function  $\ell$  at hand. Gradient boosting is a more generic procedure. The main idea behind it is to consider the addition of a new term  $\alpha F(\cdot;\theta)$  to the sum  $\alpha_1 F(\cdot;\theta_1) + \cdots + \alpha_{t-1} F(\cdot;\theta_{t-1})$  as a perturbation. Therefore, the function to minimize in (7.2) can be approximated by

$$\frac{1}{n}\sum_{i=1}^{n}\ell(y_{i}, f_{t-1}(x_{i})) + \frac{\alpha}{n}\sum_{i=1}^{n}\partial_{z}\ell(y_{i}, f_{t-1}(x_{i}))F(x_{i}; \theta),$$

the first term not depending on  $\theta$ . In order to have some maximal decrease of the above approximate functional, it suggests that the vector  $(F(x_1; \theta), \ldots, F(x_n; \theta))$  should be proportional to the vector  $(\partial_z \ell(y_1, f_{t-1}(x_1)), \ldots, \partial_z \ell(y_n, f_{t-1}(x_n)))$ . There are two ways to making this idea practical:

• minimizing over  $\theta \in \Theta$  the second part of the above approximate functional, for  $\alpha$  fixed, namely

$$\theta_t \in \operatorname*{argmin}_{\theta \in \Theta} \left\{ \frac{1}{n} \sum_{i=1}^n \partial_z \ell(y_i, f_{t-1}(x_i)) F(x_i; \theta) \right\}.$$

Once  $\theta_t$  is determined, there are various choices for the update of the predictor, including the magnitude  $\alpha_t$  for the new term, and some convex combination with previous iterates using some mixing parameter; see the discussion in [4, Section 10.2.2] for further precisions.

• fitting  $F(\cdot;\theta)$  in order for  $(F(x_1;\theta),\ldots,F(x_n;\theta))$  to be the best possible approximation of  $-(\partial_z \ell(y_1, f_{t-1}(x_1)), \ldots, \partial_z \ell(y_n, f_{t-1}(x_n)))$ . This can be obtained through some least square regression:

$$\min_{\theta \in \Theta} \left\{ \sum_{i=1}^{n} \left| F(x_i; \theta) + \partial_z \ell(y_i, f_{t-1}(x_i)) \right|^2 \right\}.$$

One can in addition find the magnitude  $\alpha_t$  by some line search, as in usual gradient methods. The so-obtained method is close to least-square boosting when  $\ell(y, z) = (y - z)^2$ .

When using trees, the new term which is added to the prediction function is taken of the form

$$F(x;\theta_t) = \sum_{j=1}^{J_t} w_{j,t} \mathbf{1}_{x \in R_{j,t}},$$

with regions  $R_{j,t}$  which can be learned as for usual decision trees for regression, based on fitting the gradient as discussed above, but with values  $w_{j,t}$  which are fit as

$$w_{j,t} \in \operatorname*{argmin}_{w \in \mathbb{R}} \left\{ \sum_{i: x_i \in R_{j,t}} \ell(y_i, f_{t-1}(x_i) + w) \right\}.$$

Indeed, for trees, the difficult part of the learning is to find the regions  $R_{j,t}$ , but once these regions are given, it is possible to fit the values  $w_{j,t}$  somewhat directly in the boosting objective function, instead of its gradient approximation. For the squared error, the value  $w_{j,t}$  is just the empirical mean of the residuals  $(y_i - f_{t-1}(x_i))_{1 \le i \le n}$  in the region (see [40, Section 18.5.5.1] and [24, Section 10.10.2]).

**XGBoost.** A very efficient and widely used implementation of gradient boosting trees is "extreme gradient boosting" (XGBoost), proposed in [11]. It improves on the above described gradient boosting method by

- adding a regularizer to the tree complexity;
- considering a second order approximation to the functional to minimize in (7.2) instead of a first order one;
- using some random sampling of features as for random forests;
- implementing various computer science tricks to improve the scalability of the method.

See for instance [40, Section 18.5.5.2] for further precisions.

# Neural networks

8.1	Feed-forward neural networks 129		
	8.1.1	Structure of multilayer perceptrons	
	8.1.2	Mathematical results on approximation properties	
8.2	Train	ing neural networks 133	
	8.2.1	Computing the gradient	
	8.2.2	Optimization methods	
	8.2.3	Regularization	
8.3	Unsupervised learning with neural networks		
	8.3.1	Structure of autoencoders	
	8.3.2	Interpretations of the loss function	
8.4	$\mathbf{Othe}$	r types of neural networks 144	

We discuss in this chapter nonlinear models based on neural networks, which have nowadays become increasingly important and useful for many machine learning applications. Our focus is on feed-forward neural networks, whose structure is described in Section 8.1, and whose training is made precise in Section 8.2. We will mostly considered supervised learning, both for classification and regression; unsupervised learning based on autoencoders is however considered in Section 8.3. We conclude the chapter with a brief presentation of other types of neural networks in Section 8.4.

The main bibliographical resources for this chapter are [40, Chapters 13 to 15], with an emphasis on Chapter 13 which covers more basic material (see also Section 20.3 for autoencoders); as well as the book by Goodfellow, Bengio and Courville [22] (see in particular Chapters 6, 7, 8 and 11). Concerning the practical implementation, many useful contents can be found in the book [57]. Some (more adavanced) elements on the mathematical analysis of neural networks can be read in [4, Chapter 9].

# 8.1 Feed-forward neural networks

The motivation for considering neural networks is to construct nonlinear mappings to go from the inputs to the outputs, without making explicit feature functions, but by describing them using a very expressive parametrized set of functions. Somehow, neural networks used in supervised learning (either for classification or regression) can be seen as an efficient way of learning a feature function  $\phi_p(x)$  (parametrized by a vector p) to be used in standard, possibly linear models, with a prediction function of the form  $f_{\theta}(x) = w^{\top} \phi_p(x) + b$ , with  $\theta = (w, b, p)$ .

We consider in this section feed-forward neural networks, also known as multilayer perceptrons. Other types of architectures are briefly discussed in Section 8.4.

#### 130 8 Neural networks

#### 8.1.1 Structure of multilayer perceptrons

Layers and parameters. The prediction function for feed-forward neural networks is defined in an iterative manner by composing elementary functions as

$$f_{\theta}(x) = \left(g_{\theta_L} \circ g_{\theta_{L-1}} \circ \cdots \circ g_{\theta_1}\right)(x)$$

This formula corresponds to a network with L - 1 hidden layers and an output layer, with input layer  $x \in \mathbb{R}^d$ . Shallow networks have small values of L, while deep neural networks correspond to larger values of L (say, about 20 for current implementations). The input is first (nonlinearly) transformed into an element of  $\mathbb{R}^{d_1}$  with  $g_{\theta_1} : \mathbb{R}^d \to \mathbb{R}^{d_1}$ . Then, each function  $g_{\theta_L} : \mathbb{R}^{d_{\ell-1}} \to \mathbb{R}^{d_\ell}$ modifies the output of the previous layer, until the last function  $g_{\theta_L} : \mathbb{R}^{d_{L-1}} \to \mathbb{R}^{d_L}$  maps it to some output vector. The dimension of the output is  $d_L = 1$  for simple regression problems, but it can be a vector in many situations, for instance when performing regression on vector labels, or when performing multiclass classification. In the latter case,  $d_L$  is typically the number of classes to be predicted, as the output is passed through a softmax function, following the same strategy as in Section 3.2.4.

The transformation functions for each layer map an input  $a_{\ell-1} \in \mathbb{R}^{d_{\ell-1}}$  obtained from the previous layer to an input  $a_{\ell} \in \mathbb{R}^{d_{\ell}}$  to be used by the next layer, of the form

$$a_{\ell} = g_{\theta_{\ell}}(a_{\ell-1}) = \rho_{\ell} \left( W_{\ell} a_{\ell-1} + b_{\ell} \right), \qquad W_{\ell} \in \mathbb{R}^{d_{\ell} \times d_{\ell-1}}, \quad b_{\ell} \in \mathbb{R}^{d_{\ell}},$$

where  $\rho_{\ell} : \mathbb{R} \to \mathbb{R}$  is a so-called activation function, which is applied componentwise. In the sequel, we denote by  $z_{\ell} = W_{\ell}a_{\ell-1} + b_{\ell} \in \mathbb{R}^{d_{\ell}}$  the argument of the activation. Historically, perceptrons were constructed with Heaviside functions for activations, but the associated models are difficult to train. Nowadays activation functions are continuous and often smooth (see below).

**Remark 8.1.** In practice, all operations are usually done in a vectorized manner. The input of the neural network is then the design matrix  $X \in \mathbb{R}^{n \times d}$  whose *i*-th line is the line vector  $x_i \in \mathbb{R}^{1 \times d}$ ; and the output is the vector  $y = (y_1, \ldots, y_n)^\top \in \mathbb{R}^{n \times d_L}$ . For simplicity of exposition, we however consider that the input is a column vector of size d, and the output is a column vector of size  $d_L$ .

The parameters of the model are  $\theta_1 = (W_1, b_1), \theta_2 = (W_2, b_2), \dots, \theta_L = (W_L, b_L)$ , so that  $\theta = (\theta_1, \dots, \theta_L)$ . The total number of parameters is therefore

$$\sum_{\ell=1}^{L} d_{\ell} (1 + d_{\ell-1}),$$

with  $d_0 = d$  the dimension of the input, and  $d_L$  the dimension of the output. The number of parameters is usually very large, possibly (much) larger than the number of training points; see for instance Exercise 8.1 for a concrete example.

**Exercise 8.1.** What is the number of parameters for a neural network classifying the MNIST set of digits (images of  $28 \times 28$  pixels) based on a single hidden layer? What is this number for a hidden layer of size 50? Comment the magnitude of this number in view of the size of the data set.

**Correction.** There are L = 2 layers overall besides the input layer. The size of the input is  $d_0 = 28^2 = 784$ , while the output has dimension  $d_2 = 10$  since there are 10 classes to predict. Therefore, the total number of parameters is  $d_1(1 + d_0) + d_2(1 + d_1) = d_1(1 + d_0 + d_2) + d_2$ . For  $d_1 = 50$ , one obtains 39,760 parameters. This number should be compared to the size of the training data set, which is composed of 60,000 training points: there are almost as many parameters as the number of training points!

Activation functions. There are various choices for the activation functions  $\rho_{\ell}: \mathbb{R} \to \mathbb{R}$  used to pass from one layer to the other:

- the linear function  $\rho_{\ell}(z) = z$ , which can in particular be considered for the final layer, to produce a result such as  $w^{\top}\phi(x) + b$ , the first L-1 layers being used to come up with a featurization function;
- the sigmoid function  $\rho_{\ell}(z) = \sigma(z) = (1 + e^{-z})^{-1};$ the hyperbolic tangent  $\rho_{\ell}(z) = \tanh(z) = \frac{e^z e^{-z}}{e^z + e^{-z}};$
- the rectifier linear unit (ReLU) function  $\rho_{\ell}(z) = \max(0, z)$ , or leaky versions such as  $\rho_{\ell}(z) =$  $\max(0, z) + \alpha \min(0, z)$  with  $\alpha > 0$  small (the latter option ensuring that  $\rho'_{\ell}$  does not vanish for z < 0, which may cause issue for training, see Section 8.2.2).

The selection of the activation function should be made on the basis of a trial and error process, *i.e.* it can somehow be considered as a hyperparameter to tune using some (cross) validation procedure. However, many choices lead to similar performances. The general trend is that activation functions with gradients that saturate at infinity, such as sigmoid or hyperbolic tangent, are fine for shallow networks; but for deeper networks one should choose activation functions which do not saturate in order for training to be possible. ReLU and its recent variations (for instance smoothed out versions such as GELU (Gaussian Error Linear Unit) or swish) are currently common choices.

The activation function for the last layer (*i.e.* for the output) depends on the task:

- linear functions are considered for regression;
- sigmoid functions are used for binary classification in  $\{0, 1\}$ , and hyperbolic tangent for binary classification in  $\{-1, 1\}$ ;
- multiclass classification can be performed with a softmax function, which, given  $z_L$  =  $(z_{L,1},\ldots,z_{L,d_L})$  returns

$$a_L = \frac{1}{Z_L} \left( e^{z_{L,1}}, \dots, e^{z_{L,d_L}} \right), \qquad Z_L = \sum_{j=1}^{d_L} e^{z_{L,j}}$$

More advanced choices are discussed in [22, Section 6.2.2.4], for instance predictions based on sampling from a Gaussian distribution with mean  $\mu_{\theta}(x)$  and variance  $\sigma_{\theta}(x)^2$ , where  $\mu_{\theta}, \sigma_{\theta}^2$  are learned by the first L-1 layers.

Loss functions. The loss functions to consider for training neural networks are the same as for other machine learning problems, in particular the square loss  $(y - \hat{y})^2$  for regression problems, and, for classification problems, the cross entropy (3.19) in the binary case or (3.21) for multiclass classification.

Architecture design. Neural networks are quite flexible, and a practical question of interest is how to choose the number of units and the way they are connected. For fully connected feed-forward neural networks, this mean choosing the number of hidden layers, and the width of these layers. For a given number of parameters, one can choose for instance between wider and shallower networks or deeper networks with fewer units per layer in order to perform some hierarchical learning. The reader is encouraged to test strategies on the tensorflow playground<sup>1</sup> to get a feeling of various options.

Some historical landmarks. Let us conclude this section with some historical perspective on the development of neural networks. Neural networks have been considered since Rosenblatt's perceptron in the 50s. Many ideas concerning the structure and training of neural networks were already present in the 80s and 90s, but other learning models were more successful at the time. A strong surge of interest arose after 2010 due to breakthough results in image classification. The progress here is due to two factors:

<sup>&</sup>lt;sup>1</sup> See https://playground.tensorflow.org

#### 132 8 Neural networks

- (i) the increase in computing power, in particular through graphical processing units (GPUs) which can very efficiently handle matrix/vector operations, and also through dedicated hardware such as tensor processing units (TPUs);
- (ii) the size of the data sets available for training have substantially increased, which made it possible to train neural networks with more parameters, in particular deep neural networks.

#### 8.1.2 Mathematical results on approximation properties

Neural networks have the so-called "universal approximation property", which allow them to approximate arbitrary continuous functions [12], and even any Borel measurable function [26]. The function to represent is however only obtained in the limit of layers of infinite width – in fact, with a single hidden layer when the output is one-dimensional. We present here a simple argument to substantiate these claims, taken from [4, Section 9.3], and, for simplicity of exposition, do not provide convergence rates (which would require extra smoothness on the function to approximate in order to state quantitative results). The proof we write is for ReLU activation functions, the result being hower rather easy to extend to any non polynomial activation function [36].

Before writing an approximation argument, let us immediately emphasize that the results mentioned above are representation results, meaning that they state that a given function can be arbitrarily well approximated by neural networks of increasing sizes. This does not mean that, in practice, the optimization algorithms used to train neural networks will be able to find parameters which indeed allow to correctly approximate the function under consideration.

To write a precise approximation result, we consider the class of continuous functions with one-dimensional input  $x \in [-R, R]$  for some R > 0 to simplify the presentation, and output in  $\mathbb{R}$ . The model we consider is that of a single layer of size m. The prediction function associated with the neural network can then be written as

$$f(x) = \sum_{j=1}^{m} \eta_j \rho\left(w_j x + b_j\right), \qquad \eta_j, w_j, b_j \in \mathbb{R},$$

with  $\rho$  the activation function, here  $\rho(z) = \max(0, z)$ . The arguments  $w_j x + b_j$  of the activation functions are obtained by an affine transformation of the input, while the parameters  $\eta_j$  perform a linear combination of the outcomes of the hidden layer to perform the prediction (note that there is no bias for the last layer). The proof proceeds in two steps: (i) we first show that any piecewise affine function with m - 2 kinks in [-R, R] can be represented by a neural network with a single hidden layer of size m; (ii) we next conclude by an approximation argument.

Representation of piecewise affine function with m-2 kinks in [-R, R]. We denote by  $x_+ = \max(x, 0)$  the ReLU function. We first consider the case when f(-R) = 0, and denote by  $-R < X_1 < \cdots < X_{m-2} < R$  the location of the kinks. On the interval  $[-R, X_1]$ , one has  $f(x) = v_1(x+R)_+$  for some slope  $v_1 \in \mathbb{R}$ . On  $[X_1, X_2]$ , we then add a part  $v_2(x-X_1)_+$ . This extra term does not change the representation on  $[-R, X_1]$ , and allows to represent the function on  $[X_1, X_2]$  by choosing  $v_2$  such that the slope of f is  $v_1 + v_2$ . We then continue adding terms of the form  $v_{k+1}(x - X_k)_+$  for  $2 \leq k \leq m-2$ , with  $v_{k+1}$  chosen such that the slope is the correct one on the interval  $[X_k, X_{k+1}]$  (with the convention that  $X_{m-1} = R$ ). Overall, this requires m-1neurons.

When  $f(-R) \neq 0$ , we need an extra neuron. We use in fact the identity

$$\frac{1}{2R}(x+R)_{+} + \frac{1}{2R}(-x+R)_{+} = 1.$$

This amounts to changing the weight  $v_1$  of the first neuron in the construction above, and adding an extra neuron outputing  $(-x+R)_+$ . Overall, the function is then represented by a linear combination of m terms of the form  $(x - X_k)_+$  with  $0 \le k \le m - 1$  (with the convention  $X_0 = -R$ ).

Approximation argument to represent continuous functions. We use the property that continuous functions can be approximated arbitrary closely in the  $C^0$  norm by piecewise affine functions. Indeed, fix  $f \in C^0([-R, R])$ . Then f is uniformly continuous on this interval: for any  $\varepsilon > 0$ , there exists  $\eta > 0$  such that  $|f(x) - f(x')| \leq \varepsilon/2$  for  $|x - x'| \leq \eta$ . Consider then the piecewise approximation  $\tilde{f}$  which coincides with f on a regular mesh of size smaller than  $\eta$ , *i.e.* define  $\tilde{\eta} = 2R/K$  with  $K = \lceil 2R/\eta \rceil$ , and introduce the points  $X_k = -R + k\tilde{\eta}$  for  $0 \leq k \leq K$ . Then, for a given  $x \in [-R, R]$ , considering the mesh point  $X_i$  closest to x (*i.e.* such that  $|x - X_i| \leq |x - X_j|$  for all  $0 \leq j \leq K$ ),

$$\left|f(x) - \widetilde{f}(x)\right| \leq \left|f(x) - f(X_i)\right| + \left|f(X_i) - \widetilde{f}(X_i)\right| + \left|\widetilde{f}(X_i) - \widetilde{f}(x)\right|.$$

The first term on the right hand side is bounded by  $\varepsilon/2$  in view of the property of uniform continuity of f. The second term vanishes since f and  $\tilde{f}$  coincide at the mesh points. For the last term, we use the fact that  $\tilde{f}$  is affine to write

$$\left| \widetilde{f}(X_i) - \widetilde{f}(x) \right| \leq \min \left\{ \left| \widetilde{f}(X_i) - \widetilde{f}(X_{i+1}) \right|, \left| \widetilde{f}(X_i) - \widetilde{f}(X_{i-1}) \right| \right\}$$
$$= \min \left\{ \left| f(X_i) - f(X_{i+1}) \right|, \left| f(X_i) - f(X_{i-1}) \right| \right\}.$$

The right hand side is bounded by  $\varepsilon/2$  since mesh points are at a distance smaller than  $\eta$ . Overall, we obtain that  $|f(x) - \tilde{f}(x)| \leq \varepsilon$  uniformly in  $x \in [-R, R]$ .

As a conclusion, one can therefore approximate arbitrarily closely in  $C^0$  norm any continuous function on a finite interval [-R, R] by a finite neural network with ReLU activation functions.

**Remark 8.2.** The above approximation result is based on networks of infinite width. A similar result for deep neural networks (i.e. where the depth would grow instead of the width) is currently less clear (see some reference in [4, Section 9.7]).

# 8.2 Training neural networks

The training of neural networks can be nicely visualized using the playground of tensorflow. This task seems to be a rather difficult one at first sight, as it corresponds to a nonconvex optimization problem involving a very large number of parameters. It can however be made computationnally efficient thanks to

- backpropagation and more generally automatic differentiation to compute gradients with respect to parameters and inputs;
- SGD and its variations to leverage minibatching to reduce the cost of one training step;
- good software to make the implementation easy for the user; in particular Tensorflow (Google) and PyTorch (Meta).

We start by discussing how to compute the gradient in Section 8.2.1, then discuss some elements on optimization and regularization which are specific to neural networks (see respectively Sections 8.2.2 and 8.2.3).

