arXiv:1710.09230v1 [cs.LG] 25 Oct 2017

# Supervised Classification: Quite a Brief Overview

Marco Loog

Pattern Recognition Laboratory
Delft University of Technology
The Netherlands
e-mail: m.loog@tudelft.nl

The Image Section
University of Copenhagen
Denmark
http: prlab.tudelft.nl

August 11, 2017

**Abstract**

The original problem of supervised classification considers the task of automatically assigning objects to their respective classes on the basis of numerical measurements derived from these objects. Classifiers are the tools that implement the actual functional mapping from these measurements—also called features or inputs—to the so-called class label—or output. The fields of pattern recognition and machine learning study ways of constructing such classifiers. The main idea behind supervised methods is that of learning from examples: given a number of example input-output relations, to what extent can the general mapping be learned that takes any new and unseen feature vector to its correct class? This chapter provides a basic introduction to the underlying ideas of how to come to a supervised classification problem. In addition, it provides an overview of some specific classification techniques, delves into the issues of object representation and classifier evaluation, and (very) briefly covers some variations on the basic supervised classification task that may also be of interest to the practitioner.

**Keywords**: supervised learning, pattern recognition, machine learning, representation, classification, evaluation.

# Contents


**1 An Introduction** **3**
- 1.1 Learning, not Modelling . . . 3
- 1.2 An Outline . . . 4

**2 Classifiers** **4**
- 2.1 Preliminaries . . . 4
- 2.2 The Bayes Classifier . . . 6
- 2.3 Generative Probabilistic Classifiers . . . 7
- 2.4 Discriminative Probabilistic Classifiers . . . 9
- 2.5 Losses and Hypothesis Spaces . . . 10
  - 2.5.1 0-1 Loss . . . 11
  - 2.5.2 Convex Surrogate Losses . . . 11
  - 2.5.3 Particular Surrogate Losses . . . 12
- 2.6 Neural Networks . . . 13
- 2.7 Neighbors, Trees, Ensembles, and All That . . . 15
  - 2.7.1 $k$ Nearest Neighbors . . . 15
  - 2.7.2 Decision Trees . . . 15
  - 2.7.3 Multiple Classifier Systems . . . 16

**3 Representations and Classifier Complexity** **17**
- 3.1 Feature Transformations . . . 17
  - 3.1.1 The Kernel Trick . . . 18
- 3.2 Dissimilarity Representation . . . 19
- 3.3 Feature Curves and the Curse of Dimensionality . . . 20
- 3.4 Feature Extraction and Selection . . . 21

**4 Evaluation** **22**
- 4.1 Apparent Error and Holdout Set . . . 22
- 4.2 Resampling Techniques . . . 23
  - 4.2.1 Leave-one-out and $k$-Fold Cross Validation . . . 23
  - 4.2.2 Bootstrap Estimators . . . 23
  - 4.2.3 Tests of Significance . . . 24
- 4.3 Learning Curves and the Single Best Classifier . . . 24
- 4.4 Some Words about More Realistic Scenarios . . . 26

**5 Regularization** **28**

**6 Variations on Standard Classification** **29**
- 6.1 Multiple Instance Learning . . . 29
- 6.2 One-class Classification, Outliers, and Reject Options . . . 29
- 6.3 Contextual Classification . . . 30
- 6.4 Missing Data and Semi-supervised Learning . . . 30

6.5 Transfer Learning and Domain Adaptation . . . . . 31
6.6 Active Learning . . . . . 31

# 1 An Introduction

Consider the playful, yet in a way realistic problem of comparing apples and oranges. More precisely, let us consider the goal of telling them apart in an automated way, e.g. by means of a machine or, more specifically, a computer. How does one go about building such a machine? One approach could be to try and construct an accurate physiological model of both types of fruit. We gather what is known about their appearances, their cell structure, their chemical compounds, etc., together with any type of physiological laws that relate these quantities and their inner processes. Being presented by a new piece of fruit, we can then measure the quantities we think matter most and check these with our two models. A new piece of fruit for which we want to predict whether it is an apple or an orange is then best assigned to the class for which the model fits best. Clearly, the performance of this approach critically depends on issues such as how well we can build such models, how good we are at deciding what are the quantities that matter and how accurate we can measure these, how we actually decide whether we have good model fit, and so on. Such a model may, for instance, not be adequate at all when dealing with cases that are pathological from a physiological point of view, e.g. pieces of fruit that suffer from deformations or rot. Nevertheless, having enough understanding of the problem at hand, we may be able to tackle it in the way sketched.

## 1.1 Learning, not Modelling

Now, let us consider the, in a way, even more challenging problem of comparing foos and bars. Again, let us consider the goal of telling them apart in an automated way on the basis of particular measurements taken from the individual foos and bars. How are we going to go about our problem now? We might not even know exactly what we are dealing with here. Foos? Bars? So, where for the fruits we could try and build, for instance, a physiological model, we now do not even know whether physiology at all applies. Or is it physics that we need in this case to describe our objects? Chemistry maybe? Economics? Linguistics?

In the absence of any precise knowledge of our problem at hand[1], another approach to construct the asked-for machine that tells two or more object classes apart is by means of learning from examples. With this, we move away from any precisely interpretable model based on more or less factual knowledge about the object classes we are dealing with. As a substitute, we consider a different type of model—in a sense more general purpose—that can learn from examples. This is one of the premier learning settings that is studied in the

[1] Unfortunately, people may often not be able to recognize such absence. Work by Tversky, Kahneman, and others tells us, for instance, that people can suffer from systematic deviations from rationality or good judgment, i.e., cognitive biases [62].

fields of pattern recognition and machine learning[2]. The particular task is referred to as the supervised classification task: given a number of example input-output relations, can a general mapping be learned that takes any new and unseen feature vector to its correct class? That is: can we generalize from finite examples?

### 1.2 An Outline

The next section starts out with a basic introduction into classification technology and some of the underlying ideas. Classifiers are the aforementioned mappings, or functions, that take the measurements made on our objects that need assigning and aim to map these values to the correct corresponding output, or label as it is often referred to. Section 3 discusses a matter that typically warrants consideration prior even to the actual learning, or training, of the classifier: what measurements is the classifier that we are about to construct actually going to rely on when predicting the labels for new and unseen examples? In the same section, we more broadly cover the issue of object representation, which generally considers in what way to present our objects to the classifier. At this point, we can build various classifiers and have different possibilities to represent our objects, but really how good is the actual classifier built at our prediction task? Section 4 delves into the question of classifier evaluation, providing some tools to get insight into the behavior and performance of the machines that we have constructed. Two of the chapter's main conclusions are covered in this section: 1) there is no single best classifier and 2) there is an inherent tradeoff between the number of examples from which to generalize and classifier complexity. Section 5 then quickly introduces elementary regularization, i.e., another way of controlling classifier complexity. In conclusion, Section 6 mentions some variations to basic supervised classification.

## 2 Classifiers

Before we get to our selection of classifiers, we introduce some notation, make mathematically more precise what the scope is within which we operate, and describe what is the ultimate aim of a classifier in equally unambiguous terms.

### 2.1 Preliminaries

In the general classification setting considered, there are a total of $N$ labeled objects $o_i$, with $i \in \{1, \ldots, N\}$, to learn from. Every object $o_i$ is represented by a $d$-dimensional feature vector $x_i \in \mathbb{R}^d$ and has corresponding labels $y_i \in \{-1, +1\}$. So we assume we actually already have a numerical representation that makes up the $d$ measurements in every vector $x_i$. In addition, with the choice of $y_i \in \{-1, +1\}$, we limit ourselves here

[2]We do not care to elaborate on the possible distinctions between machine learning and pattern recognition. For this chapter, it is perfectly fine to consider them one and the same and we, indeed, are going to treat them as such.

to two-class classification problems for simplicity of exposition[3]. Although typically not stated explicitly, we should be aware that the pairs $(x_i, y_i)$ are assumed to be random yet i.i.d. draws from some underlying true distribution $p_{XY}$. For notational convenience, we also define $N_+$ and $N_-$ ($N_+ + N_- = N$), being the number of positive and negative samples, respectively. In general, we may use shorthand $+$ and $-$, rather than the slightly more elaborate $+1$ and $-1$, e.g. $N_+ = N_{+1}$.

A classifier $C$ is a function or mapping from the space of feature vectors $\mathbb{R}^d$ to the set of labels $\{-1, +1\}$. The input space on which $C$ is defined can be taken smaller than $\mathbb{R}^d$, for instance, if we know that some measurements can only take on continuous values between 0 and 1 or if they can only be integer. At the least, the input space is typically larger than the set of $N$ training points as the aim is for our classifier to generalize to new and unseen input vectors. Here we generally consider all of $\mathbb{R}^d$ as possible inputs. One way or the other, a classifier $C$ splits up the space in two sets. One in which points are assigned to $+1$ and one in which the class assignment is $-1$. The boundary between these two sets is generally referred to as the classification or decision boundary of $C$.

Finally, we define the objective that every classifier sets out to optimize. In a sense, this will also act as the ultimate measure to decide which of two or more classifiers performs the best. The de facto standard is the measure $\varepsilon$ that determines the fraction of misclassified objects when applying classifier $C$ to the problem described by $p_{XY}$[4]. Using Iverson brackets, we have

$$\varepsilon = \int_{\mathbb{R}^d} \sum_{y \in \{-1,+1\}} [C(x) \neq y] p_{XY}(x, y) dx \,. \tag{1}$$

This measure goes under various names like (true) error rate, classification error, 0-1 risk, and probability of misclassification. Then again, some may rather refer to it in terms of accuracy, which equals one minus the error rate. The lower $\varepsilon$ for a particular classifier the better we would say it is for the particular problem $p_{XY}$, as it will perform better in

[3]Every two-class classifier can, however, in a more or less principled way, be extended to the general multi-class case. One way is through combining multiple two-class classifiers (see, for instance, [45]). Other classifiers, in particular probabilistic ones, adapt more naturally to the multi-class case.

[4]This is not to say that other measures have not been studied. Depending on the actual goal, measures like AUC, $F$ score, or the $H$ measure have been considered as well. For both general and critical coverage of such measures refer to [8, 39, 49, 50, 68, 70, 98].

expectation[5].

## 2.2 The Bayes Classifier

One should understand that if we know $p_{XY}$, we are done. In that case, we can construct an optimal classifier $C$ that attains the minimum risk $\varepsilon^*$. But how do we get to that classifier? Let $C^*$ refer to this optimal classifier, fix the feature vector that we want to label to $x$, and consider the corresponding term within the integral of Equation (1):

$$\sum_{y\in\{-1,+1\}} [C^*(x) \neq y] p_{XY}(x,y) = [C^*(x) \neq +1] p_{XY}(x,+1) + [C^*(x) \neq -1] p_{XY}(x,-1)\,. \quad (2)$$

As $C^*$ should assign $x$ to $+1$ or $-1$, we see that the the choice that adds the least to the integral at this feature vector value is the assignment to that class for which $p_{XY}$ is largest. We reach an overall minimum if we stick to this optimal choice in every location[6] $x \in \mathbb{R}^d$. In case $p_{XY}(x,+1) = p_{XY}(x,-1)$, where $x$ is actually on the decision boundary, it does not matter what decision the classifier makes, as it will induce an equally large error. In other words, we can define $C^* : \mathbb{R}^d \to \{-1,+1\}$ as follows:

$$C^*(x) = \begin{cases} -1 & \text{if } p_{XY}(x,-1) > p_{XY}(x,+1) \\ +1 & \text{otherwise} \end{cases}\,. \quad (3)$$

Again using Iverson brackets, we could equally well write this as $C^*(x) = 1 - 2[p_{XY}(x,-1) > p_{XY}(x,+1)]$.

