

Statistical Learning

In order to analyze a learning algorithm, we must define a mathematical model of how examples (\mathbf{x}, y) are generated. In the statistical learning framework, we assume that every example (\mathbf{x}, y) is obtained through an independent draw from a fixed but unknown probability distribution on $\mathcal{X} \times \mathcal{Y}$. We write (\mathbf{X}, Y) to highlight that \mathbf{x} and y are random variables. The assumption that not all data points \mathbf{x} are equally likely is quite natural (for example, when data points are images, only a small fraction of all possible pixel configurations correspond to real-world images). Similarly, as we previously argued, labels are typically noisy. This corresponds to assuming that labels of any given datapoint are stochastic.

Assuming that every example (\mathbf{x}, y) is the realization of an independent random draw from the same joint probability distribution \mathcal{D} implies that every dataset (e.g., a training set or a test set) is a statistical sample. Note that the independence assumption is actually violated in many practical domains. Consider for example the problem of categorizing news stories. The newsfeed is clearly far from being an independent process, as the evolution of news reflects developing and related stories. Although not very realistic, the independence assumption is nevertheless convenient from the viewpoint of the analytical tractability of the problem, and works surprisingly well in practice.

In statistical learning, a problem is fully specified by a pair (\mathcal{D}, ℓ) , where \mathcal{D} is the data distribution and ℓ is a loss function. The performance of a predictor $h : \mathcal{X} \rightarrow \mathcal{Y}$ with respect to (\mathcal{D}, ℓ) is evaluated via the **statistical risk**, defined by

$$\ell_{\mathcal{D}}(h) = \mathbb{E}[\ell(Y, h(\mathbf{X}))]$$

This is the expected value of the loss function on a random example (\mathbf{X}, Y) drawn from \mathcal{D} . The best possible predictor $f^* : \mathcal{X} \rightarrow \mathcal{Y}$ given \mathcal{D} is known as **Bayes optimal predictor**, and is defined by

$$f^*(\mathbf{x}) = \operatorname{argmin}_{\hat{y} \in \mathcal{Y}} \mathbb{E}[\ell(Y, \hat{y}) \mid \mathbf{X} = \mathbf{x}]$$

The quantity $\mathbb{E}[\ell(Y, \hat{y}) \mid \mathbf{X} = \mathbf{x}]$ is the conditional risk, which is the expected loss of the prediction with respect to the distribution of the label Y conditioned on \mathbf{x} . Hence $f^*(x)$ is the prediction minimizing the conditional risk. By definition of f^* , we have that

$$\mathbb{E}[\ell(Y, f^*(X)) \mid \mathbf{X} = \mathbf{x}] \leq \mathbb{E}[\ell(Y, h(X)) \mid \mathbf{X} = \mathbf{x}]$$

for every predictor $h : \mathcal{X} \rightarrow \mathcal{Y}$ and for any $\mathbf{x} \in \mathcal{X}$. Because the above inequality holds for every $\mathbf{x} \in \mathcal{X}$, it also holds in expectation with respect to the random draw of \mathbf{X} . But since, for any predictor h ,

$$\mathbb{E}\left[\mathbb{E}[\ell(Y, h(X)) \mid \mathbf{X}]\right] = \mathbb{E}[\ell(Y, h(X))] = \ell_{\mathcal{D}}(h)$$

we have that $\ell_{\mathcal{D}}(f^*) \leq \ell_{\mathcal{D}}(h)$ for every predictor h . The risk $\ell_{\mathcal{D}}(f^*)$ of the Bayes optimal predictor is called **Bayes risk**. Typically, the Bayes risk is larger than zero because labels are stochastic.

We now compute the Bayes optimal predictor for the quadratic loss function $\ell(y, \hat{y}) = (y - \hat{y})^2$ when $\mathcal{Y} \equiv \mathbb{R}$,

$$\begin{aligned}
f^*(\mathbf{x}) &= \operatorname{argmin}_{\hat{y} \in \mathbb{R}} \mathbb{E}[(Y - \hat{y})^2 \mid \mathbf{X} = \mathbf{x}] \\
&= \operatorname{argmin}_{\hat{y} \in \mathbb{R}} \left(\mathbb{E}[Y^2 \mid \mathbf{X} = \mathbf{x}] + \hat{y}^2 - 2\hat{y} \mathbb{E}[Y \mid \mathbf{X} = \mathbf{x}] \right) \\
&= \operatorname{argmin}_{\hat{y} \in \mathbb{R}} \left(\hat{y}^2 - 2\hat{y} \mathbb{E}[Y \mid \mathbf{X} = \mathbf{x}] \right) \quad (\text{ignoring the term that does not depend on } \hat{y}) \\
&= \mathbb{E}[Y \mid \mathbf{X} = \mathbf{x}] \quad (\text{minimizing the function } F(\hat{y}) = \hat{y}^2 - 2\hat{y} \mathbb{E}[Y \mid \mathbf{X} = \mathbf{x}])
\end{aligned}$$

Thus, the Bayes optimal prediction for the quadratic loss function is the expected value of the label conditioned on the instance.