#### 8.2.1 Computing the gradient

We discuss in this section how to compute the gradient of functionals of the output of neural networks, using the celebrated backpropagation algorithm. The presentation is based on [42, Chapter 2]. We assume that the activations functions are continuously differentiable for the derivation below. Network and loss function. We consider a neural network composed of L layers, whose parameters are

$$\theta = \{W_\ell, b_\ell\}_{1 \le \ell \le L}.$$

In accordance with the notation of Section 8.1.1, we denote the weighted input at layer  $\ell$  by  $z_{\ell}$  and the activation by  $a_{\ell}$ :

$$a_{\ell} = \rho_{\ell}(z_{\ell}), \qquad z_{\ell} = W_{\ell}a_{\ell-1} + b_{\ell}, \qquad b_{\ell} \in \mathbb{R}^{d_{\ell}}, \quad W_{\ell} \in \mathbb{R}^{d_{\ell} \times d_{\ell-1}},$$

starting from an input  $a_0 = x$ . We provide the derivation here for a single input for simplicity. In practice, operations are vectorized when several inputs are considered, so that the neural networks takes a matrix as input, and returns a vector.

Let us first make precise the function whose gradient needs to be computed, namely the empirical risk

$$\widehat{\mathcal{R}}_n(\theta) = \frac{1}{n} \sum_{i=1}^n \mathscr{L}(y_i, a_{L,i}),$$
(8.1)

where  $\mathscr{L}$  is some elementary loss function<sup>2</sup> with values in  $\mathbb{R}$  (for instance,  $\mathscr{L}(y, \hat{y}) = (y - \hat{y})^2$ for least square regression), and  $a_{L,i}$  is the output of the neural network for the input  $x_i$ . It is possible to add a penalization term  $\lambda \Omega(\theta)$ , but this extra term does not pose any challenge in the evaluation of its gradient, so we do not consider it in this discussion for simplicity of exposition. Also, for simplicity, we write the derivation for a single data point (*i.e.* n = 1), and denote by  $\widehat{\mathcal{R}}$ the corresponding loss function. It is straightforward to obtain gradients for general values of nfrom the formulas for n = 1.

**Computation of gradients with respect to parameters.** We now discuss how to compute the gradient of  $\widehat{\mathcal{R}}$  with respect to  $\theta$  (a similar derivation can be performed to compute the gradient of the output of the neural network with respect to the input x, which is needed for certain applications). This is done by starting from the output, and going back to the input, with the chain rule. More precisely, define

$$\forall j \in \{1, \dots, d_\ell\}, \qquad \delta_j^\ell = \frac{\partial \widehat{\mathcal{R}}}{\partial z_{\ell,j}} \in \mathbb{R},$$

where  $\widehat{\mathcal{R}}$  is seen (with some abuse of notation) as a function of  $z_{\ell}$ . In fact, we write

$$\widehat{\mathcal{R}}(\theta) = R_{\ell}(z_{\ell})$$

in the sequel in order to understand the recursion relation between  $\delta^{\ell}$  and  $\delta^{\ell-1}$ . Let us start with the output layer, for which

$$\mathcal{R}(\theta) = R_L(z_L) = \mathscr{L}(y, \rho_L(z_L)).$$

Therefore,

$$\delta^L = \rho'_L(z_L)\partial_{\widehat{y}}\mathscr{L}(y,\rho_L(z_L)) \in \mathbb{R}^{d_L},$$

We next express  $R_{L-1}$  in terms of  $R_L$  as

$$\widehat{\mathcal{R}}(\theta) = R_{L-1}(z_{L-1}) = R_L (W_L \rho_{L-1}(z_{L-1}) + b_L),$$

and claim that

$$\delta^{L-1} = \left( W_L^\top \delta^L \right) \odot \rho'_{L-1}(z_{L-1}) \in \mathbb{R}^{d_{L-1}}.$$
(8.2)

In the latter equality,  $\odot$  is the componentwise product: for two vectors  $a, b \in \mathbb{R}^d$ ,

<sup>&</sup>lt;sup>2</sup> We do not use here the notation  $\ell$  for the elementary loss function, as in other chapters, in order to avoid confusions with the index of the layer.

$$a \odot b = \begin{pmatrix} a_1 b_1 \\ a_2 b_2 \\ \vdots \\ a_d b_d \end{pmatrix}.$$

Note that the dimensions agree in (8.2) as  $z_{L-1}, W_L^{\top} \delta^L \in \mathbb{R}^{d_{L-1}}$ , so that the componentwise product indeed makes sense. Note also that for backward operations, it is the transpose of the matrix  $W_L$  that appears (while in the forward pass  $W_L$  is being used).

**Exercise 8.2.** Prove the equality (8.2).

**Correction.** We prove the result by perturbing the argument  $z_{L-1}$  by a small quantity  $h \in \mathbb{R}^{d_{L-1}}$ , and performing Taylor expansions. By definition of  $\delta^{L-1}$ , the perturbation in the output is

$$R_L \left( W_L \rho_{L-1}(z_{L-1} + h) + b_L \right) - R_L \left( W_L \rho_{L-1}(z_{L-1}) + b_L \right) = h^\top \delta^{L-1} + \mathcal{O} \left( \|h\|^2 \right).$$

Now,

$$\rho_{L-1}(z_{L-1}+h) = \rho_{L-1}(z_{L-1}) + \rho'_{L-1}(z_{L-1}) \odot h + \mathcal{O}\left(\|h\|^2\right),$$

where  $\rho_{L-1}, \rho'_{L-1}$  are applied componentwise, so that

$$W_L \rho_{L-1}(z_{L-1}+h) + b_L = W_L \rho_{L-1}(z_{L-1}) + b_L + W_L \left(\rho'_{L-1}(z_{L-1}) \odot h\right) + O\left(\|h\|^2\right)$$

Using the definition of  $\delta^L$  as the gradient of  $R_L$ , we find

$$R_L \left( W_L \rho_{L-1}(z_{L-1}+h) + b_L \right) - R_L \left( W_L \rho_{L-1}(z_{L-1}) + b_L \right) = \delta^L \cdot W_L \left( \rho'_{L-1}(z_{L-1}) \odot h \right) + \mathcal{O} \left( \|h\|^2 \right).$$

Let us finally rewrite the first term on the right hand side of the previous equality as

$$\delta^{L} \cdot W_{L} \left( \rho_{L-1}'(z_{L-1}) \odot h \right) = \sum_{j=1}^{d_{L}} \sum_{k=1}^{d_{L-1}} \delta_{j}^{L} \left[ W_{L} \right]_{j,k} \rho_{L-1}'(z_{L-1,k}) h_{k}$$
$$= \sum_{k=1}^{d_{L-1}} h_{k} \sum_{j=1}^{d_{L}} \left[ W_{L}^{\top} \right]_{k,j} \delta_{j}^{L} \rho_{L-1}'(z_{L-1,k})$$
$$= \sum_{k=1}^{d_{L-1}} h_{k} \left[ W_{L}^{\top} \delta^{L} \odot \rho_{L-1}'(z_{L-1}) \right]_{k}$$
$$= h \cdot \left[ W_{L}^{\top} \delta^{L} \odot \rho_{L-1}'(z_{L-1}) \right].$$

The claimed expression for  $\delta^{L-1}$  comes from the fact that the latter quantity is equal to  $h \cdot \delta^{L-1}$ .

It can similarly be shown that

$$\delta^{\ell-1} = \left( W_{\ell}^{\top} \delta^{\ell} \right) \odot \rho_{\ell-1}'(z_{\ell-1}) \in \mathbb{R}^{d_{\ell-1}}$$

The latter expression makes sense as the componentwise multiplication of  $\rho'_{\ell-1}(z_{\ell-1}) \in \mathbb{R}^{d_{\ell-1}}$ and  $W_{\ell}^{\top} \delta^{\ell} \in \mathbb{R}^{d_{\ell-1}}$ . Here as well, the expression of  $\delta^{\ell-1}$  involves the transpose matrix  $W_{\ell}^{\top}$ , which encodes the transpose operation allowing to go from the layer  $\ell$  to the layer  $\ell-1$ .

Once the vectors  $(\delta^{\ell})_{\ell=L,\dots,1}$  are determined, one can compute the gradients of the loss with respect to  $b_{\ell}$  as

$$\nabla_{b_\ell}\widehat{\mathcal{R}}(\theta) = \delta^\ell$$

where we used that  $z_{\ell} = W_{\ell}\rho_{\ell-1}(z_{\ell-1}) + b_{\ell}$ . The gradient with respect to  $W_{\ell}$  is

$$\nabla_{W_{\ell}}\widehat{\mathcal{R}}(\theta) = \delta^{\ell} \rho_{\ell-1}(z_{\ell-1})^{\top} = \delta^{\ell} a_{\ell-1}^{\top}.$$
(8.3)

### 136 8 Neural networks

Note that the latter equality makes sense since  $\delta^{\ell} \in \mathbb{R}^{d_{\ell} \times 1}$  while  $a_{\ell-1}^{\top} = \rho_{\ell-1}(z_{\ell-1})^{\top} \in \mathbb{R}^{1 \times d_{\ell-1}}$ , so that  $\nabla_{W_{\ell}} \widehat{\mathcal{R}}(\theta) \in \mathbb{R}^{d_{\ell} \times d_{\ell-1}}$  has the same dimension as  $W_{\ell}$ .

**Exercise 8.3.** Prove the equality (8.3).

**Correction.** Fix all parameters in  $\theta$  except the component  $W_{\ell}$ , and consider a perturbation  $w_{\ell} \in \mathbb{R}^{d_{\ell} \times d_{\ell-1}}$  of this parameter. The associated perturbed set of parameters is denoted by  $\theta + \eta_{w_{\ell}}$ . We compute

$$\begin{aligned} \widehat{\mathcal{R}} \left( \theta + \eta_{w_{\ell}} \right) &- \widehat{\mathcal{R}} (\theta) = R_{\ell} \left( z_{\ell} + w_{\ell} \rho_{\ell-1}(z_{\ell-1}) \right) - R_{\ell}(z_{\ell}) = \delta^{\ell} \cdot w_{\ell} \rho_{\ell-1}(z_{\ell-1}) + \mathcal{O}(\|w_{\ell}\|^{2}) \\ &= \sum_{j=1}^{d_{\ell}} \sum_{k=1}^{d_{\ell-1}} \delta_{j}^{\ell} [w_{\ell}]_{j,k} \rho_{\ell-1}(z_{\ell-1,k}) + \mathcal{O}(\|w_{\ell}\|^{2}) \\ &= \sum_{j=1}^{d_{\ell}} \sum_{k=1}^{d_{\ell-1}} [w_{\ell}]_{j,k} \left[ \delta^{\ell} \rho_{\ell-1}(z_{\ell-1})^{\top} \right]_{j,k} + \mathcal{O}(\|w_{\ell}\|^{2}). \end{aligned}$$

By definition of the gradient, the first term on the right hand side of the previous equality is

$$\operatorname{Tr}\left(w_{\ell}^{\top} \nabla_{W_{\ell}} \widehat{\mathcal{R}}(\theta)\right),$$

which leads to the claimed equality.

Final backpropagation algorithm. Let us summarize the procedure for the loss based on a single data point and for parameters  $\theta = \{W_{\ell}, b_{\ell}\}_{1 \leq \ell \leq L}$ :

• one first does a forward pass through the network, starting from the input  $a_0 = x$ , and evaluating, for  $\ell = 1, \ldots, L$ :

$$z_{\ell} = W_{\ell} a_{\ell-1} + b_{\ell}, \qquad a_{\ell} = \rho_{\ell}(z_{\ell}).$$

This allows to perform the final prediction  $f_{\theta}(x) = a_L$ . The associated loss is  $\widehat{\mathcal{R}}(\theta) = \mathscr{L}(y, a_L)$ .

one then performs a backward pass, starting from the gradient of the output

$$\delta^L = \rho'_L(z_L)\partial_{\widehat{y}}\mathscr{L}(y, a_L),$$

and then evaluating, for  $\ell = L, \ldots, 2$ :

$$\delta^{\ell-1} = \left( W_{\ell}^{\top} \delta^{\ell} \right) \odot \rho_{\ell-1}'(z_{\ell-1}).$$

The gradients with respect to the parameters are finally obtained for  $\ell = 1, \ldots, L$  as

$$\nabla_{W_{\ell}}\widehat{\mathcal{R}}(\theta) = \delta^{\ell} a_{\ell-1}^{\top}, \qquad \nabla_{b_{\ell}}\widehat{\mathcal{R}}(\theta) = \delta^{\ell}.$$

Extension to more complex architectures. The most recent neural networks have structures more complicated than feedforward multilayer perceptrons (see some elements in Section 8.4), and the backpropagation algorithm as described above is no longer sufficient. One example of a new element in the architecture is the presence of skip connection, where the value of a node at some layer  $\ell$  is directly fed to a subsequent layer  $\ell + k$  with  $k \ge 2$  (*i.e.* not only to the next layer).

One resorts in these situations to automatic differentiation techniques, see [22, Section 6.5] for a more detailed presentation. In these methods, the dependence of the final output in terms in the input can be made precise with a computational graph (which is a directed acyclic graph). The determination of this dependency graph is made "just in time" in software such as JAX or PyTorch. As an example, one can construct the graph associated with the evaluation of the function

$$f(x_1, x_2) = x_2 \mathrm{e}^{x_1} \sqrt{x_1 + x_2 \mathrm{e}^{x_1}}.$$

This involves going from the input  $x_1$  to  $e^{x_1}$ , then multiplying this value by the input  $x_1$ . This quantity is then added to the input  $x_1$  using a skip connection from the input, and then passed through a square root; and also directly passed to the output using a skip connection in order to multiply the result of taking the square root.

#### 8.2.2 Optimization methods

The various optimization techniques presented in Chapter 4 are used to train neural networks. Early stopping is used throughout to monitor convergence and decide when to stop training. Let us already mention that there is a strong practical knowledge to be used to make training effective, in particular in the choice of the parameters of the algorithms (minibatch size, learning rates, etc). The algorithms which are particularly relevant to train neural networks are SGD (with or without momentum variables) and Adam. The latter algorithm is currently the most common choice. These methods work surprisingly well to minimize (regularized) empirical risks, despite the non-convexity of the problem. This fact is not well understood from a theoretical viewpoint.

Let us next discuss points specific to the training of neural networks, in particular elements which make the optimization easier, and allow to effectively train deep neural networks. The main point is to avoid gradients which degenerate, either by being too small or too large for certain parameters. When gradients are small, it is no longer possible to optimize; while, when gradients are too large, learning rates need to be severely limited, and/or the computational procedure crashes due to some overflow. As the computation of the gradient involves operations over series of layers, it is generically expected that the magnitude of the gradient either exponentially decreases or increases (this can be understood for instance when multiplying vectors by random matrices). It is possible to ensure that gradients are better behaved by

- choosing *activation functions* that do not saturate, for instance ReLU instead of sigmoids or hyperbolic tangent. The idea is that derivatives of activation functions should not be too small when the input of this function are large in absolute values, otherwise gradients will vanish;
- considering *loss functions* such as the cross entropy which involve a log and therefore allow to limit the saturation of the sigmoid function used to predict probabilities for classification problems. The use of properly chosen convex surrogate functions can have an impact on the quality of the training;
- introducing batch normalization layers (see [22, Section 8.7.1]). The idea is to add a layer to renormalize the outputs  $a_{\ell} \in \mathbb{R}^{d_{\ell}}$  for the  $\ell$ -th layer to express them as  $a_{\ell} = \gamma_{\ell} \overline{a}_{\ell} + \beta_{\ell}$ , where  $\overline{a}_{\ell}$  has mean 0 and variance 1. The two parameters  $\gamma_{\ell}$ ,  $\beta_{\ell}$  are learnt, and allow to set the scale of the outputs of the  $\ell$ -th layer. This procedure does not change the expressivity of the neural network, but renders the numerical method better behaved;
- modifying the architechture. This is where, historically, most of the benefit came from. One idea is to consider "more linear" functions, as in ResNets where the update can be written as  $a_{\ell+1} = a_{\ell} + \mathcal{F}_{\ell}(a_{\ell-1}; \theta_{\ell-1})$  for some of the blocks instead of the update  $a_{\ell+1} = \mathcal{F}_{\ell}(a_{\ell-1}; \theta_{\ell-1})$ , *i.e.* there is a dominant linear part to which a perturbation is added. Another idea is to resort to skip connections in order for certain parameters to have a shorter computational path to the output;
- clipping gradients, which means that a component  $g_k$  of the gradient is replaced by  $\min(1, c/|g_k|)g_k$  for some parameter c > 0, so that the effective gradient is  $g_k$  for  $|g_k| \leq c$ , and  $c \operatorname{sign}(g_k)$  otherwise;
- carefully *initializing the parameters* of the neural network. There are various rules here, depending on the activation function and the topology of the neural network. A typical rule is that the variance of the random variables used to initialize the parameters at a given node should be proportional to  $1/(n_{\rm in} + n_{\rm out})$ , where  $n_{\rm in}$  is the number of incoming connections and  $n_{\rm out}$  is the number of outgoing connections. In order to motivate such a rule, consider the simple case when

138 8 Neural networks

$$g = \sum_{m=1}^{n_{\rm in}} w_m x_m,$$

where  $w_m \sim \mathcal{N}(0, \sigma^2)$  (one could also consider  $w_m \sim \mathcal{U}[-\sigma, \sigma]$ ). Note that  $\mathbb{E}(g) = 0$ and  $\operatorname{Var}(g) = \mathbb{E}(g^2) = n_{\mathrm{in}}\sigma^2\gamma^2$  when the inputs  $(x_m)_{1 \leq m \leq n_{\mathrm{in}}}$  are i.i.d. with  $\mathbb{E}(x_m^2) = \gamma^2$ . The variance of g is of order 1 for  $\sigma^2$  of order  $1/n_{\mathrm{in}}$ . The above computation was motivated by the computation of properties in the forward pass of the neural network, but a similar argument can be made for the backward pass, in which case  $n_{\mathrm{in}}$  should be replaced by  $n_{\mathrm{out}}$ . Refined heuristics can be worked out depending on the activation function at hand, see [22, Section 8.4]).

In terms of computational efficiency, it is often better to consider minibatches of sizes  $2^k$  for some integer  $k \ge 1$  due to the GPU memory formats. It is also good practice to shuffle the data before starting training in order to possibly destroy spurious correlations in the data set (which can arise for instance from the way data was collected). Note that second order methods (Newton or quasi-Newton techniques) are typically not used because of their computational cost; and also because there are generically many saddle points in high dimension, and not so many local minima in proportion, which is not good for Newton methods.

#### 8.2.3 Regularization

Recall that regularization allows to trade some bias in the predictions for a limitation of the variance in these predictions, by limiting the capacity of the model. Regularization is particularly important for neural networks, which can have millions of parameters. Let us first emphasize that early stopping, which is used throughout, already provides some form of regularization. In addition to this, it is customary to consider additional penalization terms in the empirical risk function similar to the ones considered in Sections 2.3 and 2.4, namely  $\|\theta\|_2$  or  $\|\theta\|_1$ . In fact, in practice, one usually considers only a penalization of the weights  $W_{\ell}$  and not of the biases  $b_{\ell}$ , as the latter parameters are merely used to recenter the outputs from one layer to the other. A  $\ell^2$  penalization of the weights corresponds to what is known as "weight decay" (the naming coming from the update of the parameter, which is of the form  $\theta^{n+1} = (1 - \lambda \gamma)\theta^n + \gamma \widehat{F}(\theta^n)$  for some stochastic estimator of the gradient of the unregularized loss, so that the first term decreases the magnitude of the coefficients of  $\theta$ ); while a  $\ell^1$  penalization encourages some form of sparsity.

One regularization specific to neural networks is dropout. This consists in randomly switching off some connections in the neural network at training time, in order to avoid the co-adaptation of units to the data at hand, and allow for a more robust learning. In practice this is implemented by adding a Bernoulli random variable at each neuron, and removing the corresponding neuron with the associated probability<sup>3</sup> (no incoming nor outgoing connection). At validation/test time, several options are possible, such as not turning off units or averaging the predictions over an ensemble of neural networks with randomly deactivated units.

A last comment is that it is currently believed that SGD-like algorithms have some regularizing effect, as they allow to converge to flat local minima which better generalize. This is known as "implicit regularization" and is an active field of research.

# 8.3 Unsupervised learning with neural networks

Autoencoders networks are an elegant tool for dimensionality reduction, denoising and generative modeling. Autoencoders have been considered early on in the neural network literature, where they were also called auto-associative neural networks [33]. The models considered in these early works correspond to what is currently known as bottleneck autoencoders, and were rather shallow.