A possibly more instructive reformulation is by considering the conditional probabilities $p_{Y|X}$, often referred to as the posterior probabilities or simply the posteriors, instead of the full probabilities. Equivalent to checking $p_{XY}(x,-1) > p_{XY}(x,+1)$, we can verify whether $p_{Y|X}(-1|x) > p_{Y|X}(+1|x)$ and in the same vein as in Equation (3) decide to assign $x$ to

[5]Expanding a bit further on footnote 4, we remark that one of the more important settings, in which another performance measure or another way of evaluating may be appropriate, is in the case where the classification cost per class is different. Equation (1) tacitly assumes that predicting the wrong class incurs a cost of one, while predicting the right class comes at no cost. In many real-world settings, however, making the one error is not as costly as making the other. For instance, building a rotten fruit detector, classifying a fresh piece of fruit as rotten could turn out less costly than classifying a bad piece of fruit as good. When building an actual classifier, life often is even worse as one may not even know what the cost really is that will be incurred by a misclassification. This is one reason to resort to an analysis of the so-called receiver operating characteristic (ROC) curve and its related measure: the area under the ROC curve (AUC, an abbreviation mention already in the previous footnote). This curve and the related area provide, in some sense, tools to study the behavior of classifiers over all possible misclassification costs simultaneously.

Another important classification setting is the one in which there is a strong imbalance in class sizes, e.g. where we expect the one class to occur significantly much more often than the other class—a situation easily imagined in various applications. Also here analyses through ROC, AUC, and related techniques are advisable. For more on this topic, the reader is kindly referred to [39], [50], and related work.

[6]Of course, if we wish, we generally can decide otherwise on a set of measure 0 without doing any harm to the optimality of the classifier $C^*$.

$-1$ if this is indeed the case and assign it to $+1$ otherwise. The latter basically states that, given the observations made, one should assign the corresponding object to the class with the largest probability conditioned on those observations. Especially formulated like this, it seems like the obviously optimal assignment strategy.

The theoretical constructs $\varepsilon^*$ and $C^*$ are referred to as the Bayes error rate and the Bayes classifier, respectively. The former gives a lower bound on the best error we could ever achieve on the problem at hand. The latter shows us how to make optimal decisions once $p_{XY}$ is known. But these quantities are merely of theoretical importance indeed. In reality, our only hope is to approximate them, as the exact $p_{XY}$ will never be available to us. The objects we can work with are the $N$ draws $(x_i, y_i)$ from that same distribution. Based on these examples, we aim to build a classifier that generalizes well to all of $p_{XY}$. In all that follows in this chapter, this is the setting considered.

## 2.3 Generative Probabilistic Classifiers

The previous subsection showed that if we have $p_{XY}$, we can compare $p_{XY}(x, -1)$ and $p_{XY}(x, +1)$, and perform an optimal class assignment. One way to get a hold on $p_{XY}$ is by making assumptions on the form of the underlying class distributions $p_{X|Y}$ and estimate its free parameters from the training data provided. Such class-conditional probabilities describe the distribution of the underlying individual classes, i.e., they consider the distribution of $X$ given $Y$. Subsequently, these $p_{X|Y}$ can be combined with an estimate of the prior probability $p_Y$ of every class—i.e., an estimate of how often every class anyway occurs—to come to an overall estimate of $p_Y p_{X|Y} = p_{XY}$. These models are called generative, because, once they are fitted to the data, they allow us to generate new data from these class-conditional distributions, though the accuracy of this generative process of course heavily depends on the accuracy of the model fit.

A very common, and one of the more simple instantiations of a generative model, is classical linear discriminant analysis (LDA)[7], which assumes the classes to be normally distributed with different means and equal covariance matrices. Denoting the normal distribution with mean $\mu$ and covariance $\Sigma$ by $g(\cdot|\mu, \Sigma)$, the full probability model can be written as

$$p(x, y|\mu_{+1}, \mu_{-1}, \Sigma) = \begin{cases} \pi_+ g(x|\mu_+, \Sigma) & \text{if } y = +1 \\ \pi_- g(x|\mu_-, \Sigma) & \text{if } y = -1 \end{cases}, \tag{4}$$

where $\pi_+$ and $\pi_- = 1 - \pi_+$ are the class priors, which are in the $[0, 1]$ interval.

The above merely specifies the class of models that we consider, but it does not tell us how we fit it to the data that we have. A classical way to come to such parameters is to determine the maximum likelihood estimates. Other approaches, like maximum a posteriori and proper Bayesian estimation, are possible as well. Relying on the log-likelihood, the

[7]Some authors take the term linear discriminant analysis to refer to a more broadly defined class of classifiers. The normality-based classifier discussed here is also referred to by some as Fisher's linear discriminant or classical discriminant analysis.

objective function we find equals

$$\begin{aligned} L(\pi_+, \pi_-, \mu_+, \mu_-, \Sigma) &= \sum_{i=1}^{N} \log p(x_i, y_i | \mu_+, \mu_-, \Sigma) \\ &= \sum_{i:y_i=+1} (\log \pi_+ + \log g(x_i | \mu_+, \Sigma)) \\ &+ \sum_{i:y_i=-1} (\log \pi_- + \log g(x_i | \mu_-, \Sigma)) \ . \end{aligned} \tag{5}$$

Maximizing this leads to the well-known closed-form estimators

$$\hat{\pi}_+ = \frac{N_+}{N} \tag{6}$$

$$\hat{\pi}_- = \frac{N_-}{N} \tag{7}$$

$$\hat{\mu}_+ = \frac{1}{N_+} \sum_{i:y_i=+1} x_i \tag{8}$$

$$\hat{\mu}_- = \frac{1}{N_-} \sum_{i:y_i=-1} x_i \tag{9}$$

$$\hat{\Sigma} = \frac{1}{N} \sum_{i=1}^{N} (x_i - \hat{\mu}_{y_i})(x_i - \hat{\mu}_{y_i})^\top \, . \tag{10}$$

We note that, no matter what estimates one chooses, the L in the abbreviation LDA refers to the fact that the decision boundary forms a $(d-1)$-dimensional hyperplane, as can be checked by explicitly solving $\pi_+ g(x|\mu_+, \Sigma) = \pi_- g(x|\mu_-, \Sigma)$ for $x$ and finding that it takes on the form of a linear equation[8]:

$$(\hat{\mu}_+ - \hat{\mu}_-)^\top \hat{\Sigma}^{-1} x + c = 0 \, , \tag{11}$$

with $c$ a constant offset that depends on the dimensionality $d$, the determinant of $\hat{\Sigma}$, and the prior probabilities $\hat{\pi}_+$ and $\hat{\pi}_-$.

Again, these estimates are just one way to specify the free parameters in the model. For instance, $\Sigma$ could also have been estimated by an estimator such as $\frac{N}{N-2}\hat{\Sigma}$ or we could have applied Laplace smoothing to our prior estimates and, for instance, consider $(N_+ + 1)/(N + 2)$ instead of $N_+/N$. Of course, probably of a more dramatic influence is the actual choice of normal class-conditional models, which just may not be realistic. Even though having a misspecified model, as such, does not mean that the classifier will not perform well, more advanced and complex choices for these class-conditional distributions have been introduced that can potentially improve upon the use of normal distributions (see, for instance, [5, 51, 86, 101, 102], but let us also refer already to Subsection 4.3 for

[8] It actually is a nonhomogeneous linear or affine one, to be a bit more precise.

some important notes regarding classifier complexity in relation to training set size). As an example, we may substitute our relatively rigid parametric choice for a highly flexible nonparametric kernel density estimator per class, which leads to a classifier that is often referred to as the Parzen classifier.

## 2.4 Discriminative Probabilistic Classifiers

In the previous subsection, we decided to model $p_{XY}$, based on which we can then come to a decision on whether to assign $o_i$ to the $+$ or the $-$ class considering its corresponding feature vector $x_i$. Subsection 2.2, however, showed that we might as well use a model of $p_{Y|X}$ to reach a decision. Of course, from the full model $p_{XY}$, we can get to the conditional $p_{Y|X}$, while going into the other direction is not possible. For classification, however, we merely need to know $p_{Y|X}$ and so we can save the trouble of building a full model. In fact, if we are unsure about the true form of the underlying class-conditionals or the marginal $p_X$ that describes the feature vector distribution, directly modeling $p_{Y|X}$ may be wise, as we can avoid potential problems due to such model misspecification. On the other hand, if the full model is accurate enough this may have a positive effect on the classifier's performance [86, 104]. Approaches that directly model $p_{Y|X}$ are called discriminative as they aim to get straightaway to the information that matters to tell the one class apart from the other.

The classical model in this setting, and in a sense a counterpart of LDA, is called logistic regression[9]. One way to get to this model is to assume that the logarithm of the so-called posterior odds ratio takes on a linear form in $x$, i.e.,

$$\log \frac{p_{Y|X}(+1|x)}{p_{Y|X}(-1|x)} = w^\top x + w_\circ \,, \tag{12}$$

with $w \in \mathbb{R}^d$ and $w_\circ \in \mathbb{R}$. From this we derive that the posterior for the positive class takes on the following form:

$$p_{Y|X}(+1|x) = 1 - p_{Y|X}(-1|x) = \frac{\exp(w^\top x + w_\circ)}{1 + \exp(w^\top x + w_\circ)} = 1 - \frac{1}{1 + \exp(w^\top x + w_\circ)} \,. \tag{13}$$

The parameters $w$ and $w_\circ$ again are typically estimated by maximizing the log-likelihood. Formally, we have to consider the likelihood of the full model and not only of its posterior, but the choice of the necessary additional marginal model for $p_X$ is of no influence on the optimum of the parameters we are interested in [86, 87] and so we may just consider

$$(\hat{w}, \hat{w}_\circ) = \underset{(w,w_\circ)\in\mathbb{R}^{d+1}}{\operatorname{argmax}} \sum_{i=1}^{N} \log \left( \frac{y_i + 1}{2} - \frac{y_i}{1 + \exp(w^\top x_i + w_\circ)} \right) . \tag{14}$$

Note that, like LDA, this classifier is linear as well. Generally, the decision boundary is located at the $x$ for which $p_{XY}(x, +1) = p_{XY}(x, -1)$ or, similarly, for which $p_{Y|X}(+1|x) =$

[9] Again, there are various other names that refer to, at least more or less, the same approach, e.g. the logistic classifier, logistic discrimination, or logistic discriminant analysis.

$p_{Y|X}(-1|x)$. But the latter case implies that the log-odds equals 0 and so the decision boundary takes on the form

$$\hat{w}^\top x_i + \hat{w}_\circ = 0\,. \tag{15}$$

As for generative models, discriminative probabilistic ones come in all different kinds of flavors. A particularly popular and fairly general form of turning linear classifiers into nonlinear ones is discussed in Subsection 3.1. These and more variations can be found, among others, in [5, 51, 52, 86, 101, 102].

## 2.5 Losses and Hypothesis Spaces

As made explicit in Equations (11) and (15), both LDA and logistic regression lead to linear decision boundaries, i.e., $(d-1)$-dimensional hyperplanes in a $d$-dimensional feature space. The functional form of these decision boundaries is indirectly fixed through the choice of probabilistic model. Ultimately, they can be seen as mapping a feature vector $x$ in a linear way to a value on the real line. A subsequent decision to assign this point to the $+$ or $-$ class depends on whether the obtained value is smaller or larger than 0. This last operation, which basically turns the linear mapping into an actual classifier, can in essence be carried out by applying the sign function following the linear transform. This maps every number to a point in the set $\{-1, +1\}$, except for the points on the decision boundary which are mapped to 0[10]. The main difference between LDA and logistic regression is then the way the free parameters of their respective linear transformations have been obtained.

