

# Evaluation

Jianxin Wu

LAMDA Group

National Key Lab for Novel Software Technology  
Nanjing University, China  
wujx2001@gmail.com

March 16, 2017

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Accuracy and error in the simple case</b>	<b>2</b>
2.1	Training and test error . . . . .	3
2.2	Overfitting and underfitting . . . . .	4
2.3	Choose the hyperparameter(s) using validation . . . . .	6
2.4	Cross validation . . . . .	7
<b>3</b>	<b>Minimizing the cost / loss</b>	<b>9</b>
3.1	Regularization . . . . .	10
3.2	Cost matrix . . . . .	11
3.3	Bayes decision theory . . . . .	12
<b>4</b>	<b>Evaluation in imbalanced problems</b>	<b>12</b>
4.1	Rates inside individual class . . . . .	12
4.2	Area under the ROC curve . . . . .	14
4.3	Precision, Recall, and F-measure . . . . .	15
<b>5</b>	<b>Can we reach 100% accuracy?</b>	<b>17</b>
5.1	Bayes error rate . . . . .	17
5.2	Groundtruth labels . . . . .	18
5.3	Bias-Variance decomposition . . . . .	19
<b>6</b>	<b>Confidence in the evaluation results</b>	<b>22</b>
6.1	Why averaging? . . . . .	22
6.2	Why report the sample standard deviation? . . . . .	23
6.3	Compare two classifiers . . . . .	24

## 1 Introduction

It is very important to evaluate the performance of a learning algorithm or a recognition system. A systematic evaluation is often a comprehensive process, which may involve many aspects such as the speed, accuracy (or other accuracy-related evaluation metrics), scalability, reliability, etc. In this note, we focus on the evaluation of the accuracy and other accuracy-related evaluation metrics. We will also briefly discuss why perfect accuracy is impossible for many tasks, and where the errors come from.

## 2 Accuracy and error in the simple case

Accuracy is a widely used evaluation metric in classification tasks, which we have briefly discussed in the previous note. In laymen language, accuracy is the percentage of examples that are correctly classified. The error rate is the percentage of examples that are wrongly (incorrectly) classified, i.e., one minus accuracy. We use  $\text{Acc}$  and  $\text{Err}$  to denote accuracy and error rate, respectively. Note that both accuracy and error rate are between 0 and 1 (e.g.,  $0.49 = 49\%$ ). However, as aforementioned in the previous note, we have to resort to quite some assumptions to make this definition crystal clear.

Suppose we are dealing with a classification task, which means that the labels of all examples are taken from a set of categorical values  $\mathbb{Y}$ . Similarly, the features of all examples are from the set  $\mathbb{X}$ . We want to evaluate the accuracy of a learned mapping  $f : \mathbb{X} \mapsto \mathbb{Y}$ . We obey the i.i.d. assumption, that is, any example  $(\mathbf{x}, y)$  satisfy that  $\mathbf{x} \in \mathbb{X}$ ,  $y \in \mathbb{Y}$  and an example is drawn independently from the same underlying joint distribution  $p(\mathbf{x}, y)$ , written as  $(\mathbf{x}, y) \sim p(\mathbf{x}, y)$ . The mapping is learned on a training set of examples  $D_{\text{train}}$ , and we also have a set of test examples  $D_{\text{test}}$  for evaluation purposes. Under these assumptions, we can define the following error rates.

- The *generalization* error. This is the expected error rate when we consider all possible examples that follow the underlying distribution, i.e.,

$$\text{Err} = \mathbb{E}_{(\mathbf{x}, y) \sim p(\mathbf{x}, y)} \left[ \mathbb{I}(\mathbf{x} \neq y) \right], \quad (1)$$

in which  $\mathbb{I}[\cdot]$  is the indicator function. If we know the underlying distribution *and* the examples and labels can be enumerated (e.g., the joint distribution is discrete) or analytically integrated (e.g., the joint distribution is continuous), the generalization error can be explicitly computed. However, in almost all tasks, the underlying distribution is unknown. Hence, we often assume the generalization error exists as a fixed number between 0 and 1, but cannot be explicitly calculated. Our purpose is to provide a robust estimate of the generalization error.

- Approximate error. Because of the i.i.d. assumption, we can use a set of examples (denoted as  $D$ ) to *approximate* the generalization error, so long as  $D$  is sampled from the underlying distribution following the i.i.d. assumption,

$$\text{Err} \approx \frac{1}{|D|} \sum_{(\mathbf{x}, y) \in D} \llbracket f(\mathbf{x}) \neq y \rrbracket, \quad (2)$$

in which  $|D|$  is the number of elements in the set  $D$ . It is easy to verify that Equation 2 calculates the percentage of examples in  $D$  that is incorrectly classified.

- The *training* and *test* error. The training set  $D_{\text{train}}$  and the test set  $D_{\text{test}}$  are two obvious candidates for the  $D$  set in Equation 2. When  $D = D_{\text{train}}$ , the calculated error rate is called the training error; and when  $D = D_{\text{test}}$ , it is called the test error.

## 2.1 Training and test error

Why do we need a test set? The answer is simple: the training error is *not* a reliable estimate of the generalization error.

Consider the nearest neighbor classifier. It is easy to deduce that the training error will be 0, because the nearest neighbor of one training example must be itself.<sup>1</sup> However, we have good reasons to believe that NN search will lead to errors in most applications (i.e., the generalization error is larger than 0). Hence, the training error is too optimistic for estimating the generalization error of the 1-NN classifier.

The NN example is only an extreme case. However, in most tasks we will find that the training error is an *overly optimistic* approximation of the generalization error. That is, it is usually smaller than the true error rate. Because the mapping  $f$  is learned to fit the characteristics of the training set, it makes sense that  $f$  will perform better on  $D_{\text{train}}$  than on examples not in the training set.

The solution is use a separate test set  $D_{\text{test}}$ . Because  $D_{\text{test}}$  is sampled i.i.d., we expect the examples in  $D_{\text{train}}$  and  $D_{\text{test}}$  to be different from each with high probability (especially when the space  $\mathbb{X}$  is a large one). That is, it is unlikely that examples in  $D_{\text{test}}$  have been utilized in learning  $f$ . Hence, the test error (which is calculated based on  $D_{\text{test}}$ ) is a better estimate compared to the training error. In practice, we divide all the available examples into two *disjoint* sets: the training set and the test set, i.e.,  $D_{\text{train}} \cap D_{\text{test}} = \emptyset$ .

---

<sup>1</sup>If two examples  $(\mathbf{x}_1, y_1)$  and  $(\mathbf{x}_2, y_2)$  are in the training set, with  $\mathbf{x}_1 = \mathbf{x}_2$  but  $y_1 \neq y_2$  due to label noise or other reasons, the training error of the nearest neighbor classifier will not be 0. A more detailed analysis of this kind of irreducible error will be presented later in this note. For now we assume if  $\mathbf{x}_1 = \mathbf{x}_2$  then  $y_1 = y_2$ .

## 2.2 Overfitting and underfitting

Over- and under-fitting are two adverse phenomena when learning a mapping  $f$ . We use a simple regression task to illustrate these two concepts and a commonly used evaluation metric for regression.

The regression mapping we study is  $f : \mathbb{R} \mapsto \mathbb{R}$ . In this simple task, the mapping can be written as  $f(x)$ , where the input  $x$  and output  $f(x)$  are both real numbers. We assume the mapping  $f$  is a polynomial with degree  $d$ , i.e., the model is

$$y = f(x) + \epsilon = p_d x^d + p_{d-1} x^{d-1} + \dots + p_1 x + p_0 + \epsilon. \quad (3)$$

This model has  $d + 1$  parameters  $p_i \in \mathbb{R}$  ( $0 \leq i \leq d$ ), and the random variable  $\epsilon$  models the noise of the mapping, i.e.,  $\epsilon = y - f(x)$ . The noise  $\epsilon$  is independent of  $x$ . We then use a training set to learn these parameters. Note that the polynomial degree  $d$  is not considered as a parameter of  $f(x)$ , which is specified prior to the learning process. We call  $d$  a *hyperparameter*. After both the hyperparameter  $d$  and all the parameters  $p_i$  are specified, the mapping  $f$  is completely learned.

It is not difficult to find the optimal parameters for polynomial regression, which we leave as an exercise. In Figure 1, we show the fitted models for the same set of training data for different polynomial degrees. The data is generated by a degree 2 polynomial, as

$$y = 3(x - 0.2)^2 - 7(x - 0.2) + 4.2 + \epsilon, \quad (4)$$

in which  $\epsilon$  is i.i.d. sampled from a standard normal distribution  $N(0, 1)$ .

In Figure 1, the  $x$ - and  $y$ -axis represent  $x$  and  $y / f(x)$  values, respectively. One blue marker corresponds to one training example, and the purple curve is the fitted polynomial regression model  $f(x)$ . In Figure 1a,  $d = 1$  and the fitted model is a linear one. It is easy to observe that the true relationship between  $x$  and  $y$  is nonlinear. Hence, the capacity or representation power of the model (degree 1 polynomial) is lower than the complexity in the data. This phenomenon is called *underfitting*, and we expect large errors when underfitting happens. As shown by the large discrepancy between the purple curve and the blue points, large regression errors indeed occur.

Now we move on to Figure 1b (with  $d = 15$ ). The learned polynomial curve is a complex one, with many upward and downward transitions along the curve. We observe that most blue points are very close to the purple curve, meaning that the fitting error is very small *on the training set*. However, since the data are generated using a degree 2 polynomial, its generalization (or test) error will be large, too. A higher order polynomial has larger representation power than a lower order one. The extra capacity is often used to model the noise in the examples (the normal noise in this case) or some peculiar property in the training set (which is not a property for the underlying distribution). This phenomenon is called *overfitting*. Nowadays we are equipped with more computing resources, which means we can afford to use models with large capacity. Overfitting has

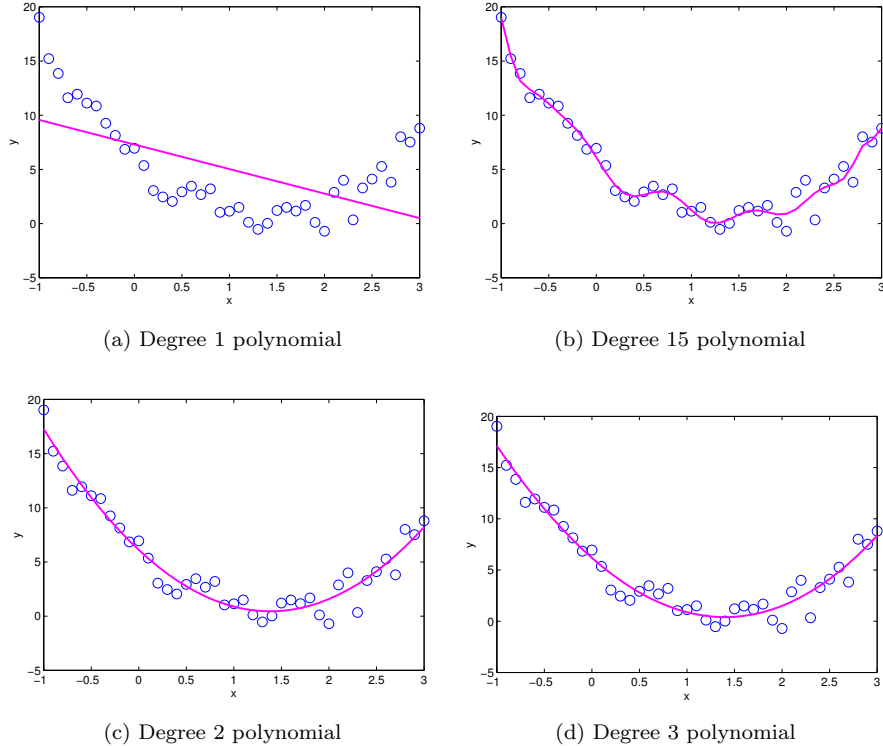


Figure 1: Fitting data with polynomials with different degrees of freedom.

become a more serious issue than underfitting in many modern algorithms and systems.

