

Assuming each internal node has two children,¹ it is not hard to show that this is a prefix-free encoding of the tree, and that the description length of a tree with n nodes is $(n+1)\log_2(d+3)$.

By Theorem 7.7 we have that with probability of at least $1 - \delta$ over a sample of size m , for every n and every decision tree $h \in \mathcal{H}$ with n nodes it holds that

$$L_{\mathcal{D}}(h) \leq L_S(h) + \sqrt{\frac{(n+1)\log_2(d+3) + \log(2/\delta)}{2m}}. \quad (18.1)$$

This bound performs a tradeoff: on the one hand, we expect larger, more complex decision trees to have a smaller training risk, $L_S(h)$, but the respective value of n will be larger. On the other hand, smaller decision trees will have a smaller value of n , but $L_S(h)$ might be larger. Our hope (or prior knowledge) is that we can find a decision tree with both low empirical risk, $L_S(h)$, and a number of nodes n not too high. Our bound indicates that such a tree will have low true risk, $L_{\mathcal{D}}(h)$.

18.2 DECISION TREE ALGORITHMS

The bound on $L_{\mathcal{D}}(h)$ given in Equation (18.1) suggests a learning rule for decision trees – search for a tree that minimizes the right-hand side of Equation (18.1). Unfortunately, it turns out that solving this problem is computationally hard.² Consequently, practical decision tree learning algorithms are based on heuristics such as a greedy approach, where the tree is constructed gradually, and locally optimal decisions are made at the construction of each node. Such algorithms cannot guarantee to return the globally optimal decision tree but tend to work reasonably well in practice.

A general framework for growing a decision tree is as follows. We start with a tree with a single leaf (the root) and assign this leaf a label according to a majority vote among all labels over the training set. We now perform a series of iterations. On each iteration, we examine the effect of splitting a single leaf. We define some “gain” measure that quantifies the improvement due to this split. Then, among all possible splits, we either choose the one that maximizes the gain and perform it, or choose not to split the leaf at all.

In the following we provide a possible implementation. It is based on a popular decision tree algorithm known as “ID3” (short for “Iterative Dichotomizer 3”). We describe the algorithm for the case of binary features, namely, $\mathcal{X} = \{0, 1\}^d$, and therefore all splitting rules are of the form $\mathbb{1}_{[x_i=1]}$ for some feature $i \in [d]$. We discuss the case of real valued features in Section 18.2.3.


The algorithm works by recursive calls, with the initial call being $\text{ID3}(S, [d])$, and returns a decision tree. In the pseudocode that follows, we use a call to a procedure $\text{Gain}(S, i)$, which receives a training set S and an index i and evaluates the gain of a split of the tree according to the i th feature. We describe several gain measures in Section 18.2.1.

¹ We may assume this without loss of generality, because if a decision node has only one child, we can replace the node by its child without affecting the predictions of the decision tree.

² More precisely, if $\text{NP} \neq \text{P}$ then no algorithm can solve Equation (18.1) in time polynomial in n, d , and m .

ID3(S, A)

INPUT: training set S , feature subset $A \subseteq [d]$
if all examples in S are labeled by 1, return a leaf 1
if all examples in S are labeled by 0, return a leaf 0
if $A = \emptyset$, return a leaf whose value = majority of labels in S
else :
 Let $j = \operatorname{argmax}_{i \in A} \operatorname{Gain}(S, i)$
 if all examples in S have the same label
 Return a leaf whose value = majority of labels in S
 else
 Let T_1 be the tree returned by ID3($\{(\mathbf{x}, y) \in S : x_j = 1\}, A \setminus \{j\}$).
 Let T_2 be the tree returned by ID3($\{(\mathbf{x}, y) \in S : x_j = 0\}, A \setminus \{j\}$).
 Return the tree:



18.2.1 Implementations of the Gain Measure

Different algorithms use different implementations of $\operatorname{Gain}(S, i)$. Here we present three. We use the notation $\mathbb{P}_S[F]$ to denote the probability that an event holds with respect to the uniform distribution over S .

Train Error: The simplest definition of gain is the decrease in training error. Formally, let $C(a) = \min\{a, 1 - a\}$. Note that the training error before splitting on feature i is $C(\mathbb{P}_S[y = 1])$, since we took a majority vote among labels. Similarly, the error after splitting on feature i is

$$\mathbb{P}_S[x_i = 1] C(\mathbb{P}_S[y = 1 | x_i = 1]) + \mathbb{P}_S[x_i = 0] C(\mathbb{P}_S[y = 1 | x_i = 0]).$$

Therefore, we can define Gain to be the difference between the two, namely,

$$\operatorname{Gain}(S, i) := C(\mathbb{P}_S[y = 1]) - \left(\mathbb{P}_S[x_i = 1] C(\mathbb{P}_S[y = 1 | x_i = 1]) + \mathbb{P}_S[x_i = 0] C(\mathbb{P}_S[y = 1 | x_i = 0]) \right).$$

Information Gain: Another popular gain measure that is used in the ID3 and C4.5 algorithms of Quinlan (1993) is the information gain. The information gain is the difference between the entropy of the label before and after the split, and is achieved by replacing the function C in the previous expression by the entropy function,

$$C(a) = -a \log(a) - (1 - a) \log(1 - a).$$

Gini Index: Yet another definition of a gain, which is used by the CART algorithm of Breiman, Friedman, Olshen, and Stone (1984), is the Gini index,

$$C(a) = 2a(1 - a).$$

Both the information gain and the Gini index are smooth and concave upper bounds of the train error. These properties can be advantageous in some situations (see, for example, Kearns & Mansour (1996)).

