Electronic Colloquium on Computational Complexity, Revision 1 of Report No. 208 (2026)

Exponential Correlation Bounds for Polynomials

Eshan Chattopadhyay* Cornell University

Pooya Hatami† Ohio State University

Chin Ho Lee‡ NC State University

Shachar Lovett§ UC San Diego

Avishay Tal¶ UC Berkeley

Emanuele Viola‖ Northeastern University

October 11, 2026

Abstract

We prove that the XOR of k majorities on disjoint blocks of ℓ bits has correlation at most $(4d/\sqrt{\ell})^k$ with every Boolean-valued degree-d polynomial over an arbitrary field. By known techniques, we obtain explicit pseudorandom generators that fool such polynomials on n Boolean inputs with seed length $d^2 \cdot \mathsf{polylog}(n/\varepsilon)$. The correlation and the generator extend to small $\mathsf{AC}^0[q]$ circuits, for every fixed prime power q. These results resolve several long-standing questions in circuit complexity and pseudorandomness.

1 Introduction

Proving correlation bounds for low-degree polynomials is a fundamental challenge in theoretical computer science. Such bounds have received a lot of attention, because they are a prerequisite for progress on a number of other long-standing challenges in computational complexity, related to matrix rigidity, multiparty communication complexity, circuit lower bounds, pseudorandomness, and more. For background on correlation bounds, see the survey [Vio22] or the book [Vio].

For two Boolean functions $f, g\colon \{0,1\}^n \to \{0,1\}$, we define their correlation by

$$\mathsf{corr}(f,g) := |\mathop{\mathbb{E}}_x[(-1)^{f(x)+g(x)}]|,$$

where x is uniform over $\{0,1\}^n$.

*eshan@cs.cornell.edu. Supported by NSF Award CCF-2514586.

†hatami.2@osu.edu.

‡chinho.lee@ncsu.edu.

§slovett@ucsd.edu. Supported by Simons Investigator Award #929894 and NSF Award CCF-2425349.

¶atal@berkeley.edu. Supported by NSF CAREER award CCF-2145474.

‖mathematicsoftheimpossible@gmail.com. Supported by NSF award 2430026.

ISSN 1433-8092

For degree-d polynomials over $\mathbb{F}_2$, the fundamental works of Razborov [Raz87] and Smolensky [Smo87, Smo93] proved correlation $O(d/\sqrt{n})$ with the Majority function on n bits, and this bound is tight [Vio21]. Remarkably, for polynomials with degree $d \geq \log n$, no explicit function (e.g. in NP) was known to have asymptotically smaller correlation. Improving this bound was a prerequisite for each of the other long-standing challenges mentioned above.

Here, we finally solve this challenge and in fact obtain exponentially small correlation bounds, $\exp(-\Omega(n/d^2))$. Our results hold over every finite field, and in fact every field. To keep this exposition focused, we consider Boolean-valued polynomials over a finite field $\mathbb{F}$, meaning polynomials that output 0 or 1 on every input in $\{0,1\}^n$. We bound their correlation with the XOR of Majority on k disjoint blocks of ℓ bits, for odd ℓ, denoted $\mathsf{Maj}_\ell^{\oplus k}$. This is a natural candidate that was considered in previous works [Vio06, BVdW07, VW08, CHH+20, Vio21].

Theorem 1. For every field $\mathbb{F}$ and every Boolean-valued polynomial $P\colon \{0,1\}^{k\ell} \to \{0,1\}$ over $\mathbb{F}$ of total degree at most d, we have

$$\mathsf{corr}(\mathsf{Maj}_\ell^{\oplus k}, P) \leq \left(\frac{4d}{\sqrt{\ell}}\right)^k.$$

Over $\mathbb{F}_2$, this bound is tight up to an absolute constant in the base. This follows by taking the XOR of degree-d approximating polynomials for Majority constructed in [Vio21] on the k disjoint blocks.[1]

Our general result also allows Majority to be replaced by any Boolean-valued function that is *versatile* over $\mathbb{F}$, such as Parity when the characteristic of $\mathbb{F}$ is different from two. For arbitrary $\mathbb{F}$-valued polynomials over a finite field $\mathbb{F}$, we obtain analogous bounds after applying any Boolean map to their outputs, with an additional dependence on $|\mathbb{F}|$. Over a prime field $\mathbb{F}_p$, we also obtain bounds for correlations with complex phases $\omega^{P(x)}$, where $\omega = e^{2\pi i/p}$, with an additional dependence on p. We give the precise statements and compare these extensions with prior work in Section 3.2.