Especially within some areas of machine learning, it is common to altogether forget about the potentially probabilistic interpretation of a classifier. Instead, one explicitly defines the set $H$ of functions that transform any feature vector into a single number and which can subsequently be turned into actual classifiers by "signing" them. Next to that, one defines a measure $\ell : \mathbb{R} \times \{-1, +1\} \to \mathbb{R}$ of how good or bad a particular function fits to any point in the training data. The set of functions is often referred to as the hypothesis space, while the measure of fit is typically referred to as the loss (which is something that one minimizes). More specifically, the loss function $\ell$ takes in a predicted value $h(x)$, with $h \in H$, and decides how good this value matches the desired output $y$. Once, $H$ and $\ell$ are fixed, we can search for a best fitting $h^* \in H$,

$$h^* = \operatorname*{argmin}_{h \in H} \sum_{i=1}^{N} \ell(h(x_i), y_i)\,, \tag{16}$$

which can then be used to classify new and unseen object $x$ by assigning it to class $\operatorname{sign}(h^*(x))$.

The foregoing is a fairly general way of formulating the training of a classifier and not every choice for $H$ and $\ell$ may be equally convenient or suitable to be used on a

[10] As this is typically a set of measure zero, where the decision boundary gets mapped to does not really influence the performance of the classifier. Nevertheless, one may of course decide to have the boundary points mapped to $-1$ or $+1$, rather than 0.

classification problem. In the following, we present some of the better-known choices and briefly introduce some of the classifiers related to the particular options.

### 2.5.1 0-1 Loss

In a way, the obvious choice is to take the loss that we are actually interested in: the fraction of misclassified observations. Equation (1) defines this fraction, i.e., the classification error or the 0-1 risk, under the true distribution. Considering our finite number $N$ of training data, the best we can do is just count the number of incorrectly assigned samples:

$$L_{0\text{-}1}(h) := \sum_{i=1}^{N} \ell_{0\text{-}1}(h(x_i), y_i)\,, \tag{17}$$

with

$$\ell_{0\text{-}1}(a, b) := [\text{sign}(a) \neq b] \tag{18}$$

being the 0-1 loss[11].

A major problem with this loss is that finding the optimal $h^*$ is for many cases computationally very hard. Take for $H$, for example, all linear functions, then finding our $h^*$ turns out to be NP-hard and even settling for approximate solutions does not necessarily help [2, 56].

### 2.5.2 Convex Surrogate Losses

The most broadly used approach to get around the problem of the computational complexity that the 0-1 loss poses is to consider a relaxation of the original problem, which turns it into a convex optimization problem that is relatively easy to solve, e.g. by gradient descent or variations of this technique.

To achieve this, first of all, one chooses $H$ to be convex. The classical choice would be the space of linear functions, but more complex choices are possible as we will see later. The second step in the relaxation of the optimization problem is to turn to so-called surrogate losses. These are losses that aim to approximate $\ell_{0\text{-}1}$ as well as possible, but have some additional benefits. For instance, one could choose a loss that is everywhere differentiable, something that does not hold for the 0-1 loss. To turn the optimization for $h$ into a convex optimization, the function in Equation (17) needs to be convex as well. To generally do so, we approximate $\ell_{0\text{-}1}$ in Equation (18) by a convex function that upper-bounds this loss everywhere. The idea of the upper bound is that if we find the minimizer to it, we know that 0-1 loss one would achieve on the training set is certainly not larger than the surrogate risk the minimizer attains.

[11] We note here that the expected loss is often referred to as the risk, which explains why the error defined in Equation (1) is also referred to as the 0-1 risk.

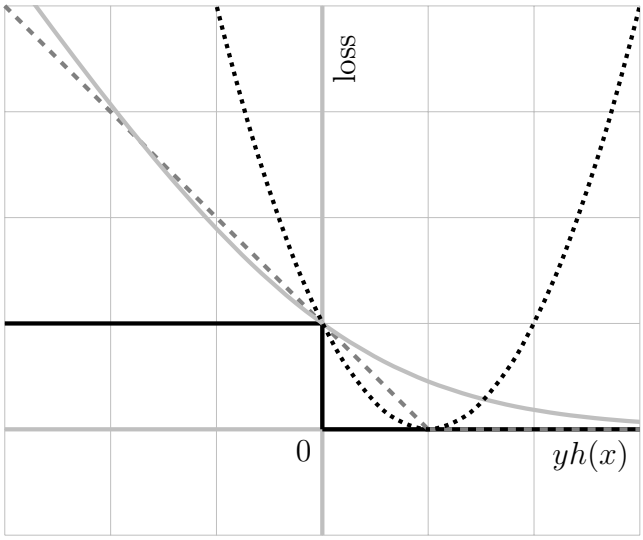


Figure 1: Plot of the 0-1 loss in terms $y$ times the predicted output $h(x)$ (in solid black) and three examples of surrogate losses that upperbound the 0-1 loss. Logistic loss is in solid light gray, hinge loss is in dashed dark gray, and squared loss is dotted black.

### 2.5.3 Particular Surrogate Losses

As we would decide on the label of a sample $x$ based on the output $h(x)$ that a trained classifier $h \in H$ provides, one typically only needs to consider what the loss function does for the value $yh(x)$. For instance, we can rewrite $\ell_{0\text{-}1}(a, b)$ as $\ell_{0\text{-}1}(a, b) = \ell_{0\text{-}1}(ba) = [b\,\mathrm{sign}(a) \neq 1]$ and achieve the same loss. In Figure 1 the shape of the 0-1 loss is plotted in these terms. The same figure shows various widely-used upper bounds for $\ell_{0\text{-}1}$.

Maybe the first one to note is the logistic loss, which is defined as

$$\ell_{\text{logistic}}(a, b) := \log_2\left(1 + e^{-ba}\right) . \tag{19}$$

The figure displays it as the solid light gray curve. Using this loss, in combination with a linear hypothesis class, leads to standard logistic regression as introduced in Subsection 2.4. So in this case, we have both a probabilistic view of the resulting classifier as well as an interpretation of logistic regression as minimizer of a specific surrogate loss. A second well-known classifier, or at least a basic form of it, is obtained by using the so-called hinge loss:

$$\ell_{\text{hinge}}(a, b) := \max(1 - ba, 0) . \tag{20}$$

This loss is at the basis of the support vector machine[12] (SVM) [6, 17, 21, 117]. A third classifier that fits the general formalism and is widely employed is obtained by using the squared loss function

$$\ell_{\text{squared}}(a, b) := (1 - ba)^2 . \tag{21}$$

Using again the set of linear hypotheses, we get basically what is, among others, referred to as the linear regression classifier, the least squares classifier, the least squares support vector machine, the Fisher classifier, or Fisher's linear discriminant [28,51,96,114]. Indeed, this classifier is a reinterpretation of the classical decision function introduced by Fisher [41] in the language of losses and hypotheses.

Finally, other losses one may encounter in the literature are the exponential loss $\exp(-ba)$, the truncated squared loss $\max(1-ba, 0)^2$, and the absolute loss $|1-ba|$. In Subsection 3.1, we introduce ways of designing nonlinear classifiers, which often rely on the same formalism as presented in this subsection.

## 2.6 Neural Networks

The use of artificial neural networks for supervised learning can be traced back at least to 1958. In that year, the perceptron was introduced [103], providing a linear classifier that could be trained using a basic iterative updating scheme for its parameters. Currently, neural networks are the dominant technique in many applications and application related areas when massively sized data sets to train from are available.

Even though the original formulation of the perceptron does not give a direct interpretation in terms of the solution to the optimizing of a particular loss within its hypothesis space of linear classifiers, such formulations are possible [23]. Neural networks that are employed nowadays readily follow the earlier sketched loss-hypothesis space paradigm. Possibly the most characteristic feature of neural networks is that the hypotheses considered are built up of relatively simple computational units called neurons or nodes. Such unit is a function $g : \mathbb{R}^q \to \mathbb{R}$ that takes in $q$ inputs and maps these to a single numerical output. Typically, $g$ takes on the form of a linear mapping followed by a nonlinear transfer or activation function $\sigma : \mathbb{R} \to \mathbb{R}$:

$$g(x) = \sigma(w^\top x + w_\circ) . \tag{22}$$

Often $\sigma$ is taken to be of sigmoidal shape, like the logistic function in Equation (13) used in logistic regression. This function squeezes the real line onto the interval $[0, 1]$, resembling a

[12]Together with the underlying theory, SVMs caused all the furore in the late 1990s and early 2000s. To many, the development of the SVM may still be one of the prime achievements of the mathematical field of statistical learning theory that started with Vapnik and Chervonenkis in the early 1970s. At least, SVMs are still one of the most widely known and used classifiers within the fields of pattern recognition and machine learning. Possibly one of the main feats of statistical learning theory was that it broke with the statistical tradition of merely studying the theory of the asymptotic behavior of estimators. Statistical learning theory is also concerned with the finite sample setting and makes, for instance, statements on the expected performance on unseen data for classifiers that have been trained on a limited set of examples [17, 23, 117].

smooth threshold function. Other choices are possible however. A choice popularized more recently, with a clearly different characteristic is the so-called rectified linear unit, which is defined as $\sigma(x) = \max(0, x)$ [89]. As for various of the previously mentioned classifiers, the free parameters are tuned by measuring how well $g$ fits the given training data. A widely used choice is the squared loss, but also likelihood based methods have been considered and links with probabilistic models have been studied [5, 52, 102].

One should realize that, whatever the choice of activation function, as long as it is monotonic using $g$ for classification will lead to a linear classifier. Nonlinear classifiers are constructed by combining various $g$s, both in parallel and in sequence. In this way, one can build arbitrarily large networks that can perform arbitrarily complex input-output mappings. This means that we are dealing with large and diverse hypothesis classes $H$. The general construction is that multiple nodes, connected in parallel, provide the inputs to subsequent nodes. Consider, for instance, the nonlinear extension where, instead of a single node $g$, as a first step, we have multiple nodes $g_1, \ldots, g_D$ that all receive the same feature vectors $x$ as input. In a second step, these $D$ outputs are collected by yet another node, $g : \mathbb{R}^D \to \mathbb{R}$ and transformed in a similar kind of way. So, all in all, we get a function $G$ of the form

$$G(x) = g\left(g_1(x), \ldots, g_D(x)\right) = \sigma\left(\sum_{i=1}^{D} w_i \sigma(w_{ij}^\top x + w_{i\circ}) + w_\circ\right). \tag{23}$$

To fully specify a particular $G$, one needs to set all the parameters in all $D + 1$ nodes. Once these are set, we can again use it to classify any $x$ to the sign of $G(x)$.

Of course, one does not have to stop at two steps. The network can have an arbitrary number of steps, or layers as they are typically referred to. Nowadays, so called deep[13] networks are being employed with hundreds of layers and millions of parameters. In addition to this, there are generally many different variations to the basic scheme we have sketched here [107]. By making different choices for the transfer function, by using multiple transfer functions, by changing the structure of the network, the number of nodes per layer, etc., one basically changes the hypothesis class that is considered. In addition, where in Subsection 2.5 the choice of $H$ and $\ell$ would typically be such that we end up with a convex optimization problem, using neural networks, we typically move away from optimization problems for which one can reasonably expect to be able to find the global optimum. As a result, to fully define the classifier, we should not only specify the loss and the hypothesis class, but also the exact optimization procedure that is employed. There are many choices possible to carry out the optimization, but most approaches rely on gradient descent or variations to this basic scheme [5, 7, 52, 122].