 $<sup>^3</sup>$  It is not very clear what happens if all the neurons in a hidden layer are deactivated... This depends on the implementation of the model.
Deep autoencoders were used later on with the advent of modern computing architectures [25]. Bottleneck autoencoders were initially introduced to provide a nonlinear generalization of PCA, as it was shown that the linear neural networks obtained by minimizing the mean-square error were essentially equivalent to PCA [9, 5]. We refer for instance to [8, Section 12.4.2], [22, Chapter 14] and [40, Section 20.3] for textbook presentations of autoencoders, which include some historical perspectives, and mention many variations and extensions that were considered.

We first discuss the structure of autoencoders in Section 8.3.1, and then more precisely interpret their loss function in Section 8.3.2. The presentation of this section is taken from [35].

#### 8.3.1 Structure of autoencoders

Autoencoders fall into the class of unsupervised machine learning methods. For a given input data point  $x \in \mathcal{X} \subset \mathbb{R}^D$ , we denote by  $f_{\theta}(x)$  the prediction of the neural network. The parameters  $\theta \in \Theta$  are chosen to minimize the expected risk

$$\mathcal{R}(\theta) = \mathbb{E}[\ell(X, f_{\theta}(X))], \tag{8.4}$$

where  $\ell$  is a given elementary loss function, and the expectation is over the realizations of the input data X distributed according to some probability measure denoted by  $p_{\text{data}}$ . The typical choice for the latter elementary loss function is the square loss  $\ell(x, y) = ||x - y||^2$ , although other choices, such as the mean absolute loss  $\ell(x, y) = ||x - y||$  could also be considered in order to give less weight to outliers. In practice, the risk  $\mathcal{R}$  is replaced by the empirical risk over a training set of n given input data points  $\mathcal{D} = \{x_1, \ldots, x_n\}$ :

$$\widehat{\mathcal{R}}(\theta) = \frac{1}{n} \sum_{i=1}^{n} \ell(x_i, f_{\theta}(x_i)).$$

**Families of autoencoders.** There are various classes of autoencoders. It is useful to distinguish between undercomplete and overcomplete models. Undercomplete models have a limited capacity that prevents them from achieving zero training loss. The most prominent example is provided by bottleneck autoencoders for which

$$f_{\theta} = f_{\text{dec},\theta_2} \circ f_{\text{enc},\theta_1},\tag{8.5}$$

where the parameters  $\theta = (\theta_1, \theta_2)$  have been decomposed into parameters used in the encoder and decoder parts, respectively (see Figure 8.1 below). The limitation in the capacity of the autoencoder arises from the fact that the encoding function  $f_{\text{enc},\theta_1}$  has values in a latent space  $\mathcal{Z} \subset \mathbb{R}^D$  of dimension D strictly smaller than the dimension d of the input/output space  $\mathcal{X}$ , usually much smaller in fact. Overcomplete models can on the other hand achieve zero training error. These models are of course useless as such since they would simply copy the input without extracting the salient features explaining the data at hand. This is why some regularization process should be considered to limit the capacity of the neural network. Regularization is however also useful for undercomplete models. Standard examples of regularization mechanisms include:

- resorting to regularization strategies commonly used for neural networks in general, in particular early stopping, dropout, and standard weight decay to name a few options (see Section 8.2.3);
- sparse autoencoders where a penalization term is added to the loss function to prevent too many neurons to be active [41];
- denoising autoencoders [54], where outputs should be predicted from inputs corrupted by some noise, which forces networks to learn the structure of the distribution of the data [2];
- contractive autoencoders, where the Jacobian of the encoder is penalized in order to ensure smoother variations of this function, and limit its sensitivity to small variations in the input [45];
- variational autoencoders (VAEs) [31, 32, 21] can also be interpreted as some regularized version of the usual autoencoder framework.

#### 140 8 Neural networks

Other variations/extensions of autoencoders were also considered for more specific purposes, and are still studied, in particular for manifold learning, where the preservation of neighborhood relationships and/or geometric information in the dimensionality reduction process are important [28, 16, 34].



Fig. 8.1: Schematic representation of a symmetric autoencoder. Blue neurons correspond to hidden layers, while the input and output layers are respectively in green and red. Activation functions are hyperbolic tangents, except for the bottleneck and output layers, for which linear activation functions are considered in order not to restrict the range of values.

In some applications, for instance molecular dynamics, overcomplete models in particular are not directly interesting in terms of dimensionality reduction.

Autoencoders architectures. Autoencoders are often symmetric in their structures. In some cases, tied weights are being used, *i.e.* the weights  $\theta_2$  are the transpose of the weights  $\theta_1$  when writing the prediction function as (8.5). This choice reproduces at the level of autoencoders the symmetric structure of PCA, where the decoder matrix is the transpose of the encoder matrix (recall Lemma 5.1). In this case, a regularization on the encoder part only, as in contractive autoencoders, in fact also regularizes the decoder part. However, there is no particular motivation to use symmetric architectures, and various works, such as [50], studied the impact of an asymmetric architecture on the performance of the model. This point is further discussed below, where we motivate that decoders should be rather expressive in order to appropriately reproduce conditional expectations in the context of dimensionality reduction.

Another important consideration, more specific to bottleneck autoencoders, is the choice of the bottleneck dimension. In PCA, the number of dimensions to retain usually corresponds to some "knee" in the plot of eigenvalues of the empirical covariance matrix. Similar plots can be drawn for bottleneck autoencoders, for which the reconstruction error can be reported as a function of the bottleneck dimension (the remainder of the architecture being fixed). The optimal dimension should ideally coincide with the intrinsic dimension of the data set in manifold learning [14], which can be quantified using the Frechet inception distance (see for instance the method described in [27] in the context of VAEs).

#### 8.3.2 Interpretations of the loss function

We discuss here various reformulations and reinterpretations of the loss function (8.4) for bottleneck autoencoders (8.5) when the loss function is the square loss, and discuss in particular the relationship with principal curves/manifolds [23, 52].

Three viewpoints on the loss function. We consider an ideal setting where we minimize upon all measurable functions  $f_{enc} : \mathcal{X} \to \mathcal{Z}$  and  $f_{dec} : \mathcal{Z} \to \mathcal{X}$ . We denote by  $\mathcal{F}_{enc}$  and  $\mathcal{F}_{dec}$  the sets of measurable functions from  $\mathcal{X}$  to  $\mathcal{Z}$  and from  $\mathcal{Z}$  to  $\mathcal{X}$ , respectively; and by  $\mathcal{F}$  the set of measurable functions obtained by composing functions of  $\mathcal{F}_{enc}$  with functions of  $\mathcal{F}_{dec}$ :

$$\mathcal{F} = \{ f = f_{\text{dec}} \circ f_{\text{enc}}, f_{\text{enc}} \in \mathcal{F}_{\text{enc}}, f_{\text{dec}} \in \mathcal{F}_{\text{dec}} \}.$$

Minimizing the reconstruction error over the set of functions in  $\mathcal{F}$  can be then rewritten as

$$\inf_{f \in \mathcal{F}} \mathbb{E}\left[ \left\| X - f(X) \right\|^2 \right].$$
(8.6)

Note that we do not consider a regularization term here, so that overfitting may occur in practice when considering empirical risks (for instance, even with Z of dimension 1,  $f_{dec}$  can parametrize a space-filling curve).

As discussed in [18] (which complements [20] which was already hinting at autoencoders), the unsupervised least-square problem (8.6) can be thought of in various ways. In particular, there is some duality in the way the minimization over  $f \in \mathcal{F}$  is performed, as one can decide to either

- (i) simultaneously minimize over  $f_{enc}$  and  $f_{dec}$ , which is the standard way to proceed when training neural networks;
- (ii) minimize first over the encoder part, which allows to reformulate the minimization as the well-known problem of finding principal manifolds;
- (iii) minimize first over the decoder part, which is natural when thinking of the reconstruction error as some total variance to be decomposed using a conditioning on the values of the encoder.

The chosen numerical approach has a natural impact on the topology of the networks which are considered: in situation (i), encoders and decoders are treated on an equal footing, and it is therefore natural to consider them to be of a similar complexity; whereas options (ii) and (iii) suggest to consider asymmetric autoencoders. For instance, in option (iii), the minimization over the decoder part, which is performed first, could be done more carefully, with more expressive networks in order to better approximate the optimal decoder for a given encoder.

**Principal manifold reformulation.** We start by minimizing the reconstruction error (8.6) over the encoder part for a given decoder:

$$\inf_{f \in \mathcal{F}} \mathbb{E}\left[ \|X - f(X)\|^2 \right] = \inf_{f_{dec} \in \mathcal{F}_{dec}} \left\{ \inf_{f_{enc} \in \mathcal{F}_{enc}} \mathbb{E}\left[ \|X - f_{dec} \circ f_{enc}(X)\|^2 \right] \right\}$$
$$= \inf_{f_{dec} \in \mathcal{F}_{dec}} \mathbb{E}\left[ \|X - f_{dec} \circ h^{\star}_{f_{dec}}(X)\|^2 \right],$$
(8.7)

where the optimal encoder  $h_{f_{\text{dec}}}^{\star}: \mathcal{X} \to \mathcal{Z}$  for a given decoder  $f_{\text{dec}}: \mathcal{Z} \to \mathcal{X}$  is defined pointwise as

$$h_{f_{\text{dec}}}^{\star}(x) \in \operatorname*{argmin}_{z \in \mathcal{Z}} \|x - f_{\text{dec}}(z)\|_{z \in \mathcal{Z}}$$

provided that this minimization problem admits a solution. When  $f_{\text{dec}}$  is smooth and has an invertible Jacobian, the principal manifold is then the set  $f_{\text{dec}}(\mathcal{Z}) \subset \mathcal{X}$ . For any  $x \in \mathcal{X}$ ,  $h_{f_{\text{dec}}}^*(x) \in \mathcal{Z}$  gives the coordinates in the latent space of the projection of x on the principal manifold (the so-called projection index).

The reformulation of the minimization problem in terms of the decoder function leads to a minimization problem similar to the one encountered when searching for principal curves and manifolds. These mathematical concepts generalize in some sense PCA to curves and surfaces rather

#### 142 8 Neural networks

than lines and hyperplanes, as already discussed in the work introducing principal curves [23]. The minimization problems associated with finding principal manifolds are in general difficult to solve as these manifolds correspond to saddle-points of the loss functional [15]. In practice, this means that it is possible to move away from a critical point without increasing the test loss. This leads to overfitting issues and prevents the use of traditional cross-validation procedures to tune regularization hyperparameters.

**Reformulating autoencoders with conditional expectations.** We discuss here how to reformulate the training of autoencoders with conditional expectations, and provide alternative interpretations to the reconstruction error. In contrast to (8.7), we minimize here the reconstruction error (8.6) by first minimizing over the decoder part for a given encoding function, as already considered in [19]. This approach is natural in molecular dynamics. From a mathematical viewpoint, it corresponds to introducing conditional averages associated with fixed values of the encoder.

The loss function for unsupervised least-squares can be rewritten as

$$\inf_{f \in \mathcal{F}} \mathbb{E}\left[ \|X - f(X)\|^2 \right] = \inf_{f_{\text{enc}} \in \mathcal{F}_{\text{enc}}} \left\{ \inf_{f_{\text{dec}} \in \mathcal{F}_{\text{dec}}} \mathbb{E}\left[ \|X - f_{\text{dec}} \circ f_{\text{enc}}(X)\|^2 \right] \right\}$$
$$= \inf_{f_{\text{enc}} \in \mathcal{F}_{\text{enc}}} \mathbb{E}\left[ \|X - g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X)\|^2 \right],$$
(8.8)

where the ideal decoder  $g_{f_{enc}}^{\star}$  for a given encoder  $f_{enc}$  is the Bayes predictor associated with the least square regression problem (similarly to what is done in Exercise 1.6; see Exercise 8.4):

$$g_{f_{\text{enc}}}^{\star}(z) = \mathbb{E}[X \mid f_{\text{enc}}(X) = z].$$

$$(8.9)$$

Let us recall that, in all these expressions, expectations are taken with respect to the probability distribution  $p_{\text{data}}$  of the input data. Equations (8.8)-(8.9) show that the question of finding the best autoencoder can be reduced to finding the best encoding function, provided that one is able to compute good approximations of the conditional expectation.

**Exercise 8.4.** Prove the equality (8.9).

**Correction.** We start by introducing the quantity  $g_{f_{\text{enc}}}^* \circ f_{\text{enc}}(x)$  in the reconstruction error: for any  $f_{\text{dec}} \in \mathcal{F}_{\text{dec}}$  and  $f_{\text{enc}} \in \mathcal{F}_{\text{enc}}$ ,

$$\mathbb{E}\left[\left\|X - f_{\mathrm{dec}} \circ f_{\mathrm{enc}}(X)\right\|^{2}\right] = \mathbb{E}\left[\left\|\left[X - g_{f_{\mathrm{enc}}}^{\star} \circ f_{\mathrm{enc}}(X)\right] + \left[g_{f_{\mathrm{enc}}}^{\star} \circ f_{\mathrm{enc}}(X) - f_{\mathrm{dec}} \circ f_{\mathrm{enc}}(X)\right]\right\|^{2}\right] \\ = \mathbb{E}\left[\left\|X - g_{f_{\mathrm{enc}}}^{\star} \circ f_{\mathrm{enc}}(X)\right\|^{2}\right] + \mathbb{E}\left[\left\|g_{f_{\mathrm{enc}}}^{\star} \circ f_{\mathrm{enc}}(X) - f_{\mathrm{dec}} \circ f_{\mathrm{enc}}(X)\right\|^{2}\right],$$
(8.10)

where we used the following identity, obtained by conditioning on the values of the random variable  $Z = f_{enc}(X)$ :

$$\mathbb{E}\left[\left(X - g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X)\right) \cdot \left(g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X) - f_{\text{dec}} \circ f_{\text{enc}}(X)\right)\right] \\ = \mathbb{E}\left[\left(\mathbb{E}\left[X|Z\right] - g_{f_{\text{enc}}}^{\star}(Z)\right) \cdot \left(g_{f_{\text{enc}}}^{\star}(Z) - f_{\text{dec}}(Z)\right)\right] = 0,$$

in view of the definition (8.9) of  $g_{f_{enc}}^{\star}$ . It is clear from (8.10) that the decoding function which minimizes the reconstruction error for a given encoder  $f_{enc} \in \mathcal{F}_{enc}$  is indeed  $g_{f_{enc}}^{\star}$ , as defined by (8.9).

The reconstruction error (8.8) can be reinterpreted in terms of variances. Indeed, on the one hand,

$$\mathbb{E}\left[\left\|X - g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X)\right\|^{2}\right] = \mathbb{E}\left[\left\|X - \mathbb{E}\left[X|f_{\text{enc}}(X)\right]\right\|^{2}\right]$$
$$= \mathbb{E}\left[\mathbb{E}\left(\left\|X - \mathbb{E}\left[X|f_{\text{enc}}(X)\right]\right\|^{2}\right|f_{\text{enc}}(X)\right)\right]$$
$$= \mathbb{E}\left[\operatorname{Var}(X|f_{\text{enc}}(X))\right].$$
(8.11)

On the other hand,

$$\mathbb{E}\left[\left\|X - g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X)\right\|^{2}\right] = \mathbb{E}\left[\left\|X - \mathbb{E}\left[X|f_{\text{enc}}(X)\right]\right\|^{2}\right]$$
$$= \mathbb{E}\left(\left\|X\right\|^{2}\right) - \mathbb{E}\left(\mathbb{E}\left[X|f_{\text{enc}}(X)\right]^{2}\right)$$
$$= \operatorname{Var}(X) - \operatorname{Var}\left[\mathbb{E}(X|f_{\text{enc}}(X))\right].$$
(8.12)

These two equalities yield the well-known formula for the total variance decomposition, namely  $\operatorname{Var}(X) = \mathbb{E}\left[\operatorname{Var}(X|f_{\operatorname{enc}}(X))\right] + \operatorname{Var}\left[\mathbb{E}(X|f_{\operatorname{enc}}(X))\right].$ 

A consequence of (8.12) is that the minimization problem (8.8) can be reformulated as the following equivalent maximization problem:

$$\sup_{f_{\rm enc} \in \mathcal{F}_{\rm enc}} \operatorname{Var}\left[\mathbb{E}(X|f_{\rm enc}(X))\right].$$
(8.13)

In words, this reformulation translates the equivalence between (the "classes" referring here to the level sets of  $f_{enc}$ )

- minimizing the intraclass dispersion (8.11): the distribution of configurations  $x \in \mathcal{X}$  for a fixed value z of  $f_{\text{enc}}$  should concentrate around the mean value  $g^{\star}_{f_{\text{enc}}}(z)$  by having a variance as small as possible;
- maximizing the interclass dispersion (8.13): the values of the conditional averages of X for fixed values of  $f_{enc}$  should be as spread out as possible over the range of  $f_{enc}$ .

Let us conclude for providing a formal characterization of the optimal encoding function. A key equality to establish (8.12), namely

$$\mathbb{E}\left[\left\|X - g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X)\right\|^{2}\right] = \mathbb{E}\left[\left\|X\right\|^{2}\right] - \mathbb{E}\left[\left\|g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X)\right\|^{2}\right],$$

shows that the minimization of the reconstruction error is equivalent to the maximization of the second moment of the conditional expectation:

$$\sup_{f_{\text{enc}} \in \mathcal{F}_{\text{enc}}} \mathbb{E}\left[ \left\| g_{f_{\text{enc}}}^{\star} \circ f_{\text{enc}}(X) \right\|^2 \right].$$
(8.14)

This alternative viewpoint allows to characterize the optimal encoding function  $f_{enc}$  by some orthogonality condition similar to the self-consistency condition of principal curves, see [20, Section 2]. In fact, it can be formally shown that critical points of (8.14) satisfy

$$\forall j \in \{1, \dots, d\}, \quad \forall x \in \operatorname{Supp}(p_{\text{data}}), \qquad \left[x - g_{f_{\text{enc}}}^{\star}(f_{\text{enc}}(x))\right]^{\top} \partial_{z_j} g_{f_{\text{enc}}}^{\star}(f_{\text{enc}}(x)) = 0, \quad (8.15)$$

where  $\text{Supp}(p_{\text{data}})$  is the support of the probability measure  $p_{\text{data}}$ . The derivation of this condition, which can be read in [35, Appendix A], can be seen as a variation of derivations of optimality conditions for principal curves, as written already in [23]; see also [20] where (8.15) is used to construct a new objective function to minimize in order to find  $f_{\text{enc}}$ .

An interesting implication of (8.15) is that the intersection of  $\text{Supp}(p_{\text{data}})$  and the submanifold

$$\Sigma_z = f_{\text{enc}}^{-1}\{z\} = \{x \in \mathcal{X} \mid f_{\text{enc}}(x) = z\}$$

is in fact included in the (d-D)-dimensional hyperplane containing the point  $g_{f_{enc}}^{\star}(z)$  and orthogonal to the vectors  $\partial_{z_1}g_{f_{enc}}^{\star}(z), \ldots, \partial_{z_d}g_{f_{enc}}^{\star}(z)$  (recalling that  $\mathcal{X}$  and  $\mathcal{Z}$  have dimensions d and D, respectively). As these hyperplanes generally have a non-empty intersection, finding a regular function  $f_{enc}$  which satisfies (8.15) is only possible for distributions  $p_{data}$  which have a support sufficiently concentrated around the principal manifold. The issue of having hyperplanes intersecting can be seen as the counterpart in the context we consider here of the concept of ambiguity points for principal curves [23].

# 8.4 Other types of neural networks

We (very) briefly mention in this section neural network architectures which go beyond the simple feedforward multilayer perceptrons. These modern architectures, which can be rather complicated, are key to the current successes of machine learning, for classification and regression, and also more and more for generative tasks such as those related to natural language processing.

**Convolution neural networks.** Convolutional neural networks (CNNs) are a default choice for image processing. Their motivation is threefold:

- they allow to treat inputs of variable sizes (depending on the resolution of the image under consideration for instance);
- they can handle inputs of large dimension, coming in the form width × height × input channel. For images, the input channel is typically RGB, *i.e.* the input for color images is composed of three matrices of pixels intensities, for the colors red, green and blue;
- they naturally encode some translation invariance and various symmetries. This is important for image classification for instance, as the main element in the image can be rotated, translated, etc.

From a structural viewpoint, matrix multiplications in CNN are replaced by convolution operations, which work with filters that take local averages of pictures. In practice, CNNs alternate between convolutional layers and pooling layers (see for instance [40, Figure 14.13]). Pooling layers take the minimum or maximum of an input over certain regions; this allows to reduce the sensitivity to the location of elements in the image.

Famous examples of CNNs include (see [40, Section 14.3]): LeNet, AlexNet (2012; this is the network which reintroduced neural networks in machine learning, by its dramatic improvement in classificiation performance), GoogLeNet (2015), ResNet (2015, Microsoft), DenseNet, ... All these networks are constructed on the similar building blocks, which are rearranged in various manners.

