On Lower Bounds for Algebraic Decision Trees over the Complex Numbers

Peter Scheiblechner
Department of Mathematics
Purdue University
West Lafayette, IN 47907-2067, USA,
Email: pscheibl@math.purdue.edu

***Abstract*—We prove a new lower bound for the decision complexity of a complex algebraic set X in terms of the sum of its (compactly supported) Betti numbers $b_c(X)$, which is for the first time better than logarithmic in $b_c(X)$.**

We apply this result to subspace arrangements including some well studied problems such as the knapsack and element distinctness problems.

I. Introduction

It came as a surprise in the 1980's that there exist linear decision trees of polynomial depth for subspace arrangements in $\mathbb{R}^n$, some of whose restrictions to binary inputs are NP-complete. A famous example of such a problem is the knapsack problem, which is an arrangement of 2^n hyperplanes. In particular, Meyer auf der Heide [1] proved that one can solve the knapsack problem by linear decision trees of polynomial depth. This result is generalized by Clarkson and Meiser [2], [3] to arbitrary subspace arrangements. More precisely, their results imply that one can decide membership to an arrangement of m subspaces of $\mathbb{R}^n$ by a linear decision tree of depth polynomial in $n \log m$ (see also [4, §3.4]).

We study the corresponding problem over the complex numbers. More generally, we consider the problem of testing membership of a point to an algebraic set in $\mathbb{C}^n$ in the computational model of algebraic decision trees of degree d. It is easy to see by a generic path argument that the results of Meyer auf der Heide etc. are impossible over $\mathbb{C}$ [5], [6]. We provide a second proof for this fact by proving a general lower bound for the decision tree complexity of a complex algebraic set X in terms of its (compactly supported) Betti numbers. There is a tradition of lower bounds for the decision complexity in terms of topological invariants [7], [8], [9], [10], [11], [12], see also the survey [13]. All known lower bounds are *logarithmic*. For instance, Yao's result [11] bounds the complexity of X by $\Omega(\log b_c(X))$, where $b_c(X)$ denotes the sum of the (compactly supported) Betti numbers of X. Our bound is the first non-logarithmic one. In particular, we prove that the decision tree complexity of an algebraic set X in $\mathbb{C}^n$ is bounded below by $(b_c(X)/\deg(X))^{\Omega(\frac{1}{n})}$, where $\deg(X)$ denotes the cumulative degree of X, i.e., the sum of the degrees of all irreducible components of X. Although this looks like a big improvement, it yields better lower bounds only for quite large Betti numbers. Note that we need $b_c(X)$ to be at least of order $2^{n \log^{1+\varepsilon} n}$ for some $\varepsilon > 0$ to have $b_c(X)^{\frac{1}{n}}$ substantially larger than $\log b_c(X)$. So for instance, our result yields a worse lower bound than Yao's for the element distinctness or k-equal problem. An examples on which our lower bound performs better is the knapsack problem. However, the generic path argument mentioned above yields the same result. We are currently looking for examples of higher codimension, where this argument is no longer applicable.

II. Preliminaries

An *algebraic decision tree* of degree d over $\mathbb{C}$ is a rooted binary tree, whose inner nodes (*branching nodes*) are labeled with polynomials in $\mathbb{C}[X_1, \ldots, X_n]$ of degree at most d, and whose leaves are labeled with either "yes" or "no". The tree encodes a program that on input $(x_1, \ldots, x_n) \in \mathbb{C}^n$ parses the tree from the root to some leaf by testing at each branching node labeled with f, whether $f(x) = 0$, and continuing to the left or right according to the outcome of the test. The program answers the label of the leaf it reaches. In this way, an algebraic decision tree over $\mathbb{C}$ decides a *constructible set* in $\mathbb{C}^n$, i.e., a set which is definable by a Boolean combination of polynomial equations.