18.2.2 Pruning

The ID3 algorithm described previously still suffers from a big problem: The returned tree will usually be very large. Such trees may have low empirical risk, but their true risk will tend to be high – both according to our theoretical analysis, and in practice. One solution is to limit the number of iterations of ID3, leading to a tree with a bounded number of nodes. Another common solution is to *prune* the tree after it is built, hoping to reduce it to a much smaller tree, but still with a similar empirical error. Theoretically, according to the bound in Equation (18.1), if we can make n much smaller without increasing $L_S(h)$ by much, we are likely to get a decision tree with a smaller true risk.

Usually, the pruning is performed by a bottom-up walk on the tree. Each node might be replaced with one of its subtrees or with a leaf, based on some bound or estimate of $L_D(h)$ (for example, the bound in Equation (18.1)). A pseudocode of a common template is given in the following.

Generic Tree Pruning Procedure

input:

function $f(T, m)$ (bound/estimate for the generalization error of a decision tree T , based on a sample of size m),
tree T .

foreach node j in a bottom-up walk on T (from leaves to root):
 find T' which minimizes $f(T', m)$, where T' is any of the following:
 the current tree after replacing node j with a leaf 1.
 the current tree after replacing node j with a leaf 0.
 the current tree after replacing node j with its left subtree.
 the current tree after replacing node j with its right subtree.
 the current tree.
 let $T := T'$.

18.2.3 Threshold-Based Splitting Rules for Real-Valued Features

In the previous section we have described an algorithm for growing a decision tree assuming that the features are binary and the splitting rules are of the form $\mathbb{1}_{[x_i=1]}$. We now extend this result to the case of real-valued features and threshold-based splitting rules, namely, $\mathbb{1}_{[x_i < \theta]}$. Such splitting rules yield decision stumps, and we have studied them in Chapter 10.

The basic idea is to reduce the problem to the case of binary features as follows. Let $\mathbf{x}_1, \dots, \mathbf{x}_m$ be the instances of the training set. For each real-valued feature i , sort the instances so that $x_{1,i} \leq \dots \leq x_{m,i}$. Define a set of thresholds $\theta_{0,i}, \dots, \theta_{m+1,i}$ such that $\theta_{j,i} \in (x_{j,i}, x_{j+1,i})$ (where we use the convention $x_{0,i} = -\infty$ and $x_{m+1,i} = \infty$). Finally, for each i and j we define the binary feature $\mathbb{1}_{[x_i < \theta_{j,i}]}$. Once we have constructed these binary features, we can run the ID3 procedure described in the previous section. It is easy to verify that for any decision tree with threshold-based splitting rules over the original real-valued features there exists a decision tree over the constructed binary features with the same training error and the same number of nodes.

If the original number of real-valued features is d and the number of examples is m , then the number of constructed binary features becomes dm . Calculating the Gain of each feature might therefore take $O(dm^2)$ operations. However, using a more clever implementation, the runtime can be reduced to $O(dm \log(m))$. The idea is similar to the implementation of ERM for decision stumps as described in Section 10.1.1.

18.3 RANDOM FORESTS

As mentioned before, the class of decision trees of arbitrary size has infinite VC dimension. We therefore restricted the size of the decision tree. Another way to reduce the danger of overfitting is by constructing an ensemble of trees. In particular, in the following we describe the method of *random forests*, introduced by Breiman (2001).

A random forest is a classifier consisting of a collection of decision trees, where each tree is constructed by applying an algorithm A on the training set S and an additional random vector, θ , where θ is sampled i.i.d. from some distribution. The prediction of the random forest is obtained by a majority vote over the predictions of the individual trees.

To specify a particular random forest, we need to define the algorithm A and the distribution over θ . There are many ways to do this and here we describe one particular option. We generate θ as follows. First, we take a random subsample from S with replacements; namely, we sample a new training set S' of size m' using the uniform distribution over S . Second, we construct a sequence I_1, I_2, \dots , where each I_i is a subset of $[d]$ of size k , which is generated by sampling uniformly at random elements from $[d]$. All these random variables form the vector θ . Then, the algorithm A grows a decision tree (e.g., using the ID3 algorithm) based on the sample S' , where at each splitting stage of the algorithm, the algorithm is restricted to choosing a feature that maximizes Gain from the set I_i . Intuitively, if k is small, this restriction may prevent overfitting.

18.4 SUMMARY

Decision trees are very intuitive predictors. Typically, if a human programmer creates a predictor it will look like a decision tree. We have shown that the VC dimension of decision trees with k leaves is k and proposed the MDL paradigm for