They can be used for other tasks than classification, such as image tagging, object detection, segmentation, ...

**Graph neural networks.** These networks use connections between nodes which are not organized in a strictly sequential manner, but rather on a graph. This can be useful for various tasks, for instance the prediction of atomic forces in molecular models; in this situation, the nodes of the graph correspond to an atom, and the connections between the nodes, which provide some transfer of information, mimick the interactions between atoms.

**Recurrent neural networks.** These networks are used to generate sequences or trajectories  $y_{1:T}$ , for instance for language modeling. They include long short term memory (LSTM) blocks to have better gradients for training and ensure that the memory in sequences is better taken into account.

Attention networks. These networks use a dictionary lookup where a query/input is compared to various keys, and the output depends on the key.

**Transformers.** Many successful current large language models, such as ChatGPT, are based on networks of GPT type ("generative pre-trained transformers"). They include attention layers – self-attention in fact. They also make use of positional encoding. Finally, they are made to scale efficiently with the data size.

Kolmogorov-Arnold networks.

Complete the paragraphs below, and also possibly the description of CNNs

# Clustering methods

9.1	Aims and scope of clustering 145
9.2	K-means and center-based clustering methods 146
9.3	Hierarchical clustering 149
9.4	Clustering with mixture models 150
9.5	Density based clustering 155
9.6	Spectral clustering 155

We discuss in this chapter a set of unsupervised techniques aiming at grouping together data points which are similar in a common class. We discuss more precisely the aim of clustering in Section 9.1, and then successively present various techniques, starting in Section 9.2 with K-means clustering, turning next in Section 9.3 to hierarchical clustering, then in Section 9.4 to clustering with mixture models, and concluding with density based clustering in Section 9.5 and spectral clustering in Section 9.6. Our presentation is based on [40, Chapter 20], [49, Chapter 22 and Section 24.4], [24, Sections 14.3 and 6.8], as well as [38, Section 13].

# 9.1 Aims and scope of clustering

We consider a set  $\mathcal{D}$  of n data points  $\{x_1, \ldots, x_n\} \subset \mathcal{X}$ , which are unlabelled. The informal aim of clustering is to group data points into clusters – similar elements should be in the same class, while dissimilar elements should be in different classes. In fact, one should talk about segmentation rather than clustering when the data is not well separated. Applications include grouping families of genes in biology, customers depending on their behavior for marketing, communities in social media, atomic configurations in molecular dynamics, etc.

There are two major difficulties with clustering:

- (1) The precise definition of clustering (*i.e.* a quantitative definition) is rather difficult to make rigorous as the notion of similarity is not a transitive relation, while cluster assignment is. To illustrate the issue at stake, think of two parallel lines of data points regularly spaced, the separation between the two lines being quite larger than the separation between the points in each line. If one emphasizes grouping close-by points, then each line would be a cluster, each data point being close to its neighbors, but the extremal points on each line can be quite different;
- (2) Another difficulty is the lack of ground truth, which is a usual problem in unsupervised learning. There may be many ways to cluster a given data set. In particular, the distance or metric used for clustering can have a dramatic impact on the results. Think for example of grouping movies by main actor, topic, rating, year, etc. The quality of the clustering can be assessed if labelled data is available, by evaluating the purity of each cluster; or if some reference clustering is available (through the Rand index computed from the number of true/false positives/negatives; see [40, Section 21.1.1.2]).

#### 146 9 Clustering methods

On the practical side, let us mention two points:

- (i) normalizing the data set is not a good idea in general as the notion of closedness is not invariant by rescaling. Normalization could however still be relevant for some problems. Whether to normalize or not is problem specific, requires some domain knowledge, and possibly some trial and error procedure;
- (ii) clustering in high dimension can be challenging as data points are generically all far away from each other, see Exercise ?? below. It makes sense to perform some dimensionality reduction as a pre-processing step, using for instance PCA (see Chapter 5) or autoencoders (see Section 8.3).

**Exercise 9.1.** Consider a data composed of two clusters of points in  $\mathbb{R}^d$ , independently drawn from Gaussian distributions with means 0 and  $\mu$  respectively, and identity covariance in both cases. Compute the average distance between points within one of the clusters, and the average distance between two points in different clusters. How easy is it to distinguish these quantities as d increases?

**Correction.** Consider  $x_1, x_2 \sim \mathcal{N}(0, \mathrm{Id}_d)$  and  $x_3 \sim \mathcal{N}(\mu, \mathrm{Id}_d)$ , which are all independent. Then, the average distance between points within one of the clusters is

$$\mathbb{E}\left[\|x_1 - x_2\|_2^2\right] = \mathbb{E}\left[\|x_1\|_2^2\right] + \mathbb{E}\left[\|x_2\|_2^2\right] = 2d,$$

while the average distance between two points in different clusters is

$$\mathbb{E}\left[\|x_1 - x_3\|_2^2\right] = \mathbb{E}\left[\|x_1\|_2^2\right] + \mathbb{E}\left[\|x_3\|_2^2\right] = 2d + \mu^2,$$

since  $\mathbb{E}\left[\|x_3\|_2^2\right] = \mu^2 + \operatorname{Var}(x_3)$ . As *d* increases, the shift between the two clusters becomes more and more difficult to detect. For high dimensional systems, the average distance between points within one of the clusters is very close to the the average distance between two points in different clusters, and in fact both are large.

### 9.2 K-means and center-based clustering methods

The aim of K-means clustering is to group the data points into K clusters. Using a combinatorial algorithm to this end is impossible because the number of possible partitions into K classes is too large (see [24, Equation (14.30)] for the formula giving the number of classes). We consider here a method to perform an approximate clustering based on some cost function, with an optimization procedure to find the assignment function  $\mathcal{X} \to \{1, \ldots, K\}$ . The cost function is itself based on a distance function  $d: \mathcal{X} \times \mathcal{X} \to \mathbb{R}_+$ , for instance the Euclidean distance.

**Loss function.** Let us write the loss function for the K-means algorithm when  $\mathcal{X} = \mathbb{R}^d$ , the chosen distance is the Euclidean norm, *i.e.*  $\mathbf{d} = \|\cdot\|_2$ , and the number of classes  $K \ge 2$  is fixed/predetermined. The algorithm updates a partition  $C_1, \ldots, C_K$  of  $\mathcal{D}$ , and considers the loss function

$$J(C_1, \cdots, C_K) = \min_{(m_1, \dots, m_K) \in (\mathbb{R}^d)^K} \left\{ \sum_{k=1}^K \sum_{x \in C_k} \|x - m_k\|_2^2 \right\} = \sum_{k=1}^K \sum_{x \in C_k} \|x - \overline{x}_k\|_2^2, \quad (9.1)$$

where

$$\overline{x}_k = \frac{1}{|C_k|} \sum_{x \in C_k} x$$

is the empirical mean of the data points in the cluster  $C_k$ .

**Exercise 9.2.** Prove that the second equality in (9.1) holds.

Correction. Note that

$$\min_{(m_1,\dots,m_K)\in(\mathbb{R}^d)^K}\sum_{k=1}^K\sum_{x\in C_k}\|x-m_k\|_2^2 = \sum_{k=1}^K \left(\min_{m_k\in\mathbb{R}^d}\sum_{x\in C_k}\|x-m_k\|_2^2\right).$$

The minimization problem on the right hand side is classical and has been solved various times in these lecture notes, the unique minimizer being  $m_k^* = \overline{x}_k$  (see for instance Exercise 7.1).

#### Remark 9.1 (Formulation using an assignment function/matrix). Note that

$$\sum_{k=1}^{K} \sum_{x \in C_k} \|x - m_k\|_2^2 = \sum_{i=1}^{n} \|x_i - m_{z(i)}\|_2^2 = \|X - ZM^{\top}\|^2,$$

where  $X \in \mathbb{R}^{n \times d}$  is the matrix whose *i*-th line is  $x_i$ , the matrix  $M = [m_1|m_2| \dots |m_K] \in \mathbb{R}^{d \times K}$ has the vectors  $m_k$  as columns, and  $Z \in \{0,1\}^{n \times K}$  is the assignment matrix, with entries  $Z_{i,k} =$  $\mathbf{1}_{z(i)=k}$  for

$$z(i) \in \operatorname*{argmin}_{1 \le k \le K} \|x_i - m_k\|_2$$

the assignement function.

K-means algorithm. The precise algorithm is the following.<sup>1</sup> Starting from some initial partition  $(C_1^{(0)}, \cdots, C_K^{(0)})$ , and given a partition  $(C_1^{(t)}, \cdots, C_K^{(t)})$  at step  $t \ge 0$ :

- compute the centers \$\overline{x}\_k^{(t)}\$ of the cluster \$C\_k^{(t)}\$;
  for every data point \$x\_i\$ with \$1 ≤ i ≤ n\$, find the index \$z\_{t+1}(i) ∈ {1,...,K}\$ of the centers \$z\_{t+1}(i) ∈ {1 ter  $(\overline{x}_k^{(t)})_{1 \leq k \leq K}$  closest to  $x_i$  (with some rule to break ties):

$$z_{t+1}(i) \in \operatorname*{argmin}_{1 \leq k \leq K} \left\| x_i - \overline{x}_k^{(t)} \right\|_2;$$

• update the clusters as  $C_k^{(t+1)} = \{x_j \in \mathcal{D} : z_{t+1}(j) = k\}.$ 

The algorithm is iterated until convergence  $(i.e. (C_1^{(t+1)}, \dots, C_K^{(t+1)}) = (C_1^{(t)}, \dots, C_K^{(t)}))$ . Note that the computational cost of a single iteration is O(n). The assignment step is based on a Voronoi tesselation of the space based on the cluster centers.

It is good practice to run several times the algorithm with random initializations and retain the clustering leading to the smallest value of the objective function (9.1). A typical way to start with a random partition is to draw at random K cluster centers, either fully at random (all points are uniformly sampled over the full space) or by emphasizing more uniform clusters in order to better cover the data. This can be achieved by choosing the random centers in a sequential manner, and emphasizing points which are further apart from the current centers (the so-called K-means++ method, which is the default initialization method in scikit-learn for instance); see [40, Section 21.3.4].

See the website http://www.bytemuse.com/post/k-means-clustering-visualization/ for a nice graphical illustration.

148 9 Clustering methods

**Convergence of the** *K*-means algorithm. We present a result on the evolution of the loss function (9.1), which however does not give insight on the number of iterations of the algorithm. This also does not provide indication on the quality of the so-obtained clustering – for instance because the minimization problem is non convex, and there is no guarantee that the minimal loss is obtained.

**Lemma 9.1.** Each iteration of the K-means algorithm does not increase the objective function (9.1):

$$\forall t \ge 0, \qquad J\left(C_1^{(t+1)}, \cdots, C_K^{(t+1)}\right) \le J\left(C_1^{(t)}, \cdots, C_K^{(t)}\right). \tag{9.2}$$

*Proof.* By definition of  $z_{t+1}(i)$ , it holds

$$\left\|x_i - \overline{x}_k^{(t)}\right\|_2 \ge \left\|x_i - \overline{x}_{z_{t+1}(i)}^{(t)}\right\|_2$$

Therefore, making first use of the reassignement of  $x_i$  (previous inequality) and then of the update of the clusters,

$$J\left(C_{1}^{(t)}, \cdots, C_{K}^{(t)}\right) \ge \sum_{k=1}^{K} \sum_{x_{i} \in C_{k}^{(t)}} \left\|x_{i} - \overline{x}_{z_{t+1}(i)}^{(t)}\right\|_{2}^{2} = \sum_{k=1}^{K} \sum_{x_{i} \in C_{k}^{(t+1)}} \left\|x_{i} - \overline{x}_{k}^{(t)}\right\|_{2}^{2}.$$

Now,

$$\sum_{x_i \in C_k^{(t+1)}} \left\| x_i - \overline{x}_k^{(t)} \right\|_2^2 = \sum_{x_i \in C_k^{(t+1)}} \left( \left\| x_i - \overline{x}_k^{(t+1)} \right\|_2^2 + \left\| \overline{x}_k^{(t+1)} - \overline{x}_k^{(t)} \right\|_2^2 \right) \ge \sum_{x_i \in C_k^{(t+1)}} \left\| x_i - \overline{x}_k^{(t+1)} \right\|_2^2.$$

Therefore,

$$J\left(C_{1}^{(t)}, \cdots, C_{K}^{(t)}\right) \geqslant \sum_{k=1}^{K} \sum_{x_{i} \in C_{k}^{(t+1)}} \left\|x_{i} - \overline{x}_{k}^{(t+1)}\right\|_{2}^{2} = J\left(C_{1}^{(t+1)}, \cdots, C_{K}^{(t+1)}\right),$$

which allows to conclude.

The proof of Lemma 9.1 shows that if equality holds in (9.2), then the partitions  $(C_1^{(t+1)}, \cdots, C_K^{(t+1)})$ and  $(C_1^{(t)}, \cdots, C_K^{(t)})$  coincide<sup>2</sup> as the centers of the clusters are unchanged.

Variations of the K-means algorithm. We present two variations of the K-means algorithm:

• *K*-medoids, for which the empirical average in a cluster is replaced by some element of the dataset. This allows to apply *K*-means type algorithms to data which is not in  $\mathbb{R}^d$ , for which averages may be undefined (think of categorical variables, or matrices with positivity or norm constraints). The objective function to be considered instead of (9.1) reads

$$\min_{(m_1,\ldots,m_K)\in\mathcal{D}}\sum_{k=1}^K\sum_{x\in C_k}\mathsf{d}(x,m_k)^2.$$

There are various approaches to try and solve this problem. A common one is "partitioning around medoids" [29], but it is computationally rather expensive. A more recent and less expensive option is the Voronoi iteration of [44], which has a cost O(nK) per iteration and works by reassigning points and updating the medoids, very similarly to what is done for *K*-means.

 $<sup>^{2}</sup>$  Except maybe for points equidistant from two centers or more.

• K-median builds upon K-medoids but relies on a loss function which uses a distance rather than a squared distance as in (9.1), namely

$$\min_{(m_1,\ldots,m_K)\in\mathcal{D}}\sum_{k=1}^K\sum_{x\in C_k}\mathsf{d}(x,m_k).$$

Considering unsquared norms may be better for problems such as facility location (think of houses as data points and look for places where to position K firestations).

• *K*-autoencoders, introduced in [43], consider the loss function

$$\sum_{i=1}^{n} \min_{1 \le k \le K} \|x_i - \Phi_{\theta_k}(x_i)\|^2,$$

which is minimized with respect to the parameters  $\theta_1, \ldots, \theta_K$  for each autoencoder. This can therefore be seen as some extension of K-means where  $\Phi_{\theta_k}(x_i)$  replaces the cluster empirical average. The cluster assignment is performed based on the index of the autoencoder for which the reconstruction loss is minimized, namely

$$C_{k} = \left\{ x_{i} \in \mathcal{X} \left| k = \underset{1 \leq k' \leq K}{\operatorname{argmin}} \left\| x_{i} - \Phi_{\theta_{k'}}(x_{i}) \right\|^{2} \right\}.$$

**Choosing the number of clusters** K**.** There is a monotonic decrease of the minimal value for the objective function (9.1) when K is increased, so that (as in PCA) it is not possible to use some form of cross-validation to identify the best value of K. There are two main ways to find an appropriate value for K:

- rely on probabilistic models (such as the Gaussian mixture models consider in Section 9.4) and use the tools of Bayesian model selection to determine K;
- look for a kink/elbow in the within cluster distances (see [24, Section 14.3.11]) or in silhouette coefficients (see [40, Section 21.3.7.3]).

## 9.3 Hierarchical clustering

There are two types of hierarchical clustering:

- *agglomerative*, in which case one starts from n clusters composed of single data points and merges them;
- *divisive*, in which case one starts from a single cluster with all data points, and splits it. This approach is not as well studied and less common (see [24, Section 14.3.12]).

A nice property of hierarchical clustering is its nested property, and the fact that one does not need to specify the desired number of clusters. Partitions of the data are obtained by transforming dendrograms representing the hierarchical structure using two possible termination criteria: cutting when a fixed number of clusters is attained, or when a certain upper bound in distance is reached.

**Hierarchical agglomerative clustering.** In order to make these statements more precise, we present agglomerative clustering in more detail. The key element to specify the method is a distance D between clusters, which allows to determine the closest clusters, which should be merged. The computational cost of the algorithm is a priori  $O(n^3)$  (as picking the two most similar clusters has a cost  $O(n^2)$ , and this operation should be repeated O(n) times). The precise algorithm is the following. Starting from the *n* clusters  $C_i^{(0)} = \{x_i\}$ , iterate on  $0 \le t \le n-2$ :

• select among the remaining n - t clusters the pair  $(C_k^{(t)}, C_\ell^{(t)})$  which minimizes the distance between clusters:

$$(k,\ell) \in \mathop{\mathrm{argmin}}_{(a,b):1 \leqslant a < b \leqslant n-t} \mathsf{D}\left(C_a^{(t)}, C_b^{(t)}\right);$$

• form a new partition  $C_1^{(t+1)}, \ldots, C_{n-t-1}^{(t+1)}$  by keeping the previous clusters  $C_a^{(t)}$  for  $a \notin \{1, \ldots, n-t\} \setminus \{k, \ell\}$ , and adding the merged cluster  $C_k^{(t)} \cup C_\ell^{(t)}$ .

#### 150 9 Clustering methods

**Choice of the distance.** The key element in this algorithm is the choice of the distance D. The result crucially depends on it. We present three choices here:

• single link is based on the distance between the two closest elements of  $C_k$  and  $C_{\ell}$ :

$$\mathsf{D}(C_k, C_\ell) = \min_{\substack{x_i \in C_k \\ x_j \in C_\ell}} \mathsf{d}(x_i, x_j).$$

A possible issue with this choice is chaining, which corresponds to grouping together data points related by a series of close intermediate points, the overall cluster being however not very compact (somehow elongated, hence the name);

• complete link is based on the distance between the two farthest elements of  $C_k$  and  $C_\ell$ :

$$\mathsf{D}(C_k, C_\ell) = \max_{\substack{x_i \in C_k \\ x_j \in C_\ell}} \mathsf{d}(x_i, x_j).$$

This metric emphasizes compact clusters but may lead to "closedness losses" (two elements of the same cluster can be separated because they are somewhat far away from each other);

• average link considers the average distance between the elements of  $C_k$  and  $C_{\ell}$ :

$$\mathsf{D}(C_k, C_\ell) = \frac{1}{|C_k| |C_\ell|} \sum_{x_i \in C_k} \sum_{x_j \in C_\ell} \mathsf{d}(x_i, x_j).$$

It somewhat interpolates between the single and complete link, and hence allows to form relatively compact clusters with elements that are relatively close. It is however sensitive to monotonic transformations of the baseline distance d.

There are many other possible choices, in particular Ward's distance. The main issue is the algorithmic complexity of the resulting algorithm. Hierarchical clustering is suitable for small data sets, or should be initialized with K-means with a large value of K, however sufficiently small compared to the number of data points (*i.e.*  $1 \ll K \ll n$ ).

Hierarchical clustering can be nicely visualized using dendrograms (see for instance the presentation in [24, Section 14.3.12]). In this representation, the data points are placed on the horizontal axis, at zero height. When two points are merged to form a cluster, say  $x_1$  and  $x_2$ , then one draws a vertical line starting from  $x_1$  and  $x_2$ , until the height corresponding to the distance  $D({x_1}, {x_2})$ where a horizontal line is drawn. This is representation, the height of each node in the resulting binary tree is proportional to the distance between the two daughter clusters. The tree can indeed be plotted as such due to the monotone increasing character of the distances between clusters.

#### 9.4 Clustering with mixture models

Mixture models make use of a latent/hidden variable in the context of a generative model. Here, this latent variable is the cluster identity of a data point, the overall distribution of data points being a mixture of Gaussian distributions with different means and variances. This requires making an assumption on the form of the data distribution, in accordance with the "no free lunch theorem" (see Section 10.3).