For lower complexity bounds one needs invariants which are subadditive. It was Yao's idea to use the compactly supported Betti numbers (or equivalently, Borel-Moore Betti numbers). These may be defined for locally closed semialgebraic sets via the cohomology with compact supports. A *semialgebraic set* in $\mathbb{R}^n$ is a set which can be defined by a Boolean combination of polynomial inequalities $f \geq 0$. A *locally closed* set in $\mathbb{R}^n$ is the intersection of an open with a closed set. We recall the definition of the cohomology with compact support [13]. Let $A \subseteq X$ be locally closed semialgebraic sets in $\mathbb{R}^n$. Denote with $H_c^k(X, A)$ the *cohomology group with compact support* of the pair (X, A), which is defined as the direct limit of the singular cohomology groups $H^k(X, X \setminus K)$ over all compact subsets $K \subseteq X$ with $A \subseteq X \setminus K$, directed by inclusion. The k-th *Betti number with compact supports* $b_{c,k}(X)$ is defined as the rank of $H_c^k(X) := H_c^k(X, \emptyset)$. We write $b_c(X)$ for the sum $\sum_{k \in \mathbb{N}} b_{c,k}(X)$.

The important subadditivity property of these Betti numbers reads as follows. Let $A \subseteq X$ be a pair of locally closed sets, where A is closed in X. Then excision implies $H_c^k(X, A) = H_c^k(U)$, where $U := X \setminus A$. It follows that the long exact sequence of the pair (X, A) takes the form

$$\cdots \to H_c^k(U) \to H_c^k(X) \to H_c^k(A) \to H_c^{k+1}(U) \to \cdots,$$

where $U := X \setminus A$. It follows

$$b_{c,k}(X) \le b_{c,k}(U) + b_{c,k}(A) \quad \text{for closed } A \subseteq X. \quad (1)$$

It is clear that for compact X the numbers $b_{c,k}(X)$ coincide with the k-th Betti numbers $b_k(X)$ with respect to the usual singular cohomology. In general, one can compute the compactly supported Betti numbers of X via the Alexandrov compactification $\dot{X} = X \cup \{\infty\}$ as follows:

$$b_{c,k}(X) = \begin{cases} b_0(\dot{X}) - 1 & \text{if } k = 0, \\ b_k(\dot{X}) & \text{if } k > 0. \end{cases} \quad (2)$$

As an example we note that the Alexandrov compactification of $\mathbb{R}^n$ is the n-sphere S^n, hence

$$b_{c,k}(\mathbb{R}^n) = \begin{cases} 0 & \text{if } k \ne n, \\ 1 & \text{if } k = n. \end{cases} \quad (3)$$

III. The General Lower Bound

Our result is based on the fundamental Oleinik-Petrovski/Milnor/Thom-Bound. The following is a version for the compactly supported Betti-numbers which is proved in [13, Corollary 4.5].

Theorem 3.1: Let $X = \{f_1 = 0, \dots, f_r = 0\} \subseteq \mathbb{R}^n$ be an algebraic set defined by the real polynomials f_i of degree at most d. Then $b_c(X) \le d(2d-1)^n$.

We need it for locally closed sets over the complex numbers.

Corollary 3.2: Let X be the locally closed subset of $\mathbb{C}^n$ defined by $f_1 = 0, \dots, f_r = 0, g_1 \ne 0, \dots, g_s \ne 0$, where $s \ge 1$ and $f_i, g_j \in \mathbb{C}[X_1, \dots, X_n]$ have degree at most d. Then $b_c(X) \le (sd+1)(2sd+1)^{2n+2}$.

Proof: The set $X = \{f_1 = 0, \dots, f_r = 0, \prod_i g_i \ne 0\}$ is homeomorphic to the closed set

$$\{f_1 = 0, \dots, f_r = 0, Y \prod_i g_i = 1\} \subseteq \mathbb{C}^{n+1},$$

whose sum of Betti numbers is bounded by $(sd+1)(2(sd+1)-1)^{2(n+1)}$ according to Theorem 3.1. ■

To state the main result of this section, we introduce the following notation. For a constructible set $X \subseteq \mathbb{C}^n$ the *degree d decision complexity* $C_d(X)$ is the minimal depth of a degree d algebraic decision tree deciding X.