[13]There is, of course, not clear-cut definition at what number of layers makes a network deep. This is similar to wondering what amount of data makes it big or what number of features makes a problem high-dimensional.

## 2.7 Neighbors, Trees, Ensembles, and All That

There are a few approaches to classification that do not really fit the aforementioned settings, but that we feel do need brief mentioning either for historical reasons or for completeness.

### 2.7.1 *k* Nearest Neighbors

The nearest neighbor rule [22,42] is maybe the classifier with the most intuitive appeal. It is a widely used and classical decision rule and one of the earliest nonparametric classifiers proposed. In order to classify a new and unseen object, one simply determines the distances between its describing feature vector and the feature vectors in the training set and decides to assign the object to the same class the closest feature vector in that training set has. Most often, the Euclidean distance is used to determine the nearest neighbor in the training data set, but in principle any other, possibly even expert-designed or learned, distance measure can be employed.

A direct, yet worthwhile extension is to not only consider the closest sample in the training set, i.e., the first nearest neighbor, but to consider the closest $k$ and assign any unseen object to the class that occurs most often among these $k$ nearest neighbors. The $k$ nearest neighbor classifier, has various nice and interesting properties [22, 23, 42]. One of the more interesting ones may be the result that roughly states that with increasing numbers of training data, the $k$ nearest neighbor classifier converges to the Bayes classifier $C^*$, given $k$ increases at the appropriate rate.

### 2.7.2 Decision Trees

Classification trees or decision trees are classifiers that can be visualized by a hierarchical tree structure [11, 100]. Every observation that is classified traverses the tree's nodes to arrive at a leave that contains the final decision, i.e., assign label $+1$ or $-1$. In every node, a basic operation decides what next node, at a level deeper, will be visited.

Probably the simplest and most extensively used form of decisions that are made at every node are those that just check whether or not a single particular feature value is larger than a chosen threshold value. As there are only two outcomes possible at every step, there will only be two nodes available at every next level that can be reached. Using training data, there are various ways to come to an automated construction of the exact tree hierarchy, together with the decisions to make at the different nodes.

One of the possible benefits of these kinds of classifiers, especially of the last mentioned type of decision tree, given the tree is not too deep, is that one can easily trace back and interpret how the decision to assign a sample to class $-1$ or $+1$ was reached. One can retrace the steps through the tree and see what subsequent decisions lead to the class label provided.

### 2.7.3 Multiple Classifier Systems

The terms multiple classifier systems, classifier combining, and ensemble methods all refer to roughly the same idea: potentially more powerful classifiers can be built by combining two or more of them [63, 65, 97]. The latter are often referred to as base[14] classifiers. So, these techniques are not classifiers as such, but ways to compile base classifiers into classifiers that in some sense befit the data better. There can be various reasons to combine classifiers.

Sometimes a classifier turns out to be overly flexible and one may wish to stabilize the base classifier (see also Section 5). One way to do so is by a well-known combining technique called bagging [10], which trains various classifiers based on bootstrap samples of the same data set and assigns any new sample based on the average output of this often large set of base classifiers. Another way to construct different base classifiers is to consider random subspaces by sampling a set of different features for every base learner. This technique has been extensively exploited in random forests and the like [53, 54].

Combining classifiers can also be exploited when dealing with a problem where, in some sense, essentially different sets of features play a role. For instance, in the analysis of patient data, one might want to use different classifiers for high-dimensional genetic measurements and low-dimensional clinical data, as these sets may behave rather differently from each other. Once the two or more specialized classifiers have been trained, various forms of so-called fixed and trained combiners can be applied to come to a final decision rule [30, 65].

At times, the base classifiers can already be quite complex, possibly being a multiple classifier system in itself. Nice examples are available from medical image analysis [90] and recommender systems [59][15]. In these cases, advanced systems have been developed independently from each other. As a result, there is a fair chance that every systems has its own strengths and weaknesses and even the best performing individual system cannot be expected to perform the best in every part of feature space. Hence combining such systems can result in significantly improved overall performance.

Another reason to employ classifier combining is to integrate contextual features into the classification process. Such approaches can especially be beneficial when integrating contextual information into image and signal analysis tasks [18, 78, 82]. These techniques can be seen as a specific form of stacked generalization or stacking [124] and are becoming relevant again these days in the connection with deep learning (see, for instance, [44, 73, 112]).

Finally, we should mention boosting approaches to multiple classifier systems and in particular adaboost [43]. Boosting was initially studied in a more theoretical setting to show that so-called weak learners, i.e., classifiers that barely reach better performance than an error rate equal to the a priori probability of the largest class, could be combined into a strong learner to significantly improve performance over the weak ones. This research culminated in the development of a combining technique that sequentially adds base classifiers to the ensemble that has already been constructed, where the next base classifier

[14] That is, \'bās\rather than \'bāz-\.

[15] Though illustrative, strictly speaking, this work does not report on a classification task.

focuses especially on samples that previous base learners were unable to correctly classify. This last feature is the adaptive characteristic of this particular combining scheme that warrants the prefix ada-.

# 3 Representations and Classifier Complexity

In all of the foregoing, we tacitly assumed that we already had decided on what feature set to use for every object or, more generally, on how to represent every object before constructing our classification rule. In actual applications, there are, however, choices to be made.

To start with, we may be lucky and even have influence on what kind of raw measurements are going to be taken. Are we making color images with a standard camera of our oranges? Do we make a CT or an MR scan? Do we measure weight, diameter, plasticity? Will we sequence tissue samples? Clearly, the actual measurements that are most informative depend on the classification task at hand. More importantly, one should realize that once key measurements have been omitted[16], one can never recover the lost information through the building of a classifier, no matter how advanced or clever the classifier one designs is.

In reality, the practitioner often has little influence on what precise measurements are going to be made and s/he has to work with a predetermined set of initial features. But even in this setting there are considerations to be made. Are we going to use all of these features? Are we using them as is, or do we construct derivative features. Such decisions can hinge, say, on one's own insight into the problem or on possible external expert knowledge that one has acquired. Do we really need the RGB values of every pixel when describing our oranges or is an average value per channel enough? Do we at all need RGB or is a measure of orangeness sufficient? And do we need that weight measurement or does our consulting expert suggest that it does not contain any relevant information and can we do without?

We give a brief overview of ways to represent objects and some tools with which we can partially automate the process of selecting features and create derivatives. In addition, we discuss some of the—initially maybe counterintuitive—effects of using more and more measurements and how this roughly relates to the complexity of the classifier. We also introduce a tool possibly valuable in the analysis of the effect of the number of features: the so-called feature curve.

## 3.1 Feature Transformations

In the previous section, we saw various linear classifiers like logistic regression, SVMs, and Fisher's discriminant. At first sight, the linearity of a classifier may seem like a

[16]Such situation may of course arise from the fact that particular measurements simply cannot be realized for a number of reasons. For instance, measurements can be too expensive to consider as part of a realistic solution or they can be invasive to an unacceptable degree.

limitation, but this is actually easily removed. Nonlinear classifiers can be created by including nonlinear transformations of the original features into the original feature vector and subsequently train a linear classifier in this extended space. This essentially extends the hypothesis space $H$. For example, if $x_i \in \mathbb{R}^2$, then we can form new and nonlinear features by transforming the original individual features, e.g. we can transform the first feature $x_{i1}$ into $x_{i1}^3$ or $\sin x_{i1}$, and by combining individual features, e.g. we can form the product feature $x_{i1}x_{i2}$. Clearly, the possibilities are endless.

Let $\varphi : \mathbb{R}^d \to \mathbb{R}^D$ be the transformation that maps the original feature vector to its extended representation. Once we learned a classifier $C_D$ in the extended space, it induces a classifier in the original space through $C_d = C_D \circ \varphi$: we simply take every feature vector to be classified, transform it by $\varphi$, and apply the classifier trained in the higher-dimensional space. Using this construct, even if we would limit ourselves to linear classifiers in $\mathbb{R}^D$, we would typically find nonlinear decision boundaries in the original $d$-dimensional feature space.

### 3.1.1 The Kernel Trick

SVMs not only received a lot of attention as a result of statistical learning theory, the SVM literature also introduced what has become widely known as the kernel trick or kernel method [6], which has its roots in the 1960s [1]. The kernel trick allows one to extend many inherently linear approaches to nonlinear ones in a computationally simple way. At its basis is that—following the representer theorem [108,120]—many solutions for the type of optimization problems for linear classifiers that we have considered in Subsection 2.5 can be expressed in terms of a weighted combination of inner products of training feature vectors and the $x$ that is being classified, i.e.,

$$h^*(x) = w^\top x + w_\circ = \sum_{i=1}^{N} a_i x_i^\top x + a_\circ \,, \tag{24}$$

with $a_i \in \mathbb{R}$. Therefore, finding $h^*$ becomes equivalent to finding the optimal coefficients $a_i$.

After mapping the original feature vectors with $\varphi$, we would be optimizing the equivalent in the $D$-dimensional space to get to a possibly nonlinear classifier:

$$h^*(x) = w^\top x + w_\circ = \sum_{i=1}^{N} a_i \varphi(x_i)^\top \varphi(x) + a_\circ \,. \tag{25}$$

It becomes apparent that the only thing that matters in these settings is that we know how to compute inner products $k(z, x) := \varphi(z)^\top \varphi(x)$ between any two mapped feature vectors $x$ and $z$. The function $k$ is also referred to as a kernel function or simply a kernel. Of course, once we have explicitly defined $\varphi$, we can always construct the corresponding kernel functions, but the power of the kernel trick is that in many settings this can be avoided. This is interesting in at least two ways.

The first one is that if one wants to construct highly nonlinear classifiers, the explicit expansion $\varphi$ could grow inconveniently large. Take a simple expansion in which we consider all (unique) second degree monomials, which number equals $\binom{d}{2}$. So the dimensionality $D$ of the feature space in which we have to take the inner product grows as $O(d^2)$. By a direct calculation, one can however show that the inner product in this larger space can be expressed in terms of a much simpler $k$. In this case particularly, we have that[17]

$$\varphi(z)^\top \varphi(x) = (z^\top x)^2 \,. \tag{26}$$

As one can imagine, moving to nonlinearities of even higher degree, the effect becomes more pronounced[18]. At some point, explicitly expanding the feature vector nonlinearly becomes prohibitive, while calculating the induced inner product may still be easy to do. An extreme example is the radial basis function or Gaussian kernel defined by

$$k(z, x) = \exp\left(-\frac{1}{\sigma^2}||x - z||^2\right) , \tag{27}$$

which corresponds to a mapping that takes the original $d$-dimensional space to an infinite dimensional expansion [117].

A second reason why the formulation in terms of inner products is of interest is that it, in principle, allows us to forget about an explicit feature representation altogether. Going back to our original objects $o_i$, if we can construct a function $k(\cdot, \cdot) \mapsto \mathbb{R}_0^+$ that takes in two objects and fulfils all the criteria of an kernel, we can directly use $k(o_i, o)$ (with $o$ the object that we want to classify) as a substitute for $\varphi(x_i)^\top \varphi(x)$ in Equation (25). Once such a kernel function $k$ has been constructed—whether it is through an explicit feature space or not, one can use it to build classifiers.

All in all, kernel methods define a very general, powerful, and flexible formalism, which allows the design of problem specific kernels. Research into this direction has spawned a massive amount of publications about such approaches (see, for instance, [17, 110]).

## 3.2 Dissimilarity Representation

Any kernel $k$ provides, in a way, a similarity measure between two feature observations $x$ and $z$ (or possible directly between two objects): the larger the value is the more similar the two observations are. As $k$ has to act like an inner product that, at least implicitly, corresponds to some underlying feature space, limitations apply. In many settings, one might actually have an idea of a proper way to measure the similarity or the, in some sense equivalent, dissimilarity between two objects[19]. Possibly, such measure is provided

[17] This can be demonstrated by explicitly writing out both sides of the equation.