**Gaussian mixture model.** The algorithm consists in determining the most likely value of the latent variable  $z_i$  for a given data point  $x_i$ . In fact, one determines the vector of discrete probabilities to belong to a cluster. The precise model for the density of the data distribution is

$$p_{\theta}(x) = \sum_{k=1}^{K} \pi_k \mathscr{G}_{\mu_k, \varSigma_k}(x), \qquad \mathscr{G}_{\mu, \varSigma}(x) = (2\pi)^{-d/2} \det(\varSigma)^{-1/2} \exp\left(-\frac{1}{2}(x-\mu)^{\top} \varSigma^{-1}(x-\mu)\right),$$

with parameters

$$\theta = (\pi_1, \ldots, \pi_K, \mu_1, \ldots, \mu_K, \Sigma_1, \ldots, \Sigma_K) \in \Theta,$$

where

$$\Theta = \left\{ \theta \in [0,1]^K \times \left( \mathbb{R}^d \right)^K \times \left( \mathbb{R}^{d \times d} \right)^K \ \left| \ \sum_{k=1}^K \pi_k = 1, \ \Sigma_k = \Sigma_k^\top > 0 \right\} \right\}$$

This models corresponds to having a prior distribution  $(\pi_1, \ldots, \pi_K)$  on the latent variable  $z_i$ , and a conditional distribution  $p_{\theta}(x_i | z_i = k) = \mathscr{G}_{\mu_k, \Sigma_k}(x_i)$ . The likelihood of a data point x is obtained by marginalizing out the latent variable, *i.e.* 

$$p_{\theta}(x) = \sum_{k=1}^{K} p_{\theta}(z_i = k) p_{\theta}(x_i \mid z_i = k).$$

Bayes' rule is used to obtain the responsibility for each data point, *i.e.* the posterior membership probability:

$$r_{i,k} := p_{\theta}(z_i = k \mid x_i) = \frac{p_{\theta}(z_i = k)p_{\theta}(x_i \mid z_i = k)}{\sum_{k'=1}^{K} p_{\theta}(z_i = k')p_{\theta}(x_i \mid z_i = k')} = \frac{\pi_k \mathscr{G}_{\mu_k, \Sigma_k}(x_i)}{\sum_{k'=1}^{K} \pi_{k'} \mathscr{G}_{\mu_{k'}, \Sigma_{k'}}(x_i)}.$$

The latter formula can be seen as an application of the general Bayes formula to the case of Gaussian mixture models (as the previous equation can be rewritten as  $p_{\theta}(z_i = k | x_i) = p_{\theta}(z_i = k)p_{\theta}(x_i | z_i = k)/p_{\theta}(x_i)$ ). Assignment can be made in various ways, for instance through the "hard assignment" rule

$$z_i \in \operatorname*{argmax}_{1 \leqslant k \leqslant K} r_{i,k}.$$

Estimation of the parameters of the model. The difficulty in this approach is that the vector of parameter  $\theta$  needs to be estimated. This can be done using some maximum likelihood method, in which case one obtains

$$\widehat{\theta} \in \operatorname*{argmax}_{\theta \in \Theta} \left\{ \log p_{\theta}(x_1, \dots, x_n) \right\} = \operatorname*{argmax}_{\theta \in \Theta} \left\{ \sum_{i=1}^n \log \left( \sum_{k=1}^K \pi_k \mathscr{G}_{\mu_k, \Sigma_k}(x_i) \right) \right\}.$$
(9.3)

Gradient-like methods could be used to find (local) maxima. A more standard strategy in this context is however to rely on Expectation-Maximization (EM) algorithms, which ensure that the likelihood is non decreasing.

The precise EM algorithm is the following. Starting from an initial guess  $\theta^0$ , iterate on  $t \ge 0$ :

• (*expectation step*) update the responsibilities as

$$r_{i,k}^{(t)} = \frac{\pi_k^{(t)} \mathscr{G}_{\mu_k^{(t)}, \Sigma_k^{(t)}}(x_i)}{\sum_{k'=1}^K \pi_{k'}^{(t)} \mathscr{G}_{\mu_{k'}^{(t)}, \Sigma_{k'}^{(t)}}(x_i)};$$

• (maximization step) update  $\theta^{t+1}$  as

$$\pi_{k}^{(t+1)} = \frac{1}{n} \sum_{i=1}^{n} r_{i,k}^{(t)} \in [0,1], \qquad \mu_{k}^{(t+1)} = \frac{\sum_{i=1}^{n} r_{i,k}^{(t)} x_{i}}{\sum_{i=1}^{n} r_{i,k}^{(t)}} \in \mathbb{R}^{d},$$

$$\Sigma_{k}^{(t+1)} = \frac{\sum_{i=1}^{n} r_{i,k}^{(t)} \left(x_{i} - \mu_{k}^{(t+1)}\right) \left(x_{i} - \mu_{k}^{(t+1)}\right)^{\top}}{\sum_{i=1}^{n} r_{i,k}^{(t)}} \in \mathbb{R}^{d \times d}.$$
(9.4)

#### 1529 Clustering methods

In the previous expressions, the data points  $x_i$  are considered as column vectors  $\mathbb{R}^{d \times 1}$ . This update corresponds to recomputing the first moments of the clusters based on the weights given by the updated responsibilities.

Motivation for the update rules in the EM algorithm. We first claim that the update rules (9.4) correspond to the maximization problem

$$\max_{\theta \in \Theta} \left\{ \sum_{i=1}^{n} \sum_{k=1}^{K} r_{i,k}^{(t)} \left( \log \pi_{k} + \log \mathscr{G}_{\mu_{k}, \Sigma_{k}}(x_{i}) \right) \right\}$$

$$= \max_{\theta \in \Theta} \left\{ \sum_{i=1}^{n} \sum_{k=1}^{K} r_{i,k}^{(t)} \left( \log \pi_{k} - \frac{1}{2} (x_{i} - \mu_{k})^{\top} \Sigma_{k}^{-1} (x_{i} - \mu_{k}) - \frac{1}{2} \log \det(\Sigma_{k}) \right) \right\};$$
(9.5)

and next motivate why this maximization problem is considered instead of (9.3). We assume that the maximization problem is well posed, with a maximizer  $\theta^{t+1} \in \Theta$ .

**Exercise 9.3.** Prove that  $\theta^{t+1}$  defined in (9.4) is the unique critical point of (9.5).

Correction. Define

$$F_t(\theta) = \sum_{i=1}^n \sum_{k=1}^K r_{i,k}^{(t)} \left( \log \pi_k - \frac{1}{2} (x_i - \mu_k) \Sigma_k^{-1} (x_i - \mu_k) - \frac{1}{2} \log \det(\Sigma_k) \right),$$

which should be maximized under the constraint

$$\Phi(\theta) = \sum_{k=1}^{K} \pi_k - 1 = 0$$

We do not consider the positivity constraints for  $\pi_k$  and  $\Sigma_k$ , but check a posteriori that these are satisfied.

A necessary condition of optimality is that  $\partial_{\pi_k} F_t(\theta^{t+1}) + \lambda \partial_{\pi_k} \Phi(\theta^{t+1}) = 0$  for some  $\lambda \in \mathbb{R}$ , which leads to

$$\sum_{i=1}^{n} \frac{r_{i,k}^{(t)}}{\pi_k^{(t+1)}} + \lambda = 0.$$

Necessarily,  $\lambda < 0$ . Therefore,

$$\pi_k^{(t+1)} = -\frac{1}{\lambda} \sum_{i=1}^n r_{i,k}^{(t)}.$$

The value of  $\lambda$  is adjusted so that  $\Phi(\theta^{(t+1)}) = 0$ , which leads to the claimed expression for  $\pi_k^{(t+1)}$ . Let us next consider the optimality condition involving  $\mu_k$ , which reads  $\nabla_{\mu_k} F_t(\theta^{t+1}) = 0$ . This leads to

$$\sum_{i=1}^{n} r_{i,k}^{(t)} \left( \Sigma_k^{(t+1)} \right)^{-1} \left( x_i - \mu_k^{(t+1)} \right) = 0 = \left( \Sigma_k^{(t+1)} \right)^{-1} \left[ \sum_{i=1}^{n} r_{i,k}^{(t)} \left( x_i - \mu_k^{(t+1)} \right) \right],$$

which vanishes with the claimed formula for  $\mu_k^{(t+1)}$ . The computation of the necessary optimality condition  $\nabla_{\Sigma_k} F_t(\theta^{(t+1)}) = 0$  requires more care. We consider to this end a perturbation  $M_k = M_k^{\top}$  of  $\Sigma_k$ , all other parameters being fixed. Denoting by  $\eta_{M_k} \in \Theta$  the vector of parameters whose single non zero component is  $M_k$ , we can write

 $\operatorname{Tr}\left[M_{k}^{\top} \nabla_{\Sigma_{k}} F\left(\theta\right)\right] = F_{t}\left(\theta + \eta_{M_{k}}\right) - F_{t}(\theta) + \operatorname{O}(\|M_{k}\|^{2}).$ 

Now,  $(\Sigma_k + M_k)^{-1} = \Sigma_k^{-1} - \Sigma_k^{-1} M_k \Sigma_k^{-1} + O(||M_k||^2)$  when  $\Sigma_k$  and  $\Sigma_k + M_k$  are both positive definite (hence invertible), so that

$$(x_i - \mu_k)(\Sigma_k + M_k)^{-1}(x_i - \mu_k) - (x_i - \mu_k)\Sigma_k^{-1}(x_i - \mu_k)$$
  
=  $-(x_i - \mu_k)\Sigma_k^{-1}M_k\Sigma_k^{-1}(x_i - \mu_k) + O(||M_k||^2)$   
=  $-\mathrm{Tr}\left[\Sigma_k^{-1}M_k\Sigma_k^{-1}(x_i - \mu_k)(x_i - \mu_k)^{\top}\right] + O(||M_k||^2)$ 

Moreover, using det(Id<sub>d</sub> + A) = 1 + Tr(A) + O(||A||<sup>2</sup>),

$$\log \det(\Sigma_k + M_k) = \log \left[ \det(\Sigma_k) \det \left( \operatorname{Id}_d + \Sigma_k^{-1} M_k \right) \right]$$
  
= log det( $\Sigma_k$ ) + log  $\left[ 1 + \operatorname{Tr} \left( \Sigma_k^{-1} M_k \right) \right] + O(||M_k||^2)$   
= log det( $\Sigma_k$ ) + Tr  $\left( \Sigma_k^{-1} M_k \right) + O(||M_k||^2)$ .

Therefore,

$$F_t(\theta + \eta_{M_k}) - F_t(\theta) = \frac{1}{2} \sum_{i=1}^n r_{i,k}^{(t)} \operatorname{Tr} \left( \Sigma_k^{-1} M_k \left[ \Sigma_k^{-1} (x_i - \mu_k) (x_i - \mu_k)^\top - \operatorname{Id}_d \right] \right) + \mathcal{O}(\|M_k\|^2)$$
  
$$= \frac{1}{2} \sum_{i=1}^n r_{i,k}^{(t)} \operatorname{Tr} \left( M_k^\top \left[ \Sigma_k^{-1} (x_i - \mu_k) (x_i - \mu_k)^\top - \operatorname{Id}_d \right] \Sigma_k^{-1} \right) + \mathcal{O}(\|M_k\|^2).$$

This allows to conclude that

$$\sum_{i=1}^{n} r_{i,k}^{(t)} \left[ \left( \Sigma_{k}^{(t+1)} \right)^{-1} \left( x_{i} - \mu_{k}^{(t+1)} \right) \left( x_{i} - \mu_{k}^{(t+1)} \right)^{\top} - \mathrm{Id}_{d} \right] = 0,$$

which leads to the claimed expression for  $\Sigma_k^{(t+1)}$ .

Let us now discuss why the maximization problem (9.5) is considered, following the presentation in [40, Section 8.7.2]. We start by rewriting the log-likelihood of the data points by summing over the cluster assignments for conditional probabilities:

$$\mathscr{L}_{n}(\theta) = \sum_{i=1}^{n} \log p_{\theta}(x_{i}) = \sum_{i=1}^{n} \log \left( \sum_{k=1}^{K} p_{\theta}(x_{i}, z_{i} = k) \right) = \sum_{i=1}^{n} \log \left( \sum_{k=1}^{K} q_{i,k} \frac{p_{\theta}(x_{i}, z_{i} = k)}{q_{i,k}} \right),$$

where we introduced a discrete probability distribution  $(q_{i,1}, \ldots, q_{i,K})$  with non zero entries for all  $1 \leq i \leq n$ . Using Jensen's inequality and the concavity of the log,

$$\mathscr{L}_{n}(\theta) \geqslant \sum_{i=1}^{n} \sum_{k=1}^{K} q_{i,k} \log\left(\frac{p_{\theta}(x_{i}, z_{i} = k)}{q_{i,k}}\right) := \sum_{i=1}^{n} \mathcal{E}(\theta, q_{i} \mid x_{i}).$$

$$(9.6)$$

The right hand side of the previous inequality is called an "evidence lower bound" (ELBO), as it provides a lower bound to the negative log-likelihood. Moreover, the terms which appear in the ELBO can be rewritten as

$$\mathcal{E}(\theta, q_i \mid x_i) = \sum_{k=1}^{K} q_{i,k} \log\left(\frac{p_{\theta}(z_i = k \mid x_i)p_{\theta}(x_i)}{q_{i,k}}\right)$$
$$= \sum_{k=1}^{K} q_{i,k} \log\left(\frac{p_{\theta}(z_i = k \mid x_i)}{q_{i,k}}\right) + \left(\sum_{k=1}^{K} q_{i,k}\right) \log p_{\theta}(x_i)$$
$$= -\mathrm{KL}(q_i \mid p_{\theta}(z_i \mid x_i)) + \log p_{\theta}(x_i), \tag{9.7}$$

where

$$\mathrm{KL}(q \mid p) = \sum_{k=1}^{K} q_k \log \frac{q_k}{p_k} \ge 0$$

#### 154 9 Clustering methods

is the Kullback–Leibler entropy between two discrete probability distributions. The quantity (9.7) is maximized by setting  $q_{i,k} = p_{\theta}(z_i = k | x_i)$ , in view of Exercise 9.4 below. This corresponds to the expectation step. On the other hand, the maximization step is recovered by maximizing the lower bound (9.6) with respect to  $\theta$ . Overall, denoting by

$$\mathscr{P}_K = \left\{ p \in [0,1]^K \; \middle| \; \sum_{k=1}^K p_k = 1 \right\}$$

the set of discrete probabilities over a finite set of cardinality K, the EM algorithm can then be rewritten as

$$q_i^{(t)} = \underset{q_i \in \mathscr{P}_K}{\operatorname{argmax}} \mathcal{E}\left(\theta^t, q_i \, \big| \, x_i\right) = p_{\theta^t}(z_i \, | \, x_i) = r_i^{(t)},$$
  
$$\theta^{t+1} = \underset{\theta \in \Theta}{\operatorname{argmax}} \left\{ \sum_{i=1}^n \mathcal{E}\left(\theta, q_i^{(t)} \, \Big| \, x_i\right) \right\} = \underset{\theta \in \Theta}{\operatorname{argmax}} \left\{ \sum_{i=1}^n \sum_{k=1}^K r_{i,k}^{(t)} \log p_{\theta}(x_i, z_i = k) \right\},$$

the latter maximization problem being (9.5) since  $p_{\theta}(x_i, z_i = k) = \pi_k \mathscr{G}_{\mu_k, \Sigma_k}$ . In addition, by (9.7),

$$\sum_{i=1}^{n} \mathcal{E}\left(\theta^{t}, q_{i}^{(t)} \middle| x_{i}\right) = \sum_{i=1}^{n} \log p_{\theta^{t}}(x_{i}) = \mathscr{L}_{n}\left(\theta^{t}\right)$$

is the log-likelihood to be maximized. Therefore, by construction of the algorithm,

$$\mathscr{L}_{n}\left(\theta^{t}\right) \leqslant \sum_{i=1}^{n} \mathcal{E}\left(\theta^{t+1}, q_{i}^{\left(t\right)} \middle| x_{i}\right) \leqslant \sum_{i=1}^{n} \mathcal{E}\left(\theta^{t+1}, q_{i}^{\left(t+1\right)} \middle| x_{i}\right) = \mathscr{L}_{n}\left(\theta^{t+1}\right),$$

so that the log-likelihood is indeed non decreasing as claimed.

**Exercise 9.4.** Consider two discrete probability distributions  $q, p \in \mathscr{P}_K$ . Prove that  $\mathrm{KL}(q \mid p) \ge 0$ , and  $\mathrm{KL}(q \mid p) = 0$  if and only if q = p. It is useful to this end to first prove the identity

$$\forall z \in [0, +\infty), \qquad z \log z - z + 1 \ge 0,$$

with equality if and only if z = 1.

**Correction.** We first prove the identity  $z \log z - z + 1 \ge 0$ , with equality if and only if z = 1. This can be proved by noting that the derivative of the function  $z \mapsto z \log z - z + 1$  is  $z \mapsto \log z$ , so that the function decreases from 1 to 0, and then increases for  $z \ge 1$ . The value at z = 1 is 0, which shows that the function remains nonnegative. Moreover, the only value of z for which the function vanishes is z = 1.

The above identity then allows to write

$$\operatorname{KL}(q \mid p) = \sum_{k=1}^{K} q_k \log \frac{q_k}{p_k} = \sum_{k=1}^{K} \left(\frac{q_k}{p_k} \log \frac{q_k}{p_k}\right) p_k \ge \sum_{k=1}^{K} \left(1 - \frac{q_k}{p_k}\right) p_k = \sum_{k=1}^{K} p_k - \sum_{k=1}^{K} q_k = 0.$$

Equality holds if and only if  $q_k/p_k = 1$  for all  $1 \leq k \leq K$ , *i.e.* q = p.

**Remark 9.2.** One could alternatively make use of the fact that  $\varphi : z \mapsto z \log z$  is convex (since its first derivative  $z \mapsto \log z$  is increasing) in order to write

$$\operatorname{KL}(q \mid p) = \sum_{k=1}^{K} p_k \varphi\left(\frac{q_k}{p_k}\right) \ge \varphi\left(\sum_{k=1}^{K} p_k \frac{q_k}{p_k}\right) = \varphi(1) = 0.$$

(Soft) K-means. The cluster assignment

$$z_i = \operatorname*{argmax}_{1 \leqslant k \leqslant K} r_{i,k}^{\infty},$$

with  $r_i^{\infty}$  the limiting responsibility obtained from the EM algorithm, allows to fall back to hard clustering. One can consider some soft clustering by using the vector  $r_i^{\infty} \in [0, 1]^K$ , which provides the probabilities to be in one of the clusters. In any case, one benefit of Gaussian mixture clustering is that anisotropic shapes for the clusters are allowed thanks to the covariance matrices  $\Sigma_k^{\infty}$ . The standard *K*-means algorithm corresponds to hard assignments  $r_i$  for which a single component is equal to 1 (at each step of the procedure), with parameters  $\pi_k = 1/K$  and  $\Sigma_k = \mathrm{Id}_d$ , so that the only quantities to estimate are the cluster centers  $\mu_k$  (see [40, Section 21.4.1.1]).

### 9.5 Density based clustering

The presentation in this section is inspired by [38, Section 13.1.3]. The idea behind density based clustering is that clusters are defined by regions of space with a higher density of data points; outliers being in regions of low density. These methods therefore allow to identify clusters and the associated points, but also outliers, which should not be assigned to any cluster.

The most famous density based clustering algorithm is DBSCAN [17] (but other methods exist, see for instance the review in [56, Section 2.2]). Its computational cost scales as  $O(n \log n)$  and it can therefore be used for large data sets. It relies on  $\varepsilon$ -neighborhoods of data points, defined as

$$\mathcal{N}_{\varepsilon}(x_i) = \left\{ x \in \mathcal{D} \, | \, \mathsf{d}(x, x_i) < \varepsilon \right\}.$$

In addition to the parameter  $\varepsilon > 0$ , the method uses the integer  $N_{\min}$ , which is the minimal number of neighbors needed for a data point to be considered a "core point". This parameter should be chosen depending on the size of the smallest cluster one expects. A point is said to be "density reachable" if it is in the neighborhood of a core point.

The precise algorithm is the following: until all points have been visited,

- pick a data point x which has not been visited, determine  $\mathcal{N}_{\varepsilon}(x)$  and mark the point as visited;
- if x is a core point (*i.e.* the cardinality of  $\mathcal{N}_{\varepsilon}(x)$  is larger or equal to  $N_{\min}$ ),
- find the set C of all points that are density reachable from x;
- C forms a cluster; mark all points in it as visited.

The algorithm then returns the cluster assignments  $C_1, \ldots, C_K$ . Note that K is not fixed in advance in this method. Note also that some points are not assigned to a cluster, and are therefore considered as noise or outliers. Each cluster contains at least one core point. Non-core points form the "edge" of a cluster. Overall, clusters are composed of one or several core points and all points (core or non-core) which are density reachable from these core points.

## 9.6 Spectral clustering

The presentation in this section is based on the review article [55]. Spectral clustering starts by constructing a similarity graph from the data points, by considering data points as vertices, linked by edges with weights  $W_{ij} = s(x_i, x_j)$  for some similarity function  $s : \mathcal{X}^2 \to \mathbb{R}_+$ . A classical choice is

$$W_{ij} = \exp\left(-\frac{1}{2\sigma^2} \|x_i - x_j\|_2^2\right),$$

or more generally  $d(x_i, x_j)$  for some distance function. The aim of spectral clustering is to find a partition of the graph such that

- edges between different groups have low weights;
- edges within a group have large weights.