Figure 1c uses the correct polynomial degree ( $d = 2$ ). It strikes a good balance between minimizing the training error and avoiding overfitting. As shown by this figure, most blue points are around the purple curve, and the curve is quite smooth. We expect it to also achieve small generalization error.

In Figure 1d, we set  $d = 3$ , which is slightly larger than the correct value ( $d = 2$ ). There are indeed differences between the curves in Figure 1c and Figure 1d, although the differences are almost indistinguishable. We also listed the true and learned parameter values in Table 1. Both learned sets of parameters match the true values quite well. The  $p_3$  parameter for the degree 3 polynomial is close to 0, which indicates that the term  $p_3x^3$  has very limited impact for predicting  $y$ . In other words, when the model capacity is only slightly higher than the data's true complexity, a learning algorithm has the potential to still properly fit the data. However, when the gap between model and data complexity is large (cf. Figure 1b), overfitting will cause serious troubles.

Table 1: True (first row) and learned parameters of polynomial regression (second and third row).

	$p_3$	$p_2$	$p_1$	$p_0$
True		3.0000	-8.2000	6.8000
$d = 2$		2.9659	-9.1970	6.1079
$d = 3$	0.0489	2.8191	-8.1734	6.1821

### 2.3 Choose the hyperparameter(s) using validation

Unlike the parameters which can be learned based on the training data, hyperparameters are usually not learnable. Hyperparameters are often related to the capacity or representation power of the mapping  $f$ . A few examples are listed as follows.

- $d$  in polynomial regression. It is obvious that a larger  $d$  is directly related to higher representation power.
- $k$  in  $k$ -NN. A large  $k$  value can reduce the effect of noise, but may reduce the discriminative power in classification. For example, if  $k$  equals the number of training examples, the  $k$ -NN classifier will return the same prediction for any example.
- Bandwidth  $h$  in kernel density estimation (KDE). We will introduce KDE in the note on probabilistic methods. Different from the above examples, there is a theoretical rule that guides the choice of optimal bandwidth.
- $\gamma$  in the RBF kernel for SVM. We will introduce the RBF kernel in the SVM note. The choice of this hyperparameter, however, is neither learnable from data nor guided by a theoretical result.

Hyperparameters are often critical in a learning method. Wrongly changing the value of one hyperparameter (e.g.,  $\gamma$  in the RBF kernel) may dramatically reduce the classification accuracy (e.g., from above 90% to less than 60%). We need to be careful about hyperparameters.

One widely adopted approach is to use a validation set. The validation set is provided in some tasks, in which we want to ensure that the training, validation and test sets are disjoint. In scenarios where this split is not provided, we can randomly split the set of all available data into three disjoint sets:  $D_{\text{train}}$ ,  $D_{\text{val}}$  (the validation set) and  $D_{\text{test}}$ .

Then, we can list a set of representative hyperparameter values, e.g.,  $k \in \{1, 3, 5, 7, 9\}$  for  $k$ -NN, and train a classifier using the training set and one of the hyperparameter value. The resulting classifier can be evaluated using the *validation* set to obtain a validation error. For example, we will obtain 5 validation  $k$ -NN error rates, one for each hyperparameter value. The hyperparameter value that leads to the smallest validation error rate will be used as the chosen hyperparameter value.

In the last step, we will train a classifier using  $D_{\text{train}}$  and the chosen hyperparameter value to train a classifier, and its performance can be evaluated using  $D_{\text{test}}$ . When there are ample examples, this strategy often works well. For example, we can use  $\frac{n}{10}$  examples for testing when there are  $n$  examples in total. In the remaining  $\frac{9n}{10}$  examples, we can use  $\frac{n}{10}$  for validation and the rest  $\frac{8n}{10}$  for training. The ratio of sizes between train/validation or train/test is not fixed, but we usually set aside a large portion of examples for training. If a split between training and test examples have been pre-specified, we can split the training set into a smaller training and a validation set when necessary.

The validation set can be used for other purposes, too. For example, the parameters of neural networks (including the convolutional neural network) are learned iteratively. The neural network model (i.e., values of its parameters) are updated in every iteration. However, it is not trivial to determine when to stop the iterative updating procedure. Neural networks have very large capacity. Hence, the training error will continue to decrease. However, after some iterations the new updates may be mainly fitting the noise, which leads to overfitting and the validation error will start to increase.<sup>2</sup>

One useful strategy is to evaluate the model's validation error. The learning process can be terminated if the validation error becomes stable (not decreasing in a few consecutive iterations) or even starts to increase, that is, when it starts to continuously overfit. Figure 2 illustrates the interaction between training and validation error rates. Note that this is only a schematic illustration. In practice, although the overall trend is that the training error is to reduce when more iterations are updated, it can have ups and downs in a short period. The validation error is also possible to be (slightly) smaller than the training error before overfitting occurs.

## 2.4 Cross validation

The validation strategy works well when there are a large number of examples. For example, if one million ( $n = 10^6$ ) examples are available, the validation set has 100,000 examples ( $\frac{n}{10}$ ), which is enough to provide a good estimate of the generalization error. However, when  $n$  is small (e.g.,  $n = 100$ ), this strategy may cease to work if we use a small portion of examples for validation (e.g.,  $\frac{n}{10} = 10$ ). If we set aside a large portion as the validation set (e.g.,  $\frac{n}{2} = 50$ ), the training set will be too small to learn a good classifier.

Cross-validation (CV) is an alternative strategy for small size tasks. A  $k$ -fold CV *randomly* divides the training set  $D_{\text{train}}$  into  $k$  (roughly) *even* and *disjoint*

---

<sup>2</sup>Note that the validation and training sets are disjoint. Hence, both the validation error and the test error rates will increase if overfitting happens.




Figure 2: Illustration of the overfitting phenomenon. One unit in the  $x$ -axis can be one or more (e.g., 1000) updating iterations.

subsets  $D_1, D_2, \dots, D_k$ , such that

$$\bigcup_{i=1}^k D_i = D_{\text{train}}, \quad (5)$$

$$D_i \cap D_j = \emptyset, \quad \forall 1 \leq i \neq j \leq k, \quad (6)$$

$$|D_i| \approx \frac{|D_{\text{train}}|}{k}, \quad \forall 1 \leq i \leq k. \quad (7)$$

The CV estimate of the error rate then operates following Algorithm 1.

In Algorithm 1, the learning algorithm is run  $k$  times. In the  $i$ -th run, all examples except those in the  $i$ -th subset  $D_i$  are used for training, while  $D_i$  is used to compute the validation error  $\text{Err}_i$ . After all the  $k$  runs, the cross-validation estimate is the average of  $k$  validation error  $\text{Err}_i$  ( $1 \leq i \leq k$ ). Note that every example is used exactly once to compute one of the validation error rates, and at the same time is used  $k - 1$  times for training classifiers. In other words, the CV strategy makes good use of the small number of examples, and hence may provide a reasonable estimate of the error rate.

Because the number of training examples is small, often a small  $k$  is used, e.g.,  $k = 10$  or  $k = 5$ . Although  $k$  classifiers have to be learned, the computational costs are still acceptable because  $|D_{\text{train}}|$  is small.

Furthermore, because there are randomness in the division of  $D_{\text{train}}$  into  $k$  folds, different  $\text{Err}_{\text{CV}}$  will be returned if we run Algorithm 1 multiple times. The difference can be large if  $|D_{\text{train}}|$  is really small (e.g., a few tens). In this case, we can run  $k$ -fold CV multiple times and return the average of their estimates. A commonly used CV schedule is 10 times average of 10-fold cross-validation.



---

**Algorithm 1** Cross-validation estimate of the error rate

---

- 1: **Input:**  $k$ ,  $D_{\text{train}}$ , a learning algorithm (with hyperparameters specified)
- 2: Randomly divide the training set into  $k$  folds satisfying Equations 5–7
- 3: **for**  $i = 1, 2, \dots, k$  **do**
- 4:   Construct a new training set, which consists of all the examples from  $D_1, \dots, D_{i-1}, D_{i+1}, \dots, D_k$
- 5:   Learn a classifier  $f_i$  using this training set and the learning algorithm
- 6:   Calculate the error rate of  $f_i$  on  $D_i$ , denoted as  $\text{Err}_i$
- 7: **end for**
- 8: **Output:** Return the error rate estimated by cross-validation as

$$\text{Err}_{\text{CV}} = \frac{1}{k} \sum_{i=1}^k \text{Err}_i. \quad (8)$$

---

When the number of available examples is small, we may use all these examples as  $D_{\text{train}}$  (i.e., do not put aside examples as a separate test set). The cross validation error rate can be used to evaluate classifiers in this case.

### 3 Minimizing the cost / loss

Based on the above analyses of error rates, it is a natural idea to formulate a learning task as an optimization problem which minimizes the error rate. We will use the classification task to illustrate this idea.

Given a training set consists of  $n$  examples  $(\mathbf{x}_i, y_i)$  ( $1 \leq i \leq n$ ), the training error rate of a mapping  $f$  is simply

$$\min_f \frac{1}{n} \sum_{i=1}^n \llbracket f(\mathbf{x}_i) \neq y_i \rrbracket. \quad (9)$$

In Equation 9, the exact form of the mapping  $f$  and its parameters are not specified. After these components are fully specified, optimizing Equation 9 will lead to an optimal (in the training error sense) mapping  $f$ .

