7 Algebraic Decision Trees (February 27 and March 4)

7.1 How Hard Is Almost Sorting?

Almost everyone is familiar with the $\Omega(n \log n)$ decision tree lower bound for sorting. Any binary decision tree that sorts must have at least $n!$ leaves, one for each possible permutation, and therefor must have depth at least $\log_2(n!) = \Omega(n \log n)$. It makes absolutely no difference what questions are asked in the internal nodes.

Now consider the following related problems.

- ELEMENT UNIQUENESS: Are any two elements of the input sequence $\langle x_1, x_2, \ldots, x_n \rangle$ equal?
- PARITY: Is the permutation of the input sequence $\langle x_1, x_2, \ldots, x_n \rangle$ even or odd?
- SET INTERSECTION: Given two sets $\{x_1, x_2, \ldots, x_n\}$ and $\{y_1, y_2, \ldots, y_n\}$, do they intersect?
- SET EQUALITY: Given two sets $\{x_1, x_2, \ldots, x_n\}$ and $\{y_1, y_2, \ldots, y_n\}$, are they equal?

We cannot apply the same purely information-theoretic argument to these problems, because there are only two possible outputs: YES and NO. And moreover, there *are* trivial decision trees with only one node that decide these problems.

Probably the most natural model to consider is the comparison tree model that we used to study sorting partial orders. It's possible to prove $\Omega(n \log n)$ lower bounds for the number of comparisons required to solve any of those four problems, using a fairly complicated adversary argument. (Essentially, we *must* sort the data to solve any of them.)

7.2 Linear Decision Trees

However, proving these lower bounds will actually be easier if we use a more powerful model of computation, introduced by David Dobkin and Richard Lipton in the late 1970s. In a *linear decision tree* (with input size n), every internal node is labeled with a vector $(a_0, a_1, \ldots, a_n)$ and has three outgoing edges labeled $-$, 0, and $+$. Given an input vector $(x_1, x_2, \ldots, x_n)$, we decide which way to branch based on the sign of the following expression:

$$a_0 + a_1 x_1 + a_2 x_2 + \cdots + a_n x_n.$$

For example, in each node in a comparison tree, we have $a_i = 1$ and $a_j = -1$ for some i and j, and $a_k = 0$ for all $k \neq i, j$.

Linear decision trees have a very simple geometric interpretation. Our generic problem can be thought of as a function $F : \mathbb{R}^n \to \{0, 1\}$, where the input is a point in n-space and the output is a single bit. Every internal node defines a hyperplane with equation

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = -a_0;$$

this equation describes a line when $n = 2$, a plane when $n = 3$, and so on. We branch at a node depending on whether the input is above, below, or on this hyperplane.[1]

Now consider the set of input points $R(v) \subseteq \mathbb{R}^n$ that reach a particular node v in a linear decision tree. The set $R(v)$ contains all the points that satisfy a set of linear equalities and inequalities; such a set is called a *convex polyhedron*. Recall that a set X is *convex* if for any two points $p, q \in X$, the entire line segment pq is contained in X. The intersection of any two convex sets is clearly convex;

[1] Linear decision trees are exactly the same as the *binary space partition trees* used in computer graphics systems.

any hyperplane divides space into three convex sets: the hyperplane itself and two *halfspaces*. Chugging through the definitions, we discover that convex polyhedra are, in fact, convex. (Whew!)

It is trivial to prove that every convex set is connected.[2] This simple observation gives us our first significant tool to prove lower bound in the linear decision tree model.

Lemma 1. *For any node v in any linear decision tree, $R(v)$ is connected.*

Now let $\#F_1$ denote the number of connected components of the set $F^{-1}(1)$ of points x such that $F(x) = 1$, and define $\#F_0$ similarly.

Lemma 2. *Any linear decision tree that computes the function $F : \mathbb{R}^n \to \{0, 1\}$ has depth at least $\lceil \log_3(\#F_0 + \#F_1) \rceil$.*

Proof: For any point $x \in \mathbb{R}^n$ such that $F(x) = 1$, there must be a 1-leaf ℓ that is reached by x. By the previous lemma, only points in the connected component of $F^{-1}(1)$ containing x can reach ℓ. it follows immediately that there are at least $\#F_1$ 1-leaves. Similarly, there must be at least $\#F_0$ 0-leaves. The lower bound now follows from the usual information-theoretic argument. □

Now we're ready to prove some lower bounds!