#### 156 9 Clustering methods

**Graph cuts.** A first attempt to partition the graph is to consider the first objective above, which leads to the so-called "mincut" problem:

$$\operatorname{Cut}(C_1,\ldots,C_K) = \sum_{k=1}^K W\left(C_k,\overline{C}_k\right), \qquad W(A,B) = \sum_{\substack{a \in A \\ b \in B}} W_{ab},$$

where  $\overline{C}_k = \mathcal{D} \setminus C_k$  is the complement of the data points  $C_k$  in the training data set. The mincut problem therefore minimizes the weight of connections between nodes of a given cluster and all other clusters. In practice, it often separates individual vertices from the rest of the graph. Think for instance of the case of a set evenly split into two subsets  $C, \overline{C}$  of sizes n/2: in this case,  $W(C, \overline{C})$ involves a summation over  $n^2/4$  terms, while  $W(\mathcal{D} \setminus \{x_i\}, \{x_i\})$  involves a summation over n-1terms only.

A simple fix is to normalize the cut in some way, by a term somewhat proportional to the size of the sets, so that the sums in the example above are of the same order of magnitude:

• unnormalized spectral clustering considers

RatioCut
$$(C_1, \ldots, C_K) = \sum_{k=1}^K \frac{1}{|C_k|} W(C_k, \overline{C}_k)$$

• normalized spectral clustering considers

$$\operatorname{NCut}(C_1, \dots, C_K) = \sum_{k=1}^K \frac{1}{\operatorname{vol}(C_k)} W\left(C_k, \overline{C}_k\right), \quad \operatorname{vol}(A) = \sum_{a \in A} d_a, \quad d_a = \sum_{i=1}^n W_{ai};$$

Normalized spectral clustering explicitly aims at maximizing the within cluster similarity  $W(C_k, C_k)$  since

$$W(C_k, C_k) = \operatorname{vol}(C_k) - W(C_k, \overline{C}_k),$$

the first term being maximized by the objective function while the second term is minimized. On the other hand, unnormalized spectral clustering emphasizes larger clusters, but these may contain many connections of low weight. See [55, Section 8.5.1] for a more detailed discussion of this point.

**Relaxed graph cuts.** It is helpful in fact to solve relaxed versions of the minimization problems NCut and RatioCut: instead of assigning points to a cluster, *i.e.* finding  $c_i \in \{0, 1\}^K$  with  $c_{ik} = 1$  if and only if  $x_i \in C_k$  (in particular,  $c_{ik} = 0$  if  $x_i \notin C_k$ ), one solves a simpler linear algebra problem; see [55, Sections 5.3 and 5.4] and Exercise 9.5 below for a precise discussion on how this relaxation is performed.

A useful matrix to consider for the relaxed problem is L = D - W, where D (the degree matrix) is a diagonal matrix with entries  $D_{ii} = d_i$ , the total weight of the edges connected to  $x_i$ . The following result is fundamental to the spectral approach to clustering (see [55, Proposition 2] and [40, Theorem 21.5.1]).

**Proposition 9.1.** The multiplicity K of the eigenvalue 0 of L is equal to the number of connected components  $C_1, \ldots, C_K$  of the graph, and the associated eigenvectors are proportional to  $\mathbf{1}_{C_1}, \ldots, \mathbf{1}_{C_K}$ .

*Proof.* Note first that L is real symmetric, so that its spectrum is contained in  $\mathbb{R}$ . For  $\xi \in \mathbb{R}^n$ , consider

$$\xi^{\top} L\xi = \sum_{i=1}^{n} d_i \xi_i^2 - \sum_{i,j=1}^{n} W_{ij} \xi_i \xi_j = \frac{1}{2} \sum_{i,j=1}^{n} W_{ij} \left(\xi_i - \xi_j\right)^2,$$

where we used  $W_{ij} = W_{ji}$  for the last equality. Therefore, L is positive semidefinite, so its spectrum is contained in  $[0, +\infty)$ .

If  $L\xi = 0$ , then  $\xi_i = \xi_j$  for all nodes for which  $W_{ij} > 0$ , *i.e.* for all nodes in the same connected component of the graph. When K = 1, there is a single connected component, and eigenvectors associated with the eigenvalue 0 are proportional to  $\mathbf{1}_{\mathcal{D}}$ . When  $K \ge 2$ , upon reordering the data points, one can write L in block diagonal form as

$$L = \begin{pmatrix} L_1 & 0 & \dots & 0 \\ 0 & L_2 & \dots & 0 \\ & \ddots & \\ 0 & 0 & \dots & L_K \end{pmatrix}.$$

Then, the spectrum of L is

$$\sigma(L) = \bigcup_{k=1}^{K} \sigma(L_k).$$

The operators  $L_k$  all admit 0 as a nondegenerate eigenvalue, with associated eigenvectors  $\mathbf{1}_{S_k}$ , which allows to conclude.

The above result suggests by a perturbative analysis that, when the connected components are weakly linked, the spectrum should consist of K small eigenvalues and then a spectral gap between the (K+1)-th and the K-th eigenvalues. This remark allows to find the relevant number of clusters K in practice by diagonalizing L and locating the spectral gap.

Moreover, the eigenvectors associated with these eigenvalues should be almost constant on the connected components (*i.e.* have values close to 1 on one of the connected components, and close to 0 on the others, when normalized). A practical algorithm to find these eigenvectors, and hence identify the clusters, is to perform K-means on the lines of the matrix

$$\mathcal{U} = \begin{pmatrix} | & | & | \\ U_1 \ U_2 \ \dots \ U_K \\ | & | & | \end{pmatrix} \in \mathbb{R}^{n \times K},$$

where  $(U_k)_{1 \leq k \leq K}$  are the first (normalized) K eigenvectors of L. This corresponds to the relaxed version of RatioCut. Some authors suggest to nomalize L as  $L_{\text{sym}} = D^{-1/2}LD^{-1/2}$  in order to take into account that some clusters are more connected to others (see [40, Section 21.5.2]) and perform K-means clustering on the first K eigenvectors of  $L_{\text{sym}}$ . This corresponds to the relaxed version of NCut. There is no consensus on which approach to prefer, see the discussions in [55, Section 8.5.3].

**Exercise 9.5.** We make precise in this exercise in which sense the spectral problem on the graph Laplacian can be understood as a relaxed version of graph cut problems.

(a) We first consider RatioCut for K = 2, i.e. consider the partition  $(C, \overline{C})$ . Introduce the vector  $h \in \mathbb{R}^n$  with components

$$h_{i} = \sqrt{\frac{|\overline{C}|}{|C|}} \mathbf{1}_{x_{i} \in C} - \sqrt{\frac{|C|}{|\overline{C}|}} \mathbf{1}_{x_{i} \in \overline{C}}.$$
(9.8)

Denoting by  $\mathbf{1}_n \in \mathbb{R}^n$  the vector for which all components are equal to 1, prove that

$$h \cdot \mathbf{1} = 0, \qquad h^{\top} h = n, \qquad \text{RatioCut}\left(C, \overline{C}\right) = \frac{1}{n} h^{\top} L h.$$

Deduce that

RatioCut 
$$(C, \overline{C}) \ge \min_{\substack{\xi \in \mathbb{R}^n \setminus \{0\}\\ \xi \cdot \mathbf{1}_n = 0}} \frac{\xi^\top L\xi}{\xi^\top \xi}.$$

What is the minimizer of the right hand side?

158 9 Clustering methods

(b) We next consider RatioCut for  $K \ge 2$ . Define the matrix  $H \in \mathbb{R}^{n \times K}$  with entries

$$H_{ik} = \frac{1}{\sqrt{|C_k|}} \mathbf{1}_{x_i \in C_k}.$$
(9.9)

Prove that  $H^{\top}H = \mathrm{Id}_K$  and that

$$\frac{1}{|C_k|} \operatorname{Cut} \left( C_k, \overline{C}_k \right) = \left( H^\top L H \right)_{kk}$$

Deduce that

RatioCut
$$(C_1, \ldots, C_K) \ge \min_{H \in \mathbb{R}^{n \times K}} \left\{ \operatorname{Tr} \left( H^\top L H \right) \mid H^\top H = \operatorname{Id}_K \right\}$$

What is the minimizer of the right hand side?

(c) We finally consider NCut. By proceeding as in the previous question, and introducing the matrix  $\widetilde{H} \in \mathbb{R}^{n \times K}$  with entries

$$\widetilde{H}_{ik} = \frac{1}{\sqrt{\operatorname{vol}(C_k)}} \mathbf{1}_{x_i \in C_k},$$

show that the relaxed version of NCut is

$$\min_{H \in \mathbb{R}^{n \times K}} \left\{ \operatorname{Tr} \left( H^{\top} L_{\operatorname{sym}} H \right) \mid H^{\top} H = \operatorname{Id}_{K} \right\}.$$

#### Correction.

(a) The normalization conditions on h are readily verified:

$$h \cdot \mathbf{1} = \sum_{i=1}^{n} h_i = \sum_{x_i \in C} \sqrt{\frac{|\overline{C}|}{|C|}} - \sum_{i \in \overline{C}} \sqrt{\frac{|C|}{|\overline{C}|}} = 0,$$

and

$$h^{\top}h = \sum_{i=1}^{n} h_i^2 = \sum_{x_i \in C} \frac{|\overline{C}|}{|C|} + \sum_{i \in \overline{C}} \frac{|C|}{|\overline{C}|} = |\overline{C}| + |C| = n.$$

Moreover, since  $h_i = h_j$  when  $(x_i, x_j) \in C^2$  or  $(x_i, x_j) \in \overline{C}^2$ ,

$$\frac{1}{n}h^{\top}Lh = \frac{1}{2n}\sum_{i,j=1}^{n}W_{ij}(h_i - h_j)^2 = \frac{1}{2n}\left(\sum_{\substack{x_i \in C \\ x_j \in \overline{C}}} + \sum_{\substack{x_i \in \overline{C} \\ x_j \in C}}\right)W_{ij}\left(\sqrt{\frac{|\overline{C}|}{|C|}} + \sqrt{\frac{|C|}{|\overline{C}|}}\right)^2$$

$$= \frac{1}{n}\left(\frac{|\overline{C}|}{|C|} + \frac{|C|}{|\overline{C}|} + 2\right)\sum_{\substack{x_i \in C \\ x_j \in \overline{C}}}\sum_{\substack{x_j \in \overline{C}}}W_{ij}$$

$$= \frac{1}{n}\left(\frac{|\overline{C}| + |C|}{|C|} + \frac{|\overline{C}| + |C|}{|\overline{C}|}\right)\operatorname{Cut}(C,\overline{C})$$

$$= \frac{|\overline{C}| + |C|}{n}\left(\frac{1}{|C|} + \frac{1}{|\overline{C}|}\right)\operatorname{Cut}(C,\overline{C}) = \operatorname{RatioCut}(C,\overline{C}).$$

Therefore, denoting by  $h_C$  the vector with components (9.8), and by

$$\mathcal{H} = \{h_C, C \subset \mathcal{D}\} \subset \left\{\xi \in \mathbb{R}^n \mid \xi \cdot \mathbf{1}_n = 0, \ \xi^\top \xi = n\right\},\$$

it holds

$$\min_{C \subset \mathcal{D}} \operatorname{RatioCut} \left( C, \overline{C} \right) = \min_{\xi \in \mathcal{H}} \frac{\xi^\top L \xi}{\xi^\top \xi} \ge \min_{\substack{\xi \in \mathbb{R}^n \setminus \{0\} \\ \xi \cdot \mathbf{1}_n = 0}} \frac{\xi^\top L \xi}{\xi^\top \xi} = \sigma_2(L),$$

where  $\sigma_2(L)$  is the second eigenvalue of L, the first one being 0 with associated eigenvector  $\mathbf{1}_n$ . The minimal value is attained for  $\xi$  proportional to the second eigenvector of L.

(b) The fact that  $H^{\top}H = \text{Id}_K$  follows by a direct computation. Moreover, by computations similar to the ones performed in the previous question, and denoting by  $H_k$  the k-th column of H,

$$\frac{1}{|C_k|} \operatorname{Cut}\left(C_k, \overline{C}_k\right) = \frac{1}{|C_k|} \sum_{x_i \in C} \sum_{x_j \in \overline{C}} W_{ij} = H_k^\top L H_k = \left[H^\top L H\right]_{kk},$$

which gives the claimed formula. Therefore, denoting by  $H_C$  the matrix with components (9.9), and by

$$\mathcal{H} = \{ H_C, C \subset \mathcal{D} \} \subset \{ H \in \mathbb{R}^{n \times K} \mid H^\top H = \mathrm{Id}_K \},\$$

it holds

RatioCut
$$(C_1, \ldots, C_K) = \min_{H \in \mathcal{H}} \operatorname{Tr} (H^\top L H) \ge \min_{H \in \mathbb{R}^{n \times K}} \left\{ \operatorname{Tr} (H^\top L H) \mid H^\top H = \operatorname{Id}_K \right\}$$

By the Rayleigh-Ritz principle, the right hand side of the previous inequality is minimized for  $H = [U_1| \dots |U_K]$ , the matrix whose k-th column  $U_k$  is a normalized eigenvector associated with the k-th eigenvalue of L.

(c) Computations very similar to the ones of the previous question give

$$\widetilde{H}^{\top} D \widetilde{H} = \mathrm{Id}_K,$$

so that

$$\operatorname{NCut}(C_1,\ldots,C_K) \ge \min_{\widetilde{H} \in \mathbb{R}^{n \times K}} \left\{ \operatorname{Tr}\left(\widetilde{H}^\top L \widetilde{H}\right) \mid \widetilde{H}^\top D \widetilde{H} = \operatorname{Id}_K \right\}.$$

By writing  $\widetilde{H} = D^{1/2}H$ , the claimed lower bound is obtained on the right hand side of the previous inequality.

# **Complements: Elements of statistical learning**

10.1 Empi	rical risk minimization and statistical learning theory 161	
10.1.1	Framework	
10.1.2	Decomposition of the risk minimization	
10.1.3	Proof of Proposition 10.2	
10.2 Model selection 166		
10.2.1	Validation	
10.2.2	Model selection with validation	
10.3 Statistical learning theory 168		

This chapter gathers additional material which is mentioned at various places in the lecture notes: guarantees on empirical risk minimization in Section 10.1, elements on model selection in Section 10.2, and some key facts in Statistical Learning Theory in Section 10.3.

#### 10.1 Empirical risk minimization and statistical learning theory

We present here a brief introduction to some theoretical aspects of empirical risk minimization for supervised learning, in particular estimates motivating that this approach provides good predictors. The material in this section is taken from [4, Sections 2.3 to 2.5] as well as selected parts of Chapter 4 of this reference, [40, Section 5.4] and [39, Sections 4.1 and 4.2].

#### 10.1.1 Framework

Let us briefly recall here the setting of empirical risk minimization. We consider a parametrized family of predictors  $f_{\theta} : \mathcal{X} \to \mathcal{Y}$  for  $\theta \in \Theta$ . In practice, predictors can be obtained from data by minimizing with respect to  $\theta$  the empirical risk

$$\widehat{\mathcal{R}}_n(f_\theta) = \frac{1}{n} \sum_{i=1}^n \ell(y_i, f_\theta(x_i)),$$

where  $\mathcal{D}_n = \{(x_1, y_1), \ldots, (x_n, y_n)\}$  is a dataset of i.i.d. pairs  $(x_i, y_i)$  sampled with respect to some unknown distribution  $p_{data}$  (see for instance (2.1) and specific formulations such as (2.3), or (8.1), to give just two examples). The minimization of such empirical risks can sometimes be easy (as for linear regression problems, for which the optimal value of the parameter has an analytic expression; or for logistic regression, for which the empirical risk is convex), and can in any cases be considered for parameters of arbitrary dimensions. However, the minimization can also be quite difficult, for instance for non-convex empirical risks (as those encountered for supervised learning with neural networks), or when the gradient with respect to  $\theta$  is complex or expensive to compute. Moreover, classification requires parametrizing functions with values in  $\{0, 1\}$  (and/or convexifying the loss, as discussed in Section 3.1), and in general one needs capacity control to avoid overfitting. 162 10 Complements: Elements of statistical learning

#### 10.1.2 Decomposition of the risk minimization

The aim of this section is to provide estimates on the quality of the predictor found by empirical risk minimization. We decompose to this end the excess risk into an estimation error and an approximation error. We consider for this a family  $\mathcal{F}$  of predictor functions  $\mathcal{X} \to \mathcal{Y}$  (for instance  $\{f_{\theta}, \theta \in \Theta\}$  with  $\Theta$  belonging to  $\mathbb{R}^d$  or to some ball in order to have some capacity control). The predictor found by the empirical risk minimization is

$$\widehat{f} \in \underset{f \in \mathcal{F}}{\operatorname{argmin}} \widehat{\mathcal{R}}_n(f).$$
(10.1)

The question is how large the excess risk  $\mathcal{R}(\hat{f}) - \mathcal{R}^{\star}$  is. We rewrite the latter quantity as

$$\underbrace{\mathcal{R}(\widehat{f}) - \mathcal{R}^{\star}}_{\text{excess risk}} = \underbrace{\mathcal{R}(\widehat{f}) - \inf_{f' \in \mathcal{F}} \mathcal{R}(f')}_{\text{estimation error}} + \underbrace{\inf_{f' \in \mathcal{F}} \mathcal{R}(f') - \mathcal{R}^{\star}}_{\text{approximation error}}.$$
(10.2)

Let us discuss the two errors on the right hand side of the previous equality:

- The approximation error is a deterministic quantity, which depends on  $\mathcal{F}$  and the underlying (unknown) distribution  $p_{\text{data}}$ . It quantifies how well the Bayes predictor  $f^*$  can be represented in  $\mathcal{F}$ . The idea is that the larger  $\mathcal{F}$  is, the smaller the approximation error is. Precise convergence rates depend on the regularity of  $f^*$  (which is a classical result in analysis and approximation theory).
- The estimation error is a random quantity, which depends implicitly on the data set through the minimizer  $\hat{f}$  of the empirical risk (see (10.1)). This error is not easy to quantify as such. A key idea here is to upper bound it in terms of the supremum over  $f \in \mathcal{F}$  of the differences  $|\mathcal{R}(f) \widehat{\mathcal{R}}_n(f)|$ , as made precise in Proposition 10.1.

**Proposition 10.1 (Bound on the estimation error).** For any data set  $\mathcal{D}_n$ , it holds

$$0 \leqslant \mathcal{R}(\widehat{f}) - \inf_{f \in \mathcal{F}} \mathcal{R}(f) \leqslant 2 \sup_{f \in \mathcal{F}} \left| \widehat{\mathcal{R}}_n(f) - \mathcal{R}(f) \right|.$$

The main interest of this result is that it is more convenient to work with  $\widehat{\mathcal{R}}_n$  than  $\widehat{f}$ . Let us emphasize that the upper bound is a random quantity (depending on the data set). This random quantity usually decays with n (with  $\widehat{\mathcal{R}}_n(f)$  converging almost surely to  $\mathcal{R}(f)$  for f given).

*Proof.* Assume that  $\inf_{f \in \mathcal{F}} \mathcal{R}(f) = \mathcal{R}(g_{\mathcal{F}})$  is attained for some  $g_{\mathcal{F}} \in \mathcal{F}$  (otherwise the argument below needs to be modified as in the proof of [39, Proposition 4.1]). Then,

$$\begin{split} \mathcal{R}(\widehat{f}) &- \inf_{f \in \mathcal{F}} \mathcal{R}(f) = \mathcal{R}(\widehat{f}) - \mathcal{R}(g_{\mathcal{F}}) \\ &= \left( \mathcal{R}(\widehat{f}) - \widehat{\mathcal{R}}_n(\widehat{f}) \right) + \underbrace{\widehat{\mathcal{R}}_n(\widehat{f}) - \widehat{\mathcal{R}}_n(g_{\mathcal{F}})}_{\leqslant 0} + \left( \widehat{\mathcal{R}}_n(g_{\mathcal{F}}) - \mathcal{R}(g_{\mathcal{F}}) \right) \\ &\leqslant 2 \sup_{f \in \mathcal{F}} \left| \mathcal{R}(f) - \widehat{\mathcal{R}}_n(f) \right|, \end{split}$$

which is the claimed result.

