

Almost all decision trees do not allow significant quantum speed-up

Ashley Montanaro*

October 29, 2012

Abstract: We show that, for any d , all but a doubly exponentially small fraction of decision trees of depth at most d require $\Omega(d)$ quantum queries to be computed with bounded error. In other words, most efficient classical algorithms in the query complexity model do not admit a significant quantum speed-up. The proof is based on showing that, with high probability, the average sensitivity of a random decision tree is high.

Key words and phrases: quantum query complexity, decision trees

1 Introduction

Many of the most important examples of quantum algorithms which outperform classical algorithms operate in the query complexity model, where the quantity of interest is the number of queries to the input required to compute some function [4]. Can *most* classical query algorithms be significantly accelerated by a quantum computer? In one sense, the answer is a resounding no. It was shown by Ambainis in 1999 that almost all boolean functions on n bits require $n/4 - O(\sqrt{n} \log n)$ quantum queries to be computed with bounded error [2]; using different techniques, Ambainis et al have very recently improved this bound to $n/2 - o(n)$ [3].

However, random boolean functions are arguably somewhat uninteresting in general in that they are hard both for classical and quantum algorithms to compute. One can interpret the result of [2] as saying that there are too many boolean functions, and too few efficient algorithms, whether classical or quantum. Perhaps a more interesting question is whether most functions with efficient classical algorithms in the query model – i.e. short decision trees – could be computed significantly more quickly on a quantum computer than is possible classically. In this note we show that this is also not the case.

*Supported by an EPSRC Postdoctoral Research Fellowship.

We first need to define the model of decision trees which we will use (see [4] for further background). A (boolean) decision tree T is a rooted binary tree where each internal vertex has exactly two children. Such trees are called full; only full binary trees will be considered in this paper. Each vertex is labelled with a variable x_i , $1 \leq i \leq n$, and each leaf is labelled with 0 or 1, corresponding to the output of the tree. Each edge from a node labelled with x_i to its children is labelled with 0 or 1, corresponding to the value of x_i . T computes a boolean function $T : \{0, 1\}^n \rightarrow \{0, 1\}$ in the obvious way: starting with the root, the variable labelling each vertex is queried, and dependent on whether the answer is 0 or 1 the left or right subtree is evaluated. When a leaf is reached, the tree outputs the label of that leaf. The depth of a vertex is defined as follows: the root has depth 0, and the depth of any other vertex is equal to the depth of its parent, plus 1. The depth of T is the largest depth of any vertex, i.e. the worst-case number of queries made on any input.

Following [7], we will consider the *uniform* model of random decision trees. In this model, we obtain a random depth d decision tree T on $n \geq d$ variables simply by picking T uniformly at random from the set of all non-redundant decision trees of depth at most d on n variables. A decision tree is said to be non-redundant if no variable occurs more than once on any path from the root to a leaf. Observe that choosing a boolean function by picking a tree uniformly at random does not pick the corresponding function itself uniformly at random from the set of all functions with decision trees of depth d , as many different decision trees can represent the same function, and some functions may have more trees representing them than others. Also observe that the internal structure of a uniformly random decision tree T and the values assigned to its leaves are independent. In other words, a random decision tree T of depth at most d on n variables can be obtained by a two-stage process: first, pick the structure and labels of the internal vertices of T uniformly at random from the set of all decision trees of depth at most d on n variables, then assign 0 or 1 to the leaves of T uniformly at random.

One could also consider an alternative model of random decision trees, known as the *complete* model [7]. In this model, T 's structure is always that of a complete binary tree of depth d . We randomly assign variables to T 's internal vertices, consistent with T being non-redundant, and assign 0 or 1 to T 's leaves uniformly at random. In the present work we concentrate on the uniform model, which seems to be the most natural way of producing a random decision tree, as it is precisely the model which puts equal weight on every decision tree of depth at most d . However, our results easily extend to the complete model (indeed, the proofs are somewhat simpler).

We can now state our main result.

Theorem 1. Let T be a random depth d decision tree on n variables. Then there exists a universal constant $\alpha > 0$ such that $\Pr[Q_2(f) < \alpha d] \leq 2^{-\Omega(2^{d/3})}$.

In the above theorem, $Q_2(f)$ is the bounded-error quantum query complexity of f [4], i.e. the smallest number of quantum queries required to compute f with worst-case success probability at least $2/3$. Observe that the concentration bound obtained is doubly exponentially small in d . This result thus implies that the vast majority of efficient classical algorithms in the query complexity model do not admit a significant quantum speed-up.