Substituting in the conditional risk formula $\mathbb{E}[(Y - f^*(\mathbf{X}))^2 \mid \mathbf{X} = \mathbf{x}]$ the Bayes optimal predictor $f^*(\mathbf{x}) = \mathbb{E}[Y \mid \mathbf{X} = \mathbf{x}]$ we obtain

$$\mathbb{E}[(Y - f^*(\mathbf{X}))^2 \mid \mathbf{X} = \mathbf{x}] = \mathbb{E}[(Y - \mathbb{E}[Y \mid \mathbf{x}])^2 \mid \mathbf{X} = \mathbf{x}] = \operatorname{Var}[Y \mid \mathbf{X} = \mathbf{x}] .$$

In words, the conditional risk of the Bayes optimal predictor for the quadratic loss is the variance of the label conditioned on the instance. By averaging over \mathbf{X} we obtain $\ell_{\mathcal{D}}(f^*) = \mathbb{E}[\operatorname{Var}[Y \mid \mathbf{X}]]$. Namely, the Bayes risk for the quadratic loss is the expected conditional variance of the label. Note that $\mathbb{E}[\operatorname{Var}[Y \mid \mathbf{X}]]$ is generally different from $\operatorname{Var}[Y]$. Indeed, the law of total variance says that $\operatorname{Var}[Y] - \mathbb{E}[\operatorname{Var}[Y \mid \mathbf{X}]] = \operatorname{Var}[\mathbb{E}[Y \mid \mathbf{X}]]$.

We now focus on binary classification, where $\mathcal{Y} = \{-1, 1\}$. Let $\eta(\mathbf{x})$ be the probability of $Y = 1$ conditioned on $\mathbf{X} = \mathbf{x}$. We view $\eta(\mathbf{x}) = \mathbb{P}(Y = +1 \mid \mathbf{X} = \mathbf{x})$ as the value on \mathbf{x} of a function $\eta : \mathcal{X} \rightarrow [0, 1]$.

Let $\mathbb{I}\{A\} \in \{0, 1\}$ be the indicator function of an event A ; that is, $\mathbb{I}\{A\} = 1$ if and only if A occurs. The statistical risk with respect to the zero-one loss $\ell(y, \hat{y}) = \mathbb{I}\{\hat{y} \neq y\}$ is therefore defined by

$$\ell_{\mathcal{D}}(h) = \mathbb{E}[\ell(Y, h(\mathbf{X}))] = \mathbb{E}[\mathbb{I}\{h(\mathbf{X}) \neq Y\}] = \mathbb{P}(h(\mathbf{X}) \neq Y) .$$

The Bayes optimal predictor $f^* : \mathcal{X} \rightarrow \{-1, 1\}$ for binary classification is derived as follows

$$\begin{aligned}
f^*(\mathbf{x}) &= \operatorname{argmin}_{\hat{y} \in \{-1, 1\}} \mathbb{E}[\ell(Y, \hat{y}) \mid \mathbf{X} = \mathbf{x}] \\
&= \operatorname{argmin}_{\hat{y} \in \{-1, 1\}} \mathbb{E}[\mathbb{I}\{Y = +1\} \mathbb{I}\{\hat{y} = -1\} + \mathbb{I}\{Y = -1\} \mathbb{I}\{\hat{y} = +1\} \mid \mathbf{X} = \mathbf{x}] \\
&= \operatorname{argmin}_{\hat{y} \in \{-1, 1\}} \left(\mathbb{P}(Y = +1 \mid \mathbf{X} = \mathbf{x}) \mathbb{I}\{\hat{y} = -1\} + \mathbb{P}(Y = -1 \mid \mathbf{X} = \mathbf{x}) \mathbb{I}\{\hat{y} = +1\} \right) \\
&= \operatorname{argmin}_{\hat{y} \in \{-1, 1\}} \left(\eta(\mathbf{x}) \mathbb{I}\{\hat{y} = -1\} + (1 - \eta(\mathbf{x})) \mathbb{I}\{\hat{y} = +1\} \right) \\
&= \begin{cases} -1 & \text{if } \eta(\mathbf{x}) < 1/2, \\ +1 & \text{if } \eta(\mathbf{x}) \geq 1/2. \end{cases}
\end{aligned}$$