[18] We note that first degree monomials can also be included, either by explicitly including an additional feature to the original feature vector that is constant, say $c$, or implicitly by defining the inner product as $(z^\top x + c^2)^2$.

[19] Depending on the requirements one imposes upon dissimilarities (or, proximities, distances, etc.), similarities $s$ can be turned into dissimilarities $\delta$. For instance, by taking $\delta = \frac{1}{s}$ or $\delta = -s$. Next to these very basic transforms, there are various more advanced possibilities to construct such conversions [93].

by an expert that is working in the field where you are asked to build your classifier for. It therefore may be expected to be a well thought-through quantity that captures the essential resemblance of or difference between two objects.

The dissimilarity approach [31, 93, 95] allows one to build classifiers similar to kernel-based classifiers, but without some of the restrictions. One of the core ideas is that every objects can be represented, not by what one can see as absolute measurements that can be performed on every individual object, but rather by relative measurements that tells us how (dis)similar the object of interest is with a set of $D$ representative objects. These representative objects are also referred to as the prototypes. In particular, having such a set of prototypes $p_i$ with $i \in \{1, \dots D\}$, and having our favorite dissimilarity measure $\delta$, every object $o$ can be represented by the $D$-dimensional dissimilarity vector

$$x = (\delta(p_1, o), \dots, \delta(p_D, o))^\top . \tag{28}$$

Training, for instance, a linear classifier in this space leads to a hypothesis of the form

$$w^\top x + w_\circ = \sum_{i=1}^{N} a_i \delta(p_i, o) + w_\circ , \tag{29}$$

which should be compared to Equation (25).

The linear classifier is just one example of a classifier one can use in these $D$ dimensions. In this dissimilarity space, one can of course use the full range of classifiers that have been introduced in this chapter.

## 3.3 Feature Curves and the Curse of Dimensionality

Measuring more and more features on every object seems to imply that we gather more and more useful information on them. The worst that can happen is that we measure features that are partly or completely redundant, e.g. measuring the density, while we already have measured the mass and the volume. But once we have the information present in the features, it cannot vanish anymore. In a sense this is indeed true, but the question is whether we can still extract the information relevant with growing feature numbers. All classifiers rely on some form of estimation, which is basically what we do when we train a classifier, but estimation typically becomes less and less reliable when the space in which we carry it out grows[20]. The net result of this is that, while we typically would get improved performance with every additional feature in the very beginning, this effect will gradually wear off, and in the long run even leads to a deterioration in performance as soon as the estimates become unreliable enough. This behavior is what is often referred to as the curse of dimensionality [5, 32, 51, 61].

A curve that plots the performance of a classifier against an increasing number of features is called a feature curve [32]. It can be used as a simple analytic tool to get an idea of how sensitive our classifier is to the number of measurements that each object is

[20] One can say that, in this specific sense, the complexity of the classifier increases.

described with. Possibly of equal importance is that such curves can be used to compare two or more classifiers with each other. The forms feature curves take on depends heavily on the specific problem that we are dealing with, on the complexity of the classification method, the way this complexity relates to the specific problem, and on the number $N$ of training samples we have to train our classifier. As far as it is at all possible, the exact mathematical quantification of these quantities is a real challenge. Very roughly, one can state that the more complex a classifier is, the quicker its performance starts deteriorating with an increasing number of features. On the other hand: the more training data that is available, the later the deterioration in performance sets in. Also, the one classification technique is more complex than the other if the possible decision boundaries the former can model are more flexible or, similarly, less smooth. Another way to think about this is that the hypothesis of the former classification method is larger than the latter one[21].

## 3.4 Feature Extraction and Selection

The curse of dimensionality indicates that in particular cases it can be beneficial for the performance of our classifier to lower the feature dimensionality. This may be applicable, for instance, if one has little insight in the classification problem at hand, in which case one tends to define lots of potentially useful features and/or dissimilarities in the hope that at least some of them pick up what is important to discriminate between the two classes. Carrying out a more or less systematic reduction of the dimensionality after defining such large class of features can lead to acceptable classification results.

Roughly speaking, there are two main approaches [48,51,60,102]. The first one is feature selection and the second one is feature extraction. The former reduces the dimensionality by picking a subset from the original feature set, while the latter allows the combination of two or more features into fewer new features. This combination is often restricted to linear transformations, i.e., weighted sums, of original features, meaning that one considers linear subspaces of the original feature space. In principle, however, feature extraction also encompasses nonlinear dimensionality reductions. Feature selection is, by construction, linear, where the possible subspace is even further limited to linear spaces that are parallel to the feature space axes. Lowering the feature dimensionality by feature selection can also aid in interpreting classification results. At least it can shed some light on what features seem to matter mostly and possibly we can gain some insight into their interdependencies, for instance, by studying the coefficients of a trained linear classifier. Aiming for a more interpretable classifier, we might even sacrifice some of the performance for the sake of a really limited feature set size. Feature selection can also be used to select the right prototypes when employing the dissimilarity approach [94], as in this case every feature basically corresponds to all distances to a particular prototype.

[21] Vapnik was one of the first to concern himself with quantifying classifier complexity in a mathematically rigorous fashion. Details can be found in his book on learning theory: [117].

# 4 Evaluation

So how good are all these classification approaches really? How can we decide for one feature set over the other? How to compare two classifiers? What prototypes work and which do not? In Section 2.1, we already stated that the (true) error rate as defined in Equation (1) is, in our context, the ultimate measure that decides which procedure is best: the lower the error rate, the better. A first problem we face, the same one we encountered in dealing with the Bayes error and classifier—$\varepsilon^*$ and $C^*$, is that $p_{XY}$ is not available to us and so we are unable to determine the exact classification error for any procedure. Like in building a classifier, we have to rely on the $N$ samples from $p_{XY}$ that we have and we can merely come to an estimated error rate. But the situation is worse even. Not only do we have to give an estimate of our performance measure based on these $N$ samples—the simplest of ways would be to just count the number of misclassifications on that set and divide that number by $N$. Typically, we also have to use these same $N$ samples to train our classifier. In many real-world settings there simply is not more data available or there is no time or money left to gather more.

We consider some alternatives to estimate an error rate (see also [106]), introduce so-called learning curves that give some basic insight into classifier behavior, mention overtraining and in what sense there is no single best classifier, and offer some further considerations when it comes to developing a complete classifier system.

## 4.1 Apparent Error and Holdout Set

A major mistake, which is still being made among users and practitioners of pattern recognition and machine learning, is that one simply uses all available samples to build a classifier and then estimates the error on these same $N$ samples. This estimate is called the resubstitution or apparent error, denoted $\varepsilon^{\mathrm{A}}$ . The problem with this approach is that one gets an overly optimistic estimate. The classifier has been adapting to these specific points with these specific labels and therefore performs particularly well on this set. To more faithfully estimate the actual generalization performance of a classifier one would need a training set to train the classifier and a completely independent so-called test set[22] to estimate its performance. The latter is also referred to as the holdout set.

In reality, we often have only a single set at our disposal, in which case we can construct a training and a holdout set by splitting the initial set in two. But how do we decide on the sizes of these two sets? We are dealing with two conflicting goals here. We would like to train on as much data as possible, as this would typically give us the best performing classifier[23]. So the final classifier we would deliver, say, to a client, would be trained on

[22] In particular settings, when dealing with transductive learning for instance [117], the inputs may be available and exploited in the training phase. In that case, the training is still performed independent of the labels in the test set, which is the primary requirement. The related setting of semi-supervised learning is briefly covered in Subsection 6.4.

[23] Typically, yes, but classifiers can act counterintuitively and perform structurally worse with increasing numbers of labeled data in certain settings [81].

the full initial set available. But to get an idea of the true performance of this classifier—a possible selling point if low, we at least need some independent samples. The smaller we take this set, however, the less trustworthy this estimate will be. In the extreme case of one test sample, for instance, the estimate for error rate will always be equal to 0 or 1. But adding data to the test set will reduce the amount of data in the training set, which removes us further from the setting in which we train our final classifier on the full set. The following approaches, relying on resampling the training data, provide a resolution.

## 4.2 Resampling Techniques

We consider some resampling techniques that provide us with some possibilities to evaluate classifiers that, in real-world settings, are often more practical then using a holdout set.

### 4.2.1 Leave-one-out and *k*-Fold Cross Validation

Cross validation is an evaluation technique that offers the best of both worlds and allows us to both train and test on large data sets. Moreover, when it comes to estimation accuracy, so-called leave-one-out cross validation is probably one of the best options we have.

The latter approach loops through the whole data set for all $i \in \{1, \ldots, N\}$. At step $i$ the pair $(x_i, y_i)$ is used to evaluate the classifier that has been trained on all examples from the full set, except for that single sample $(x_i, y_i)$. So we have a training set of size $N - 1$ and a test set size of 1. Given that we want at least some data to test on, this is the best training set size we can have. The test set size is almost the worst we can have, but this is just for this single step in our loop. Going through all data available, every single sample will at some point act as a test set, giving us an estimated error rates $\varepsilon_i$ (all of value 0 or 1), which we can subsequently average to get to a better overall estimate

$$\varepsilon^{\mathrm{LOO}} = \frac{1}{N} \sum_{i=1}^{N} \varepsilon_i \,. \tag{30}$$

This procedure is called leave-one-out cross validation and its resulting estimate the leave-one-out estimate [35, 66, 86].

For computational reasons, e.g. when dealing with rather large data sets or classifiers that take long to train, one can consider to settle for so-called $k$-fold cross validation instead of its leave-one-out variant. In that case, the original data set is split in $k$, preferably, equal sized sets or $k$ folds. After this, the procedure is basically the same as with leave-one-out: we loop over the $k$ folds, which we consecutively leave out during training and which we then test on. Leave-one-out is then the same as $N$-fold cross validation.

### 4.2.2 Bootstrap Estimators

Bootstrapping is a common resampling technique in statistics, the basic version of which samples from the observed empirical distribution with replacement. Various bootstrap

estimators of the error rate aim to correct the bias—that is, the overoptimism, in the apparent error.

One of the more simple approaches proceeds as follows [35, 36]. From our data set of $N$ training samples, we generate $M$ bootstrap samples of size $N$ and calculate the $M$ corresponding apparent error rates $\varepsilon_i^{\mathrm{A}}$ for our particular choice of classifier. Using every time that same classifier, we also calculate the error $\varepsilon_i^{\mathrm{T}}$ rate on the total data set. An estimate of the bias is now given by their averaged difference.

$$\beta = \frac{1}{M} \sum_{i=1}^{M} \varepsilon_i^{\mathrm{A}} - \varepsilon_i^{\mathrm{T}} \,. \tag{31}$$

The bias corrected version of the apparent error, and as such an improved estimate of the true error, is now given by $\varepsilon^{\mathrm{A}} - \beta$.

Various improvements upon and alternatives to this scheme have been suggested and studied [35–37]. Possibly the best-known is the .632 estimator $\varepsilon^{.632} = 0.368\varepsilon^{\mathrm{A}} + 0.632\varepsilon^{\mathrm{O}}$. With the first term on the right-hand side the apparent error and the second term the out-of-bootstrap error. The latter is determined by counting all the samples from the original data set that are misclassified and that are not part of the current bootstrap sample based on which the classifier is built. Adding all these mistakes over all $M$ rounds and dividing this number by the total number of out-of-bootstrap samples, gives us $\varepsilon^{\mathrm{O}}$.