However, the indicator function  $\llbracket \cdot \rrbracket$  is not smooth, even not convex for most functional form and parameterization of  $f$ . Hence, as we have discussed in the math note, solving Equation 9 will be very difficult. We may also have to require  $f(\mathbf{x}_i)$  to be categorical (or integer) values, which makes designing the form of  $f$  more difficult.

When  $\llbracket f(\mathbf{x}_i) \neq y_i \rrbracket$  is 1, the  $i$ -th training example is wrongly classified by  $f$ , which means some *cost* or *loss* has been attributed to the mapping  $f$ , and Equation 9 is minimizing the average loss on the training set. Hence, we also call this type of learning as cost minimization or loss minimization.

A widely used strategy is to replace the difficult-to-optimize loss function  $\llbracket \cdot \rrbracket$  with a new optimization friendly loss function. For example, if the classification

problem is binary, i.e.,  $y_i \in \{0, 1\}$  for all  $i$ , we can use the following *mean squared error* (MSE) loss

$$\min_f \frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2. \quad (10)$$

A few properties of the MSE loss makes it a good loss function:

1. We can treat  $y_i$  as real values, and  $f(\mathbf{x}_i)$  does not need to be categorical. During prediction, the output can be 1 if  $f(\mathbf{x}_i) \geq 0.5$  and 0 if otherwise.
2. In order to minimize the loss, MSE encourages  $f(\mathbf{x}_i)$  to be close to 0 if  $y_i = 0$ , and pushes  $f(\mathbf{x}_i)$  to be around 1 if  $y_i = 1$ . In other words, the behavior of MSE mimics some *key characteristics* of the indicator function. We expect a classifier learned based on MSE to have low training error rate too (even though it might not be the smallest possible).
3. The square function is smooth and differentiable, which makes the minimization of MSE an easy task. Hence, in this formalization of loss minimization, we can treat the MSE as a simplification of the indicator function.

MSE is also a natural loss function for regression tasks (e.g., polynomial regression). In this course, we will encounter many other loss functions.

### 3.1 Regularization

As illustrated by Figure 1, both underfitting and overfitting are harmful. As we will show later in this note, a classifier with small capacity will have large *bias*, which constitutes part of the error rate. In contrast, a classifier with large capacity or representation power will have small bias but large *variance*, which is another major contributor to the error rate.

Since we do not know the inherent complexity in our task's examples, it is never easy to choose a classifier that has suitable capacity for a specific task. One commonly used strategy is to use a mapping  $f$  with high capacity, and at the same time add a *regularization* term to it. An ideal situation is that the bias will be small (because  $f$  is complex enough) and the variance may also be small (because of the regularization term penalizes complex  $f$ ). Let  $\mathcal{R}(f)$  be a regularizer on  $f$ , the MSE minimization (Equation 10) now becomes

$$\min_f \frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2 + \lambda \mathcal{R}(f), \quad (11)$$

The hyperparameter  $\lambda > 0$  is a tradeoff parameter, which controls the balance between a small training set cost ( $\frac{1}{n} \sum_{i=1}^n (f(\mathbf{x}_i) - y_i)^2$ ) and a small regularization cost ( $\mathcal{R}(f)$ ).

The regularization term (also called the regularizer) has different forms for different classifiers. Some regularizers will be introduced in this course.

### 3.2 Cost matrix

The simple accuracy or error rate criterion may work poorly in many situations. As we have introduced, minimizing the error rate in a severely imbalanced binary classification task may cause the classifier to classify all examples into the majority class.

Most imbalanced learning tasks are also *cost-sensitive*: making different types of errors will incur losses that are also imbalanced. Let us consider the small startup IT company with 24 employees once again. Company X has signed a contract with your company, and you (with your employee ID being 24) is assigned to this project. Company X produces part Y for a complex product Z, and decides to replace the quality inspector with an automatic system. The percentage of qualified parts is around 99% in company X, hence you have to deal with an imbalanced learning task. If a qualified part is predicted as defective, this part will be discarded, which will cost 10 RMB. In contrast, if a part is in fact defective but is predicted as qualified, it will be assembled into the final product Z and disable the entire product Z. Company X has to pay 1000 RMB for one such defective part. Hence, your problem is also obviously cost-sensitive.

The cost minimization framework can be naturally extended to handle cost-sensitive tasks. Let  $(\mathbf{x}, y)$  be an example and  $f(\mathbf{x})$  be a mapping that predicts the label for  $\mathbf{x}$ .  $y$  is the groundtruth label for  $\mathbf{x}$ , which is denoted as 0 if the part  $\mathbf{x}$  is qualified and 1 if it is defective. For simplicity, we assume  $f(\mathbf{x})$  is also binary (0 if it predicts  $\mathbf{x}$  as qualified, and 1 for defective). We can define a  $2 \times 2$  *cost matrix* to summarize the cost incurred in different cases:

$$\begin{bmatrix} c_{00} & c_{01} \\ c_{10} & c_{11} \end{bmatrix} = \begin{bmatrix} 0 & 10 \\ 1000 & 10 \end{bmatrix} \quad (12)$$

In Equation 12,  $c_{ij}$  is the cost incurred when the groundtruth label  $y$  is  $i$  and the prediction  $f(\mathbf{x})$  is  $j$ . Note that even a correct prediction can incur certain cost. For example,  $c_{11} = 10$  because a defective part will cost 10 RMB even if it is correctly detected as a defective one.

With an appropriate cost matrix, Equation 9 becomes

$$\min_f \frac{1}{n} \sum_{i=1}^n c_{y_i, f(\mathbf{x}_i)}, \quad (13)$$

which handles different costs.

Cost-sensitive learning is an important topic in learning and recognition. Although we will not cover its details in this course, we want to present a few notes on the cost matrix.

1. Minimizing the error rate can be considered as a special case of cost minimization, in which the cost matrix is  $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ .
2. It is, however, difficult to determine the  $c_{ij}$  values in many applications.

3. The cost matrix can be easily extended to multi-class problems. In an  $m$  class classification problem, the cost matrix is  $m \times m$ , and  $c_{ij}$  is the cost when the groundtruth label is  $i$  but the prediction is  $j$ .

### 3.3 Bayes decision theory

The Bayes decision theory minimizes the cost function in the probabilistic sense. Assume the data  $(\mathbf{x}, y)$  is a random vector and the joint probability is  $\Pr(\mathbf{x}, y)$ , the Bayes decision theory seeks to minimize the *risk*, which is the expected loss with respect to the joint density<sup>3</sup>

$$\sum_{\mathbf{x}, y} c_{y, f(\mathbf{x})} \Pr(\mathbf{x}, y). \quad (14)$$

The risk is the average loss over the true underlying probability, which can also be succinctly written as

$$\mathbb{E}_{(\mathbf{x}, y)} [c_{y, f(\mathbf{x})}]. \quad (15)$$

By picking the mapping  $f(\mathbf{x})$  that minimizes the risk, the Bayes decision theory is optimal under these assumptions and the minimum cost sense.

## 4 Evaluation in imbalanced problems

Error rate is not a good evaluation metric in imbalanced or cost-sensitive tasks. The cost in Equation 12 also has difficulty in determining the  $c_{ij}$  values. In practice, we use another set of evaluation metrics (such as *precision* and *recall*) for these tasks.

In this section, we will only consider binary classification tasks. Following the commonly used terminology, we call the two classes *positive* and *negative*, respectively. The positive class often happen to be the minority class, that is, the class with fewer examples (e.g., defective parts). The negative class often happen to be the majority class, that is, the class with more examples (e.g., qualified parts). We use  $+1$  and  $-1$  to denote positive and negative classes, respectively.

### 4.1 Rates inside individual class

There are four possible combinations of the groundtruth label  $y$  and prediction  $f(\mathbf{x})$  for one example  $\mathbf{x}$ , which is summarized in Table 2.

We use two words to pinpoint one of these four possible cases, such as true positive or false negative. In each possible case, the second word refers to the predicted label (which is shown in red in Table 2). The first word describes whether the prediction is correct or not. For example, *false positive* indicates

---

<sup>3</sup>We are using the notation for discrete random variables here. If continuous or hybrid random variables are involved, integrations (or a combination of summations and integrations) can be used to replace the summations.

Table 2: Possible combinations for the groundtruth and predicted label.

	Prediction $f(\mathbf{x}) = +1$	Prediction $f(\mathbf{x}) = -1$
True label $y = +1$	True <b>positive</b>	False <b>negative</b>
True label $y = -1$	False <b>positive</b>	True <b>negative</b>

the predicted label is “positive” (+1) and this prediction is wrong (“false”); hence, the true label is “negative” (-1). We use the initials as abbreviations for these four cases: TP, FN, FP, TN.

Given a test set with  $TOTAL$  examples, we also use these abbreviations to denote the number of examples falling into each of the 4 cases. For example,  $TP=37$  means that 37 examples in the test set are true positives. The following quantities are defined based on TP, FN, FP, and TN.

- The total number of examples:  $TOTAL = TP + FN + FP + TN$ ;
- The total number of positive examples (whose true label is +1):  $P = TP + FN$ ;
- The total number of negative examples (whose true label is -1):  $N = FP + TN$ ;
- Accuracy:  $ACC = \frac{TP+TN}{TOTAL}$ ;
- Error rate:  $Err = \frac{FP+FN}{TOTAL} = 1 - Acc$ ;
- True positive rate:  $TPR = \frac{TP}{P}$ , which is the ratio between the number of true positives and the total number of positive examples;
- False positive rate:  $FPR = \frac{FP}{N}$ , which is the ratio between the number of false positives and the total number of negative examples;
- True negative rate:  $TNR = \frac{TN}{N}$ , which is the ratio between the number of true negatives and the total number of negative examples;
- False negative rate:  $FNR = \frac{FN}{P}$ , which is the ratio between the number of false negatives and the total number of positive examples. To help remember these four rates, we note that the denominator can be deduced from the numerator in the definitions. For example, in  $FPR = \frac{FP}{N}$ , the numerator says “false positive” (whose true label is negative), which determines the denominator as  $N$ , the total number of negative examples.

These rates have been defined in different areas (such as statistics) and have different names. For example, TPR is also called *sensitivity* and FPR is the *miss rate*. One simple way to understand these rates is to treat them as evaluation results only on examples of *single* class. For examples, if we only use examples from the positive class ( $y = +1$ ) to evaluate  $f(\mathbf{x})$ , the obtained accuracy and error rate are  $TPR$  and  $FNR$ , respectively. If only the negative class is used, the accuracy and error are  $TNR$  and  $FPR$ , respectively.