The counting technique used by Ambainis [1] to prove that most boolean functions do not have query-efficient quantum algorithms does not seem to suffice to prove Theorem 1. Indeed, Ambainis showed that there are at most $2^{O(n^{2d+3}d)}$ boolean functions on n bits that can be computed by quantum

algorithms making d queries. However, there can be $n^{\Theta(2^d)}$ decision trees of depth d . Thus this argument does not disallow, for example, the possibility that random decision trees of depth $d = O(\log^2 n)$ could be evaluated using $O(\log n)$ quantum queries.

We therefore take another approach to proving the above theorem, by giving a lower bound on the *average sensitivity* (also known as total influence) of random decision trees. This quantity is mathematically tractable and is known to lower bound quantum query complexity [10].

2 Proof of Theorem 1

The average sensitivity of a boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is defined as

$$\bar{s}(f) = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} \sum_{i=1}^n |f(x) - f(x^i)|,$$

where x^i is the n -bit string obtained by flipping the i 'th bit of x . Up to normalisation, $\bar{s}(f)$ thus counts the number of neighbours x, y such that $f(x) \neq f(y)$. Average sensitivity gives the following lower bound on quantum query complexity.

Theorem 2 (Shi [10], see [11] for the version here). Assume that $f : \{0, 1\}^n \rightarrow \{0, 1\}$ can be computed by a quantum algorithm making q queries and with worst-case failure probability $\varepsilon \leq 1/2$. Then $q \geq \frac{1}{2}(1 - 2\varepsilon)^2 \bar{s}(f)$.

Our goal will now be to upper bound the probability that the average sensitivity of a random decision tree is low. To do so, we will use the following powerful measure concentration result.

Theorem 3 (See e.g. [9] or [6, Corollary 5.2]). Fix $\eta > 0$ and assume that $g : \{0, 1\}^n \rightarrow \mathbb{R}$ satisfies $|g(x) - g(y)| \leq \eta$ for all $x, y \in \{0, 1\}^n$ such that $d(x, y) = 1$, where $d(x, y)$ is the Hamming distance between x and y . Let $x \in \{0, 1\}^n$ be picked uniformly at random. Then

$$\Pr_x [g(x) < \mathbb{E}_z [g(z)] - \delta] \leq e^{-2\delta^2/(n\eta^2)}.$$

We first calculate the expected average sensitivity of random decision trees. As discussed previously, we will choose a uniformly random decision tree by first choosing the structure of the tree, then choosing assignments to the leaves uniformly at random. So fix a decision tree T on n variables and let \mathcal{L} be the set of leaves of T , setting $L := |\mathcal{L}|$. Let $T_z : \{0, 1\}^n \rightarrow \{0, 1\}$ be the boolean function obtained by assigning $z \in \{0, 1\}^L$ to the leaves of T ; we will eventually apply Theorem 3 to the function $g : \{0, 1\}^L \rightarrow \mathbb{R}$ defined by $g(z) = \bar{s}(T_z)$. For any leaf ℓ , let $d(\ell)$ be the depth of ℓ .

Lemma 4.

$$\mathbb{E}_{z \in \{0, 1\}^L} [\bar{s}(T_z)] = \frac{1}{2} \sum_{\ell \in \mathcal{L}} d(\ell) 2^{-d(\ell)}.$$

Proof. From the definition of g ,

$$\mathbb{E}_{z \in \{0, 1\}^L} [\bar{s}(T_z)] = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} \sum_{i=1}^n \mathbb{E}_{z \in \{0, 1\}^L} [|T_z(x) - T_z(x^i)|].$$

Now observe that for each i , if T queries the i 'th bit on input x , $\mathbb{E}_{z \in \{0,1\}^L} [|T_z(x) - T_z(x^i)|] = \frac{1}{2}$. Otherwise, $\mathbb{E}_{z \in \{0,1\}^L} [|T_z(x) - T_z(x^i)|] = 0$. So we obtain

$$\mathbb{E}_{z \in \{0,1\}^L} [\bar{s}(T_z)] = \frac{1}{2^{n+1}} \sum_{x \in \{0,1\}^n} |\{i : T \text{ queries } i \text{ on } x\}| = \frac{1}{2} \mathbb{E}_{x \in \{0,1\}^n} [\text{path length on input } x].$$

As each query to the input gives 0 or 1 with equal probability on a random input x , the probability of ending up at a given leaf ℓ on a random input is just $2^{-d(\ell)}$. The claim follows. \square

We can also bound $\eta = \max_{d(w,z)=1} |\bar{s}(T_w) - \bar{s}(T_z)|$ as follows.