Theorem 3.3: Let $X \subseteq \mathbb{C}^n$ be an algebraic set. Then

$$C_d(X) \ge \frac{1}{4d} \left(\frac{b_c(X)}{\deg(X)} \right)^{\frac{1}{3n+2}}.$$

Remark 3.4: If X is a subspace arrangement, i.e., a union of m affine linear subspaces, then $\deg(X) = m$, hence

$$C_d(X) \ge \frac{1}{4d} \left(\frac{b_c(X)}{m} \right)^{\frac{1}{3n+2}}.$$

Proof: The proof of this theorem begins as the proof of Yao's lower bound. The case $X = \mathbb{C}^n$ is clear by (3), so assume $X \ne \mathbb{C}^n$. Let T be a degreee d decision tree of depth k deciding X. For a leaf ν denote by D_ν its *leaf set* consisting of those $x \in \mathbb{C}^n$ which follow the path of T to the leaf ν. We clearly have the decomposition

$$X = \bigsqcup_{\nu \in \Upsilon} D_\nu, \quad \text{where } \Upsilon = \{\text{yes-leaves of } T\}. \quad (4)$$

Furthermore, each leaf set can be written as

$$D_\nu = \{f_1 = 0, \dots, f_r = 0, g_1 \ne 0, \dots, g_s \ne 0\}, \quad (5)$$

where $\deg f_i,\ \deg g_i \le d$. Since X is an algebraic set of dimension at most $n-1$, we have $s \le k-1$. Using the subadditivity we prove

$$b_c(X) \le \sum_{\nu \in \Upsilon} b_c(D_\nu). \quad (6)$$

For a node u of T let X_u be the set of inputs $x \in X$ passing through u. Since X_u is defined by the sign-conditions on the path leading to u, it is a locally closed set. Also, if u is labeled by the polynomial f and has descendants u_0 and u_1, then X_u is the disjoint union of $X_{u_0} = X_u \cap \{f = 0\}$ and $X_{u_1} \cap \{f \ne 0\}$, so X_{u_0} is closed in X_u. Hence (1) implies $b_c(X_u) \le b_c(X_{u_0}) + b_c(X_{u_1})$. By induction on the depth of T our claim (6) follows.

Now (4) and Corollary 3.2 imply

$$b_c(X) \le \sum_{\nu \in \Upsilon} b_c(D_\nu) \le |\Upsilon| kd(2kd)^{2n+2}. \quad (7)$$

Yao now bounds $|\Upsilon|$ by the number of *all* leaves, which is in our case at most 2^k. The new idea is to improve this to a polynomial bound in k.

First note that each irreducible component Z of X must be contained in a unique $\overline{D}_{\nu(Z)}$. We set $I_1 := \{\nu(Z) \mid Z$ component of $X\}$ and $X_1 := \bigcup_{\nu \in \Upsilon \setminus I_1} \overline{D}_\nu$. Then X_1 is an algebraic subset of X of dimension $< m := \dim X$, and each component Z of X_1 is contained in a unique $\overline{D}_{\nu(Z)}$. We set $I_2 := I_1 \sqcup \{\nu(Z) \mid Z$ component of $X_1\}$ and $X_2 := \bigcup_{\nu \in \Upsilon \setminus I_2} \overline{D}_\nu$. Continuing this way we get sequences of subvarieties $X = X_0 \supseteq X_1 \supseteq \cdots \supseteq X_m \supseteq X_{m+1} = \emptyset$ and subsets $\emptyset = I_0 \subseteq I_1 \subseteq \cdots \subseteq I_m \subseteq I_{m+1} = \Upsilon$ with

$$X_i = \bigcup_{\nu \in \Upsilon \setminus I_i} \overline{D}_\nu \quad \text{for} \quad 0 \le i \le m+1.$$

By construction we have $|I_i| \le \sum_{j<i} \deg(X_j)$, since the number of irreducible components of an algebraic set is bounded by its degree. Hence

$$|\Upsilon| \le \sum_{i=0}^{m} \deg(X_i), \quad (8)$$

and we are left with the task to bound $\deg(X_i)$.