## 4.2 Area under the ROC curve

$FPR$  and  $FNR$  calculate the error rates in two classes *separately*, which shall not be affected by the class imbalance issue. However, we usually need one single number to describe the performance of one classifier, rather than two rates.

What is more, in most classifiers we can obtain many pairs of different ( $FPR, FNR$ ) values. Suppose  $f(\mathbf{x}) = 2(\llbracket \mathbf{w}^T \mathbf{x} > b \rrbracket - 0.5)$ , that is, the prediction is positive (+1) if  $\mathbf{w}^T \mathbf{x} > b$  (-1 if  $\mathbf{w}^T \mathbf{x} \leq b$ ). The parameters of this classifier include both  $\mathbf{w}$  and  $b$ . After we obtain the optimal  $\mathbf{w}$  parameters, we can gradually change the parameter  $b$ :

- When  $b = -\infty$ , the prediction is always +1. That is,  $FPR$  is 1 and  $FNR$  is 0.
- When  $b$  gradually increases, some test examples will be classified as negative. If the true label for one such example is positive, a false negative is created, hence the  $FNR$  will increase; if the true label for this example is negative, then it was a false positive when  $b = -\infty$ , hence, the  $FPR$  will be reduced. In general, when  $b$  increases,  $FPR$  will gradually decrease and  $FNR$  will gradually increase.
- Finally when  $b = \infty$ , the prediction is always -1. Hence,  $FPR = 0$  and  $FNR = 1$ .

The Receiver Operating Characteristics (ROC) curve records the process of these changes into a curve.<sup>4</sup> As shown in Figure 3, the  $x$ -axis is the false positive rate  $FPR$  and the  $y$ -axis is the true positive rate  $TPR$ . Because  $TPR = 1 - FNR$ , when  $b$  sweeps from  $\infty$  to  $-\infty$  by gradually decreasing  $b$ , the ( $FPR, TPR$ ) pair will move from the coordinate (0, 0) to (1, 1), and the curve is non-decreasing.

Because the ROC curve summarizes the classifier's performance over the entire operational range (by changing the value of  $b$ ), we can use the area under the ROC curve (called the AUC-ROC) as a single-number metric for its evaluation, which is also suitable in imbalanced tasks.<sup>5</sup> The ROC of a classifier that randomly guesses its answer should be a line segment that connects (0, 0) and (1, 1) (i.e., the diagonal), and its AUC-ROC is 0.5. Hence, we expect any reasonable classifier to have its AUC-ROC larger than 0.5.

The best ROC consists of two line segments: one vertical one (from (0, 0) to (0, 1)) and a horizontal one (from (0, 1) to (1, 1)). The classifier corresponds to this ROC always correctly classify all test examples no matter what the  $b$  value is. Its AUC-ROC is 1.

---

<sup>4</sup>The name ROC was first invented to describe the performance of finding enemy targets by radar, whose original meaning is not very important in our usage of this curve.

<sup>5</sup>In a classifier other than the linear classifier  $\mathbf{w}^T \mathbf{x} + b$ , we often can find another parameter that acts as a decision threshold (similar to  $b$ ), and then can generate the ROC by altering that parameter.




Figure 3: The Receiver Operating Characteristics (ROC) curve.

### 4.3 Precision, Recall, and F-measure

Precision and recall are two measures suitable for evaluating imbalanced tasks, too. They are defined as

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (16)$$

$$\text{Recall} = \frac{TP}{P}. \quad (17)$$

Note that the recall measure is simply another name for true positive rate.

Consider an image retrieval application. The test set has 10000 images in total, in which 100 are related to *Nanjing University* (which we treat as positive class examples). When you type “Nanjing University”, the retrieval system returns 50 images to you, which are those predicted as related to Nanjing University (aka, predicted as positive). The returned results contain some images that indeed are related to Nanjing University (TP), but some are irrelevant (FP). Now two measures are important when you evaluate the retrieval system:

- Are the returned results precise? You want most (if not all) returned images are positive / related to Nanjing University. This measure is the *precision*, the percentage of true positives among those that are predicted as positive, i.e.,  $\frac{TP}{TP+FP}$ .
- Are all positive examples recalled? You want most (if not all) positive images are included in the returned set. This measure is the *recall* (or TPR), the percentage of true positives among all positive images in the entire test set, i.e.,  $\frac{TP}{P}$ .

It is not difficult to observe that these two performance measures are still meaningful in imbalanced tasks.




Figure 4: The Precision-Recall (PR) curve.

There are two widely used ways to summarize precision and recall into one number. Similar to the ROC curve, we can alter the threshold parameter in a classifier and generate many pairs of (precision, recall) values, which form the precision-recall (PR) curve. Then, we can use the AUC-PR (area under the precision-recall curve) as the evaluation metric. Figure 4 shows one example PR curve.

One notable difference between the PR and ROC curves is that the PR curve is no longer non-decreasing. In most tasks, we will observe PR curves that are zigzagging, like the one in Figure 4.

The AUC-PR measure is more discriminating than AUC-ROC. One may observe AUC-ROC scores that are very close to each other for many classifiers. In that case, AUC-PR may be a better evaluation metric.

The second way to combine precision and recall is the F-measure, which is defined as

$$F = \frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}. \quad (18)$$

F-measure is the harmonic mean of precision and recall, which is always between 0 and 1. A higher F-measure indicates a better classifier. It is also easy to see that

$$F = \frac{2TP}{2TP + FP + FN}. \quad (19)$$

The F-measure treats precision and recall equally. An extension of the F-measure is defined for a fixed  $\beta > 0$ , as

$$F_\beta = (1 + \beta^2) \cdot \frac{\text{precision} \cdot \text{recall}}{\beta^2 \cdot \text{precision} + \text{recall}}. \quad (20)$$

Either the precision or recall is considered as more important than the other, depending on the value of  $\beta$ . Note that the F-measure is a special case of  $F_\beta$  when  $\beta = 1$ .



An intuitive way to understand the relative importance between precision and recall is to consider extreme  $\beta$  values. If  $\beta \rightarrow 0$ ,  $F_\beta \rightarrow$  precision, that is, precision is more important for small  $\beta$  values. If  $\beta \rightarrow \infty$ ,  $F_\beta \rightarrow$  recall, that is, recall is more important for large  $\beta$ . When  $\beta = 1$ , the F-measure is symmetric about precision and recall, i.e.,  $F_\beta(\text{precision}, \text{recall}) = F_\beta(\text{recall}, \text{precision})$ . Hence, they are equally important when  $\beta = 1$ . More properties and implications of  $F_\beta$  will be discussed in the exercises.

## 5 Can we reach 100% accuracy?

Till now, we have introduced ways to evaluate the system's errors in a few different tasks. What we really want is that, in fact, a learning algorithm or recognition system that has *no* error. Can we implement such a system with 100% accuracy?

The answer to this question hinges on the data in your task. In the probabilistic interpretation, the data are expressed as random variables  $(\mathbf{x}, y)$ , and the examples are sampled i.i.d. from an underlying distribution  $\Pr(\mathbf{x}, y)$  (or use the density  $p(\mathbf{x}, y)$  for continuous random variables). If there exists an instance  $\mathbf{x}^1$  and two different values of  $y$  ( $y^1$  and  $y^2$ , and  $y^1 \neq y^2$ ), such that  $\Pr(\mathbf{x}^1, y^1) > 0$  and  $\Pr(\mathbf{x}^1, y^2) > 0$ , then a classifier with 100% (generalization) accuracy is impossible.

### 5.1 Bayes error rate

We assume any classifier  $f(\mathbf{x})$  is a valid function, which can only maps  $\mathbf{x}$  to a single fixed element  $y$ , that is,  $f(\mathbf{x}^1) = y^1$  and  $f(\mathbf{x}^1) = y^2$  cannot both happen if  $y^1 \neq y^2$ . This assumption is a valid one, because in a classifier, we expect a deterministic answer. Hence, in either the training or test set, any instance  $\mathbf{x}$  will only be assigned one groundtruth label  $y$  and the prediction is also unique.

Now we can consider three cases for the above example.

- Case 1: The true label for the instance  $\mathbf{x}^1$  is set to be  $y^1$ . Then, the example  $(\mathbf{x}^1, y^2)$  is considered as an error. Because  $\Pr(\mathbf{x}^1, y^2) > 0$ , this example will contribute at least  $\Pr(\mathbf{x}^1, y^2)$  to the generalization error. That is, no matter what classifier is used,  $(\mathbf{x}^1, y^2)$  will contribute  $\Pr(\mathbf{x}^1, y^2)$  to this classifier's generalization error.
- Case 2: The true label for the instance  $\mathbf{x}^1$  is set to be  $y^2$ . Then, the example  $(\mathbf{x}^1, y^1)$  is considered as an error. Because  $\Pr(\mathbf{x}^1, y^1) > 0$ , this example will contribute at least  $\Pr(\mathbf{x}^1, y^1)$  to the generalization error. That is, no matter what classifier is used,  $(\mathbf{x}^1, y^1)$  will contribute  $\Pr(\mathbf{x}^1, y^1)$  to this classifier's generalization error.
- Case 3: The true label for the instance  $\mathbf{x}^1$  is set to be neither  $y^1$  nor  $y^2$ . Then, both examples  $(\mathbf{x}^1, y^1)$  and  $(\mathbf{x}^1, y^2)$  are errors. Because both  $\Pr(\mathbf{x}^1, y^1) > 0$  and  $\Pr(\mathbf{x}^1, y^2) > 0$ , these two examples will contribute at least  $\Pr(\mathbf{x}^1, y^1) + \Pr(\mathbf{x}^1, y^2)$  to the generalization error of any possible classifier.

These three cases exhaust all possible cases for  $(\mathbf{x}^1, y^1)$  and  $(\mathbf{x}^1, y^2)$ . If there are more than two  $y$  values make  $\Pr(\mathbf{x}^1, y) > 0$ , analyses can be made in the same way. Hence, we reach the following conclusions:

- Errors are inevitable if there exists  $\mathbf{x}^1$  and  $y^1, y^2$  ( $y^1 \neq y^2$ ) such that  $\Pr(\mathbf{x}^1, y^1) > 0$  and  $\Pr(\mathbf{x}^1, y^2) > 0$ . Note that this statement is classifier agnostic—it is true regardless of the classifier.
- If non-deterministic label is inevitable for an instance  $\mathbf{x}^1$ , we shall compare the two probabilities  $\Pr(\mathbf{x}^1, y^1)$  and  $\Pr(\mathbf{x}^1, y^2)$ . If  $\Pr(\mathbf{x}^1, y^1) > \Pr(\mathbf{x}^1, y^2)$ , we set the groundtruth label for  $\mathbf{x}^1$  to  $y^1$  so that its contribution of unavoidable error  $\Pr(\mathbf{x}^1, y^2)$  is small (compared to the case when the groundtruth label is set to  $y^2$ ). If  $\Pr(\mathbf{x}^1, y^2) > \Pr(\mathbf{x}^1, y^1)$ , we must set the groundtruth label to  $y^2$ .
- Let us consider the general situation. There are  $m$  classes ( $m \geq 2$ ), which is denoted as  $y = i$  ( $1 \leq i \leq m$ ) for the  $i$ -th class. For any instance  $\mathbf{x}$ , we should set its groundtruth label as  $y_\star = \arg \max_y \Pr(\mathbf{x}, y)$  to obtain the smallest unavoidable error, which is  $\sum_{y \neq y_\star} \Pr(\mathbf{x}, y)$  for  $\mathbf{x}$ .