Applications. Theorem 1 has several applications. First, we note that taking $\ell = 64d^2 + 1$, gives an explicit function on $n = \ell k$ variables (namely $\mathsf{Maj}_\ell^{\oplus k}$) with correlation at most $2^{-k} = 2^{-\Omega(n/d^2)}$ with all Boolean-valued degree-d polynomials over any field.

These bounds also have consequences for circuit lower bounds. For any fixed prime power $q = p^r$, the polynomial approximation method of Razborov and Smolensky [Raz87, Smo93, KS18] simulates constant-depth circuits with MOD_q gates,[2] namely $\mathsf{AC}^0[q]$ circuits, by low-degree polynomials over $\mathbb{F}_p$. Our correlation bounds for polynomials thus give an explicit function with correlation at most $2^{-n^{\Omega(1)}}$ with $\mathsf{AC}^0[q]$ circuits of size $2^{n^{\Omega(1)}}$ at any fixed depth. By the discriminator lemma [HMP+93], these correlation bounds also imply $2^{n^{\Omega(1)}}$ size lower bounds for circuits consisting of a majority gate applied to $\mathsf{AC}^0[q]$ circuits.

Next, we discuss applications to derandomization, another central goal in complexity theory. The Nisan-Wigderson generator [NW94], based on our hard function, gives an explicit

[1] An earlier version of this paper was restricted to $\mathbb{F}_2$ and used a different weight function, yielding a slightly better constant in the base. That proof remains valid. We changed the weight function (see Section 3.5) to extend the argument to arbitrary fields.

[2] A MOD_q gate outputs 1 if the sum of its input bits is nonzero modulo q, and 0 otherwise. In particular, MOD_2 is a parity gate.

pseudorandom generator (PRG) for Boolean-valued degree-d polynomials over any field on n Boolean inputs with seed length $\mathsf{poly}(d, \log(n/\varepsilon))$. Applying the polarizing random walk generator [CHHL19, CGL+21] gives an improved seed length of $d^2 \cdot \mathsf{polylog}(n/\varepsilon)$. Over $\mathbb{F}_2$, all previous constructions [Bog05, Lov09, Vio09] had an exponential dependence on the degree, and so gave nothing for $d \geq \log n$. As part of the analysis of the polarizing random walk generator, we resolve a conjecture of [CGL+21] on the growth of Fourier coefficients of low-degree polynomials over $\mathbb{F}_2$. In fact, we show that for any Boolean-valued degree-d polynomial P over *any field*, the function $(-1)^P$ has level-k Fourier coefficients whose sum has absolute value bounded by $O(d)^k$ for all k. We provide details in Sections 4.1 and 4.2.

Using the same polynomial simulation as before, we derive PRGs for $\mathsf{AC}^0[q]$ circuits for every fixed prime power q. We obtain a PRG with polylogarithmic seed length for polynomial-size circuits. The previous best seed length was $n(1 - o(1))$ [FSUV13]. We provide details in Section 4.3.

Our generator for low-degree polynomials also yields PRGs for Boolean-valued sparse polynomials over any field. We obtain seed length $\mathsf{polylog}(nS/\varepsilon)$ for polynomials with at most S monomials. Over any finite field, when $S \geq n$ and $\varepsilon \geq 1/\mathsf{poly}(S)$, this improves the $2^{O(\sqrt{\log S})}$ seed length of previous constructions [LVW93, Vio07, ST18, ST22, Kum25] to $\mathsf{polylog}(S)$.

