Information and Computation 243 (2015) 2–25

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Dual lower bounds for approximate degree and Markov–Bernstein inequalities

Mark Bun [a,*,1], Justin Thaler [a,b,2]

[a] Harvard University, School of Engineering and Applied Sciences, 33 Oxford St., Cambridge, MA 02138, United States
[b] Yahoo! Labs, 111 West 40th St. 9th Floor, New York, NY 10018, United States

ARTICLE INFO

Article history:
Received 1 September 2013
Available online 12 December 2014

ABSTRACT

The ε-*approximate degree* of a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ is the minimum degree of a real polynomial that approximates f to within error ε in the ℓ_∞ norm. We prove several lower bounds on this important complexity measure by explicitly constructing solutions to the dual of an appropriate linear program. Our first result resolves the ε-approximate degree of the two-level AND–OR tree for any constant $\varepsilon > 0$. We show that this quantity is $\Theta(\sqrt{n})$, closing a line of incrementally larger lower bounds. The same lower bound was recently obtained independently by Sherstov (Theory Comput. 2013) using related techniques. Our second result gives an explicit *dual polynomial* that witnesses a tight lower bound for the approximate degree of any symmetric Boolean function, addressing a question of Špalek (2008). Our final contribution is to reprove several Markov-type inequalities from approximation theory by constructing explicit dual solutions to natural linear programs. These inequalities underly the proofs of many of the best-known approximate degree lower bounds, and have important uses throughout theoretical computer science.

1. Introduction

Approximate degree is an important measure of the complexity of a Boolean function. It captures whether a function can be approximated by a low-degree polynomial with real coefficients in the ℓ_∞ norm, and it has diverse applications in theoretical computer science. For instance, lower bounds on approximate degree underly fundamental circuit complexity lower bounds [25,7,36] and oracle separations between complexity classes [8]. In quantum computing, many tight lower bounds on quantum query complexity have been proved via lower bounds on approximate degree [2,19,5]. Approximate degree lower bounds have also found important uses in communication complexity [39,22,13,42,10,41,38], enabling the resolution of long-standing open problems regarding both randomized and quantum formulations of bounded-error, small-bias, and multiparty communication. Meanwhile, upper bounds on approximate degree have had several important algorithmic uses. For instance, in computational learning theory, approximate degree upper bounds underly the best known algorithms for PAC learning DNF and read-once formulas, and agnostically learning disjunctions [20,3,18].

* Corresponding author.
E-mail addresses: mbun@seas.harvard.edu (M. Bun), jthaler@fas.harvard.edu (J. Thaler).

1 Supported by an NDSEG Fellowship and NSF grant CNS-1237235.
2 This work was performed while the author was a graduate student at Harvard University, School of Engineering and Applied Sciences, and supported by an NSF Graduate Research Fellowship and NSF grants CNS-1011840 and CCF-0915922.

http://dx.doi.org/10.1016/j.ic.2014.12.003
0890-5401/

In this paper, we seek to advance our understanding of this fundamental complexity measure. We focus on proving approximate degree lower bounds by specifying explicit *dual polynomials*, which are dual solutions to a certain linear program capturing the approximate degree of any function. These polynomials act as certificates of the high approximate degree of a function, and their construction is of interest because these dual objects have been used recently to resolve several long-standing open problems in communication complexity (e.g. [39,22,13,42,10,41]). See the survey of Sherstov [34] for an excellent overview of this body of literature.

Our contributions Our first result resolves the approximate degree of the function $f(x) = \bigwedge_{i=1}^{N} \bigvee_{j=1}^{N} x_{ij}$, showing this quantity is $\Theta(N)$. Known as the two-level AND–OR tree, f is perhaps the simplest function whose approximate degree was not previously characterized. A series of works spanning nearly two decades proved incrementally larger lower bounds on the approximate degree of this function, and this question was recently re-posed by Aaronson in a tutorial at FOCS 2008 [1]. Our proof not only yields a tight lower bound, but it specifies an explicit dual polynomial for the high approximate degree of f, answering a question of Špalek [42] in the affirmative.