If we repeat the above analyses for all possible instances  $\mathbf{x} \in \mathbb{X}$ , the generalization error of any classifier has a lower bound, as

$$\sum_{\mathbf{x} \in \mathbb{X}} \sum_{y \neq y_\star(\mathbf{x})} \Pr(\mathbf{x}, y) = 1 - \sum_{\mathbf{x} \in \mathbb{X}} \Pr(\mathbf{x}, y_\star(\mathbf{x})), \quad (21)$$

in which  $y_\star(\mathbf{x}) = \arg \max_y \Pr(\mathbf{x}, y)$  is the class which makes  $\Pr(\mathbf{x}, y)$  the largest for  $\mathbf{x}$ .

Equation 21 defines the *Bayes error rate*, which is a theoretical bound of the smallest error rate any classifier can attain. Comparing Equations 15 and 21, it is easy to see that the loss incurred by the Bayes decision theory is exactly the Bayes error rate if the cost matrix is  $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ .

## 5.2 Groundtruth labels

The Bayes error rate is a theoretical bound, which assumes the joint density  $p(\mathbf{x}, y)$  or joint probability  $\Pr(\mathbf{x}, y)$  is known and can be analytically computed. The underlying distribution, however, is unknown in almost all real-world tasks. Then, how shall the groundtruth labels been decided?

In some applications, labels of training examples are accessible. For example, you try to predict day  $t + 1$ 's NYSE closing index value based on previous  $n$  days' closing NYSE index values, which is a regression task. On day  $t + 2$ , you have observed the label of this regression task for day  $t + 1$ . Hence, you can collect the training instances and labels using NYSE history data. A model learned from the history data can be used to predict next day's trading trend (although that prediction may be highly inaccurate).

In most cases, these computations are replaced by human judgments. An expert in the task domain will manually specify labels for instances. For example, an experienced doctor will look at a patient's bone X-ray image, and

decide whether a bone injury exists or not. This type of human annotation is the major labeling method.

Human annotations are susceptible to errors and noises. In some difficult cases, even an expert will have difficulties in determining an accurate groundtruth label. An exhausted expert may make mistakes during labeling, especially when he or she has to label many examples. The labeling process can be extremely time-consuming, which also leads to high financial pressure—imagine you have to hire a top-notch doctor for 100 business days to label your data! In cases where both  $\Pr(\mathbf{x}^1, y^1) > 0$  and  $\Pr(\mathbf{x}^1, y^2) > 0$ , the expert has to choose from  $y^1$  or  $y^2$ , which may be difficult in some applications.

In summary, the labels can be a categorical value (for classification), a real number (for regression), a vector or matrix of real numbers, or even more complex data structures. The groundtruth labels may also contain uncertainty, noise or error. In this introductory course, we will focus on simple classification and regression tasks, and assume the labels are error and noise free *if not otherwise specified* (mainly because methods that explicitly handle label errors are advanced and beyond the scope of this course).

### 5.3 Bias-Variance decomposition

Now that the error is inevitable, it is helpful to study what components contribute to the error. Hopefully, this understanding will help us in reducing the part of error that is larger than the Bayes error rate. We will use regression as an example to illustrate the *bias-variance* decomposition. Similar decompositions also exist for classifiers, but are more complex.

We need quite many assumptions to set up the stage for bias-variance decomposition. First, we have a function  $F(\mathbf{x}) \in \mathbb{R}$ , which is the function that generates our training data. However, the training data is susceptible to noise. A training example  $(\mathbf{x}, y)$  follows

$$y = F(\mathbf{x}) + \epsilon, \tag{22}$$

in which  $\epsilon \sim N(0, \sigma^2 I)$  is a Gaussian random noise, which is *independent* of  $\mathbf{x}$ .

Next, we can sample (i.i.d.) from Equation 22 to generate different training sets, which may contain different number of training examples. We use a random variable  $D$  to represent the training set.

Third, we have a regression model  $f$ , which will generate a mapping  $f(\mathbf{x}; D)$  by learning from the training set  $D$ , and predicts  $f(\mathbf{x}; D)$  for any instance  $\mathbf{x}$ . We put  $D$  in the prediction's notation to emphasize that the prediction is based on the mapping learned from  $D$ . When different samples of  $D$  (different training sets) are used, we expect different prediction results for the same  $\mathbf{x}$ .

However, we assume that the classifier is deterministic. That is, given the same training set many times, it will produce the same mapping and predictions.

Finally, because of the i.i.d. sampling assumption, we just need to study for one specific instance  $\mathbf{x}$ , and examine where does the error come from in predicting the regression output for  $\mathbf{x}$ .

An important note about these assumptions: since  $F$ , the regression learning process, and  $\mathbf{x}$  are deterministic, the randomness mainly comes from the training set  $D$ . For example,  $F(\mathbf{x})$  is deterministic, but  $f(\mathbf{x}; D)$  is a random variable since it depends on  $D$ . For notational simplicity, we will simplify  $\mathbb{E}_D[f(\mathbf{x}; D)]$  as  $\mathbb{E}[f(\mathbf{x})]$ , but it is essential to remember that the expectation is with respect to the distribution of  $D$ , and  $f(\mathbf{x})$  means  $f(\mathbf{x}; D)$ .

Now we have all the tools to study  $\mathbb{E}[(y - f(\mathbf{x}))^2]$ , the generalization error in the squared error sense. Because we only consider one fixed example  $\mathbf{x}$ , we will write  $F(\mathbf{x})$  as  $F$  and  $f(\mathbf{x})$  as  $f$ .  $F$  is deterministic (hence  $\mathbb{E}[F] = F$ ), and  $\mathbb{E}[f]$  means  $\mathbb{E}_D[f(\mathbf{x}; D)]$  (which is also a fixed value).

The error is then  $\mathbb{E}[(y - f)^2] = \mathbb{E}[(F - f + \epsilon)^2]$ . Because the noise  $\epsilon$  is independent of all other random variable, we have

$$\mathbb{E}[(y - f)^2] = \mathbb{E}[(F - f + \epsilon)^2] \quad (23)$$

$$= \mathbb{E}[(F - f)^2 + \epsilon^2 + 2(F - f)\epsilon] \quad (24)$$

$$= \mathbb{E}[(F - f)^2] + \sigma^2. \quad (25)$$

Note that  $\mathbb{E}[\epsilon^2] = (\mathbb{E}[\epsilon])^2 + \text{Var}(\epsilon) = \sigma^2$  and  $\mathbb{E}[(F - f)\epsilon] = \mathbb{E}[F - f]\mathbb{E}[\epsilon] = 0$  (because of the independence).

We can further expand  $\mathbb{E}[(F - f)^2]$ . From the math note we know

$$\mathbb{E}[(F - f)^2] = (\mathbb{E}[F - f])^2 + \text{Var}(F - f). \quad (26)$$

For the first term in the RHS of Equation 26, because  $\mathbb{E}[F - f] = F - \mathbb{E}[f]$ , we have  $(\mathbb{E}[F - f])^2 = (F - \mathbb{E}[f])^2$ .

For the second term, because  $F$  is deterministic, we have  $\text{Var}(F - f) = \text{Var}(-f) = \text{Var}(f) = \mathbb{E}[(f - \mathbb{E}[f])^2]$ , i.e., the variance of  $f(\mathbf{x}; D)$ .

Putting all these results together, we have

$$\mathbb{E}[(y - f)^2] = (F - \mathbb{E}[f])^2 + \mathbb{E}[(f - \mathbb{E}[f])^2] + \sigma^2, \quad (27)$$

which is the bias-variance decomposition for regression. This decomposition states that the generalization error for any example  $\mathbf{x}$  comes from three parts: the squared *bias*, the *variance*, and the noise.

- $F - \mathbb{E}[f]$  is called the bias. Because an expectation is taken on  $f$ , bias is not depend on a training set. Hence, it is determined by the regression model, e.g., are you using a degree 2 or degree 15 polynomial? After we fix the form of our regression model, the bias is fixed.
- $\mathbb{E}[(f - \mathbb{E}[f])^2]$  is the variance of the regression *with respect to the variation in the training set*.
- $\sigma^2$  is the variance of the noise, which is irreducible (cf. the Bayes error rate in classification). Even if we know the underlying function  $F(\mathbf{x})$  exactly and set  $f = F$ , the generalization error is still  $\sigma^2$ .

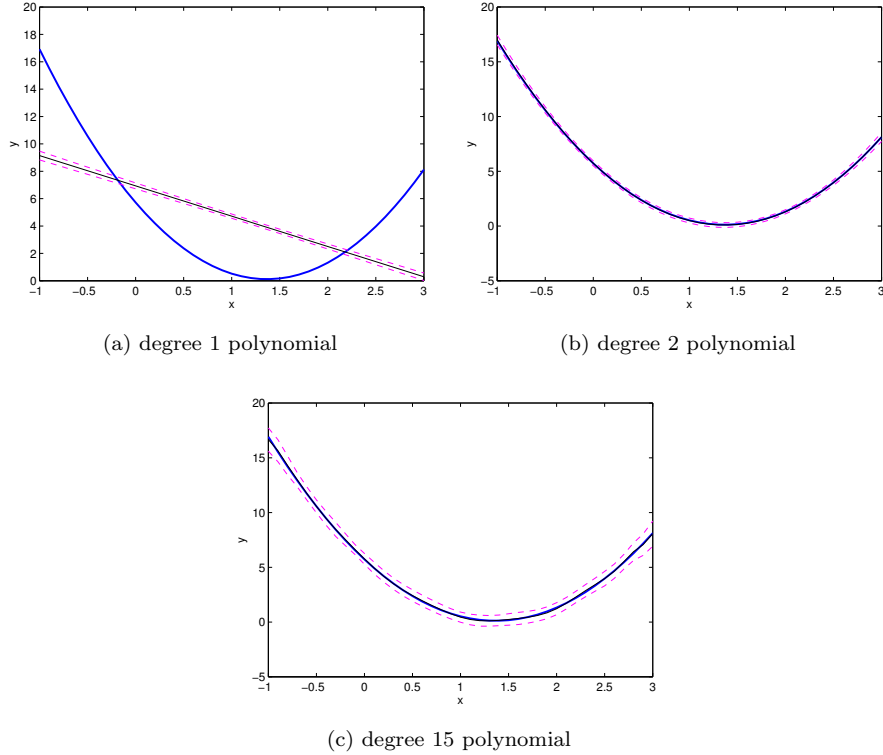


Figure 5: Illustration of bias and variance in a simple polynomial regression task. The blue curve is the groundtruth, the black curve is the mean of 100 regression models learned from different training sets, and the two purple curves are the mean plus/minus one standard deviation of the regression models.

