Chapter 11

Decision Trees

A *decision tree* is a model of computation used to study the number of bits of an input that need to be examined in order to compute some function on this input. Consider a function $f : \{0,1\}^n \to \{0,1\}$. A decision tree for f is a tree for which each node is labelled with some x_i, and has two outgoing edges, labelled 0 and 1. Each tree leaf is labelled with an output value 0 or 1. The computation on input $x = x_1 x_2 \ldots x_n$ proceeds at each node by inspecting the input bit x_i indicated by the node's label. If $x_i = 1$ the computation continues in the subtree reached by taking the 1-edge. The 0-edge is taken if the bit is 0. Thus input x follows a path through the tree. The output value at the leaf is $f(x)$. An example of a simple decision tree for the majority function is given in Figure 11.1

Figure unavailable in pdf file.

Figure 11.1: A decision tree for computing the majority function $Maj(x_1, x_2, x_3)$ on three bits. Outputs 1 if at least two input bits are 1, else outputs 0.

Recall the use of decision trees in the proof of the lower bound for comparison-based sorting algorithms. That study can be recast in the above framework by thinking of the input —which consisted of n numbers — as consisting of $\binom{n}{2}$ bits, each giving the outcome of a pairwise comparison between two numbers.

We can now define two useful decision tree metrics.

Definition 11.1

The cost of tree t on input x, $cost(t, x)$, is the number of bits of x examined by t.

Definition 11.2
The **decision tree complexity** of function f, $D(f)$, is defined as follows, where T below refers to the set of decision trees that decide f.

$$D(f) = \min_{t \in T} \max_{x \in \{0,1\}^n} cost(t, x) \tag{1}$$

The decision tree complexity of a function is the number of bits examined by the most efficient decision tree on the worst case input to that tree. We are now ready to consider several examples.

Example 11.3
(*Graph connectivity*) Given a graph G as input, in adjacency matrix form, we would like to know how many bits of the adjacency matrix a decision tree algorithm might have to inspect in order to determine whether G is connected. We have the following result.

Theorem 11.4
Let f be a function that computes the connectivity of input graphs with m vertices. Then $D(f) = \binom{m}{2}$.

The idea of the proof of this theorem is to imagine an adversary that constructs a graph, edge by edge, in response to the queries of a decision tree. For every decision tree that decides connectivity, the strategy implicitly produces an input graph which requires the decision tree to inspect each of the $\binom{m}{2}$ possible edges in a graph of m vertices.

Adversary Strategy:
Whenever the decision tree algorithm asks about edge e_i,
answer "no" unless this would force the graph to be disconnected.

After i queries, let N_i be the set of edges for which the adversary has replied "no", Y_i the set of edges for which the adversary has replied "yes". and E_i the set of edges not yet queried. The adversary's strategy maintains the invariant that Y_i is a disconnected forest for $i < \binom{m}{2}$ and $Y_i \cup E_i$ is connected. This ensures that the decision tree will not know whether the graph is connected until it queries every edge.

Example 11.5

(*OR Function*) Let $f(x_1, x_2, \ldots x_n) = \bigvee_{i=1}^{n} x_i$. Here we can use an adversary argument to show that $D(f) = n$. For any decision tree query of an input bit x_i, the adversary responds that x_i equals 0 for the first $n-1$ queries. Since f is the OR function, the decision tree will be in suspense until the value of the nth bit is revealed. Thus $D(f)$ is n.

Example 11.6

Consider the AND-OR function, with $n = 2^k$. We define f_k as follows.

$$f_k(x_1, \ldots, x_n) = \begin{cases} f_{k-1}(x_1, \ldots x_{2^{k-1}-1}) \wedge f_{k-1}(x_{2^{k-1}}, \ldots x_{2^k}) & \text{if } k \text{ is even} \\ f_{k-1}(x_1, \ldots x_{2^{k-1}-1}) \vee f_{k-1}(x_{2^{k-1}}, \ldots x_{2^k}) & \text{if } k > 1 \text{ and is odd} \\ x_i & \text{if } k = 1 \end{cases} \tag{2}$$

A diagram of a circuit that computes the AND-OR function is shown in Figure 11.2. It is left as an exercise to prove, using induction, that $D(f_k) = 2^k$.