Lemma 5. Let $w, z \in \{0,1\}^L$ satisfy $d(w, z) = 1$. Then

$$|\bar{s}(T_w) - \bar{s}(T_z)| \leq \max_{\ell \in \mathcal{L}} d(\ell) 2^{1-d(\ell)}.$$

Proof. The idea is essentially the same as the proof of Lemma 4. We have

$$\bar{s}(T_w) - \bar{s}(T_z) = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} \sum_{i=1}^n (|T_w(x) - T_w(x^i)| - |T_z(x) - T_z(x^i)|).$$

As w and z only differ on one leaf ℓ , $|T_w(x) - T_w(x^i)| = |T_z(x) - T_z(x^i)|$ unless precisely one of x and x^i leads to ℓ . Thus

$$|\bar{s}(T_w) - \bar{s}(T_z)| \leq \frac{1}{2^n} \sum_{x, x \rightarrow \ell} \sum_{i, x^i \rightarrow \ell} 1 + \frac{1}{2^n} \sum_{x, x \not\rightarrow \ell} \sum_{i, x^i \rightarrow \ell} 1 = \frac{1}{2^{n-1}} \sum_{x, x \rightarrow \ell} |\{i : x^i \rightarrow \ell\}|,$$

where we use the notation $x \rightarrow \ell$ and $x \not\rightarrow \ell$ to mean that evaluation of the tree does or does not end up in leaf ℓ on input x , respectively. For each x , there can only be at most $d(\ell)$ variables i such that $x \rightarrow \ell$ and $x^i \not\rightarrow \ell$. Therefore

$$|\bar{s}(T_w) - \bar{s}(T_z)| \leq 2d(\ell) \Pr_{x \in \{0,1\}^n} [x \rightarrow \ell] = d(\ell) 2^{1-d(\ell)}.$$

\square

Note that this bound is quite weak when there exists a leaf ℓ in T of low depth. Nevertheless, its dependence on $d(\ell)$ is essentially tight: consider the tree T that computes

$$T(x) = \alpha x_1 + (1 - x_1)x_2x_3 \dots x_n,$$

for $\alpha \in \{0,1\}$. Then changing α from 0 to 1 changes $\bar{s}(T)$ from close to 0 to close to 1. Luckily, it turns out that in fact $\min_{\ell \in \mathcal{L}} d(\ell)$ will be high with high probability; roughly speaking, a random tree is unlikely to have any low-hanging fruit. We formalise this as the following lemma, which we prove at the end.

Lemma 6. Let T be picked uniformly at random from the set of decision trees of depth at most d . Then, for any $h \leq d - \log_2 d - 2$, the probability that T has a leaf with depth at most h is at most $2^{1-2^{d-h-2}}$.

Assume that $d(\ell) \geq 2d/3$ for all $\ell \in \mathcal{L}$. Then, using Theorem 3 and Lemma 5, we have

$$\Pr_z[\bar{s}(T_z) < \mathbb{E}[\bar{s}(T_z)] - \delta] \leq e^{-2\delta^2/(L\eta^2)} \leq e^{-(9/8)2^{d/3}\delta^2/d^2}.$$

Also, if $d(\ell) \geq 2d/3$ for all $\ell \in \mathcal{L}$, by Lemma 4 $\mathbb{E}[\bar{s}(T_z)] \geq d/3$. Thus

$$\Pr_z[\bar{s}(T_z) < (1 - \gamma)d/3] \leq e^{-2^{d/3-3}\gamma^2}.$$

Now let T be picked uniformly at random from the set of decision trees of depth at most d . Taking $h = 2d/3$ in Lemma 6 and using a union bound over the two bad events that $\min_{\ell \in \mathcal{L}} d(\ell) \leq 2d/3$ and $\bar{s}(T_z) < (1 - \gamma)d/3$, we obtain the overall bound that

$$\Pr_T[\bar{s}(T) < (1 - \gamma)d/3] \leq 2^{1-2^{d/3-2}} + e^{-2^{d/3-3}\gamma^2},$$

which is clearly of order $2^{-\Omega(2^{d/3}\gamma^2)}$. Taking $\varepsilon = 1/3$ in Theorem 2 we have $Q_2(T) \geq \bar{s}(T)/18$, and choosing γ to be an arbitrary constant such that $0 < \gamma < 1$ completes the proof of Theorem 1.

Observe that, if one is content with looser concentration bounds, it is possible to show that $\Pr_T[\bar{s}(T) < d/2 - O(\log d)] = 2^{-\Omega(d)}$, by taking $h = d - O(\log d)$ in Lemma 6.

2.1 Remaining lemmas

In order to prove Lemma 6 we will need a simple bound on the number of low depth binary trees.