- This decomposition can be applied to any  $\mathbf{x}$ . Because we omitted  $\mathbf{x}$  from the notations in Equation 27, the bias and variance will have different values when  $\mathbf{x}$  changes.

For the regression task in Equation 4, we can compute the bias and variance for any  $x$ , and the results are shown in Figure 5. Three regression models are considered: degree 1, 2, and 15 polynomials. In Figure 5,  $F$  is shown as the blue curve;  $\mathbb{E}[f]$  is the black curve; the two purple curves are above and below  $\mathbb{E}[f]$  by one standard deviation of  $f$ , respectively. Hence, the difference between the black and blue curves equals the bias, and the squared difference between the black and purple curves equals the variance at every  $x$  coordinate. In order to compute  $\mathbb{E}[f]$  and  $\text{Var}(f)$ , we i.i.d. sampled 100 training sets with the same size.

When the model has enough capacity (e.g., degree 2 or 15 polynomials), the bias can be very small. The black and blue curves in Figures 5b and 5c are

almost identical. However, the linear (degree 1 polynomial) model is too simple for this task, and its bias is huge: the distance between the black and blue curves is very large at most points. Also note that the bias and variance is not constant when  $x$  changes. In Figure 5a, the bias term changes quickly when  $x$  changes; while in Figure 5c, the variance is the largest when  $x$  is close to both ends of the exhibited  $x$  range.

However, we do not want the model to be too complex. Although Figures 5a and 5b both exhibit small distances between the purple and black curves (variances), the degree 15 polynomial model in Figure 5c shows quite large variance. That is, a slight change in the training set can lead to large variations in the learned regression results. We call this type of models *unstable*. In short, we want a learning model that has enough capacity (to have small bias) and is stable (to have small variance).

These requirements are contradictory to each other in most cases. For example, the linear model (cf. Figure 5a) is stable but its capacity is low. Adding regularization to complex models is one viable approach. Averaging many models (i.e., model ensemble) is also widely used.

## 6 Confidence in the evaluation results

At the end of this note, we briefly discuss the following question: after you have obtained an estimate of your classifier's error rate, how confident are you about this estimate?

One factor that affects the confidence is the size of your test or validation set. In some tasks such as the ImageNet image classification challenge, the test or validation set size is very large (150,000 images in its validation set).<sup>6</sup> Researchers use the test or validation error rate with high confidence.

In some other problems with moderate number of examples, we can randomly divide all available examples into a training set and a test set. This split will be repeated  $k$  times (with  $k = 10$  being a typical value), and the error rates evaluated in all  $k$  splits are averaged. The sample standard deviation (standard deviation computed from the  $k$  error rates) are almost always reported along with the average error rate.

When there are even fewer (e.g., a few hundreds) examples, within each of the  $k$  splits, cross-validation is used to calculate the cross-validation error rate in each split. The average and sample standard deviation of the  $k$  splits are reported.

### 6.1 Why averaging?

As shown by the bias-variance decomposition, the bias will not change after the classifier has been fixed. Hence, the error rate's variation are caused by different training and test sets. Hence, by reducing this variation we have more confidence in the error rate estimate.

---

<sup>6</sup><http://www.image-net.org/>




Figure 6: The standard normal probability density function and the 1- and 2-sigma range. The one sigma range has an area of 0.6827, and the two sigma range has an area of 0.9545.

Let  $E$  be a random variable corresponding to the error rate. If  $E_1, E_2, \dots, E_k$  are  $k$  samples of  $E$ , computed from i.i.d. sampled training and test sets (with the same training and test set size). The error rate is often modeled as a normal distribution, i.e.,  $E \sim N(\mu, \sigma^2)$ . It is easy to verify that the average

$$\bar{E} = \frac{1}{k} \sum_{j=1}^k E_j \sim N\left(\mu, \frac{\sigma^2}{k}\right), \quad (28)$$

i.e., the average of multiple *independent* error rates will reduce the variance to  $\frac{1}{k}$ . This fact means that averaging can reduce the variance of the error rate estimates. However, if we split a set of examples  $k$  times, the  $k$  training and  $k$  test sets cannot be independent because they are split from the same set of examples. Hence, we expect the average  $\bar{E}$  to have a smaller variance than  $E$ , but not as small as  $\frac{\sigma^2}{k}$ . The smaller variance is still useful in providing a better estimate than using a single train / test split.

## 6.2 Why report the sample standard deviation?

If we know both  $\bar{E}$  and  $\sigma$  (the population / true standard deviation, not the sample standard deviation), we can deduce how confident we are about  $\bar{E}$ .

Because  $\bar{E} \sim N\left(\mu, \frac{\sigma^2}{k}\right)$ , we have  $\frac{\bar{E}-\mu}{\sigma/\sqrt{k}} \sim N(0, 1)$ . In Figure 6 we show the p.d.f. of the standard normal  $N(0, 1)$ . The area of the green region is the probability  $\Pr(|X| \leq 1)$  if  $X \sim N(0, 1)$ , which is 0.6827; the area of the green plus two blue regions is the probability  $\Pr(|X| \leq 2)$  if  $X \sim N(0, 1)$ , which is 0.9545.

Because  $\frac{\bar{E}-\mu}{\sigma/\sqrt{k}} \sim N(0, 1)$ , we have  $\Pr\left(\left|\frac{\bar{E}-\mu}{\sigma/\sqrt{k}}\right| \leq 2\right) = 0.9545$ , or,

$$\Pr\left(\bar{E} - \frac{2\sigma}{\sqrt{k}} \leq \mu \leq \bar{E} + \frac{2\sigma}{\sqrt{k}}\right) > 0.95. \quad (29)$$

That is, although we do not know the generalization error  $\mu$ , a pretty confident (95%) estimate can be provided by  $\bar{E}$  and  $\sigma$ . We know it is very likely (95% confident) that  $\mu$  will be in the interval  $\left[\bar{E} - \frac{2\sigma}{\sqrt{k}}, \bar{E} + \frac{2\sigma}{\sqrt{k}}\right]$ . A smaller variance  $\sigma^2$  means the confidence interval is smaller. Hence, it is useful to report the variance (standard deviation).

However, two caveats exist about this interpretation. First, we do not know the true population standard deviation  $\sigma$ , which will be replaced by the sample standard deviation computed from the  $k$  splits. Hence, Equation 29 will not be correct anymore. Fortunately, the distribution after using the sample standard deviation has a closed form, the Student's  $t$ -distribution. We will discuss more of the  $t$ -distribution soon.

Second,  $\bar{E}$  is a random variable in Equation 29. In practice, we need to replace it with the sample mean  $\bar{e}$  computed from the  $k$  splits in one experiment. However, the interval  $\left[\bar{e} - \frac{2\sigma}{\sqrt{k}}, \bar{e} + \frac{2\sigma}{\sqrt{k}}\right]$  is deterministic, which does *not* have a probability associated with it. A 95% confidence means the following:

*In one experiment, we will split the data  $k$  times, compute one  $\bar{e}$  and one confidence interval accordingly. We will obtain 100  $\bar{e}$  and 100 confidence intervals in 100 experiments. Then, among the 100 experiments, around 95 times  $\mu$  will fall into the confidence intervals, but around 5 times it will not.*

### 6.3 Compare two classifiers

If we have two classifier  $f_1$  and  $f_2$ , and want to evaluate which one is better in a problem (with a set of examples  $D$ ). We can evaluate  $f_1$  and  $f_2$ , and estimate  $f_1$ 's error rate as  $0.08 \pm 0.002$  (meaning the sample mean and standard deviation for the error of  $f_1$  is 0.08 and 0.002, respectively). Similarly, the evaluation result for  $f_2$  is  $0.06 \pm 0.003$ . Now that  $\frac{0.08-0.06}{0.002+0.003} = 4$ , we are confident enough (> 99%) to say that  $f_2$  is better than  $f_1$  on this problem.

However, if  $f_1$  is  $0.08 \pm 0.01$  and  $f_2$  is evaluated as  $0.06 \pm 0.012$ , we have  $\frac{0.08-0.06}{0.012+0.01} = 0.91$ , that is, the one standard deviation confidence interval (whose confidence is < 0.6827) of  $f_1$  and  $f_2$ 's estimates overlap with each other. Hence, there is not enough confidence to say which one of the two classifiers is better.

The Student's  $t$ -test<sup>7</sup> is useful for such comparisons. Full details of the  $t$ -test involves a lot of different cases and details. In this note, we only introduce an application of paired  $t$ -test when we compare two classifiers on a wide range of problems.

Suppose we evaluate  $f_1$  and  $f_2$  on  $n$  datasets, yielding errors  $E_{ij}$  where  $i \in \{1, 2\}$  is the classifier index, and  $j$  ( $1 \leq j \leq n$ ) is the index to the datasets.

<sup>7</sup>William Sealy Gosset designed this test's statistics, and published his results under the pen name "Student".



Individually,  $E_{1j}$  and  $E_{2j}$  may not be confidently compared for most of the  $j$  values / datasets. However, this is a *paired* comparison because for any  $j$ , the same dataset is used by  $f_1$  and  $f_2$ . Hence, we can study the properties of  $E_{1j} - E_{2j}$  for all  $1 \leq j \leq n$ . Furthermore, we assume the datasets are not dependent on each other, thus different  $j$  values lead to *independent*  $E_{1j} - E_{2j}$  random variables. Because the sum of two normal is once again a normal distribution, we also assume  $E_{1j} - E_{2j}$  is *normally distributed*.

For notational simplicity, we denote  $X_j = E_{1j} - E_{2j}$  ( $1 \leq j \leq n$ ). Under the above assumptions, we know they are i.i.d. samples from  $X \sim N(\mu, \sigma^2)$ , but the parameters  $\mu$  and  $\sigma$  are unknown. The average  $\bar{X} = \frac{1}{n} \sum_{j=1}^n X_j \sim N(\mu, \frac{\sigma^2}{n})$ . Hence,  $\frac{\bar{X} - \mu}{\sigma/\sqrt{n}} \sim N(0, 1)$ .