Figure unavailable in pdf file.

Figure 11.2: A circuit showing the computation of the AND-OR function. The circuit has k layers of alternating gates, where $n = 2^k$.

11.1 Certificate Complexity

We now introduce the notion of *certificate complexity*, which, in a manner analogous to decision tree complexity above, tells us the minimum amount of information needed to be convinced of the value of a function f on input x.

Definition 11.7

Consider a function $f : \{0,1\}^n \to \{0,1\}$. If $f(x) = 0$, then a **0-certificate** for x is a sequence of bits in x that proves $f(x) = 0$. If $f(x) = 1$, then a **1-certificate** is a sequence of bits in x that proves $f(x) = 1$.

Definition 11.8
The **certificate complexity** $C(f)$ of f is defined as follows.

$$C(f) = \max_{x:input} \{\text{number of bits in the smallest 0- or 1- certificate for x}\} \tag{3}$$

Example 11.9
If f is a function that decides connectivity of a graph, a 0-certificate for an input must prove that some cut in the graph has no edges, hence it has to contain all the possible edges of a cut of the graph. When these edges do not exist, the graph is disconnected. Similarly, a 1-certificate is the edges of a spanning tree. Thus for those inputs that represent a connected graph, the minimum size of a 1-certificate is the number of edges in a spanning tree, $n-1$. For those that represent a disconnected graph, a 0 certificate is the set of edges in a cut. The size of a 0-certificate is at most $(n/2)^2 = n^2/4$, and there are graphs (such as the graph consisting of two disjoint cliques of size $n/2$) in which no smaller 0-certificate exists. Thus $C(f) = n^2/4$.

Example 11.10
We show that the certificate complexity of the AND-OR function f_k of Example 11.6 is $2^{\lceil k/2 \rceil}$. Recall that f_k is defined using a circuit of k layers. Each layer contains only OR-gates or only AND-gates, and the layers have alternative gate types. The bottom layer receives the bits of input x as input and the single top layer gate outputs the answer $f_k(x)$. If $f(x) = 1$, we can construct a 1-certificate as follows. For every AND-gate in the tree of gates we have to prove that both its children evaluate to 1, whereas for every OR-gate we only need to prove that *some* child evaluates to 1. Thus the 1-certificate is a subtree in which the AND-gates have two children but the OR gates only have one each. Thus the subtree only needs to involve $2^{\lceil k/2 \rceil}$ input bits. If $f(x) = 0$, a similar argument applies, but the role of OR-gates and AND-gates, and values 1 and 0 are reversed. The result is that the certificate complexity of f_k is $2^{\lceil k/2 \rceil}$, or about $\sqrt{n}$.

The following is a rough way to think about these concepts in analogy to Turing machine complexity as we have studied it.

$$\text{low decision tree complexity} \leftrightarrow \quad \mathbf{P} \tag{4}$$
$$\text{low 1-certificate complexity} \leftrightarrow \quad \mathbf{NP} \tag{5}$$
$$\text{low 0-certificate complexity} \leftrightarrow \quad \mathbf{coNP} \tag{6}$$

The following result shows, however, that the analogy may not be exact since in the decision tree world, $\mathbf{P} = \mathbf{NP} \cap \mathbf{coNP}$. It should be noted that the result is tight, for example for the AND-OR function.

Theorem 11.11
For function f, $D(f) \leq C(f)^2$.

Proof: Let S_0, S_1 be the set of minimal 0-certificates and 1-certificates, respectively, for f. Let $k = C(f)$, so each certificate has at most k bits.

Remark 11.12
Note that every 0-certificate must share a bit position with every 1-certificate, and furthermore, assign this bit differently. If this were not the case, then it would be possible for both a 0-certificate and 1-certificate to be asserted at the same time, which is impossible.

The following decision tree algorithm then determines the value of f in at most k^2 queries.

Algorithm: Repeat until the value of f is determined: Choose a remaining 0-certificate from S_0 and query all the bits in it. If the bits are the values that prove the f to be 0, then stop. Otherwise, we can prune the set of remaining certificates as follows. Since all 1-certificates must intersect the chosen 0-certificate, for any $c_1 \in S_1$, one bit in c_1 must have been queried here. Eliminate c_1 from consideration if the certifying value of c_1 at at location is different from the actual value found. Otherwise, we only need to consider the remaining $k - 1$ bits of c_1.