Fact 2.1. *Let N_d be the number of binary trees of depth at most d . For $d \geq 1$, $N_d = N_{d-1}^2 + 1$, and $N_0 = 1$. Thus $N_d \geq 2^{2^{d-1}}$ for all $d \geq 1$.*

Proof. Every binary tree of depth at most d either consists of a single leaf or of two independent binary trees of depth at most $d - 1$; the claimed recurrence for N_d is immediate. The second part easily follows from this recurrence by induction. \square

Lemma 7. *Let T be picked uniformly at random from the set of binary trees of depth at most d . Then, for any $h \leq d - \log_2 d - 2$, the probability that T has a leaf with depth at most h is at most $2^{1-2^{d-h-2}}$.*

Proof. Using a union bound and Fact 2.1, the probability that T has a leaf with depth at most h is upper bounded by

$$\sum_{k=0}^h 2^k \Pr[\text{a vertex at depth } k \text{ is a leaf}] = \sum_{k=0}^h \frac{2^k}{N_{d-k}} \leq \sum_{k=0}^h 2^{k-2^{d-k-1}} \leq \sum_{k=0}^h 2^{-2^{d-k-2}} \leq 2^{1-2^{d-h-2}}.$$

\square

Picking a decision tree T uniformly at random from the set of decision trees of depth at most d is similar to, but not quite the same as, making T 's structure a uniformly random binary tree of depth at most d , then assigning variables to the internal vertices at random, and values to the leaves at random. However, the above conclusions also apply to uniformly random decision trees (in other words, Lemma 7 implies Lemma 6). The reason is that choosing a tree from a distribution where trees with different variables assigned to the internal vertices are considered to be distinct can only decrease the probability that an arbitrary vertex at a given depth is a leaf.

3 Conclusions

We have shown that, in the query complexity model, most short decision trees do not admit a quantum speed-up greater than a constant factor. An interesting open question is to determine the extent of the quantum speed-up which *is* possible. It is known that, for all boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$, $Q_2(f) \leq n/2 + O(\sqrt{n})$ [5]. Could it be the case that all decision trees of depth d can be computed with bounded error using $d/2 + O(\sqrt{d})$ quantum queries? Indeed, one can show that a non-zero speed-up is possible for essentially all decision trees: if f has a decision tree of depth $d \geq 2$, then $Q_2(f) \leq d - 1$. The reason is that every boolean function on 2 bits can be computed with success probability $9/10$ using only one quantum query [8], so the quantum algorithm can simply follow the classical decision tree, replacing the last two classical queries with one quantum query.

Acknowledgements

I would like to thank Tony Short for his interpretation of Lemma 6, and Ronald de Wolf and an anonymous referee for helpful comments on a previous version.

References

- [1] A. AMBAINIS: A better lower bound for quantum algorithms searching an ordered list. In *Proc. 40th Annual Symp. Foundations of Computer Science*, pp. 352–357. IEEE, 1999. [quant-ph/9902053](#). [2](#)
- [2] A. AMBAINIS: A note on quantum black-box complexity of almost all boolean functions. *Inf. Proc. Lett.*, 71:5–7, 1999. [quant-ph/9811080](#). [1](#)
- [3] A. AMBAINIS, A. BAČKURS, J. SMOTROVS, AND R. DE WOLF: Optimal quantum query bounds for almost all Boolean functions, 2012. [arXiv:1208.1122](#). [1](#)
- [4] H. BUHRMAN AND R. DE WOLF: Complexity measures and decision tree complexity: a survey. *Theoretical Computer Science*, 288:21–43, 2002. [1](#), [2](#)
- [5] W. VAN DAM: Quantum oracle interrogation: Getting all information for almost half the price. In *Proc. 39th Annual Symp. Foundations of Computer Science*, pp. 362–367. IEEE, 1998. [quant-ph/9805006](#). [6](#)
- [6] D. DUBHASHI AND A. PANCONESI: *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009. [3](#)
- [7] J. JACKSON AND R. SERVEDIO: Learning random log-depth decision trees under uniform distribution. *SIAM J. Comput.*, 34(5):1107–1128, 2005. [2](#)
- [8] I. KERENIDIS AND R. DE WOLF: Exponential lower bound for 2-query locally decodable codes via a quantum argument. *J. Comput. Syst. Sci.*, 69(3):395–420, 2004. [quant-ph/0208062](#). [6](#)

- [9] M. LEDOUX: *The concentration of measure phenomenon*. AMS Mathematical Surveys and Monographs 89. American Mathematical Society, 2001. [3](#)
- [10] Y. SHI: Lower bounds of quantum black-box complexity and degree of approximating polynomials by influence of boolean variables. *Inf. Proc. Lett.*, 75(1–2):79–83, 2000. [quant-ph/9904107](#). [3](#)
- [11] R. DE WOLF: A brief introduction to Fourier analysis on the boolean cube. *Theory of Computing Library Graduate Surveys*, 1:1–20, 2008. [3](#)

AUTHOR

Ashley Montanaro,
Centre for Quantum Information and Foundations,
Department of Applied Mathematics and Theoretical Physics,
University of Cambridge, UK.
am994@cam.ac.uk.