Theorem 3. *Any linear decision tree that computes the* Element Uniqueness *function has depth $\Omega(n \log n)$.*

Proof: Let $x = (x_1, x_2, \ldots, x_n)$ and $y = (y_1, y_2, \ldots, y_n)$ be two vectors with distinct coordinates that are sorted in two different orders. In particular, for some pair of indices i and j, we have $x_i < x_j$ and $y_i > y_j$. Any continuous path from x to y must contain a point z such that $z_i = z_j$, by the intermediate value theorem, but then we have $F(z) = 0$. Thus, points with different permutations are in different connected components of $F^{-1}(1)$, so $\#F_1 \geq n!$. The lower bound now follows immediately from Lemma 2. □

Theorem 4. *Any linear decision tree that computes the* Set Intersection *function has depth $\Omega(n \log n)$.*

Proof: For this problem, the input is a point $(x, y) = (x_1, x_2, \ldots, x_n; y_1, y_2, \ldots, y_n) \in \mathbb{R}^{2n}$, where the x- and y-coordinates represent the two sets X and Y. If X and Y are disjoint, then we can partition X and Y into disjoint subsets $X_1, X_2, \ldots, X_k$ and $Y_1, Y_2, \ldots, Y_k$ that satisfy the partial order

$$X_1 < Y_1 < X_2 < Y_2 < \cdots < X_k < Y_k.$$

Here, $A < B$ means that every element of A is less than every element of B, but the elements within A are incomparable. Either X_1 or Y_k or both could be empty.

There are exactly $n!^2$ such partial orders where the x- and y- elements alternate, that is, where each set X_i and Y_i is a singleton, and $k = n$. As in the previous theorem, any two points that obey different partial orders lie in different connected components of $F^{-1}(0)$. Thus, $\#F_0 \geq n!^2$, and the theorem follows immediately from Lemma 2. □

Finally, let's consider a function whose complexity is still open.

- 3Sum: Do any three elements in the input set $\{x_1, x_2, \ldots, x_n\}$ sum to zero?

[2] Careful with those vacuous cases: the empty set is both convex and connected!

The fastest algorithm known for this problem runs in $O(n^2)$ time, and this is conjectured to be optimal.[3] However, the following argument gives the strongest bound known in any general model of computation:

Theorem 5. *Any linear decision tree that computes the* 3SUM *function has depth* $\Omega(n \log n)$.

Proof: As usual, we prove the lower bound by counting the connected components of $F^{-1}(0)$. There are $\binom{n}{3}$ possible triples of input elements that could sum to zero. Each one defines a hyperplane of the form $x_i + x_j + x_k = 0$.

Suppose the Little Birdie tells us that

$$-2 < x_i < -1 \text{ for all } i < n/2, \qquad x_{n/2} = 0, \qquad 1 < x_i < 2 \text{ for all } i > n/2.$$

This set of inequalities defines a convex polyhedron Π. The Little Birdie Principle implies that the complexity of 3SUM is at least the complexity of 3SUM restricted to Π. Since Π is convex, we can apply the same counting arguments as when the input space is unrestricted.

A point $x \in \Pi$ has three coordinates that sum to zero if and only if $x_i = -x_j$ for some pair of indices $i < n/2$ and $j > n/2$. In other words, this restriction of 3SUM is exactly the same as the set intersection problem, for which we already have an $\Omega(n \log n)$ lower bound. In particular, $F^{-1}(0) \cap \Pi$ has at least $(n/2)!^2$ connected components. □

It may come as a surprise that this is the best lower bound we can prove for 3SUM using this method. Any set H of hyperplanes in $\mathbb{R}^n$ defines a cell complex called the *arrangement*. The full-dimensional cells are the connected components of $\mathbb{R}^n \setminus H$. A relatively straightforward inductive argument (in any computational geometry book) implies that the number of full-dimensional cells in an arrangement of N hyperplanes in $\mathbb{R}^n$ is at most

$$\sum_{d=0}^{n} \binom{N}{d} < N^n.$$

(The summation bound is exact if the hyperplanes are in *general position*: the intersection of any $d \leq n$ hyperplanes has dimension $n - d$.) The 0-set for 3SUM consists of the full-dimensional cells in the arrangement of $\binom{n}{3}$ hyperplanes in $\mathbb{R}^n$, so its has *less than* $\binom{n}{3}^n = O(n^{3n})$ connected components. Thus, we have no hope of proving a $\omega(n \log n)$ lower bound by counting connected components.[4]

7.3 Algebraic Decision Trees

Now let's consider the following obvious generalization of linear decision trees, first proposed by Guy Steele and Andy Yao. In a *dth-order algebraic decision tree*, every node v is labeled with a polynomial $q_v \in \mathbb{R}[x_1, x_2, \ldots, x_n]$ of degree at most d. As in the linear case, each node has three branches labeled $-$, 0, and $+$, and computation at node v branches according to the sign of the polynomial expression $q_v(x)$. A 1st-order algebraic decision tree is just a linear decision tree.