This algorithm can repeat at most k times. For each iteration, the unfixed lengths of the uneliminated 1-certificates decreases by one. This is because once some values of the input have been fixed due to queries, for any 0-certificate, it remains true that all 1-certificates must intersect it in at least one location that has not been fixed, otherwise it would be possible for both a 0-certificate and a 1-certificate to be asserted. With at most k queries for at most k iterations, a total of k^2 queries is used. ■

11.2 Randomized Decision Trees

There are two equivalent ways to look at randomized decision trees. We can consider decision trees in which the branch taken at each node is determined by the query value and by a random coin flip. We can also consider probability distributions over deterministic decision trees. The analysis that follows uses the latter model.

We will call $\mathcal{P}$ a probability distribution over a set of decision trees $\mathcal{T}$ that compute a particular function. $\mathcal{P}(t)$ is then the probability that tree t is chosen from the distribution. For a particular input x, then, we define $c(\mathcal{P}, x) = \sum_{tin\mathcal{T}} \mathcal{P}(t)cost(t, x)$. $c(\mathcal{P}, x)$ is thus the expected number of queries a tree chosen from $\mathcal{T}$ will make on input x. We can then characterize how well randomized decision trees can operate on a particular problem.

Definition 11.13
The **randomized decision tree complexity**, $\mathcal{R}(f)$, of f, is defined as follows.

$$\mathcal{R}(f) = \min_{\mathcal{P}} \max_{x} c(\mathcal{P}, x) \tag{7}$$

The randomized decision tree complexity thus expresses how well the best possible probability distribution of trees will do against the worst possible input for a particular probability distribution of trees. We can observe immediately that $\mathcal{R}(f) \geq C(f)$. This is because $C(f)$ is a minimum value of $cost(t, x)$. Since $\mathcal{R}(f)$ is just an expected value for a particular probability distribution of these *cost* values, the minimum such value can be no greater than the expected value.

Example 11.14
Consider the majority function, $f = Maj(x_1, x_2, x_3)$. It is straightforward to see that $D(f) = 3$. We show that $\mathcal{R}(f) \leq 8/3$. Let $\mathcal{P}$ be a uniform distribution over the (six) ways of ordering the queries of the three input bits. Now if all three bits are the same, then regardless of the order chosen, the decision tree will produce the correct answer after two queries. For such x, $c(\mathcal{P}, x) = 2$. If two of the bits are the same and the third is different, then there is a 1/3 probability that the chosen decision tree will choose the two similar bits to query first, and thus a 1/3 probability that the cost will be 2. There thus remains a 2/3 probability that all three bits will need to be inspected. For such x, then, $c(\mathcal{P}, x) = 8/3$. Therefore, $\mathcal{R}(f)$ is at most 8/3.

How can we prove *lowerbounds* on randomized complexity? For this we need another concept.

11.3 Lowerbounds on Randomized Complexity

NEEDS CLEANUP NOW

To prove lowerbounds on randomized complexity, it suffices by Yao's Lemma (see Section 11.6) to prove lowerbounds on *distributional complexity*. Where randomized complexity explores distributions over the space of decision trees for a problem, distributional complexity considers probability distributions on inputs. It is under such considerations that we can speak of "average case analysis."

Let $\mathcal{D}$ be a probability distribution over the space of input strings of length n. Then, if A is a deterministic algorithm, such as a decision tree, for a function, then we define the distributional complexity of A on a function f with inputs distributed according to $\mathcal{D}$ as the expected cost for algorithm A to compute f, where the expectation is over the distribution of inputs.

DEFINITION 11.15
The **distributional complexity** $d(A, \mathcal{D})$ of algorithm A given inputs distributed according to $\mathcal{D}$ is defined as:

$$d(A, \mathcal{D}) = \sum_{x:input} \mathcal{D}(x) cost(A, x) = \mathbf{E}_{x \in \mathcal{D}}[cost(A, x)] \tag{8}$$

From this we can characterize distributional complexity as a function of a single function f itself.

DEFINITION 11.16
The **distributional** decision tree complexity, $\Delta(f)$ of function f is defined as:

$$\Delta(f) = \max_{\mathcal{D}} \min_{A} d(A, \mathcal{D}) \tag{9}$$

Where A above runs over the set of decision trees that are deciders for f.