Now we prove

$$\deg(X_i) \le \deg(X)(kd)^i \quad \text{for} \quad 0 \le i \le m+1, \quad (9)$$

which is trivial for $i = 0$. Assume it is proved for some $i \ge 0$. We write (5) in the form $D_\nu = V_\nu \cap U_\nu$, where V_ν is the closed set defined by the equations, and U_ν the open set defined by the inequalities. According to the construction of X_{i+1}, for each irreducible component Z of X_i there is a unique $\nu(Z)$ such that $Z \cap D_{\nu(Z)}$ is dense in Z. It follows $Z \subseteq V_{\nu(Z)}$, hence $Z \cap D_{\nu(Z)} = Z \cap U_{\nu(Z)}$. Denote by G_Z the product of the polynomials g_j from (5) for $\nu(Z)$. Note that $\deg G_Z \leq kd$. Then we have

$$Z \setminus D_{\nu(Z)} = Z \cap \mathcal{Z}(G_Z),$$

in particular, this set is closed. Hence

$$\begin{aligned} X_{i+1} &= \overline{\bigcup_{\nu \in \Upsilon \setminus I_{i+1}} D_\nu} = \overline{X_i \setminus \bigcup_{\nu \in I_{i+1} \setminus I_i} D_\nu} \\ &= \bigcup_Z (Z \setminus D_{\nu(Z)}) = \bigcup_Z (Z \cap \mathcal{Z}(G_Z)), \end{aligned}$$

where the union runs over all irreducible components Z of X_i. We conclude with Bézout's Theorem

$$\begin{aligned} \deg(X_{i+1}) &\leq \sum_Z \deg(Z \cap \mathcal{Z}(g_Z)) \leq \sum_Z \deg(Z) \deg(g_Z) \\ &\leq \deg(X_i) kd, \end{aligned}$$

which completes the proof of (9).

From (8) and (9) we conclude

$$\begin{aligned} |\Upsilon| &\leq \deg(X) \sum_{i=0}^{m} (kd)^i \leq 2 \deg(X) (kd)^m \\ &\leq 2 \deg(X) (kd)^{n-1}, \end{aligned}$$

since $k \geq 2$ and $m \leq n-1$. Plugging this into (7) and solving for k yields our Theorem. ■

IV. Applications

In order to bound the Betti numbers of subspace arrangements, we collect some well known facts. A nice exposition about this topic is [14]. We write $b(X)$ for the sum of the singular Betti numbers of a space X. Let $X = \bigcup_i A_i \subseteq \mathbb{R}^n$ be a real subspace arrangement. Its *complexification* is defined to be $X^{\mathbb{C}} := \bigcup_i A_i^{\mathbb{C}}$, where $A_i^{\mathbb{C}} \subseteq \mathbb{C}^n$ is defined by the same equations as A_i. It is a remarkable fact that the cohomology of the complement of an arrangement X is a combinatorial invariant, i.e., it depends only on the intersection semilattice of X. A consequence is that

$$b(\mathbb{R}^n \setminus X) = b(\mathbb{C}^n \setminus X^{\mathbb{C}}),$$

see [14, §8.3]. Furthermore, we have

Lemma 4.1: For any closed subset $X \subseteq \mathbb{R}^n$ we have

$$b_c(X) = b(\mathbb{R}^n \setminus X) - 1.$$

Proof: Using (2) and Alexander duality in $S^n = \mathbb{R}^n \cup \{\infty\}$ we conclude

$$b_c(X) = b(\dot{X}) - 1 = b(S^n \setminus \dot{X}) - 1,$$

but $S^n \setminus \dot{X} = \mathbb{R}^n \setminus X$. ■

A. Element Distinctness

Consider the element-distinctness problem

$$\mathrm{Dist}_n = \bigcup_{i<j} \{x_i = x_j\}.$$

It is shown in [4, Corollary (11.10)] that $b_0(\mathbb{R}^n \setminus \mathrm{Dist}_n)) \geq n!$, hence Lemma 4.1 and Theorem 3.3 only imply $C_d(\mathrm{Dist}_n) \geq \Omega((n/e)^{\frac{1}{3}(1-\varepsilon)})$ for all $\varepsilon > 0$, whereas Yao's result shows the better bound $C_d(\mathrm{Dist}_n) \geq \Omega(n \log n)$.

B. Knapsack

For the knapsack problem

$$\mathrm{KN}_n = \bigcup_{I \subseteq [n]} \left\{ \sum_{i \in I} x_i = 1 \right\},$$

Lemma 4.1 together with [4, Lemma (11.14)] shows $b_c(\mathrm{KN}_n) \geq 2 \cdot 2^{\binom{n}{2}} - 1$. Hence Theorem 3.3 yields $C_d(\mathrm{KN}_n) \geq 2^{\Omega(n)}$, now better than Yao's quadratic bound. However, the aforementioned generic path argument implies $C_d(X) \geq \frac{m}{d}$ for a hyperplane arrangement X with m hyperplanes, which also yields the single exponential lower bound for the knapsack problem.