Hence, the Bayes optimal classifier predicts the label whose probability is the highest when conditioned on the instance. Finally, it is easy to verify that the Bayes risk in this case is $\ell_{\mathcal{D}}(f^*) = \mathbb{E}[\min\{\eta(\mathbf{X}), 1 - \eta(\mathbf{X})\}]$.

Bounding the risk. Next, we study the problem of bounding the risk of a predictor. From now on, we assume $\ell(y, \hat{y}) \in [0, 1]$. However, keep in mind that our analysis continues to hold also when $\ell(y, \hat{y}) \in [0, M]$ for any $M > 0$.

It should be clear that, given an arbitrary predictor h , we cannot directly compute its risk $\ell_{\mathcal{D}}(h)$ with respect to \mathcal{D} because \mathcal{D} is typically unknown (if we knew \mathcal{D} , we could directly construct the Bayes optimal predictor). We thus consider the problem of estimating the risk of a given predictor h . In order to compute this estimate, we can use the **test set** $S' = \{(\mathbf{x}'_1, y'_1), \dots, (\mathbf{x}'_n, y'_n)\}$. We can then estimate $\ell_{\mathcal{D}}(h)$ with the **test error**, which is the average loss of h on the test set,

$$\ell_{S'}(h) = \frac{1}{n} \sum_{t=1}^n \ell(y'_t, h(\mathbf{x}'_t)) .$$

Under the assumption that the test set is generated through independent draws from \mathcal{D} , the test error corresponds to the **sample mean** of the risk. Indeed, for each $t = 1, \dots, n$ the example (\mathbf{X}'_t, Y'_t) is an independent draw from \mathcal{D} . Therefore,

$$\mathbb{E} \left[\ell(Y'_t, h(\mathbf{X}'_t)) \right] = \ell_{\mathcal{D}}(h) \quad t = 1, \dots, n$$

Note that the above equalities rely on the assumption that h does not depend on the test set. If it did, then the above equalities would not be necessarily true. This fact is important in the analysis of learning algorithms.

In order to compute how good is the test error as an estimate for the risk, we can use the following result about the law of large numbers.

Lemma 1 (Chernoff-Hoeffding). *Let Z_1, \dots, Z_n be independent and identically distributed random variables with expectation μ and such that $0 \leq Z_t \leq 1$ for each $t = 1, \dots, n$. Then, for any given $\varepsilon > 0$,*

$$\mathbb{P} \left(\frac{1}{n} \sum_{t=1}^n Z_t > \mu + \varepsilon \right) \leq e^{-2\varepsilon^2 n} \quad \text{and} \quad \mathbb{P} \left(\frac{1}{n} \sum_{t=1}^n Z_t < \mu - \varepsilon \right) \leq e^{-2\varepsilon^2 n} .$$

In the rest of this course, we repeatedly use the following facts:

- For any two events A and B , if $A \Rightarrow B$, then $\mathbb{P}(A) \leq \mathbb{P}(B)$
- (Union bound) For any collection A_1, \dots, A_n of (not necessarily disjoint) events,

$$\mathbb{P}(A_1 \cup \dots \cup A_n) \leq \sum_{i=1}^n \mathbb{P}(A_i)$$

If the events A_1, \dots, A_n are pairwise disjoint, then the union bound holds with equality.

Using the Chernoff-Hoeffding bound with $Z_t = \ell(y_t, h(x_t)) \in [0, 1]$ we can compute a confidence interval for the risk as follows (where the test error is written as ℓ instead of $\ell_{S'}$),

$$\begin{aligned} \mathbb{P} \left(|\ell_{\mathcal{D}}(h) - \ell(h)| > \varepsilon \right) &= \mathbb{P} \left(\ell_{\mathcal{D}}(h) - \ell(h) > \varepsilon \cup \ell(h) - \ell_{\mathcal{D}}(h) > \varepsilon \right) \\ &= \mathbb{P} \left(\ell_{\mathcal{D}}(h) - \ell(h) > \varepsilon \right) + \mathbb{P} \left(\ell(h) - \ell_{\mathcal{D}}(h) > \varepsilon \right) \leq 2e^{-2\varepsilon^2 n} \end{aligned} \quad (1)$$

where in the last step we applied the union bound to the disjoint events $\ell_{\mathcal{D}}(h) - \ell(h) > \varepsilon$ and $\ell(h) - \ell_{\mathcal{D}}(h) > \varepsilon$. Note that the probability is computed with respect to the random draw of the test set. This inequality shows that the probability that a test set gives a test error $\ell_{S'}(h)$ differing from the true risk $\ell_{\mathcal{D}}(h)$ for more than ε quickly decreases with the size n of the test set.