Let us give a few comments on the above result:

- the statistical dependence between  $\widehat{f}$  and  $\widehat{\mathcal{R}}_n$  is removed at the end of the proof by considering a uniform bound over all possible functions in  $\mathcal{F}$ . This uniform deviation bound typically grows as  $\mathcal{F}$  becomes larger;
- it would be possible to take into account an additional optimization error (quantifying the fact that the minimizer of the empirical risk is obtained by a numerical procedure) by writing that  $\widehat{\mathcal{R}}_n(\widehat{f}) \widehat{\mathcal{R}}_n(g_{\mathcal{F}}) \leqslant \varepsilon$  instead of saying that this quantity is nonnegative.

In order to instantiate the bound on the estimation error provided by Proposition 10.1, we consider the simple setting when  $\mathcal{F}$  is a set of finite cardinality  $|\mathcal{F}| < +\infty$ , and the elementary loss  $\ell$  is bounded.

**Proposition 10.2.** Suppose that  $\mathcal{F}$  is finite, and the elementary loss  $\ell$  is bounded: there exists  $L < +\infty$  such that

$$\forall (y,z) \in \mathcal{Y}^2, \qquad 0 \leqslant \ell(y,z) \leqslant \mathsf{L} < +\infty.$$

Then, for any  $\delta \in (0,1)$ , the following bound holds with probability larger than  $1-\delta$ :

$$0 \leq \mathcal{R}(\widehat{f}) - \inf_{f \in \mathcal{F}} \mathcal{R}(f) \leq \mathsf{L}\sqrt{\frac{2}{n} \left(\log|\mathcal{F}| + \log\frac{2}{\delta}\right)}.$$
(10.3)

This result, which is based on Hoeffding's inequality, is proved in Section 10.1.3. Let us make a few comments on it:

- the error is "with high probability": the upper bound becomes larger as the probability  $1 \delta$  comes closer to 1;
- the decay rate of the upper bound is  $1/\sqrt{n}$ , consistent with the scaling  $\widehat{\mathcal{R}}(f) \mathcal{R}(f) = O(n^{-1/2})$  obtained from the central limit theorem;
- the size  $|\mathcal{F}|$  of the set of functions for the approximation appears only logarithmically. The result can be extended to set which are not discrete, under some conditions, in which case the term  $\log |\mathcal{F}|$  should be replaced by the dimension of the set.

The result on the estimation error of Propositions 10.1 and 10.2, combined with the decomposition of the excess risk (10.2), suggests that a trade-off should be reached in the complexity of  $\mathcal{F}$ : the approximation error is smaller for  $\mathcal{F}$  larger, while the estimation error increases in this case; and conversely. The increase of the estimation error is related to the generalization capability of the model: a model with a small estimation error generalizes better.

A typical scenario, as the size of  $\mathcal{F}$  (or the set  $\Theta$  of parameters for the function  $f_{\theta}$  in the set  $\mathcal{F}$ ) increases, is the following:

- for small sizes of  $\mathcal{F}$ , the train and validation (test) errors are both large. This corresponds to the regime of *underfitting*, where the estimation error is small but the approximation error is large ("large bias but small variance");
- for large sizes of  $\mathcal{F}$ , the train error is small but the validation error is large. This corresponds to the regime of *overfitting*, where the approximation error is small but the estimation error is large since the predictor learnt from the train dataset fits too closely the training data and hence generalizes poorly ("small bias but large variance");

#### 10.1.3 Proof of Proposition 10.2

A key tool in the proof of Proposition 10.2 is Hoeffding's inequality (proved in Exercise 10.4 below).

**Proposition 10.3 (Hoeffding's inequality).** If  $Z_1, \ldots, Z_n$  are independent random variables such that  $Z_i \in [0,1]$  almost surely, then, for any  $t \ge 0$ ,

$$\mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n}Z_{i}-\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_{i}] \ge t\right) \le e^{-2nt^{2}}.$$
(10.4)

Note that the random variables  $Z_i$  need not be identically distributed. The result above can be considered as a quantitative version of the Law of Large Numbers.

**Exercise 10.1.** Under the same assumptions as in Proposition 10.3, prove that

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_{i=1}^{n}Z_{i}-\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_{i}]\right| \ge t\right) \le 2\mathrm{e}^{-2nt^{2}}.$$

164 10 Complements: Elements of statistical learning

**Correction.** The result directly follows from the fact that  $\{|X| \ge t\} = \{X \ge t\} \cup \{-X \ge t\}$  together with the union bound. Indeed, the inequality (10.4) holds with  $Z_i$  replaced by  $-Z_i$  since it suffices to apply Proposition 10.3 to the family of random variables  $Z'_i = 1 - Z_i$ .

**Remark 10.1.** The bound (10.4) can be compared with the asymptotic bound given by the Central Limit Theorem when the random variables  $Z_i$  are *i.i.d.* Indeed, as  $n \to +\infty$ ,

$$\lim_{n \to +\infty} \mathbb{P}\left(\frac{1}{n} \sum_{i=1}^{n} Z_i - \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}[Z_i] \geqslant \frac{t}{\sqrt{n}}\right) = \sqrt{\frac{1}{2\pi\sigma^2}} \int_{t}^{+\infty} e^{-z^2/(2\sigma^2)} dz, \quad (10.5)$$

for  $\sigma^2 = \operatorname{Var}(Z_1)$ . The integral in the right-hand side of the above equality can be upper bounded as

$$\begin{split} \int_{t}^{+\infty} \mathrm{e}^{-z^{2}/(2\sigma^{2})} \, dz &= \mathrm{e}^{-t^{2}/(2\sigma^{2})} \int_{0}^{+\infty} \mathrm{e}^{-y^{2}/(2\sigma^{2})} \mathrm{e}^{-ty/\sigma^{2}} \, dy \leqslant \mathrm{e}^{-t^{2}/(2\sigma^{2})} \int_{0}^{+\infty} \mathrm{e}^{-y^{2}/(2\sigma^{2})} \, dy \\ &= \frac{\sqrt{2\pi\sigma^{2}}}{2} \mathrm{e}^{-t^{2}/(2\sigma^{2})}, \end{split}$$

so that the right-hand side of (10.5) is smaller than  $e^{-t^2/(2\sigma^2)}/2$ . Now, random variables with values in [0, 1] have a variance smaller than 1/4 (see Exercise 10.2), so that  $e^{-t^2/(2\sigma^2)}/2 \leq e^{-2t^2}/2$ , which is itself smaller than the bound  $e^{-2t^2}$  obtained from (10.4). The result (10.5) is however only an asymptotic result, whereas (10.4) holds for any  $n \ge 1$ .

**Exercise 10.2.** Consider a random variable Y which takes almost surely values in [0,1]. Prove that  $Var(Y) \leq 1/4$ .

Correction. Note that

$$\operatorname{Var}(Y) = \min_{a \in [0,1]} \mathbb{E}\left(|Y-a|^2\right) \leq \mathbb{E}\left(\left|Y-\frac{1}{2}\right|^2\right) = \frac{1}{4}\mathbb{E}\left(|2Y-1|^2\right).$$

Now,  $|2Y - 1| \in [0, 1]$  almost surely, so that  $\mathbb{E}(|2Y - 1|^2) \leq 1$ , which leads to the claimed result.

**Exercise 10.3.** Which bounds holds instead of (10.4) when  $Z_i \in [a, b]$  almost surely with  $-\infty < a < b < +\infty$ ?

Correction. The random variable

$$\overline{Z}_i = \frac{Z_i - a}{b - a}$$

has values in [0, 1] almost surely, so that Proposition 10.3 implies that

$$\mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n}Z_{i}-\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_{i}] \ge (b-a)t\right) \le e^{-2nt^{2}}.$$

This shows that

$$\mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n}Z_{i}-\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_{i}] \ge t\right) \leqslant e^{-2nt^{2}/(b-a)^{2}}.$$

We are now in position to prove Proposition 10.2 with Hoeffding's inequality. For  $t \ge 0$ , using the upper bound on the estimation error provided by Proposition 10.1, then the union bound,

$$\mathbb{P}\left(\mathcal{R}(\widehat{f}) - \inf_{f \in \mathcal{F}} \mathcal{R}(f) \ge t\right) \leqslant \mathbb{P}\left(2\sup_{f \in \mathcal{F}} \left|\widehat{\mathcal{R}}_n(f) - \mathcal{R}(f)\right| \ge t\right)$$
$$\leqslant \sum_{f \in \mathcal{F}} \mathbb{P}\left(2\left|\widehat{\mathcal{R}}_n(f) - \mathcal{R}(f)\right| \ge t\right).$$

For f given, it holds, in view of Exercises 10.3 and 10.1 and since  $Z_i = \ell(f(x_i), y_i) \in [0, L]$ ,

$$\mathbb{P}\left(2\left|\widehat{\mathcal{R}}_{n}(f) - \mathcal{R}(f)\right| \ge t\right) \le 2\exp\left(-\frac{2n}{\mathsf{L}^{2}}\left(\frac{t}{2}\right)^{2}\right).$$

Therefore,

$$\mathbb{P}\left(\mathcal{R}(\widehat{f}) - \inf_{f \in \mathcal{F}} \mathcal{R}(f) \ge t\right) \leqslant 2|\mathcal{F}| \mathrm{e}^{-nt^2/(2\mathsf{L}^2)}.$$

We set  $\delta = 2|\mathcal{F}|e^{-nt^2/(2\mathsf{L}^2)}$ , *i.e.* 

$$t = \mathsf{L}\sqrt{\frac{2}{n}\log\left(\frac{2|\mathcal{F}|}{\delta}\right)},$$

from which the result follows.

**Exercise 10.4 (Proof of Hoeffding's inequality).** Consider a random variable Z which almost surely has values in [0,1], and introduce

$$\varphi(s) = \log \mathbb{E}\left[e^{s(Z - \mathbb{E}(Z))}\right]$$

We start by proving that  $e^{\varphi(s)} \leq e^{s^2/8}$  for any  $s \geq 0$ , and then derive the claimed bound.

- (1) Compute  $\varphi'(s)$  and  $\varphi''(s)$ , and prove that  $\varphi''(s) \ge 0$  for any  $s \ge 0$ .
- (2) Show that  $\varphi''(s) \leq 1/4$  by relying on Exercise 10.2, and deduce that

$$\mathbb{E}\left[\mathrm{e}^{s(Z-\mathbb{E}(Z))}\right] \leqslant \mathrm{e}^{s^2/8}.$$
(10.6)

(3) Use Markov's inequality  $P(X \ge a) \le \mathbb{E}(X)/a$  for a nonnegative random variable X and a > 0 to prove

$$\mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n}Z_{i}-\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_{i}] \ge t\right) \leqslant e^{-st+s^{2}/(8n)}$$

Conclude.

**Correction.** (1) Denote by  $X_s = e^{s(Z - \mathbb{E}(Z))}$ . A simple computation shows that

$$\varphi'(s) = \frac{\mathbb{E}\left[(Z - \mathbb{E}(Z))X_s\right]}{\mathbb{E}(X_s)}, \qquad \varphi''(s) = \frac{\mathbb{E}\left[(Z - \mathbb{E}(Z))^2 X_s\right]}{\mathbb{E}(X_s)} - \left(\frac{\mathbb{E}\left[(Z - \mathbb{E}(Z))X_s\right]}{\mathbb{E}(X_s)}\right)^2.$$

A Cauchy–Schwarz inequality implies that  $\mathbb{E}\left[(Z - \mathbb{E}(Z))^2 X_s\right] \mathbb{E}(X_s) \ge \mathbb{E}\left[(Z - \mathbb{E}(Z)) X_s\right]^2$ , so that  $\varphi''(s) \ge 0$ .

(2) Introduce the random variable  $Y_s$  with values in [0, 1], defined by the following equality: for any measurable set  $A \subset [0, 1]$ ,

$$\mathbb{P}(Y_s \in A) = \frac{\mathbb{E}(\mathbf{1}_A X_s)}{\mathbb{E}(X_s)}.$$

#### 166 10 Complements: Elements of statistical learning

Alternatively, one can define  $Y_s$  as a random variable with law proportional to  $e^{-s(z-\mathbb{E}(Z))}\mu(dz)$ , where  $\mu$  is the law of Z. In any case,  $\operatorname{Var}(Y_s) = \varphi''(s)$ , so that  $\varphi''(s) \leq 1/4$  by Exercise 10.2. Since  $\varphi(0) = \varphi'(0) = 0$ , a Taylor formula with exact remainder gives, for some  $\theta_s \in [0, 1]$ ,

$$\varphi(s) = \varphi(0) + s\varphi'(0) + \frac{s^2}{2}\varphi''(\theta_s s) = \frac{s^2}{2}\varphi''(\theta_s s) \leqslant \frac{s^2}{8}$$

which gives the desired upper bound by exponentiation.

(3) Introduce

$$\overline{Z}_n = \frac{1}{n} \sum_{i=1}^n Z_i.$$

Then, using first Markov's inequality, then the independence of  $Z_1, \ldots, Z_n$ , and finally (10.6),

$$\mathbb{P}\left(\frac{1}{n}\sum_{i=1}^{n}Z_{i}-\frac{1}{n}\sum_{i=1}^{n}\mathbb{E}[Z_{i}] \ge t\right) = \mathbb{P}\left(\exp\left[s\left(\overline{Z}_{n}-\mathbb{E}[\overline{Z}_{n}]\right)\right] \ge e^{st}\right)$$
$$\leqslant e^{-st}\mathbb{E}\left(\exp\left[s\left(\overline{Z}_{n}-\mathbb{E}[\overline{Z}_{n}]\right)\right]\right)$$
$$\leqslant e^{-st}\prod_{i=1}^{n}\mathbb{E}\left(\exp\left[\frac{s}{n}\left(Z_{i}-\mathbb{E}[Z_{i}]\right)\right]\right)$$
$$\leqslant e^{-st}\prod_{i=1}^{n}e^{s^{2}/(8n^{2})} = \exp\left(-st+\frac{s^{2}}{8n}\right)$$

The right-hand side of the previous inequality is minimized for s/(4n) = t, *i.e.* s = 4nt. The argument of the exponential function for this value of s is  $-2nt^2$ , which allows to conclude.

#### 10.2 Model selection

This section discusses the mathematical foundations behind the techniques to estimate the expected risk and choose hyperparameters of models. More precisely, we rely on model selection, whose aim is to find the right balance between the estimation and approximation errors. We present the approach when looking for the best hyperparameters (*e.g.* the number of neighbors for KNN), but this can be readily generalized to choosing a regularization strength (which corresponds to the so-called structural risk minimization). The presentation is based on [49, Chapter 11] and [39, Sections 4.4 and 4.5].

#### 10.2.1 Validation

The principle of cross-validation is introduced in Section 1.3.2, when studying KNN. It corresponds to using some part of the training set as validation set to select hyperparameters. To formalize the discussion, introduce a dataset  $\mathcal{D} = \{(x_1, y_1), \ldots, (x_n, y_n)\}$  is a dataset of i.i.d. pairs  $(x_i, y_i)$ sampled with respect to some unknown probability distribution. This dataset is randomly decomposed into a training data set  $\mathcal{D}_{\text{train}}$  of size  $(1 - \alpha)n$  and a validation set  $\mathcal{D}_{\text{val}}$  of size  $\alpha n$ , for some fraction  $\alpha \in [0, 1]$  (chosen such that  $\alpha n$  is an integer; a typical value would be  $\alpha = 0.2$ ). This is equivalent to independently sampling a training and validation set of i.i.d. data points, of respective sizes  $(1 - \alpha)n$  and  $\alpha n$ .

**Proposition 10.4.** Assume that the elementary loss function is bounded, namely  $0 \leq \ell(y, z) \leq L < +\infty$  for any  $y, z \in \mathcal{Y}$ . Then,

10.2 Model selection 167

$$\mathbb{P}\left(\left|\mathcal{R}(\widehat{f}_{\mathcal{D}_{\text{train}}}) - \widehat{\mathcal{R}}_{\text{val}}(\widehat{f}_{\mathcal{D}_{\text{train}}})\right| \ge \varepsilon\right) \le 2e^{-2\alpha n\varepsilon^2/\mathsf{L}^2},\tag{10.7}$$

where  $\hat{f}_{\mathcal{D}_{\text{train}}}$  is a minimizer of the training loss (defined on  $\mathcal{D}_{\text{train}}$ ). Therefore, with probability larger that  $1 - \delta$  over the choice of the train and validation sets,

$$\left| \mathcal{R}(\widehat{f}_{\mathcal{D}_{\text{train}}}) - \widehat{\mathcal{R}}_{\text{val}}(\widehat{f}_{\mathcal{D}_{\text{train}}}) \right| \leq \mathsf{L}_{\sqrt{\frac{1}{2\alpha n}}} \log\left(\frac{2}{\delta}\right).$$
(10.8)

**Remark 10.2.** Note that the bound (10.8) is tighter than bounds such as (10.3) (in the sense that it involves a single term with  $\log(2/\delta)$ ) because a "fresh" sample, independent of the training set, is considered.

*Proof.* The first inequality is obtained by fixing  $\hat{f}_{\mathcal{D}_{\text{train}}}$  (which is determined by the training set) and then using Hoeffding's inequality (10.4) with respect to realizations of the validation set. The choice  $\delta = 2e^{-2\alpha n\varepsilon^2/L^2}$ .

$$\varepsilon = \mathsf{L}\sqrt{\frac{1}{2\alpha n}\log\left(\frac{2}{\delta}\right)},$$

directly leads to the second inequality.

#### 10.2.2 Model selection with validation

Consider the training of models on a training set  $\mathcal{D}_{\text{train}}$  of size  $(1 - \alpha)n$ , for r different choices of hyperparameters (*e.g.* order of polynomials for polynomial regression, topology of neural networks, etc), henceforth leading to associated predictors  $\hat{f}_1, \ldots, \hat{f}_r$ . A natural question is to choose a single predictor out of these r candidates, to be used on the test set.

A standard procedure is to characterize the performance of the prediction on a validation set  $\mathcal{D}_{val}$  of size  $\alpha n$ ; more precisely, to choose the predictor which minimizes the validation error:

$$\min_{1 \leqslant i \leqslant r} \mathcal{R}_{\mathrm{val}}(\widehat{f_i}),$$

with

$$\mathcal{R}_{\mathrm{val}}(\widehat{f}_i) = \frac{1}{\alpha n} \sum_{(x_j, y_j) \in \mathcal{D}_{\mathrm{val}}} \ell\left(\widehat{f}_i(x_j), y_j\right).$$

This is similar to learning from a finite set  $\mathcal{F}$ , but this set  $\mathcal{F}$  is not fixed ahead in time as it depends on the training dataset  $\mathcal{D}_{\text{train}}$ . However, the sets  $\mathcal{D}_{\text{train}}$  and  $\mathcal{D}_{\text{val}}$  are independent, so that  $\mathcal{D}_{\text{val}}$  is independent of the set of predictors being considered. We can therefore use the same techniques as the ones used to obtain learning guarantees for finite sets  $\mathcal{F}$ .

**Theorem 10.1.** Consider an arbitrary set of predictors  $\mathcal{F} = \{f_1, \ldots, f_r\}$  (possibly depending on  $\mathcal{D}_{\text{train}}$ ), and assume that the elementary loss function is bounded, namely  $0 \leq \ell(y, z) \leq L < +\infty$  for any  $y, z \in \mathcal{Y}$ . Consider a validation set  $\mathcal{D}_{\text{val}}$  of size  $\alpha n$ , sampled independently of  $\mathcal{F}$ . Then, with probability at least  $1 - \delta$  over the choice of  $\mathcal{D}_{\text{val}}$ ,

$$\forall h \in \mathcal{F}, \qquad |\mathcal{R}(h) - \mathcal{R}_{\mathrm{val}}(h)| \leq \mathsf{L}\sqrt{\frac{1}{2\alpha n}\log\left(\frac{2|\mathcal{F}|}{\delta}\right)}.$$

*Proof.* The estimate (10.7) for a given element  $h \in \mathcal{F}$ , combined with the union bound, leads to

$$\mathbb{P}\left(\sup_{h\in\mathcal{F}}\left|\mathcal{R}(h)-\widehat{\mathcal{R}}_{\mathrm{val}}(h)\right| \ge \varepsilon\right) \le \sum_{h\in\mathcal{F}} \mathbb{P}\left(\left|\mathcal{R}(h)-\widehat{\mathcal{R}}_{\mathrm{val}}(h)\right| \ge \varepsilon\right) \le 2|\mathcal{F}|\mathrm{e}^{-2\alpha n\varepsilon^{2}/\mathsf{L}^{2}},$$

from which the proof is concluded similarly to the proof of Proposition 10.4.

#### 168 10 Complements: Elements of statistical learning

The interpretation of Theorem 10.1 is that the error on the validation set approximates the true error as long as  $\mathcal{F}$  is not too large. If one tries too many methods or values of the parameters (*i.e.* the number r of predictors  $\hat{f}_1, \ldots, \hat{f}_r$  is too large), then there is no good guarantee on the performance of the predictor. This corresponds to a situation of overfitting, where the training loss is typically small while the test loss is large.