[3] And in fact, that bound *is* optimal in a weak special case of the linear decision tree model of computation; see my PhD thesis!

[4] In fact, there is no general method to derive $\omega(n \log n)$ lower bounds in any of these decision/computation tree models, as long as the problem is defined by a polynomial number of equalities and inequalities. There are several similar techniques for proving lower bounds that use different notions of the "complexity" of a semi-algebriac set—the number of components in an intersection with a subspace, the volume, the Euler characteristic, the number of boundary features of each dimension, various Betti numbers, etc. The POTM Theorem and its generalizations imply that the complexity of the set 3SUM, for any reasonable notion of "complexity", is only $n^{O(n)}$.

Unfortunately, we can't use the same connectedness argument for algebraic decision trees as we did in the linear case. It's quite easy to come up with polynomials that divide $\mathbb{R}^n$ into more than three connected components, or pairs of polynomials p and q such that the sets $\{x \mid p(x) > 0\}$ and $\{x \mid q(x) > 0\}$ are connected, but their intersection is not.

However, the following theorem gives us a bound on the number of connected components of any semi-algebraic set.[5]

Theorem 6 (Petrovskiĭ, Oleĭnik, Thom, Milnor). *Let X be a semi-algebraic subset of $\mathbb{R}^n$ defined by m polynomial equations and h polynomial inequalities, each of degree at most $d \geq 2$. Then X has at most $d(2d-1)^{n+h-1}$ connected components.*

Corollary 7. *Any dth order algebraic decision tree that computes a function $F : \mathbb{R}^n \to \{0, 1\}$ has depth $\Omega(\log_d(\#F_0 + \#F_1) - n)$.*

Proof: Suppose F can be computed by a dth order algebraic decision tree with depth h. By the POTM Theorem, the set of points that reaches any leaf has at most $d(2d+1)^{n+h-1}$ connected components. Since there are less than 3^h leaves, we have the inequality

$$3^h d(2d+1)^{n+h-1} \geq \#F_0 + \#F_1.$$

Solving for h completes the proof.

$$\begin{aligned}
(6d+3)^h &\geq \frac{\#F_0 + \#F_1}{d(2d+1)^{n-1}}, \\
h &\geq \log_{6d+3} \frac{\#F_0 + \#F_1}{d(2d+1)^{n-1}} \\
&= \frac{\ln(\#F_0 + \#F_1)}{\ln(6d+3)} - (n-1)\frac{\ln(2d+1)}{\ln(6d+3)} - \frac{\ln d}{\ln(6d+3)} \\
&= \Omega(\log_d(\#F_0 + \#F_1) - n)
\end{aligned}$$

□

Corollary 8. *Any algebraic decision tree that computes the* Element Uniqueness *function has depth $\Omega(n \log n)$.*

Corollary 9. *Any algebraic decision tree that computes the* Set Intersection *function has depth $\Omega(n \log n)$.*

Corollary 10. *Any algebraic decision tree that computes the* 3sum *function has depth $\Omega(n \log n)$.*

7.4 Algebraic Computation Trees

Consider the following alternative algorithm for solving the Element Uniqueness problem. Given the input vector $(x_1, x_2, \ldots, x_n)$, we compute its *discriminant*

$$\prod_{1 \leq i < j \leq n} (x_i - x_j)$$

[5]This used to be called "Milnor's theorem", since John Milnor proved it in the late 1960s. Then some time in the 1980's, someone noticed that René Thom had proved the same theorem a few years earlier. Then in the late 1990s, some Russians pointed out that two Russian mathematicians, Petrovskiĭ and Oleĭnik, had proved the theorem several years before Thom. Finally, in the early 2000s, someone noticed that Milnor's paper actually cited Petrovskiĭ and Oleĭnik's earlier paper.

and compare it to zero. Clearly, the discriminant is zero if and only if some pair of elements is equal. It's possible to compute the discriminant in $O(n \log^2 n)$ time using Fast Fourier Transforms. In other words, we can solve the ELEMENT UNIQUENESS problem in near-linear time *without* sorting the input first, using a straight-line program without branches.

A further generalization of algebraic decision trees, proposed by Michael Ben-Or, captures algorithms of this type. An *algebraic computation tree* is a tree with two types of internal nodes.