More specifically: if we set to $\delta \in (0, 1)$ the right-hand side of (1) and then solve for ε , we get that

$$|\ell_{\mathcal{D}}(h) - \ell_{S'}(h)| \leq \sqrt{\frac{1}{2n} \ln \frac{2}{\delta}}$$

holds with probability at least $1 - \delta$ with respect to the random draw of the test set.

The inequality (1) is telling us how to use a test set to estimate the risk of a classifier. More precisely, the inequality shows that the test set, which is how we measure in practice the performance of a classifier on unseen data, is close to the statistical risk with high probability.

Overfitting and underfitting. Fix a learning problem (\mathcal{D}, ℓ) and consider a generic learning algorithm A . Let \mathcal{H}_A be the set of predictors generated by A : $h \in \mathcal{H}_A$ if and only if there exists a finite training set S such that $A(S) = h$. For example, A could be an algorithm for training a neural network and \mathcal{H}_A is the set of predictors obtained by training the neural network using arbitrary training sets of finite size. Let h^* be any predictor with minimum risk $\ell_{\mathcal{D}}(h^*)$ in \mathcal{H}_A . That is,

$$\ell_{\mathcal{D}}(h^*) \leq \min_{h \in \mathcal{H}_A} \ell_{\mathcal{D}}(h)$$

Fix a training set S and let $h_S = A(S)$. The following is called the **bias-variance** decomposition:

$$\begin{aligned} \ell_{\mathcal{D}}(h_S) &= \ell_{\mathcal{D}}(h_S) - \ell_{\mathcal{D}}(h^*) && \text{estimation error (large when overfitting)} \\ &+ \ell_{\mathcal{D}}(h^*) - \ell_{\mathcal{D}}(f^*) && \text{approximation error (large when underfitting)} \\ &+ \ell_{\mathcal{D}}(f^*) && \text{Bayes error (unavoidable)} \end{aligned}$$

where f^* is the Bayes optimal predictor for (\mathcal{D}, ℓ) . Note that:

- the Bayes error is not controllable because it only depends on \mathcal{D} and the loss ℓ ,
- the approximation (or bias) error, which is large when the algorithm underfits, arises because \mathcal{H}_A does not necessarily contain the Bayes optimal predictor,
- the estimation (or variance) error, which is large when the algorithm overfits, arises because the risk of h_S is generally different from the risk of h^* .

We now use the bias-variance decomposition to balance overfitting and underfitting in the ERM algorithm run over a finite class \mathcal{H} . Recall that ERM minimizes the training error in \mathcal{H} ,

$$h_S = \operatorname{argmin}_{h \in \mathcal{H}} \ell_S(h)$$

where $\ell_S(h)$ is the training error of h on the training set S . Similarly to before, the best predictor in the class \mathcal{H} is any predictor h^* satisfying

$$\ell_{\mathcal{D}}(h^*) = \min_{h \in \mathcal{H}} \ell_{\mathcal{D}}(h) .$$

Thanks to the law of large numbers, we know that the training error $\ell_S(h^*)$ is close to $\ell_{\mathcal{D}}(h^*)$ with high probability with respect to the random draw of the training set S . Unfortunately, we cannot directly apply the Chernoff-Hoeffding bound to h_S to show that $\ell_{\mathcal{D}}(h_S)$ is close to $\ell_S(h_S)$. The reason is that h_S is a function of the training set, and thus a random variable. Chernoff-Hoeffding ensures that $\widehat{\ell}_S(h)$ is close to $\ell_{\mathcal{D}}(h)$ for any fixed h , whereas h_S is not fixed as it depends on the sample.