So the distributional decision tree complexity measures the expected efficiency of the most efficient decision tree algorithm works given the worst case distribution of inputs.

The following theorem follows from Yao's lemma.

THEOREM 11.17
$\mathcal{R}(f) = \Delta(f)$.

So in order to find a lower bound on some randomized algorithm, it suffices to find a lower bound on $\Delta(f)$. Such a lower bound can be found by postulating an input distribution $\mathcal{D}$ and seeing whether every algorithm has expected cost at least equal to the desired lower bound.

Example 11.18
We return to considering the majority function, and we seek to find a lower bound on $\Delta(f)$. Consider a distribution over inputs such that inputs in which all three bits match, namely 000 and 111, occur with probability 0. All other inputs occur with probability 1/6. For any decision tree, that is, for any order in which the three bits are examined, there is exactly a 1/3 probability that the first two bits examined will be the same value, and thus there is a 1/3 probability that the cost is 2. There is then a 2/3 probability that the cost is 3. Thus the overall expected cost for this distribution is 8/3. This implies that $\Delta(f) \geq 8/3$ and in turn that $\mathcal{R}(f) \geq 8/3$. So $\Delta(f) = \mathcal{R}(f) = 8/3$.

11.4 Some techniques for decision tree lowerbounds

Definition 11.19 (Sensitivity)
If $f:\{0,1\}^n \to \{0,1\}$ is a function and $x \in \{0,1\}^n$ then the *sensitivity of f on x*, denoted $s_x(f)$, is the number of bit positions i such that $f(x) \neq f(x^i)$, where x^i is x with its ith bit flipped. The *sensitivity* of f, denoted $s(f)$, is $\max_x \{s_x(f)\}$.

The *block sensitivity of f on x*, denoted $bs_x(f)$, is the maximum number b such that there are disjoint blocks of bit positions $B_{1,2}, \ldots, B_b$ such that $f(x) \neq f(x^{B_i})$ where x^{B_i} is x with all its bits flipped in block B_i. The *block sensitivity* of f denoted $bs(f)$ is $\max_x \{bs_x(f)\}$.

It is conjectured that there is a constant c (as low as 2) such that $bs(f) = O(s(f)^c)$ for all f but this is wide open. The following easy observation is left as an exercise.

Lemma 11.20
For any function, $s(f) \leq bs(f) \leq D(f)$.

Theorem 11.21 (Nisan)
$C(f) \leq s(f)bs(f)$.

PROOF: For any input $x \in \{0,1\}^n$ we describe a certificate for x of size $s(f)bs(f)$. This certificate is obtained by considering the largest number of disjoint blocks of variables $B_1, B_2, \ldots, B_b$ that achieve $b = bs_x(f) \leq bs(f)$. We claim that setting these variables according to x constitutes a certificate for x.

Suppose not, and let x' be an input that is consistent with the above certificate. Let B_{b+1} be a block of variables such that $x' = x^{B_{b+1}}$. Then B_{b+1} must be disjoint from $B_1, B_2, \ldots B_b$, which contradicts $b = bs_x(f)$.

Note that each of $B_1, B_2, \ldots, B_b$ has size at most $s(f)$ by definition of $s(f)$, and hence the size of the certificate we have exhibited is at most $s(f)bs(f)$. ■

Recent work on decision tree lowerbounds has used *polynomial representations* of boolean functions. Recall that a multilinear polynomial is a polynomial whose degree in each variable is 1.

DEFINITION 11.22
An n-variate polynomial $p(x_1, x_2, \ldots, x_n)$ *represents* $f : \{0,1\}^n \to \{0,1\}$ if $p(x) = f(x)$ for all $x \in \{0,1\}^n$.

The *degree* of f, denoted $deg(f)$, is the degree of the multilinear polynomial that represents f.

(The exercises ask you to show that the multilinear polynomial representation is unique, so $deg(f)$ is well-defined.)

EXAMPLE 11.23
The AND of n variables $x_1, x_2, \ldots, x_n$ is represented by the multilinear polynomial $\prod_{i=1}^n x_i$ and OR is represented by $1 - \prod_{i=1}^n (1 - x_i)$.