Note that a particular parameter value  $\mu = 0$  is of special interest to us. When  $\mu = 0$ , we have  $\mathbb{E}[\bar{X}] = \mathbb{E}[X_j] = \mathbb{E}[E_{1j} - E_{2j}] = 0$ , that is, on average  $f_1$  and  $f_2$  has the same error rate. When  $\mu > 0$ ,  $f_1$ 's error is higher than that of  $f_2$ , and  $f_1$ 's error is smaller than that of  $f_2$  if  $\mu < 0$ .

The  $t$ -test answers the following question: does  $f_1$  and  $f_2$  has different error? The *null hypothesis* is that  $\mu = 0$ , meaning there is no significant difference between  $f_1$  and  $f_2$ . If  $f_1$  is a new algorithm proposed by you and  $f_2$  is a method in the literature, you may want the null hypothesis to be *rejected* (i.e., with enough evidence to believe it is not true), because you hope your new algorithm is better.

How can we confidently reject the null hypothesis? We can define a *test statistics*  $T$ , whose distribution can be derived *by assuming the null hypothesis is true* (i.e., when  $\mu = 0$ ). For example, assuming  $\mu = 0$ ,  $\sigma$  is known and  $T = \frac{\bar{X}}{\sigma/\sqrt{n}}$ , we know the distribution of  $T$  is the standard normal  $N(0, 1)$ .

In the next step, you can compute the value of the statistics  $T$  based on your data, denoted as  $t$ , and say  $t = 3.1$ . Because  $T$  is a standard normal, we have  $\Pr(|T| > 3) = 0.0027$ , or 0.27%, it is a *small probability event* to observe  $t = 3.1$  in one experiment—something must have went wrong if you are not extremely unlucky.

The only thing can be wrong is the assumption “the null hypothesis is true”. Hence, when we observe unusual or extreme values for  $t$ , we have confidence to reject the null hypothesis. Precisely speaking, we can specify a *significance level*  $\alpha$ , which is a small number between 0 and 1. When assuming the null hypothesis is true, but the probability of observing an extreme value  $T = t$  is smaller than  $\alpha$  (i.e.,  $\Pr(|T| > t) < \alpha$ ), we conclude that we are rejecting the null hypothesis at confidence level  $\alpha$ .

In practice,  $\sigma$  is unknown. Let  $e_{ij}$  and  $x_j$  be samples for  $E_{ij}$  and  $X_j$  in one experiment, respectively. We can compute the sample mean and sample standard deviation for  $X$  as

$$\bar{x} = \frac{1}{n} \sum_{j=1}^n x_j, \quad s = \sqrt{\frac{1}{n-1} \sum_{j=1}^n (x_j - \bar{x})^2}, \quad (30)$$

which corresponds to random variables  $\bar{X}$  and  $S$ . Note that  $\frac{1}{n-1}$  instead of  $\frac{1}{n}$




Figure 7: Probability density function of the Student’s  $t$ -distribution with different degrees of freedom.

is used.

The new paired  $t$ -test statistics is

$$T = \frac{\bar{X}}{S/\sqrt{n}}, \quad (31)$$

which replaces  $\sigma$  with  $S$ . Consequently,  $T$  is no longer the standard normal distribution. Fortunately, we know that  $T$  follows a Student’s  $t$ -distribution with degree of freedom  $\nu = n - 1$ . Note that the degree of freedom is *not*  $n$ .

A  $t$ -distribution has one integer parameter: the degree of freedom  $\nu$ . Its p.d.f. is

$$p(t) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\nu\pi}\Gamma(\frac{\nu}{2})} \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}}, \quad (32)$$

in which  $\Gamma$  is the gamma function defined as  $\Gamma(t) = \int_0^\infty x^{t-1}e^{-x} dx$ .

As shown in Figure 7, the  $t$ -distribution is symmetric about 0 and looks like the standard normal distribution. However, the  $t$ -distribution has more density at its tails than the standard normal distribution. When the degree of freedom  $\nu$  grows,  $t$ -distribution gets closer to  $N(0, 1)$ . When  $\nu \rightarrow \infty$ , the  $t$ -distribution converges to the standard normal distribution.

Often the degree of freedom is small. For one fixed degree of freedom  $\nu$  and a given significance level  $\alpha$ , a critical value  $c_{\nu, \alpha/2} > 0$  can be found in statistical tables, satisfying  $\Pr(|T| > c_{\nu, \alpha/2}) = \alpha$ . Because the  $t$ -distribution is symmetric and the region of extremal values can be at either side with equal probability,  $\alpha/2$  is used in finding the critical value. Hence, for one experiment, we can compute the sample statistics  $t = \frac{\bar{x}}{s/\sqrt{n}}$ , and reject the null hypothesis if  $|t| > c_{\nu, \alpha/2}$ . The paired  $t$ -test for comparing two classifiers is summarized in Algorithm 2.

---

**Algorithm 2** The paired  $t$ -test for comparing two classifiers

---

- 1: **Input:** Two classifiers  $f_1$  and  $f_2$ , whose error rate estimates on  $n$  datasets are  $e_{ij}$  ( $i \in \{1, 2\}$ ,  $1 \leq j \leq n$ )
  - 2: Choose a significance level  $\alpha$  (widely used values are 0.05 and 0.01)
  - 3: Find the critical value  $c_{n-1, \alpha/2}$
  - 4:  $x_j \leftarrow e_{1j} - e_{2j}$ , for  $1 \leq j \leq n$
  - 5:  $\bar{x} \leftarrow \frac{1}{n} \sum_{j=1}^n x_j$ ,  $s \leftarrow \sqrt{\frac{1}{n-1} \sum_{j=1}^n (x_j - \bar{x})^2}$
  - 6:  $t \leftarrow \frac{\bar{x}}{s/\sqrt{n}}$
  - 7: **if**  $|t| > c_{n-1, \alpha/2}$  **then**
  - 8:   We believe that  $f_1$  and  $f_2$  have different error rates. (The null hypothesis is rejected at confidence level  $\alpha$ .)
  - 9: **else**
  - 10:   We do not believe there is significant difference between the error rates of  $f_1$  and  $f_2$
  - 11: **end if**
- 

Note that the paired  $t$ -test requires: 1) the errors  $e_{ij}$  are paired; 2) the errors are independent for different  $j$ ; and, 3) the difference  $x_j$  is normally distributed. Although slight or moderate violation of the third assumption is usually acceptable, the first two (paired and independent) assumptions cannot be violated.

Algorithm 2 specifies a *two-tailed* test, which only cares about whether  $f_1$  and  $f_2$  has same / similar error rate or not. In many scenarios we need a *one-tailed* version of the test, for example, if you want to show that your new algorithm  $f_1$  is better than  $f_2$ .

In the one-tailed test, if you want to show the error rate of  $f_1$  is smaller than that of  $f_2$ , you need to change the critical value to  $c_{n-1, \alpha}$  and further require  $t < -c_{n-1, \alpha}$ . If you want to show that  $f_2$  is better than  $f_1$ , you need to have  $t > c_{n-1, \alpha}$ .

If statistical tables are not handy but you have a computer at your disposal (which happens to have some suitable software packages installed on it), you can compute the critical value by one line in your favorite software. For example, let compute the two-tailed critical value for the paired  $t$ -test, with  $\nu = 6$  and  $\alpha = 0.05$ . Because the  $t$ -distribution is symmetric,  $\Pr(T > c_{\nu, \alpha/2}) = \alpha/2$ . Let  $\Phi$  denote the c.d.f. of a  $t$ -distribution with  $\nu$  degrees of freedom, we have  $\Phi(c_{\nu, \alpha/2}) = 1 - \alpha/2$ . Hence,  $c_{\nu, \alpha/2} = \Phi^{-1}(1 - \alpha/2)$ , in which  $\Phi^{-1}(\cdot)$  is the inverse c.d.f. function. The Matlab function `tinv(p, nu)` calculates the inverse c.d.f. value for probability `p` and degree of freedom  $\nu$ . One single Matlab command `tinv(1-0.05/2, 6)` tells us the critical value for a two-tailed, 6 degree of freedom, significance level 0.05 paired  $t$ -test is 2.4469. Similarly, for the one-tailed test, we have the critical value as `tinv(1-0.05, 6)`, which is 1.9432.

The  $t$ -test is a relatively simple statistical test which can be adopted to compare classification methods (or other learning algorithms / systems). Many other tests (e.g., various rank tests) are useful in this aspect too. We hope

the introduction in this note will help readers in understanding other statistical tests, but will not go into details of these tests.

One final note: whenever you want to apply a statistical test to your data, the first thing to check is whether your data satisfy the assumptions of that particular test!

## Exercises

1. In a binary classification problem, we know  $P = N = 100$  (i.e., there are 100 positive and 100 negative examples in the test set). If  $FPR = 0.3$  and  $TPR = 0.2$ , then what is the precision, recall,  $F_1$  score? What is its accuracy and error rate?
2. (Linear regression) Consider a set of  $n$  examples  $(\mathbf{x}_i, y_i)$  ( $1 \leq i \leq n$ ) where  $\mathbf{x}_i \in \mathbb{R}^d$  and  $y_i \in \mathbb{R}$ . A *linear regression* model assumes  $y = \mathbf{x}^T \boldsymbol{\beta} + \epsilon$  for any example  $(\mathbf{x}, y)$ , where  $\epsilon$  is a random variable modeling the regression error and  $\boldsymbol{\beta} \in \mathbb{R}^d$  are the parameters of this model. For the  $i$ -th example, we have  $\epsilon_i = y_i - \mathbf{x}_i^T \boldsymbol{\beta}$ .
  - (a) Express the linear regression task as an optimization problem over the training set, using the training examples, the parameters  $\boldsymbol{\beta}$  and the squared error ( $\sum_{i=1}^n \epsilon_i^2$ , which is the MSE times the number of examples).
  - (b) We can organize the training examples  $\mathbf{x}_i$  into a  $n \times d$  matrix  $X$ , whose  $i$ -th row is the vector  $\mathbf{x}_i^T$ . Similarly, we can organize  $y_i$  into a vector  $\mathbf{y} \in \mathbb{R}^n$ , with  $y_i$  in the  $i$ -th row. Rewrite the optimization problem in (a) using  $X$  and  $\mathbf{y}$ .
  - (c) Find the optimal values for  $\boldsymbol{\beta}$ . For now, assume  $X^T X$  is invertible. This solution is called the *ordinary linear regression* solution.
  - (d) When there are more dimensions than examples, i.e., when  $d > n$ , will  $X^T X$  be invertible?
  - (e) If we add a regularizer  $\mathcal{R}(\boldsymbol{\beta}) = \boldsymbol{\beta}^T \boldsymbol{\beta}$  (with tradeoff parameter  $\lambda$ , and  $\lambda > 0$ ) to linear regression, what effect will that regularizer have? Linear regression with this regularizer is called the *ridge regression*, and this regularizer is a special case of the *Tikhonov regularization*.
  - (f) Express the optimization problem in ridge regression using  $X$ ,  $\mathbf{y}$ ,  $\boldsymbol{\beta}$  and  $\lambda$ . Find the solution.
  - (g) Ordinary linear regression will encounter difficulties when  $X^T X$  is not invertible. How will ridge regression help in this aspect?
  - (h) What will be the ridge regression solution if  $\lambda = 0$ ? What if  $\lambda = \infty$ ?
  - (i) Can we learn a good  $\lambda$  value by treating  $\lambda$  as a regular parameter (instead of a hyperparameter), that is, by minimizing the ridge regression loss function jointly over  $\lambda$  and  $\boldsymbol{\beta}$  on the training set (without using a validation set)?