C. Higher Codimension

The generic path argument does not apply to arrangements of higher codimension. To get such a subspace arrangement with sufficiently large Betti numbers, one can take

$$\mathrm{KN}_n^2 = \bigcup_{I \subseteq [n], J \subseteq [n]} \left\{ \sum_{i \in I} x_i = \sum_{j \in J} y_j = 1 \right\} \subseteq \mathbb{C}^n \times \mathbb{C}^n,$$

which has codimension 2. Using the Künneth-Theorem one can show $b_c(X \times Y) = b_c(X) b_c(Y)$ [15], hence Theorem 3.3 implies $C_d(\mathrm{KN}_n^2) \geq 2^{\Omega(2n)}$. However, this lower bound may also be derivable by reducing to the knapsack problem. An interesting problem for our bound arises, if we require the two subsets I and J to be the same:

$$X = \bigcup_{I \subseteq [n]} \left\{ \sum_{i \in I} x_i = \sum_{i \in I} y_i = 1 \right\} \subseteq \mathbb{C}^n \times \mathbb{C}^n,$$

Unfortunately, at present we do not know a bound for the Betti numbers of this arrangement.

Acknowledgment

This paper has also been presented at the conference "Logical Approaches to Barriers in Computing and Complexity" in February 2010 in Greifswald, Germany. This work was partially supported by DFG grant SCHE 1639/1-1. The author would like to thank the anonymous referees of this and an earlier version for valuable comments.

References

[1] F. Meyer auf der Heide, "A polynomial linear search algorithm for the n-dimensional knapsack problem," *J. ACM*, vol. 31, no. 3, pp. 668–676, 1984.

[2] K. L. Clarkson, "New applications of random sampling in computational geometry," *Discrete Comput. Geom.*, vol. 2, pp. 195–222, 1987.

[3] S. Meiser, "Point location in arrangements of hyperplanes," *Inf. Comput.*, vol. 106, no. 2, pp. 286–303, 1993.

[4] P. Bürgisser, M. Clausen, and M. Shokrollahi, *Algebraic complexity theory*, ser. Grundlehren der Mathematischen Wissenschaften. Berlin: Springer-Verlag, 1997, vol. 315.

[5] L. Blum, F. Cucker, M. Shub, and S. Smale, *Complexity and Real Computation*. Springer, 1998.

[6] P. Koiran, "Computing over the reals with addition and order," in *Selected papers of the workshop on Continuous algorithms and complexity*. New York, NY, USA: Elsevier Science Inc., 1994, pp. 35–47.

[7] M. Ben-Or, "Lower bounds for algebraic computation trees," in *STOC '83: Proceedings of the fifteenth annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 1983, pp. 80–86.

[8] A. Björner, L. Lovász, and A. Yao, "Linear decision trees: volume estimates and topological bounds," in *STOC '92: Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 1992, pp. 170–177.

[9] A. Yao, "Algebraic decision trees and Euler characteristics," in *FOCS '92: Proceedings of the 33rd Annual Symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 1992, pp. 268–277.

[10] M. Ben-Or, "Algebraic computation trees in characteristic $p > 0$," in *FOCS '94: Proceedings of the 35th Annual Symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 1994, pp. 534–539.

[11] A. Yao, "Decision tree complexity and Betti numbers," in *STOC '94: Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 1994, pp. 615–624.

[12] T. Lickteig and M. Roy, "Semi-algebraic complexity of quotients and sign determination of remainders," *J. Compl.*, vol. 12, no. 4, pp. 545–571, 1996.

[13] P. Bürgisser, "Lower bounds and real algebraic geometry," in *Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science*, 2001, pp. 35–54.

[14] A. Björner, "Subspace arrangements," in *Proc. of 1st European Congress of Mathematics*. Paris, 1992: Birkhäuser, 1992, pp. 321–370.

[15] A. Dold, *Lectures on Algebraic Topology*, ser. Grundlehren der Mathematischen Wissenschaften. Berlin: Springer-Verlag, 1980, vol. 200.