Our second result gives an explicit dual polynomial witnessing the high approximate degree of any *symmetric* Boolean function, recovering a well-known result of Paturi [28]. Our solution builds on work of Špalek [42], who gave an explicit dual polynomial for the OR function, and addresses an open question from that work.

Our final contribution is to reprove several classical Markov-type inequalities from approximation theory. These inequalities bound the derivative of a polynomial in terms of its degree. Combined with the well-known symmetrization technique (see e.g. [25,1]), Markov-type inequalities have traditionally been the primary tool used to prove approximate degree lower bounds on Boolean functions (e.g. [2,4,26,40]). Our proofs of these inequalities specify explicit dual solutions to a natural linear program (that differs from the one used to prove our first two results). While these inequalities have been known for over a century [9,23,24], to the best of our knowledge our proof technique is novel, and we believe it sheds new light on these results.

2. Preliminaries

We work with Boolean functions $f : \{-1,1\}^n \to \{-1,1\}$ under the standard convention that 1 corresponds to logical false, and -1 corresponds to logical true. We let $\|f\|_\infty = \max_{x\in\{-1,1\}^n} |f(x)|$ denote the ℓ_∞ norm of f. The ε-approximate degree of a function $f : \{-1,1\}^n \to \{-1,1\}$, denoted $\deg_\varepsilon(f)$, is the minimum (total) degree of any real polynomial p such that $\|p - f\|_\infty \le \varepsilon$, i.e., $|p(x) - f(x)| \le \varepsilon$ for all $x \in \{-1,1\}^n$. We use $\widetilde{\deg}(f)$ to denote $\deg_{1/3}(f)$, and use this to refer to the *approximate degree* of a function without qualification. The choice of 1/3 is arbitrary, as $\widetilde{\deg}(f)$ is related to $\deg_\varepsilon(f)$ by a constant factor for any constant $\varepsilon \in (0,1)$. We let OR_n and AND_n denote the OR function and AND function on n variables respectively, and we let $\mathbf{1}_n \in \{-1,1\}^n$ denotes the n-dimensional all-ones vector. Define $\widetilde{\mathrm{sgn}}(x) = -1$ if $x < 0$ and 1 otherwise.

In addition to approximate degree, *block sensitivity* is also an important measure of the complexity of a Boolean function. We introduce this measure because functions with low block sensitivity are an "easy case" in the analysis of Theorem 2 below. The block sensitivity $\mathrm{bs}_x(f)$ of a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ at the point x is the maximum number of pairwise disjoint subsets $S_1, S_2, S_3, \dots \subseteq \{1,2,\dots,n\}$ such that $f(x) \ne f(x^{S_1}) = f(x^{S_2}) = f(x^{S_3}) = \dots$. Here, x^S denotes the vector obtained from x by negating each entry whose index is in S. The block sensitivity $\mathrm{bs}(f)$ of f is the maximum of $\mathrm{bs}_x(f)$ over all $x \in \{-1,1\}^n$.

2.1. A dual characterization of approximate degree

For a subset $S \subseteq \{1,\dots,n\}$ and $x \in \{-1,1\}^n$, let $\chi_S(x) = \prod_{i\in S} x_i$. Given a Boolean function f, let $p(x) = \sum_{|S|\le d} c_S \chi_S(x)$ be a polynomial of degree d that minimizes $\|p - f\|_\infty$, where the coefficients c_S are real numbers. Then p is an optimum of the following linear program.

min	ε	
such that	$\lvert f(x) - \sum_{\lvert S\rvert\le d} c_S \chi_S(x)\rvert \le \varepsilon$	for each $x \in \{-1,1\}^n$
	$c_S \in \mathbb{R}$	for each $\lvert S\rvert \le d$
	$\varepsilon \ge 0$	

The dual LP is as follows.

max	$\sum_{x\in\{-1,1\}^n} \phi(x) f(x)$	
such that	$\sum_{x\in\{-1,1\}^n} \lvert\phi(x)\rvert = 1$	
	$\sum_{x\in\{-1,1\}^n} \phi(x)\chi_S(x) = 0$	for each $\lvert S\rvert \le d$
	$\phi(x) \in \mathbb{R}$	for each $x \in \{-1,1\}^n$