- **Computation:** A computation node v has an associated value f_v determined by one of the instructions

$$f_v \leftarrow f_u + f_w \qquad f_v \leftarrow f_u - f_w \qquad f_v \leftarrow f_u \cdot f_w \qquad f_v \leftarrow f_u / f_w \qquad f_v \leftarrow \sqrt{f_u}$$

  where f_u and f_w are either values associated with ancestors of v, input values x_i, or arbitrary real constants.[6] Every computation node has one child.

- **Branch:** A branch node u contains one of the test instructions

$$f_u > 0 \qquad f_u \geq 0 \qquad f_u = 0$$

  where f_u is either a value associated with an ancestor of v or an input value x_i. Every branch node has two children.

Given an input $(x_1, x_2, \ldots, x_n)$, we follow a path from the root of the tree down to a leaf. At each computation node, we perform the corresponding arithmetic operation; at each branch node, we branch according to the result of the corresponding test. When we reach a leaf, its value is returned as algorithm's output. As usual, the running time of the algorithm is the length of the path traversed, and the worst-case running time is the depth of the tree.

The algebraic computation tree model can be equivalently described using a *real random access machine* or *real RAM*. A real RAM is pretty close to an actual (abstract) computer: it has random access memory, instructions, control flow, an execution stack, and so forth. The big difference is that the main memory in a real RAM stores *arbitrary real numbers*. Real arithmetic operations $+, -, \cdot, /, \sqrt{}$ and comparisons between real numbers all take constant time. A real RAM may also have integer variables, but **we are not allowed to convert between integer variables and real variables**.[7] In particular, the real RAM model does not allow use of the floor function $\lfloor x \rfloor$.

Given any algorithm written for the real RAM model, we can extract an algebraic computation tree by recording all possible branches and real arithmetic operations in any execution of the algorithm. Thus, any lower bound derived in the algebraic computation tree model immediately applies in the real RAM model as well.

Lemma 11. *Any algebraic computation tree that computes a function $F : \mathbb{R}^n \to \{0, 1\}$ has depth $\Omega(\log(\#F_0 + \#F_1) - n)$.*

Proof: Suppose F can be computed by an algebraic computation tree T with depth h. I claim that $2^{h+1}3^{n+h-1} \geq \#F_0 + \#F_1$; the lemma follows immediately from this claim by the usual arguments.

Let ℓ be a leaf of T with depth at most h. We can describe the set $R(\ell)$ of points that reach ℓ as a semi-algebraic set by manipulating the arithmetic and test instructions on the path from the

[6] Obviously, dividing by zero or taking the square root of a negative number is never allowed.

[7] If we could freely convert between integers and reals, and still do exact real arithmetic in constant time, we could solve some NP-hard problems in linear time.

root to ℓ. If v is a computation node, we obtain an algebraic equation as follows:

Operation	Equation
$f_v \leftarrow f_u + f_w$	$f_v = f_u + f_w$
$f_v \leftarrow f_u - f_w$	$f_v = f_u - f_w$
$f_v \leftarrow f_u \cdot f_w$	$f_v = f_u \cdot f_w$
$f_v \leftarrow f_u / f_w$	$f_u = f_v \cdot f_w$
$f_v \leftarrow \sqrt{f_u}$	$f_u = f_v^2$

Similarly, for any branch node v, we obtain either a new equation or a new inequality, depending on the outcome of the test.

Test	TRUE	FALSE
$f_u > 0$	$f_u > 0$	$-f_u \geq 0$
$f_u \geq 0$	$f_u \geq 0$	$-f_u > 0$
$f_u = 0$	$f_u = 0$	$f_u \cdot f_v = 1$

Note that an unsuccessful equality test introduces a new variable.

The points $(x_1, x_2, \ldots, x_n, f_1, f_2, \ldots, f_h)$ that satisfy this list of polynomial equalities and inequalities describe a semi-algebraic set $U(\ell) \subseteq \mathbb{R}^{n+h}$. Since every polynomial in this list has degree at most 2, the POTM Theorem implies that $U(\ell)$ has at most $2 \cdot 3^{n+h-1}$ connected components.

A point $x \in \mathbb{R}^d$ reaches leaf ℓ if and only if there exists a point $f = (f_1, f_2, \ldots, f_h) \in \mathbb{R}^h$ such that $(x, f) \in U(\ell)$. In other words, we can obtain $R(\ell)$ by projecting $U(\ell)$ onto its first n coordinates. Since projection can only decrease the number of connected components, we conclude that $R(\ell)$ has at most $2 \cdot 3^{n+h-1}$ components. The claim, and thus the lemma, now follows from the fact that T has at most 2^h leaves. □

The lower bound argument requires only that we pay for multiplications, divisions, roots, and branches; additions and subtractions could be performed for free. In fact, the lower bound holds in a more general model where each computation node computes an arbitrary bilinear function of its ancestor values. At the cost of a factor of d, we can also include an operation that computes the roots of an arbitrary polynomial of degree d, whose coefficients are ancestor values.[8] We can even allow computations with complex numbers, by representing each $z = x + yi$ as a pair (x, y), or equivalently, allowing the projection operators $\Re$ and $\Im$ to be performed at no cost.