### 4.2.3 Tests of Significance

In all of the foregoing, it is of course important to remember that our estimates are based on random samples and so we will only observe an instantiation of a random variable. As a result, to decide on anything like a significant difference in performances of two or more classifiers, we have to resort to statistical tests [24]. To get a rough idea of the standard deviation of any of our error estimates $\varepsilon^{\mathrm{X}}$ based on a test size of size $N_{\mathrm{test}}$, we can for instance use a simple estimator that can be derived as the result of an averaging of $N_{\mathrm{test}}$ Bernoulli trials:

$$\mathrm{std}[\varepsilon^{\mathrm{X}}] = \sqrt{\frac{\varepsilon^{\mathrm{X}}(1-\varepsilon^{\mathrm{X}})}{N_{\mathrm{test}}}} \,. \tag{32}$$

## 4.3 Learning Curves and the Single Best Classifier

In Subsection 3.3, we briefly introduced feature curves, which give us an impression of how the error rate evolves with an increasing numbers of features. We discussed the curse of dimensionality in this context. It should be clear by now that also these feature curves can, like the error rate, only be estimated and to do so, one would typically apply the estimation techniques described in the foregoing. Another, maybe more important curve that provides us with insight into the behavior of a classification method is the so-called learning curve [20, 32, 69]. The learning curve plots the (estimated) true error rate against

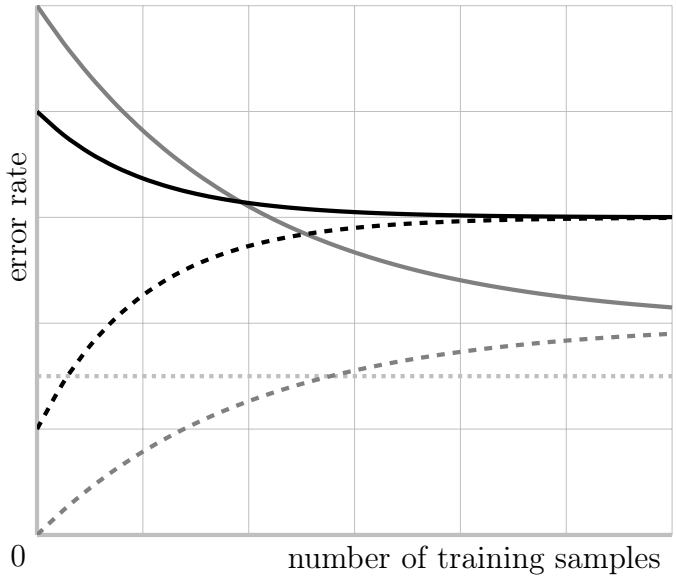


Figure 2: Stylized plots of the learning curves of two different classifiers. Solid lines are the (estimated) true errors for different training set sizes, while the dashed lines sketch the apparent errors. Black lines below to a classifier with relatively low complexity, while the gray lines illustrate the behavior of a more complex classifier. The light gray, dotted, horizontal line is the Bayes error for the problem considered.

the number of training samples. To complete the picture, one typically also plots the apparent error in the same figure.

Figure 2 displays stylized learning curves for two classifiers of different complexity. There are various characteristics of interest that we can observe in these plots and that reflect the typical behavior for many a classifier on many classification problems. To start with, with growing sample size, the classifier is expected to perform better in terms of the error rate[24]. In addition, for the apparent error we would typically observe the opposite behavior: the curve increases as it becomes more and more difficult to solve the classification problem for the growing training set[25]. In the limit of an infinite amount of data points, both curves come together[26]: the more training data one has, the better it describes the general data that we may encounter at test time and the closer to each other true error and apparent error get. In fact, the gap that we see is an indication that the trained classifier focuses too much on specifics that are in the training but not in the test set. This is called overtraining or overfitting. The larger the gap between true and apparent error is, the more overtraining has occurred.

[24]See, however, footnote 23.

[25]Clearly, there are exceptions. For instance, the nearest neighbor classifier generally gives a zero error rate on the training set.

[26]Again exceptions apply. Again the nearest neighbor is an example. Statistical learning theory formally studies settings for which such consistency does apply.

From the way that both learning curves for one classifier come together, one can also glean some insight. Classifiers that are less complex typically drop off more quickly, but also level out earlier than more complex ones. In addition, the former converges to an error rate that is most often above the limiting error rate of the latter: given enough data, one can get closer to the Bayes error when employing a more complex classifier[27]. As a result, it often is the case that one classifier is not uniformly better than another, even if we consider the same classification problem. It really matters what training set size we are dealing with and, when benchmarking the one classification method against the other, this should really be taken into account. Generally, the smaller the training data set is, the better it is to stick with simple classifiers, e.g. using a linear hypothesis class and few features.

The fact that the best choice of classifier may depend not only on the type of classification problem we need to tackle, but also on the number of training samples that we have at our disposal, may lead one to wonder what generally can be said about the superiority of one classifier over the other. Wolpert [123,125] (see also [27]) made this question mathematically precise and demonstrated that for this and several variations of this question the answer is that, maybe surprisingly, there does not exist such distinctions between learning algorithms. This so-called no free lunch theorem states, very roughly, that averaged over all classification problems possible, there is no one classification method that outperforms any other. Though the result is certainly gripping, one should interpret it with some care. It should be realized, for instance, that among all possible classification problems that one can construct, there probably are many that do not reflect any kind of realistic setting. What we can say, nevertheless, is that generally there is no single best classifier.

Finally, a learning curve may give us an idea of whether gathering more training data may improve the performance. In Figure 2, the classifier corresponding to the black curves can hardly be improved, even if we add enormous amounts of additional data. The other classifier, the gray curves, can probably improve a bit, reaching a slightly lower error rate when enlarging the traning set.

## 4.4 Some Words about More Realistic Scenarios

In real-world applications, designing and building a full classifier system will often be a process in which one may consider many feature representations, in which one will try various feature reduction schemes, and in which one will compare many different types of classifiers. On top of all that, there might be all kinds of preprocessing steps that are applied to the data (and that are not explicitly covered in this chapter). Working with images or signals for instance, one can perform various types of enhancement, smoothing, and normalization techniques that may have positive or negative effect on the performance of our final classifier.

A real problem in all this is that it is difficult to dispose of a truly independent test set. Unless one has a massive amount of labeled training data, one easily gets into the situation

[27]This also shows that complex classifiers typically have a smaller bias than simple classifiers.

that data that is also going to be used for evaluation leaks into the training phase. The estimated test errors are therefore overly optimistic and more so for complex classifiers than for simple ones. In the end, the result of this is that we may end up with a wrongly trained classifier, together with an overly optimistic estimate of its performance.

Let us consider some examples where things go wrong.

- A very simple instance is where one has decided, at some point, to use the $k$ nearest neighbor classifier. The only thing that remains to be done is finding the best value for $k$ and one decides to determine it on the basis of the performance for every $k$ on the test set. It may seem like a minor offense, but often there are many of such choices: the best number of features, the number of nodes in a layer of a neural network, the free parameters in some of the kernels, etc. (cf. [72] and, in particular point 7 in the list).

- Here is an example where it is maybe more difficult to see that one may have gone wrong. We decide to set up everything in a seemingly clean way. We prefix all classifiers that we want to study, all the feature selection schemes that we want to try, decide beforehand on all the kernels we want to consider, and all classifier combining schemes that we may want to employ. This gives a finite number of different classification schemes that we then compare based on cross validation. In the end, we then pick the scheme that provides the best performance. Even though this approach is actually fairly standard, again something does go wrong here. If the number of different classification schemes that we try out in this way gets out of hand, and it easily does, we still run the risk that we pick an overtrained solution with a badly biased estimate for its true error rate, especially when dealing with small training sets (cf. [9, 58, 105]).

- Even more complicated issues arise when multiple groups work on a large and challenging classification task. Nowadays, there are various research initiatives in which labeled data is provided publicly by a third party on which researchers can work simultaneously and collaboratively, but also in competition with each other. The flow of information and, in particular, the possibly indirect leakage of test data becomes difficult to oversee, let alone that we can easily correct for it when providing error estimates and corresponding confidence intervals or the like. How does one, for instance, correct for the fact that one's own method is inspired by some of the results by another group one has read about in the research literature? Though some statistical approaches are available that can alleviate particular problems [33, 34], it is safe to say that there currently is no generally applicable solution—if such at all exists.

Now the above primarily pertains to evaluation. In real scenarios, we of course also have to worry about the reproducibility and replicability of our findings. Otherwise, what kind of science would this be? Clearly, these are all issues that in one way or the other also play a significant role in other areas of research. In general, it turns out, however, that it

is difficult to control all of these aspects and that mistakes are made, mostly unwittingly but in some case possibly even knowingly. For some potential, more or less dramatic consequences, we refer to the following good reads: [29, 38, 57, 71, 88, 91].

# 5 Regularization

Regularization is actually a rather important yet relatively advanced topic in supervised learning [17, 96, 110, 117, 119] and unfortunately we are going to be fairly brief about it here.

The main idea of regularization is to have a means of performing complexity control. As we have seen already, classifier complexity can be controlled by the number of features that are used or through the complexity of the hypothesis class and, in a way, regularization is related to both of these. One of the well-known ways of regularizing a linear classifier is by constraining the already limited hypothesis space further. This is typically done by restricting the admissible weights $w$ of the linear classifier to a sphere with radius $t > 0$ around the origin of the hypothesis space, which means we solve the constraint optimization problem

$$\begin{aligned} h^* = & \underset{(w,w_\circ)\in\mathbb{R}^{d+1}}{\operatorname{argmin}} \sum_{i=1}^{N} \ell(w^\top x_i + w_\circ, y_i) \\ & \text{subject to } ||w||^2 \leq t\,. \end{aligned} \tag{33}$$

A formulation that is essentially equivalent is constructed by including the constraint directly into the objective function:

$$h^* = \underset{(w,w_\circ)\in\mathbb{R}^{d+1}}{\operatorname{argmin}} \sum_{i=1}^{N} \ell(w^\top x_i + w_\circ, y_i) + \lambda||w||^2\,, \tag{34}$$

where $\lambda > 0$ is known as the regularization parameter. The regularization is stronger with larger $\lambda$.

This procedure is the same as the one used in classical ridge regression [55] and effectively stabilizes the solution that is obtained. The effect of regularization is that the bias of our classification method increases, as we cannot reach certain linear classifiers anymore due to the added constraint. At the same time, the variance in our classifier estimates decreases due to the constraint (which is another way of saying that the classifier becomes more stable). In the average, with a small to moderate parameter $\lambda$, the worsening in performance we may get because of the increased bias is amply compensated with an improvement in performance due to the reduced variance, in which case regularization will lead to an improved classifier. If, however, we regularize too strongly, the bias will start to dominate and pull our model too far away from any reasonable solution at which point the true error rate will start to increase again.

A basic explanation of the effects of this so-called bias-variance tradeoff can already be found in the earlier mentioned work of Hoerl and Kennard [55]. The phenomenon can

be seen in various guises and its importance has been acknowledged early on in statistics and data analysis [118,119]. A more explicit dissection of the bias-variance tradeoff, in the context of learning methods, was published in [46]. The more complex a classifier is, the higher the variance we are faced with when training such model, and the more important some form of regularization becomes.

Equations (33) and (34) only consider the most basic form of regularization. There are many more variations on this theme. Among others, there are regularizers with built-in feature selectors [116] and regularizers that have deep connections to our earlier discussed kernels [47,113] .

# 6 Variations on Standard Classification

We have introduced and discussed the main aspects of supervised classification. In this last section, we like to review some slight variations to this basic learning problem. Though basic, there are very many decision problems that can actually be cast into a classification problem. Nevertheless, in reality, one may often be confronted with problems that still do not completely fit this restricted setting. Some of these we cover here.