Next, we bound the variance error. For every given training set S of size m , we have that

$$\begin{aligned} \ell_{\mathcal{D}}(h_S) - \ell_{\mathcal{D}}(h^*) &= \ell_{\mathcal{D}}(h_S) - \ell_S(h_S) + \ell_S(h_S) - \ell_{\mathcal{D}}(h^*) \\ &\leq \ell_{\mathcal{D}}(h_S) - \ell_S(h_S) + \ell_S(h^*) - \ell_{\mathcal{D}}(h^*) \\ &\leq |\ell_{\mathcal{D}}(h_S) - \ell_S(h_S)| + |\ell_S(h^*) - \ell_{\mathcal{D}}(h^*)| \\ &\leq 2 \max_{h \in \mathcal{H}} |\ell_S(h) - \ell_{\mathcal{D}}(h)| \end{aligned}$$

where we used the assumption that h_S minimizes $\ell_S(h)$ among all $h \in \mathcal{H}$. Therefore, for all $\varepsilon > 0$,

$$\ell_{\mathcal{D}}(h_S) - \ell_{\mathcal{D}}(h^*) > \varepsilon \quad \Rightarrow \quad \max_{h \in \mathcal{H}} |\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2} \quad \Rightarrow \quad \exists h \in \mathcal{H} : |\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}.$$

Since the above chain of implications holds for any realization of the training set of size m , we can write

$$\mathbb{P}(\ell_{\mathcal{D}}(h_S) - \ell_{\mathcal{D}}(h^*) > \varepsilon) \leq \mathbb{P}\left(\exists h \in \mathcal{H} : |\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}\right).$$

We now study the case $|\mathcal{H}| < \infty$, that is when the model space contains a finite number of predictors. Note that the event

$$\exists h \in \mathcal{H} : |\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}$$

is the union over $h \in \mathcal{H}$ of the (not necessarily disjoint) events $|\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}$. Using the union bound we get

$$\begin{aligned} \mathbb{P}\left(\exists h \in \mathcal{H} : |\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}\right) &= \mathbb{P}\left(\bigcup_{h \in \mathcal{H}} \left(|\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}\right)\right) \\ &\leq \sum_{h \in \mathcal{H}} \mathbb{P}\left(|\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}\right) \\ &\leq |\mathcal{H}| \max_{h \in \mathcal{H}} \mathbb{P}\left(|\ell_S(h) - \ell_{\mathcal{D}}(h)| > \frac{\varepsilon}{2}\right) \\ &\leq |\mathcal{H}| 2e^{-m\varepsilon^2/2} \end{aligned} \tag{2}$$

where in the last step we used the Chernoff-Hoeffding bound.

In conclusion, we have that

$$\mathbb{P}(\ell_{\mathcal{D}}(h_S) - \ell_{\mathcal{D}}(h^*) > \varepsilon) \leq 2|\mathcal{H}|e^{-m\varepsilon^2/2}. \tag{3}$$

Setting the right-hand side of (3) equal to δ and solving for ε we obtain that

$$\ell_{\mathcal{D}}(h_S) \leq \ell_{\mathcal{D}}(h^*) + \sqrt{\frac{2}{m} \ln \frac{2|\mathcal{H}|}{\delta}}$$

holds with probability at least $1 - \delta$ with respect to the random draw of a training set of cardinality m .

For a given cardinality m of the training set, in order to decrease our bound on the variance error of ERM, $\sqrt{\frac{2}{m} \ln \frac{2|\mathcal{H}|}{\delta}}$ we must decrease $|\mathcal{H}|$. But decreasing $|\mathcal{H}|$ might cause an increase of $\ell_{\mathcal{D}}(h^*)$, which produces a corresponding increase of the bias error. In light of this statistical analysis, we conclude that the ERM algorithm generates predictors with high risk (compared to Bayes risk) when there is an unbalance between the variance error and the bias error. In particular, overfitting occurs when the variance error dominates the bias error, and underfitting occurs when the bias error dominates the variance error.

In the proof of the bound on the variance error, we have also shown in (2) that

$$\forall h \in \mathcal{H} \quad |\ell_S(h) - \ell_{\mathcal{D}}(h)| \leq \sqrt{\frac{1}{2m} \ln \frac{2|\mathcal{H}|}{\delta}}$$

with probability at least $1 - \delta$ with respect to the random draw of the training set. This implies that when the cardinality of the training set is sufficiently large with respect to $\ln |\mathcal{H}|$, then the training error $\ell_S(h)$ becomes a good estimate for the statistical risk $\ell_{\mathcal{D}}(h)$ *simultaneously* for all predictors $h \in \mathcal{H}$. This is sufficient to prevent overfitting, as it tells us the ranking the predictors in \mathcal{H} according to their training error approximately corresponds to ranking them according to their risk.