Corollary 12. *Any algebraic computation tree that computes the* ELEMENT UNIQUENESS *function has depth* $\Omega(n \log n)$.

Corollary 13. *Any algebraic computation tree that computes the discriminant* $\prod_{1 \leq i < j \leq n} (x_i - x_j)$ *has depth* $\Omega(n \log n)$.

Proof: Once we compute the resultant, we can solve the ELEMENT UNIQUENESS problem with just one more branch. □

More accurately, we have an $\Omega(n \log n)$ lower bound on the number of multiplications required to compute the resultant. This lower bound is actually tight; the resultant can be computed using $O(n \log n)$ multiplications, in $O(n \log^2 n)$ time. (Most of the additional time is additions and subtractions.)

[8] More accurately, the natural cost of computing a root of a polynomial $P(x)$ is the minimum time required to evaluate $P(x)$. In particular, the cost of computing $\sqrt[d]{f_u}$ is $\Theta(\log d)$, by repeated squaring.

Corollary 14. *Any algebraic computation tree that computes the* SET INTERSECTION *function has depth* $\Omega(n \log n)$.

Corollary 15. *Any algebraic computation tree that computes the* 3SUM *function has depth* $\Omega(n \log n)$.

7.5 Generic Width

Now consider the problem of computing the maximum element in a set of n numbers. It's not hard to prove a lower bound of $n-1$ in the comparison tree model using an adversary argument, but this argument doesn't generalize to arbitrary algebraic decision or computation trees. Perhaps there is a faster algorithm to compute the maximum element using higher-degree polynomials! Alas, the following argument of Montaña, Pardo, and Recio implies that there is no such algorithm.[9]

Any closed semi-algebraic set X in $\mathbb{R}^n$ can be written in the canonical form

$$X = \bigcup_{i=1}^{t} \bigcap_{j=1}^{r} \{p_{ij} \geq 0\},$$

where each p_{ij} is a polynomial with n variables. The *width* of X is defined as the minimum r for which such a representation is possible. By convention, the empty set and $\mathbb{R}^n$ both have width zero. The *generic width* of a (not necessarily closed) semi-algebraic set X is defined as

$$\overline{w}(X) = \min\{\text{width}(S) \mid \dim(S \oplus X) \leq n\},$$

where the minimum is taken over all closed semi-algebraic sets S. The generic width of X never exceeds the width of X; take $S = X$.

Lemma 16. *Any algebraic decision or computation tree that decides whether a point $x \in \mathbb{R}^n$ lies inside a fixed semi-algebriac set X has depth at least $\overline{w}(X)$.*

Proof: Let T be an algebraic decision tree, and let ℓ be an arbitrary leaf with depth h. Let $R(\ell)$ be the set of possible inputs that reach ℓ. If the path to ℓ contains any $=$-branches, the dimension of $R(\ell)$ is less than n, so $\overline{w}(R(\ell)) = 0$. Otherwise, $R(\ell)$ is the intersection h *open* algebraic halfspaces, which implies that $\overline{w}(R(\ell)) \leq h$.

Since T correctly decides membership in X, we must have

$$X = \bigcup_{\text{1-leaves } \ell} R(\ell),$$

which implies that

$$\overline{w}(X) \leq \max_{\text{1-leaves } \ell} \overline{w}(R(\ell)), \leq \max_{\text{1-leaves } \ell} \text{depth}(\ell) \leq \text{depth}(T).$$

Unlike most results about algebriac decision trees, this argument does *not* require a constant upper bound on the degree of the query polynomials. Thus, it immediately applies to algebaric computation trees as well. □

Theorem 17. *Any algebriac decision or computation tree that computes the largest element of an n-element set has depth at least $n-1$.*

[9] This result is often credited to Rabin, but his proof has a bug.

Proof (sketch): In fact, we can prove an $n-1$ lower bound for the simpler problem of *verifying* the largest element, or equivalently, determining membership in the polyhedron

$$X = \bigcap_{i=2}^{n} \{x_1 - x_i \geq 0\}.$$

It is a fairly tedious exercise to prove that $\overline{w}(X) \geq n-1$. (It is hopefully obvious that $\overline{w}(X) \leq n-1$; the lower bound is considerbly less trivial.) □