## 6.1 Multiple Instance Learning

In particular settings, it is more appropriate or it simply is easier to describe every object $o_i$, not with a single feature vector $x_i$, but with a set of such feature vectors. This approach is, for example, common in various image analysis tasks, in which a set of so-called descriptors, i.e., feature vectors that capture the local image content at various locations in the image, act as the representation of that image. Every image, in both the training and the test set, is represented by such a set of descriptors and the goal is to construct a classifier for such sets. The research area that studies approaches applicable to this setting, in which every object can be described with sets of feature vectors having different sizes, but where the feature vectors are from the same measurement space, is called multiple instance learning. A large number of classification routines have been developed for this specific problem, which range from basic extensions of classifiers from the supervised classification domain by means of combining techniques, via dissimilarity-based approaches, to approaches specifically designed for the purpose of set classification [12, 15, 75, 85, 127]. The classical reference, in which the initial problem has been formalized, is [25].

## 6.2 One-class Classification, Outliers, and Reject Options

There are various problems where it is difficult to find sufficient examples of one of the classes, because they are very difficult to find or simply occur very seldom. In that case, one-class classification might be of use. Instead of trying to solve the two-class problem straightaway, it aims to model the distribution or support of the oft-occurring class accurately and based on that decides which points really do not belong to that class and,

therefore, will be assigned to the class of which little is known [109, 115]. Such techniques have direct relations to approaches that perform outlier or novelty detection in data and data streams [13,84] in which one aims to identify objects that are, in some sense, far away from the bulk of the data.

The more a test data point is an outlier, the less training data will be present in its vicinity and, therefore, the less certain a classifier will be in assigning the corresponding object to one or the other class. Consequently, outlier detection and related techniques are also used to implement so-called reject options [16]. These aim to identify points for which, say, $p_X$ is small and any automated decision by the classifier at hand is probably unreliable. In such case, the ultimate decision may be better left to a human expert. We might, for instance, be dealing with a sample from a third class; something that our classifier never saw examples of. This kind of rejection is also referred to as the distance reject option [26]. A second option is ambiguity rejection, in which case the classifier rather looks at $p_{Y|X}$ and leaves the final decision to a human expert if (in the two-class case) the two posteriors are very close to each other, i.e., approximately $\frac{1}{2}$ [26]. For ambiguity and distance rejection, one should realize that both an erroneous decision by the classifier and deploying a human expert come with their own costs. One of the main challenges in the use of a reject option is then to trade these two costs off in an optimal way.

## 6.3 Contextual Classification

Contextual classification has already been mentioned in Subsection 2.7.3 on multiple classifier systems. In these contextual approaches, samples are not classified in isolation, but they may have various types of neighborhood relations that can be exploited to improve the overall performance. The classical approach to this employs Markov random fields [3, 74] and specific variations to those techniques like conditional random fields [67]. The earlier mentioned methods using classifier combining techniques [18, 78, 82] are often more easily applicable and can leverage the full potential of more general classification methodologies. As already indicated, in Subsection 2.7.3 as well, the latter class of techniques seems to become relevant again in the context of nowadays deep learning approaches.

## 6.4 Missing Data and Semi-supervised Learning

In many real-world setting, missing data is a considerable and reoccurring problem. In the classification setting this means that particular features and/or class labels have not been observed. Missing features can occur because of the failure of a measurement apparatus, because of human non-response, or because the data was not recorded or got accidentally erased. There are various ways to deal with such deletions, which is a topic thoroughly studied in statistics [77].

The case of missing labels can have additional causes. It may simply have been too expensive to label more data or additional input data has been collected afterwards to extend the already available data, but the collector is not a specialist that can provide the necessary annotation. The case of missing label data is known within pattern recognition

and machine learning as semi-supervised learning [14, 128]. Also for this problem, which has been studied for over 50 years already, many different techniques have been developed. Though maybe more in a theoretical sense, there is still no completely satisfactory and practicable solution to the problem[28]. One of the major issues is the question to what extent one can guarantee that a supervised classifier can indeed be improved by taking all unlabeled data into account as well [64, 76, 79, 80].

## 6.5 Transfer Learning and Domain Adaptation

For various kinds of reasons, the distribution of the data at training time can be rather different from that at test time. Examples are medical devices that are trained on samples (subjects) from one country, while the machine is also deployed in another country. More generally, machines and sensors suffer from wear and tear and, as a result, measurement statistics at a later point in time may not really match their distribution at time of training. Depending on what can be assumed about the difference of the two domains or, as they are often referred to more specifically, the source and target, particular approaches can be employed that can alleviate the discrepancy between them [92, 99]. The areas of domain adaptation and transfer learning study techniques for these challenging settings. Depending on the actual transfer that has to be learned, more or less successful approaches can be identified to tackle these problems.

## 6.6 Active Learning

The final variation on supervised classification is actually concerned with regular supervised classification. The difference, however, with the main setting discussed throughout this chapter is that active learning sets out to improve the data collection process. It tries to answer various related questions, one of which is as follows. Given that we have a large number of unlabeled samples and a budget to label $N$ of these samples, what instances should we consider for labeling to enable us to train a better classifier than we would be able to in case we would rely on random sampling[29]? So can we in a more systematic way collect data to be labeled, such that we quicker come to a well-trained classifier?

The problem formulation has direct relations to sequential analysis [121] and optimal experimental design [40]. Overviews of current techniques can be found in [19], [111], and [126]. One of the major issues in active learning is that the systematic collection of labeled training data typically leads to a systematic bias as well. Correcting for this seems essential [4] (see also [83]). In a way, it points to a problem one will more generally encounter in practical settings and which directly relates to some of the issues indicated in Subsection 6.5: one of the key assumptions in supervised classification is that the training and test set consist of i.i.d. samples from the same underlying problem defined by the density $p_{XY}$. In reality, this assumption is most probably violated and care should be taken.

[28] In a way, this can probably also be said about many of the other problems discussed.

[29] The latter of which is the de facto standard.

## Acknowledgements

Many thanks go to Mariusz Flasiński (Jagiellonian University, Poland) and a second, anonymous reviewer for their critical, yet encouraging appraisal of this overview. Their comments helped to improve the exposition of the chapter and to drastically reduce the number of spelling and grammar glitches.

## References

[1] A Aizerman, Emmanuel M Braverman, and LI Rozoner. Theoretical foundations of the potential function method in pattern recognition learning. *Automation and Remote Control*, 25:821–837, 1964.

[2] Shai Ben-David, Nadav Eiron, and Philip M Long. On the difficulty of approximately maximizing agreements. *Journal of Computer and System Sciences*, 66(3):496–514, 2003.

[3] Julian Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 259–302, 1986.

[4] Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance weighted active learning. In *Proceedings of the 26th annual International Conference on Machine Learning*, pages 49–56. ACM, 2009.

[5] Christopher M Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.

[6] Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth Annual Workshop on Computational Learning Theory*, pages 144–152. ACM, 1992.

[7] Léon Bottou. Stochastic gradient learning in neural networks. In *Proceedings of Neuro-Nîmes 91*, Nimes, France, 1991. EC2.

[8] Andrew P Bradley. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997.

[9] Ulisses M Braga-Neto and Edward R Dougherty. Is cross-validation valid for small-sample microarray classification? *Bioinformatics*, 20(3):374–380, 2004.

[10] Leo Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

[11] Leo Breiman, Jerome Friedman, Charles J Stone, and Richard A Olshen. *Classification and Regression Trees*. CRC Press, 1984.

[12] Marc-André Carbonneau, Veronika Cheplygina, Eric Granger, and Ghyslain Gagnon. Multiple instance learning: A survey of problem characteristics and applications. *arXiv preprint arXiv:1612.03365*, 2016.

[13] Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 41(3):15, 2009.

[14] O Chapelle, B Schölkopf, and A Zien. *Semi-Supervised Learning*. MIT Press, 2006.

[15] Veronika Cheplygina, David MJ Tax, and Marco Loog. Multiple instance learning with bag dissimilarities. *Pattern Recognition*, 48(1):264–275, 2015.

[16] C Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1):41–46, 1970.

[17] Nello Christianini and John Shawe-Taylor. An introduction to support vector machines and other kernel-based learning methods. 2000.

[18] William W Cohen and Vitor R Carvalho. Stacked sequential learning. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, pages 671–676, 2005.

[19] David A Cohn, Zoubin Ghahramani, and Michael I Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, 4(1):129–145, 1996.

[20] Corinna Cortes, LD Jackel, Sara A Solla, Vladimir Vapnik, and John S Denker. Learning curves: Asymptotic values and rate of convergence. In *Proceedings of the 6th International Conference on Neural Information Processing Systems*, pages 327–334. Morgan Kaufmann Publishers Inc., 1993.

[21] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.

[22] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13(1):21–27, 1967.

[23] Luc Devroye, László Györfi, and Gábor Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag, 1996.

[24] Thomas G Dietterich. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7):1895–1923, 1998.

[25] Thomas G Dietterich, Richard H Lathrop, and Tomás Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artificial Intelligence*, 89(1):31–71, 1997.

[26] Bernard Dubuisson and Mylene Masson. A statistical decision rule with incomplete knowledge about classes. *Pattern Recognition*, 26(1):155–165, 1993.

[27] Richard O Duda, Peter E Hart, and David G Stork. *Pattern Classification*. Wiley-Interscience, 2001.

[28] RO Duda and PE Hart. *Pattern Classification and Scene Analysis*. John Wiley, 1973.

[29] Robert PW Duin. Superlearning and neural network magic. *Pattern Recognition Letters*, 15(3):215–217, 1994.

[30] Robert PW Duin. The combining classifier: To train or not to train? In *Proceedings of the 16th International Conference on Pattern Recognition*, volume 2, pages 765–770. IEEE, 2002.

[31] Robert PW Duin, Dick de Ridder, and David MJ Tax. Experiments with a featureless approach to pattern recognition. *Pattern Recognition Letters*, 18(11):1159–1166, 1997.

[32] RPW Duin and E Pękalska. *Pattern Recognition: Introduction and Terminology*. 37 Steps, 2015.

[33] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toni Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. In *Advances in Neural Information Processing Systems*, pages 2350–2358, 2015.

[34] Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248):636–638, 2015.

[35] Bradley Efron. *The Jackknife, the Bootstrap and Other Resampling Plans*. SIAM, 1982.

[36] Bradley Efron. Estimating the error rate of a prediction rule: Improvement on cross-validation. *Journal of the American Statistical Association*, 78(382):316–331, 1983.

[37] Bradley Efron and Robert Tibshirani. Improvements on cross-validation: The 632+ bootstrap method. *Journal of the American Statistical Association*, 92(438):548–560, 1997.

[38] Daniele Fanelli. Negative results are disappearing from most disciplines and countries. *Scientometrics*, 90(3):891–904, 2011.

[39] Tom Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.

[40] Valerii Vadimovich Fedorov. *Theory of optimal experiments*. Elsevier, 1972.

[41] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936.

[42] Evelyn Fix and Joseph L Hodges Jr. Discriminatory analysis-nonparametric discrimination: Consistency properties. Technical report, DTIC Document, 1951.

[43] Yoav Freund and Robert E Schapire. A desicion-theoretic generalization of on-line learning and an application to boosting. In *European Conference on Computational Learning Theory*, pages 23–37. Springer, 1995.

[44] Huan Fu, Chaohui Wang, Dacheng Tao, and Michael J Black. Occlusion boundary detection via deep exploration of context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 241–250, 2016.