3. (Polynomial regression) The polynomial regression model  $y = f(x) + \epsilon$  assumes the mapping  $f$  is a polynomial. A degree  $d$  polynomial is of the form

$$f(x) = \sum_{i=0}^d p_i x^i, \quad (33)$$

with  $d + 1$  parameters  $p_i$  ( $0 \leq i \leq d$ ). Use the ordinary linear regression to find the optimal parameters for polynomial regression. (Hint: Set the parameters of linear regression as  $\beta = (p_0, p_1, \dots, p_d)^T$ .)

4. ( $F_\beta$  measure) Answer the following two questions about the  $F_\beta$  measure.  
 (a) Prove that  $0 \leq F_\beta \leq 1$  for any  $\beta \geq 0$ .

(b) When  $\beta$  takes different values, the  $F_\beta$  measure places different relative importance between the precision and recall. Which one (precision or recall) is more important if  $\beta > 1$ ? Which one is more important if  $0 \leq \beta < 1$ ? (Hint: What is the speed of  $F_\beta$ 's change when the precision or recall changes?)

5. (AUC-PR and AP) We have not discussed the details of how the AUC-PR measurement is calculated. For a binary classification task, we assume every example  $\mathbf{x}$  has a score  $f(\mathbf{x})$ , and sort the test examples in the descending order of these scores. Then, for every example, we set the classification threshold as the current example's score (i.e., only this example and examples before it are classified as positive). A pair of precision and recall values are computed at this threshold. The PR curve is drawn by connecting nearby points using line segments. Then, AUC-PR is the area under the PR curve. Let  $(r_i, p_i)$  denote the  $i$ -th recall and precision rates ( $i = 1, 2, \dots$ ). When computing the area, the contribution between  $r_i$  and  $r_{i-1}$  is calculated using the *trapezoidal* interpolation  $(r_i - r_{i-1}) \frac{p_i + p_{i-1}}{2}$ , in which  $r_i - r_{i-1}$  is the length on the  $x$ -axis, and  $p_i$  and  $p_{i-1}$  are the lengths of two vertical lines in the  $y$ -axis. Summing over all  $i$  values, we obtain the AUC-PR score. Note that we assume the first pair  $(r_0, p_0) = (0, 1)$ , which is a pseudo-pair corresponding to the threshold  $+\infty$ .

(a) For the test set with 10 examples (indexed from 1 to 10) in Table 3, calculate the precision ( $p_i$ ) and recall ( $r_i$ ) when the threshold is set as the current example's  $f(x_i)$  value. Use class 1 as positive, and fill these values in Table 3. Fill the trapezoidal approximation  $(r_i - r_{i-1}) \frac{p_i + p_{i-1}}{2}$  in the "AUC-PR" column for the  $i$ -th row; and fill their sum to the last row.

(b) *Average precision* (AP) is another way to summarize the PR curve into one number. Similar to AUC-PR, AP approximates the contribution between  $r_i$  and  $r_{i-1}$  using a rectangle, as  $(r_i - r_{i-1})p_i$ . Fill in this approximation into the "AP" column for the  $i$ -th row; and fill their sum to the last row. Both AUC-PR and AP are summarizing the PR curve, hence they should be similar to each other. Are they?

(c) Both AUC-PR and AP are sensitive to the order of labels. If the label

Table 3: Calculation of AUC-PR and AP.

index	label	score	precision	recall	AUC-PR	AP
0			1.0000	0.0000	-	-
1	1	1.0				
2	2	0.9				
3	1	0.8				
4	1	0.7				
5	2	0.6				
6	1	0.5				
7	2	0.4				
8	2	0.3				
9	1	0.2				
10	2	0.1				
					(?)	(?)

of the 9-th and the 10-th rows are exchanged, what is the new AUC-PR and AP?

(d) Write a program to calculate both AUC-PR and AP based on the labels, scores, and the positive class. Validate your program's correctness using the example test set in Table 3.

- We can use the  $k$ -NN method for regression. Let  $D = \{\mathbf{x}_i, y_i\}_{i=1}^n$  be a training set, where the labels  $y \in \mathbb{R}$  is generated by  $y = F(\mathbf{x}) + \epsilon$ , in which the true regression function  $F$  is contaminated by noise  $\epsilon$  to generate the labels  $y$ . We assume the random noise  $\epsilon$  is independent of anything else,  $\mathbb{E}[X] = 0$  and  $\text{Var}(\epsilon) = \sigma^2$ . For any testing example  $\mathbf{x}$ , the  $k$ -NN method finds its  $k$  ( $k$  is a positive integer) nearest neighbors in  $D$ , denoted as  $\mathbf{x}_{nn(1)}, \mathbf{x}_{nn(2)}, \dots, \mathbf{x}_{nn(k)}$ , where  $1 \leq nn(i) \leq n$  is the index of the  $i$ -th nearest neighbor. Then, the prediction for  $\mathbf{x}$  is  $f(\mathbf{x}; D) = \frac{1}{k} \sum_{i=1}^k y_{nn(i)}$ .

  - What is the bias-variance decomposition for  $\mathbb{E}[(y - f(\mathbf{x}; D))^2]$ , in which  $y$  is the label for  $\mathbf{x}$ ? Do not use abbreviations (Equation 27 uses abbreviations, e.g.,  $\mathbb{E}[f]$  should be  $\mathbb{E}_D[f(\mathbf{x}; D)]$ .) Use  $\mathbf{x}$ ,  $y$ ,  $F$ ,  $f$ ,  $D$  and  $\sigma$  to express the decomposition.
  - Use  $f(\mathbf{x}; D) = \frac{1}{k} \sum_{i=1}^k y_{nn(i)}$  to compute  $\mathbb{E}[f]$  (abbreviation can be used from now on).
  - Replace the  $f$  term in the decomposition by using  $\mathbf{x}$  and  $y$ .
  - What is the variance term? How will it change when  $k$  changes?
  - What is the squared bias term? How will it change with  $k$ ? (Hint: consider  $k = n$ )?
- (Bayes decision theory) Consider a binary classification task, in which the label  $y \in \{1, 2\}$ . If an example  $x \in \mathbb{R}$  belongs to class 1, it is generated by the class conditional p.d.f.  $p(x|y = 1) = N(-1, 0.25)$ , and a class 2

example is sampled from the class conditional distribution  $p(x|y = 2) = N(1, 0.25)$ . Suppose  $\Pr(y = 1) = \Pr(y = 2) = 0.5$ .

- (a) What is the joint p.d.f.  $p(x, y)$ ?
- (b) Let us use the cost matrix  $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ . Show that for any  $x$ , if we choose  $f(x) = \arg \max_y p(y|x)$  to be our prediction for  $x$ , the cost  $\mathbb{E}_{(x,y)}[c_{y,f(x)}]$  is minimized, and hence it is the optimal solution. Is this rule optimal if  $y \in \{1, 2, \dots, C\}$  ( $C > 2$ ) (i.e., in a multi-class classification problem)?
- (c) Using the cost matrix  $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$  and the Bayes decision theory, which classification strategy will be optimal for this task? What is the Bayes risk in this example?
- (d) If the cost matrix is  $\begin{bmatrix} 0 & 10 \\ 1 & 0 \end{bmatrix}$  (i.e., when the true label is 1 but the prediction is 2, the cost is increased to 10). What is the new decision rule?

8. (Stratified sampling) Let  $D$  be a training set with only 10 examples, whose labels are 1, 1, 2, 2, 2, 2, 2, 2, 2, 2, respectively. This dataset is both small in size and imbalanced. We need cross-validation during evaluation, and 2-fold CV seems a good choice.

- (a) Write a program to *randomly* split this dataset into two subsets, with 5 examples in each subset. Repeat this random split 10 times. The histogram of class 1 examples in these two subsets can be (0, 2) or (1, 1)—one subset has zero (two) and the other has two (zero) class 1 examples, or every subset has exactly one class 1 example). In your 10 splits, how many times (0, 2) appears? (Note: this number can be different if you perform the experiments multiple times.)
- (b) What is the probability that (0, 2) will appear in one random split of these 10 examples?
- (c) In your 2-fold CV evaluation, if the split's class 1 distribution in the two subsets are (0, 2), how will it affect the evaluation?
- (d) One commonly used way to avoid this issue to use *stratified sampling*. In stratified sampling, we perform the train/test split for *every* class separately. Show that if stratified sampling is used, the distribution of class 1 examples will always be (1, 1).

9. (Confusion matrix) In a classification problem with  $K$  classes, the cost matrix is of size  $K \times K$ :  $C = [c_{ij}]$  (in which  $c_{ij}$  is the cost when one example belongs to class  $i$  but is predicted as in class  $j$ ). Similarly, a confusion matrix is a  $K \times K$  matrix:  $A = [a_{ij}]$ , in which  $a_{ij}$  is the number of class  $i$  examples that are classified as belonging to class  $j$ . Let the confusion matrix be computed based on a test set with  $N$  examples. We often normalize the confusion matrix to obtain  $\hat{A}$ , by  $\hat{a}_{ij} = \frac{a_{ij}}{\sum_{k=1}^K a_{ik}}$ . Hence, the sum of all elements in any row of  $\hat{A}$  equals 1. We call  $\hat{A}$  the

normalized confusion matrix.

(a) Prove that the total cost on the test set equals  $\text{tr}(C^T A)$ .

(b) In an imbalanced classification problem, do you prefer the confusion matrix or the normalized one? Why?

10. (McNemar's test) Find resources about McNemar's test.<sup>8</sup> *Carefully* read through these resources until you believe you have understood when this test can be applied and how to apply it to compare two classifiers.

---

<sup>8</sup>For example, in your school's library or search in the Internet.