In practice, one should avoid testing too many parameters – both because it is dangerous from the above discussion on the interpretation of Theorem 10.1, and also because it is time consuming. To this end, one should

- start with a rough coverage of parameter values (for instance using a coarse logscale for positive parameters);
- refine in a second step the parameter grid in the most relevant regions to fine tune the choice of the parameters;
- potentially, it may be a good idea not to optimize all hyperparameters and models at the same time, and proceed sequentially (for instance, set the hyperparameters of the training procedure in order to have an efficient training for extreme values of the other hyperparameters, and then look for optimal values of these other hyperparameters).

**Remark 10.3 (Learning curves).** The behavior of the training and validation losses as a function of the size of the training set is a useful criterion to determine how to improve the learning when the validation error is large, in particular when the training error is small but the validation error is large. In essence, does the issue come from n being too small (large estimation error) or the class  $\mathcal{F}$  being too small (large approximation error)? When the approximation error is large ( $\mathcal{F}$  is too small), the validation error does not decrease much as n increases. On the other hand, when the approximation error is small ( $\mathcal{F}$  is large enough), the validation error can be large for n small but then decreases as n is increased since the estimation error decreases, and one expects that the difference between the train and validation errors vanishes as  $n \to +\infty$ .

### 10.3 Statistical learning theory

The goal of learning theory is to provide guarantees of performance of unseen data. From a technical viewpoint, this can be described as relating the in-sample and out-of-sample errors. Recall that we work in these lecture notes in the statistical framework: the data points are assumed to be i.i.d. with respect to some unknown distribution  $p_{\text{data}}(dx \, dy)$ . This corresponds to considering a random set  $\mathcal{D}_{\text{train}} = \{(x_1, y_1), \ldots, (x_n, y_n)\} := \mathcal{D}_n(p_{\text{data}})$ . In this context, an algorithm  $\mathscr{A}$  is a mapping which associates to a dataset  $\mathcal{D}_n(p_{\text{data}})$  a function  $\mathcal{X} \to \mathcal{Y}$ . The associated expected risk is  $\mathcal{R}_{p_{\text{data}}}(f)$  with  $f = \mathscr{A}(\mathcal{D}_n(p_{\text{data}}))$ . The aim is to find  $\mathscr{A}$  such that

$$\mathcal{R}_{p_{\text{data}}}\left(\mathscr{A}(\mathscr{D}_n(p_{\text{data}}))\right) - \mathcal{R}^{\star} \tag{10.9}$$

is small. The challenge is that  $p_{\text{data}}$  is unknown and should be considered as arbitrary (so that ideally one should make minimalistic assumptions on it); and that the the risk is random since  $\mathscr{D}_n(p_{\text{data}})$  is random.

Metrics of performance. To quantify the performance, there are two main metrics:

• minimize the average error, namely the expectation of (10.9) with respect to realizations of  $\mathscr{D}_n(p_{data})$ :

$$\mathbb{E}\left[\mathcal{R}_{p_{\text{data}}}\left(\mathscr{A}(\mathscr{D}_{n}(p_{\text{data}}))\right)\right] - \mathcal{R}^{\star}; \tag{10.10}$$

• in the probably approximately correct (PAC) framework, one wants to ensure that (10.9) is smaller than  $\varepsilon$  with probability larger than  $1 - \delta$ :

$$\mathbb{P}\left[\mathcal{R}_{p_{\text{data}}}\left(\mathscr{A}(\mathscr{D}_n(p_{\text{data}}))\right) - \mathcal{R}^{\star} \leqslant \varepsilon\right] \ge 1 - \delta.$$

Typically, one fixes  $\delta \in (0,1)$  and seeks to minimize  $\varepsilon > 0$ , or fixes  $\varepsilon > 0$  and seeks to minimize  $\delta \in (0,1)$ .

An algorithm is consistent in expectation if (10.10) converges to 0 as  $n \to +\infty$ . It is PAC consistent if, for any  $\varepsilon > 0$ , the inequality (10.10) holds for any  $n \ge 1$  with  $\delta_n \to 0$  as  $n \to +\infty$ .

**Consistency.** Consistency can also be assessed over classes of problems, with an analysis that can be asymptotic or not. An algorithm is said to be universally consistent if it is consistent in expectation for all distributions  $p_{data}$  (but the rate of convergence will generically depend on  $p_{data}$  in view of the "no free lunch theorem"; see Theorem 10.2 below). Consistency can also be established for classes of distributions with some regularity properties (for instance with compact support, or leading to a Bayes predictor which is Lipschitz, etc). For such a given class  $\mathscr{P}$ , this amounts to finding an algorithm  $\mathscr{A}$  such that

$$\sup_{p \in \mathscr{P}} \mathbb{E} \left[ \mathcal{R}_p \left( \mathscr{A}(\mathscr{D}_n(p)) \right) \right] - \mathcal{R}^*$$

is as small as possible.

One can also consider minimax risks, which correspond to taking infima over  $\mathscr{A}$  in the previous metrics of convergence. Upper bounds on the minimax risk can be derived from the results obtained by studying one particular algorithm. Lower bounds are more difficult to establish.

No free lunch theorems. The spirit of various results going under the name "no free lunch" theorem is that there is no algorithm that works optimally for all distributions – *i.e.* learning is not possible without assumptions. We present here one possible result, which shows that, for any algorithm and any fixed number of samples  $n \ge 1$ , there is a probability distribution which makes the algorithm useless – *i.e.* no better than guessing at random ("chance level").

**Theorem 10.2 (No free lunch).** Consider binary classification with the elementary loss function  $\ell(y, z) = \mathbf{1}_{y \neq z}$  and a space of inputs  $\mathcal{X}$  of infinite cardinality  $(|\mathcal{X}| = +\infty)$ . Denote by  $\mathscr{P}$  the set of probability measures on  $\mathcal{X} \times \{0, 1\}$ . Then, for any learning algorithm  $\mathscr{A}$  and any integer  $n \ge 1$ ,

$$\sup_{p \in \mathscr{P}} \mathbb{E} \left[ \mathcal{R}_p \left( \mathscr{A}(\mathscr{D}_n(p)) \right) \right] - \mathcal{R}^* \ge \frac{1}{2}.$$

There are versions of such results with some uniformity in n, for instance statements such as: for any decreasing sequence  $(a_n)_{n \ge 1}$  with  $a_n \to 0$  as  $n \to +\infty$ , there exists  $p_{\text{data}} \in \mathscr{P}$  such that  $\mathbb{E}\left[\mathcal{R}_{p_{\text{data}}}\left(\mathscr{A}(\mathscr{D}_n(p_{\text{data}}))\right)\right] - \mathcal{R}^* \ge a_n$ .

*Proof.* We write the result for  $\mathcal{X} = \mathbb{N}$ . Fix  $k \in \mathbb{N}$  (this number will ultimately be sent to infinity). The proof proceeds in two steps: we first construct a probability distribution supported on k elements of  $\mathbb{N}$ , with  $k \gg n$  such that the knowledge of the n labels does not imply doing well on all k elements; and next choose the parameters of the distribution (characterized by the vector r below) by comparing the performance of the algorithm to the one of random guesses.

For the first step, we introduce a vector of labels  $r \in \{0,1\}^k$  and define a joint distribution  $p_r$ on (x, y) by

$$\forall j \in \{1, \dots, k\}, \quad \mathbb{P}(x = j, y = r_j) = \frac{1}{k}, \quad \forall j \in \{k + 1, \dots, n\}, \quad \mathbb{P}(x = j, \{0, 1\}) = 0.$$

This means that x is uniformly distributed over  $\{1, \ldots, k\}$  with labels  $r_x$  deterministically obtained from x. Since the relationship from inputs to labels is deterministic,  $\mathcal{R}^* = 0$ .

We next denote by

$$S(r) = \mathbb{E}\left[\mathcal{R}_{p_r}\left(\widehat{f}_{\mathscr{D}_n}\right)\right]$$

the average (over realizations of the data set) of the expected risk, with the predictor

$$\widehat{f}_{\mathscr{D}_n} = \mathscr{A}(\mathscr{D}_n(p_r)).$$

In order to maximize the quantity S with respect to  $r \in \{0, 1\}^k$ , we consider in fact a distribution q of possible elements in  $\{0, 1\}^k$  (for instance the uniform distribution). The maximum over all elements is larger that the average with respect to q. More precisely,

170 10 Complements: Elements of statistical learning

$$\begin{aligned} \max_{r \in \{0,1\}^k} S(r) &\geq \mathbb{E}_{r \sim q}[S(r)] = \mathbb{P}\left(\widehat{f}_{\mathscr{D}_n}(x) \neq r_x\right) \\ &= \mathbb{E}\left[\mathbb{P}\left(\widehat{f}_{\mathscr{D}_n}(x) \neq r_x \mid (x_1, r_1), \dots, (x_n, r_n)\right)\right] \\ &\geq \mathbb{E}\left[\mathbb{P}\left(\widehat{f}_{\mathscr{D}_n}(x) \neq r_x \text{ and } x \in \{x_1, \dots, x_n\} \mid (x_1, r_1), \dots, (x_n, r_n)\right)\right] \\ &= \frac{1}{2}\mathbb{E}\left[\mathbb{P}\left(x \in \{x_1, \dots, x_n\} \mid (x_1, r_1), \dots, (x_n, r_n)\right)\right] \\ &= \frac{1}{2}\prod_{i=1}^n \mathbb{P}(x_i \neq x \mid x) = \frac{1}{2}\left(1 - \frac{1}{k}\right)^n.\end{aligned}$$

where the second line is obtained by conditioning on  $\mathscr{D}_n$  and r, so that the conditional probability is over realizations of x; while in the third line the lower bound corresponds to realizations of x not in the data set, for which the label  $r_x$  was not observed, and so it has probability 1/2 to be 0 or 1 when averaging over  $r \sim q$ . Note that the random variables  $x_1, \ldots, x_n$  (realizations of the data set  $\mathscr{D}_n$ ), r (realizations of the labels) and x are independent. The desired lower bound is finally obtained in the limit  $k \to +\infty$ .

# Implementing and debuging machine learning programs

We refer for instance to [22, Chapter 11] for various practical advice on how to implement and debug machine learning programs. When predictions on the test set are poor, several options should be considered:

- increase or decrease the model capacity, depending on whether underfitting or overfitting is the issue. This can be achieved by playing with the parameters of the model (for instance increase or decrease the size of a neural network) and/or the regularization used;
- improve the quality of the optimization of the parameters of the model if the training error is large;
- debug the software implementation by looking at the worst errors, for instance wrong predictions obtained with a good confidence (think of a binary classification problem, where the probability to observe a class label would be close to 1 for an incorrect label);
- gather more data and/or clean it up. This can be assessed and quantified with learning curves, which report the predictive performance on the training and test sets as a function of the number of data points, possibly with a logarithmic scale on the number of data points;
- carefully select the hyperparameters of the method with refined grid searches and cross validations.

# References

- [1] Scikit-learn user guide. https://scikit-learn.org/stable/user\_guide.html.
- [2] G. ALAIN ET Y. BENGIO, What regularized auto-encoders learn from the data-generating distribution, J. Mach. Learn. Res 15(110) (2014) 3743–3773.
- [3] E. ALPAYDIN, *Introduction to Machine Learning*, 3 ed., Adaptive Computation and Machine Learning (MIT Press, 2014).
- [4] F. BACH, Learning Theory from First Principles (MIT Press, 2023). To appear.
- [5] P. BALDI ET K. HORNIK, Neural networks and principal component analysis: Learning from examples without local minima, *Neural Networks* 2 (1989) 53–58.
- [6] D. BARBER, Bayesian Reasoning and Machine Learning (Cambridge University Press, 2012).
- [7] G. BIAU ET E. SCORNET, A random forest guided tour, TEST 25 (2016) 197–227.
- [8] C. M. BISHOP, Pattern Recognition and Machine Learning (Information Science and Statistics) (Springer-Verlag, 2006).
- [9] H. BOURLARD ET Y. KAMP, Auto-association by multilayer perceptrons and singular value decomposition, *Biol. Cybern.* 59(4-5) (1988) 291–294.
- [10] C.-C. CHANG ET C.-J. LIN, LIBSVM: A library for support vector machines, ACM Transactions on Intelligent Systems and Technology 2 (2011) 1-27. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.
- [11] T. CHEN ET C. GUESTRIN. XGBoost: A scalable tree boosting system. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2016), KDD '16, Association for Computing Machinery, p. 785–794.
- [12] G. CYBENKO, Approximation by superpositions of a sigmoidal function, Math. Control Signal Systems 2 (1989) 303–314.
- [13] A. BELOTTO DA SILVA ET M. GAZEAU, A general system of differential equations to model first-order adaptive algorithms, *Journal of Machine Learning Research* **21** (2020) 1–42.
- [14] B. DAI ET D. WIPF. Diagnosing and enhancing VAE models. In International Conference on Learning Representations (2019).
- [15] T. DUCHAMP ET W. STUETZLE, Extremal properties of principal curves in the plane, Ann. Statist. 24(4) (1996) 1511–1520.
- [16] A. F. DUQUE, S. MORIN, G. WOLF, ET K. MOON. Extendable and invertible manifold learning with geometry regularized autoencoders. In 2020 IEEE International Conference on Big Data (Big Data), Atlanta, GA, USA (2020), pp. 5027–5036.
- [17] M. ESTER, H.-P. KRIEGEL, J. SANDER, ET X. XU. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining* (1996), KDD'96, AAAI Press, p. 226–231.
- [18] S. GERBER, Saddlepoints in unsupervised least squares, arXiv preprint 2104.05000 (2021).
- [19] S. GERBER, T. TASDIZEN, ET R. WHITAKER. Dimensionality reduction and principal surfaces via kernel map manifolds. In 2009 IEEE 12th International Conference on Computer Vision (2009), pp. 529–536.

#### 174 REFERENCES

- [20] S. GERBER ET R. WHITAKER, Regularization-free principal curve estimation, J. Mach. Learn. Res. 14 (2013) 1285–1302.
- [21] P. GHOSH, M. S. M. SAJJADI, A. VERGARI, M. BLACK, ET B. SCHOLKOPF. From variational to deterministic autoencoders. In *International Conference on Learning Representations* (2020).
- [22] IAN GOODFELLOW, YOSHUA BENGIO, ET AARON COURVILLE, Deep Learning (MIT Press, 2016). http://www.deeplearningbook.org.
- [23] T. HASTIE ET W. STUETZLE, Principal curves, J. Amer. Statist. Assoc. 84(406) (1989) 502-516.
- [24] T. HASTIE, R. TIBSHIRANI, ET J. FRIEDMAN, *The Elements of Statistical Learning*, 2nd ed., Springer Series in Statistics (Springer, 2009).
- [25] G. E. HINTON ET R. R. SALAKHUTDINOV, Reducing the dimensionality of data with neural networks, *Science* **313**(5786) (2006) 504–507.
- [26] K. HORNIK, M. STINCHCOMBE, ET H. WHITE, Multilayer feedforward networks are universal approximators, *Neural Networks* 2(5) (1989) 359–366.
- [27] K. IYER, R. BHALODIA, ET S. ELHABIAN, RENs: Relevance encoding networks, arXiv preprint 2205.13061 (2022).
- [28] K. JIA, L. SUN, S. GAO, Z. SONG, ET B. E. SHI, Laplacian auto-encoders: An explicit learning of nonlinear data manifold, *Neurocomputing* 160 (2015) 250–260.
- [29] L. KAUFMAN ET P. ROUSSEEUW. Clustering by means of Medoids. In Statistical Data Analysis Based on the L1-norm and Related Methods (1987), Y. Dodge, Ed., North-Holland, pp. 405–416.
- [30] D.P. KINGMA ET J. BA, Adam: A method for stochastic optimization, 3rd International Conference for Learning Representations (2014).
- [31] D. P. KINGMA ET M. WELLING. Auto-encoding variational Bayes. In 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings (2014), Y. Bengio and Y. LeCun, Eds.
- [32] D. P. KINGMA ET M. WELLING, An introduction to variational autoencoders, *Foundations* and *Trends®* in Machine Learning **12**(4) (2019) 307–392.
- [33] M. A. KRAMER, Nonlinear principal component analysis using autoassociative neural networks, AIChE Journal 37(2) (1991) 233-243.
- [34] Y. LEE, S. YOON, M. SON, ET F. C. PARK. Regularized autoencoders for isometric representation learning. In *International Conference on Learning Representations* (2022).
- [35] T. LELIÈVRE, T. PIGEON, G. STOLTZ, ET W. ZHANG, Analyzing multimodal probability measures with autoencoders, J. Phys. Chem. B (2024).
- [36] M. LESHNO, V. Y. LIN, A. PINKUS, ET S. SCHOCKEN, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, *Neural Networks* 6(6) (1993) 861–867.
- [37] G. LOUPPE, Understanding Random Forests: From Theory to Practice (PhD thesis, Université de Liège, 2014).
- [38] PANKAJ MEHTA, MARIN BUKOV, CHING-HAO WANG, ALEXANDRE G.R. DAY, CLINT RICHARDSON, CHARLES K. FISHER, ET DAVID J. SCHWAB, A high-bias, low-variance introduction to machine learning for physicists, *Physics Reports* 810 (2019) 1–124.
- [39] M. MOHRI, A. ROSTAMIZADEH, ET A. TALWALKAR, Foundations of Machine Learning, 2 ed. (MIT Press, 2018).
- [40] K. P. MURPHY, Probabilistic Machine Learning: An introduction (MIT Press, 2022).
- [41] A. NG. Sparse autoencoder. In CS294A Lecture notes (2011), volume 72, pp. 1–19.
- [42] M. A. NIELSEN, Neural Networks and Deep Learning (Determination Press, 2015).
- [43] Y. OPOCHINSKY, S.E. CHAZAN, S GANNOT, ET J. GOLDBERGER. k-Autoencoders deep clustering. In ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (2020), pp. 4037–4041.
- [44] H.-S. PARK ET C.-H. JUN. A simple and fast algorithm for K-medoids clustering. In Expert Systems with Applications (2009), volume 36.2, pp. 3336–3341.
- [45] S. RIFAI, P. VINCENT, X. MULLER, X. GLOROT, ET Y. BENGIO. Contractive auto-encoders: Explicit invariance during feature extraction. In *ICML* (2011), L. Getoor and T. Scheffer, Eds., Omnipress, pp. 833–840.
- [46] H. ROBBINS ET S. MONRO, A stochastic approximation method, Ann. Math. Statist. 22(3) (1951) 400–407.
- [47] B. SCHÖLKOPF, A. SMOLA, ET K.-R. MÜLLER, Nonlinear component analysis as a kernel eigenvalue problem, *Neural Comput.* 10(5) (1998) 1299–1319.
- [48] I. SEKKAT ET G. STOLTZ, Mini-batching error and adaptive Langevin dynamics, Journal of Machine Learning Research 24(329) (2023) 1–58.
- [49] S. SHALEV-SHWARTZ ET S. BEN-DAVID, Understanding Machine Learning: From Theory to Algorithms (Cambridge University Press, 2014).
- [50] Y. SUN, H. MAO, Q. GUO, ET Y. ZHANG, Learning a good representation with unsymmetrical auto-encoder, *Neural Comput. Appl.* 27 (2016) 1361–1367.
- [51] I. SUTSKEVER, J. MARTENS, G. DAHL, ET G. HINTON. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning* (Atlanta, Georgia, USA, 17–19 Jun 2013), S. Dasgupta and D. McAllester, Eds., volume 28 de *Proceedings of Machine Learning Research*, PMLR, pp. 1139–1147.
- [52] R. TIBSHIRANI, Principal curves revisited, Stat. Comput. 2 (1992) 183–190.
- [53] R. TIBSHIRANI, Regression shrinkage and selection via the Lasso, Journal of the Royal Statistical Society. Series B (Methodological) 58(1) (1996) 267–288.
- [54] P. VINCENT, H. LAROCHELLE, I. LAJOIE, Y. BENGIO, ET P.-A. MANZAGOL, Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion, J. Mach. Learn. Res 11(110) (2010) 3371–3408.
- [55] U. VON LUXBURG, A tutorial on spectral clustering, Stat. Comput. 17 (2007) 395–416.
- [56] X. XU, H. HOU, ET S. DING, Semi-supervised deep density clustering, Applied Soft Computing 148 (2023) 110903.
- [57] A. ZHANG, Z. C. LIPTON, M. LI, ET A. J. SMOLA, *Dive into Deep Learning* (Cambridge University Press, 2023).
- [58] J. H. ZHAO, P. L. H. YU, ET Q. JIANG, ML estimation for factor analysis: EM or non-EM?, Stat. Comput. 18 (2008) 109–123.