[45] Mikel Galar, Alberto Fernández, Edurne Barrenechea, Humberto Bustince, and Francisco Herrera. An overview of ensemble methods for binary classifiers in multi-class problems: Experimental study on one-vs-one and one-vs-all schemes. *Pattern Recognition*, 44(8):1761–1776, 2011.

[46] Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural Computation*, 4(1):1–58, 1992.

[47] Federico Girosi, Michael Jones, and Tomaso Poggio. Regularization theory and neural networks architectures. *Neural Computation*, 7(2):219–269, 1995.

[48] Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.

[49] David Hand and Peter Christen. A note on using the F-measure for evaluating record linkage algorithms. *Statistics and Computing*, pages 1–9, 2017.

[50] David J Hand and Christoforos Anagnostopoulos. A better Beta for the H measure of classification performance. *Pattern Recognition Letters*, 40:41–46, 2014.

[51] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer, 2001.

[52] Geoffrey E Hinton. Connectionist learning procedures. *Artificial Intelligence*, 40(1-3):185–234, 1989.

[53] Tin Kam Ho. Random decision forests. In *Proceedings of the Third International Conference on Document Analysis and Recognition*, volume 1, pages 278–282. IEEE, 1995.

[54] Tin Kam Ho. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):832–844, 1998.

[55] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

[56] Klaus-Uwe Hoffgen, Hans-Ulrich Simon, and Kevin S Vanhorn. Robust trainability of single neurons. *Journal of Computer and System Sciences*, 50(1):114–125, 1995.

[57] John PA Ioannidis. Why most published research findings are false. *PLoS Med*, 2(8):e124, 2005.

[58] Anders Isaksson, Mikael Wallman, Hanna Göransson, and Mats G Gustafsson. Cross-validation and bootstrapping are unreliable in small sample classification. *Pattern Recognition Letters*, 29(14):1960–1965, 2008.

[59] Michael Jahrer, Andreas Töscher, and Robert Legenstein. Combining predictions for accurate recommender systems. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 693–702. ACM, 2010.

[60] Anil Jain and Douglas Zongker. Feature selection: Evaluation, application, and small sample performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):153–158, 1997.

[61] Anil K Jain, Robert PW Duin, and Jianchang Mao. Statistical pattern recognition: A review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(1):4–37, 2000.

[62] Daniel Kahneman, Paul Slovic, and Amos Tversky, editors. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, 1982.

[63] Josef Kittler and Fabio Roli, editors. *Multiple Classifier Systems: First International Workshop*, volume 1857 of *Lecture Notes in Computer Science*, Cagliari, June 2000. Springer.

[64] Jesse H Krijthe and Marco Loog. Projected estimators for robust semi-supervised classification. *Machine Learning*, 106(7):993–1008, 2017.

[65] Ludmila I Kuncheva. *Combining Pattern Classifiers: Methods and Algorithms*. John Wiley & Sons, 2004.

[66] Peter A Lachenbruch and M Ray Mickey. Estimation of error rates in discriminant analysis. *Technometrics*, 10(1):1–11, 1968.

[67] John Lafferty, Andrew McCallum, Fernando Pereira, et al. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the eighteenth International Conference on Machine Learning*, volume 1, pages 282–289, 2001.

[68] Thomas CW Landgrebe, Pavel Paclik, and Robert PW Duin. Precision-recall operating characteristic (P-ROC) curves in imprecise environments. In *18th International Conference on Pattern Recognition*, volume 4, pages 123–127. IEEE, 2006.

[69] Pat Langley. Machine learning as an experimental science. *Machine Learning*, 3(1):5–8, 1988.

[70] Nada Lavrač, Peter Flach, and Blaz Zupan. Rule evaluation measures: A unifying view. In *International Conference on Inductive Logic Programming*, pages 174–185. Springer, 1999.

[71] JT Leek and RD Peng. Statistics: p values are just the tip of the iceberg. *Nature*, 520(7549):612–612, 2015.

[72] Deborah Levine, Alexander A Bankier, and Elkan F Halpern. Submissions to radiology: Our top 10 list of statistical errors. *Radiology*, (2), 2009.

[73] Ke Li, Bharath Hariharan, and Jitendra Malik. Iterative instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3659–3667, 2016.

[74] Stan Z Li. *Markov Random Field Modeling in Image Analysis*. Springer Science & Business Media, 2009.

[75] Yan Li, David MJ Tax, Robert PW Duin, and Marco Loog. Multiple-instance learning as a classifier combining problem. *Pattern Recognition*, 46(3):865–874, 2013.

[76] Yu-Feng Li and Zhi-Hua Zhou. Towards making unlabeled data never hurt. In *Proceedings of the 28th ICML*, pages 1081–1088, 2011.

[77] Roderick JA Little and Donald B Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, 2014.

[78] Marco Loog. *Supervised Dimensionality Reduction and Contextual Pattern Recognition in Medical Image Processing*. Ph.D. thesis, Utrecht University, 2004.

[79] Marco Loog. Constrained parameter estimation for semi-supervised learning: the case of the nearest mean classifier. *European Conference on Machine Learning*, pages 291–304, 2010.

[80] Marco Loog. Contrastive pessimistic likelihood estimation for semi-supervised classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(3):462–475, 2016.

[81] Marco Loog and Robert PW Duin. The dipping phenomenon. In *Proceedings of the 2012 Joint IAPR International Conference on Structural, Syntactic, and Statistical Pattern Recognition*, pages 310–317. Springer-Verlag, 2012.

[82] Marco Loog and Bram van Ginneken. Supervised segmentation by iterated contextual pixel classification. In *16th International Conference on Pattern Recognition*, volume 2, pages 925–928. IEEE, 2002.

[83] Marco Loog and Yazhou Yang. An empirical investigation into the inconsistency of sequential active learning. In *23rd International Conference on Pattern Recognition*, pages 210–215. IEEE, 2016.

[84] Markos Markou and Sameer Singh. Novelty detection: A review—part 1: Statistical approaches. *Signal Processing*, 83(12):2481–2497, 2003.

[85] Oded Maron and Tomás Lozano-Pérez. A framework for multiple-instance learning. *Advances in Neural Information Processing Systems*, pages 570–576, 1998.

[86] Geoffrey McLachlan. *Discriminant Analysis and Statistical Pattern Recognition*. John Wiley & Sons, 2004.

[87] Tom Minka. Discriminative models, not discriminative training. Technical Report Technical Report MSR-TR-2005-144, Microsoft Research, 2005.

[88] Ramal Moonesinghe, Muin J Khoury, and A Cecile JW Janssens. Most published research findings are false—but a little replication goes a long way. *PLoS Med*, 4(2):e28, 2007.

[89] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning*, pages 807–814, 2010.

[90] Meindert Niemeijer, Marco Loog, Michael David Abramoff, Max A Viergever, Mathias Prokop, and Bram van Ginneken. On combining computer-aided detection systems. *IEEE Transactions on Medical Imaging*, 30(2):215–223, 2011.

[91] Silas Boye Nissen, Tali Magidson, Kevin Gross, and Carl T Bergstrom. Publication bias and the canonization of false facts. *Elife*, 5:e21451, 2016.

[92] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.

[93] Elżbieta Pękalska and Robert PW Duin. *The Dissimilarity Representation for Pattern Recognition: Foundations and Applications*, volume 64. World Scientific, 2005.

[94] Elżbieta Pękalska, Robert PW Duin, and Pavel Paclík. Prototype selection for dissimilarity-based classifiers. *Pattern Recognition*, 39(2):189–208, 2006.

[95] Elżbieta Pękalska, Pavel Paclik, and Robert PW Duin. A generalized kernel approach to dissimilarity-based classification. *Journal of Machine Learning Research*, 2:175–211, 2001.

[96] Tomaso Poggio and Steve Smale. The mathematics of learning: Dealing with data. *Notices of the AMS*, 50(5):537–544, 2003.

[97] Robi Polikar. Ensemble based systems in decision making. *IEEE Circuits and Systems Magazine*, 6(3):21–45, 2006.

[98] Foster Provost, Tom Fawcett, and Ron Kohavi. The case against accuracy estimation for comparing induction algorithms. In *Fifteenth International Conference on Machine Learning*, 1997.

[99] J Quiñonero-Candela, M Suyiyama, A Schwaighofer, and ND Lawrence. *Dataset Shift in Machine Learning*. MIT Press, 2009.

[100] J Ross Quinlan. Induction of decision trees. *Machine Learning*, 1(1):81–106, 1986.

[101] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[102] Brian D Ripley. *Pattern Recognition and Neural Networks*. Cambridge University Press, 2007.

[103] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386, 1958.

[104] Y Dan Rubinstein and Trevor Hastie. Discriminative vs informative learning. In *Third International Conference on Knowledge Discovery and Data Mining*, pages 49–53, 1997.

[105] Cullen Schaffer. Selecting a classification method by cross-validation. *Machine Learning*, 13(1):135–143, 1993.

[106] Rosa A Schiavo and David J Hand. Ten more years of error rate research. *International Statistical Review*, 68(3):295–310, 2000.

[107] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[108] Bernhard Schölkopf, Ralf Herbrich, and Alex J Smola. A generalized representer theorem. In *International Conference on Computational Learning Theory*, pages 416–426. Springer, 2001.

[109] Bernhard Schölkopf, John C Platt, John Shawe-Taylor, Alex J Smola, and Robert C Williamson. Estimating the support of a high-dimensional distribution. *Neural Computation*, 13(7):1443–1471, 2001.

[110] Bernhard Schölkopf and Alexander J Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2002.

[111] Burr Settles. Active learning literature survey. Computer Sciences TR 1648, University of Wisconsin, 2010.

[112] Abhinav Shrivastava and Abhinav Gupta. Contextual priming and feedback for faster R-CNN. In *European Conference on Computer Vision*, pages 330–348. Springer, 2016.

[113] Alexander J Smola and Bernhard Schölkopf. On a kernel-based method for pattern recognition, regression, approximation, and operator inversion. *Algorithmica*, 22(1):211–231, 1998.

[114] Johan AK Suykens and Joos Vandewalle. Least squares support vector machine classifiers. *Neural Processing Letters*, 9(3):293–300, 1999.

[115] David MJ Tax and Robert PW Duin. Support vector domain description. *Pattern Recognition Letters*, 20(11):1191–1199, 1999.

[116] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.

[117] Vladimir Naumovich Vapnik. *Statistical Learning Theory*. Wiley, 1998.

[118] G Wahba. Smoothing and ill-posed problems. In *Solution Methods for Integral Equations*, pages 183–194. Springer, 1979.

[119] Grace Wahba. Constrained regularization for ill posed linear operator equations, with applications in meteorology and medicine. Technical report, DTIC Document, 1981.

[120] Grace Wahba. *Spline Models for Observational Data*. SIAM, 1990.

[121] Abraham Wald. *Sequential Analysis*. Courier Corporation, 1973.

[122] Halbert White. Learning in artificial neural networks: A statistical perspective. *Neural Computation*, 1(4):425–464, 1989.

[123] David H Wolpert. On the connection between in-sample testing and generalization error. *Complex Systems*, 6(1):47–94, 1992.

[124] David H Wolpert. Stacked generalization. *Neural Networks*, 5(2):241–259, 1992.

[125] David H Wolpert. The lack of a priori distinctions between learning algorithms. *Neural Computation*, 8(7):1341–1390, 1996.

[126] Yazhou Yang and Marco Loog. A benchmark and comparison of active learning for logistic regression. *arXiv preprint arXiv:1611.08618*, 2016.

[127] Zhi-Hua Zhou. Multi-instance learning: A survey. Technical report, Department of Computer Science & Technology, Nanjing University, 2004.

[128] X Zhu. Semi-supervised learning literature survey. Computer Sciences TR 1530, University of Wisconsin, 2008.