The degree of AND and OR is n, and so is their decision tree complexity. There is a similar connection for other problems too, but it is not as tight. The first part of the next theorem is an easy exercise; the second part is nontrivial.

THEOREM 11.24

1. *$deg(f) \leq D(f)$.*
2. *(Nisan-Smolensky) $D(f) \leq deg(f)^2 bs(f) \leq O(deg(f)^4)$.*

11.5 Comparison trees and sorting lowerbounds

TO BE WRITTEN

11.6 Yao's MinMax Lemma

This section presents Yao's minmax lemma, which is used in a variety of settings to prove lowerbounds on randomized algorithms. Therefore we present it in a very general setting.

Let $\mathcal{X}$ be a finite set of inputs and $\mathcal{A}$ be a finite set of algorithms that solve some computational problem on these inputs. For $x \in \mathcal{X}, a \in \mathcal{A}$, we denote by $\mathrm{cost}(A, x)$ the *cost* incurred by algorithm A on input x. A *randomized algorithm* is a probability distribution $\mathcal{R}$ on $\mathcal{A}$. The *cost* of $\mathcal{R}$ on input x, denoted $\mathrm{cost}(\mathcal{R}, x)$, is $E_{A\in\mathcal{R}}[\mathrm{cost}(A, x)]$. The *randomized complexity* of the problem is

$$\min_{\mathcal{R}} \max_{x\in\mathcal{X}} \mathrm{cost}(\mathcal{R}, x). \tag{10}$$

Let $\mathcal{D}$ be a distribution on inputs. For any deterministic algorithm A, the cost incurred by it on $\mathcal{D}$, denoted $\mathrm{cost}(A, \mathcal{D})$, is $E_{x\in\mathcal{D}}[\mathrm{cost}(A, x)]$. The *distributional complexity* of the problem is

$$\max_{\mathcal{D}} \min_{A\in\mathcal{A}} \mathrm{cost}(A, \mathcal{D}). \tag{11}$$

Yao's Lemma says that these two quantitities are the same. It is easily derived from von Neumann's minmax theorem for zero-sum games, or with a little more work, from linear programming duality.

Yao's lemma is typically used to lowerbound randomized complexity. To do so, one defines (using some insight and some luck) a suitable distribution $\mathcal{D}$ on the inputs. Then one proves that *every* deterministic algorithm incurs high cost, say C, on this distribution. By Yao's Lemma, it follows that the randomized complexity then is at least C.

Exercises

§1 Suppose f is any function that depends on all its bits; in other words, for each bit position i there is an input x such that $f(x) \neq f(x^i)$. Show that $s(f) = \Omega(\log n)$.

§2 Consider an f defined as follows. The n-bit input is partitioned into $\lfloor\sqrt{n}\rfloor$ blocks of size about $\sqrt{n}$. The function is 1 iff there is at least one block in which two consecutive bits are 1 and the remaining bits in the block are 0. Estimate $s(f), bs(f), C(f), D(f)$ for this function.

§3 Show that there is a unique multilinear polynomial that represents $f:\{0,1\}^n \to \{0,1\}$. Use this fact to find the multilinear representation of the PARITY of n variables.

§4 Show that $deg(f) \leq D(f)$.

Chapter notes and history

The result that the decision tree complexity of connectivity and many other problems is $\binom{n}{2}$ has motivated the following conjecture (atributed variously to Anderaa, Karp, Yao):

Every monotone graph property has $D(\cdot) = \binom{n}{2}$.

Here "monotone" means that adding edges to the graph cannot make it go from having the property to not having the property (e.g., connectivity). "Graph property" means that the property does not depend upon the vertex indices (e.g., the property that vertex 1 and vertex 2 have an edge between them). This conjecture is known to be true up to a $O(1)$ factor; the proof uses topology and is excellently described in Du and Ko [**?**]. A more ambitious conjecture is that even the randomized decision tree complexity of monotone graph properties is $\Omega(n^2)$ but here the best lowerbound is close to $n^{4/3}$.

The polynomial method for decision tree lowerbounds is surveyed in Buhrman and de Wolf [**?**]. The method can be used to lowerbound randomized decision tree complexity (and more recently, *quantum* decision tree complexity) but then one needs to consider polynomials that *approximately* represent the function.