Finally, our PRG for $\mathsf{AC}^0[2]$ yields unconditional pseudorandom functions (PRFs) that are ε-secure against $\mathsf{AC}^0[2]$ oracle circuits. For any constant-depth, size-s oracle circuits making n-bit queries, our PRF has key length $\mathsf{polylog}(ns/\varepsilon)$. No such result was known, even for ordinary AC^0 circuits. The construction combines hashing with a pseudorandom table and relies on our PRG for circuits with parity gates. We provide details in Section 4.4.

AI methodology. This work began with several authors prompting AI tools about correlation bounds for the XOR of Majority. One of these interactions eventually produced a new correlation bound, albeit much weaker than the one presented in this paper, and with a much more complicated proof. The authors then undertook a research project to improve the bounds and simplify the argument through repeated interactions with AI tools, interspersed with their own insights and simplifications. Some intermediate proofs were extremely complex; the final proof emerged through many rounds of refinement. The applications were obtained without AI assistance.

Acknowledgment. This work was done while most of the authors were at the Simons Institute for Theory of Computing, attending the Fall 2026 Pseudorandomness & High-Dimensional Expansion program, which provided an ideal working environment. We thank Gil Cohen for helpful discussions, and Eric Allender and Ronald de Wolf for useful feedback.

Paper organization. We prove Theorem 1 in Section 3 and discuss applications in Section 4. We conclude with open questions in Section 5.

2 Proof Overview

In this section, we provide an overview of the proof of Theorem 1.

Let $\mathbb{F}$ be a field. Let $n = k\ell$, and partition the input into k blocks of ℓ bits each, where $\ell = 2m+1$. Throughout, inputs are Boolean strings, and polynomials are evaluated on $\{0,1\}^n$. Let $h_1, \ldots, h_k$ denote the majority function on the respective blocks, and let $E \subseteq \{0,1\}^n$ be the set of inputs on which a polynomial $P\colon \{0,1\}^n \to \mathbb{F}$ (not necessarily Boolean-valued) of total degree at most d agrees with $H \coloneqq h_1 \oplus \cdots \oplus h_k$.

Our main goal is to show that $|E|$ is not much larger than 2^{n-1}. Since the space of functions $f\colon E \to \mathbb{F}$ has dimension $|E|$, we do this by showing that every function on E is the restriction of a function from a suitably chosen low-dimensional space.

Smolensky in the language of bases. We first review Smolensky's correlation bound [Smo87, Smo93] for $k = 1$ majority, phrased in the language we use later for multiple blocks. Let h denote the majority function on ℓ bits, and let E be the set of inputs on which a polynomial $P\colon \{0,1\}^\ell \to \mathbb{F}$ of degree at most d agrees with h. The vector space of functions $f\colon E \to \mathbb{F}$ has dimension exactly $|E|$, so to show that $|E|$ is not much larger than $2^{\ell-1}$ it suffices to bound the dimension of this space. We do so by picking a basis for the space of functions on the full cube, and then showing that only slightly more than half of the basis elements are needed once the domain is restricted to E.

The basis comes from the interpolation underlying Smolensky's argument. Let $H_0 = \{x \in \{0,1\}^\ell : |x| \le m\}$ and $H_1 = \{x \in \{0,1\}^\ell : |x| \ge m+1\}$ be the two majority halves, so that h is zero on H_0 and one on H_1, and let $U_{\le m}$ be the space of multilinear polynomials of degree at most m. The monomials x^A with $|A| \le m$ form a basis of $U_{\le m}$, and

$$\dim(U_{\le m}) = \sum_{j=0}^{m} \binom{\ell}{j} = 2^{\ell-1} = |H_0| = |H_1|.$$

The interpolation fact is that an arbitrary assignment of values on either H_0 or H_1 extends uniquely to a polynomial in $U_{\le m}$. For H_0 this can be seen directly: if $u(x) = \sum_{|A|\le m} c_A x^A$, its value at the input, $\mathbf{1}_B$, whose set of ones is B is $\sum_{A\subseteq B} c_A$. To interpolate the coefficients c_A from the evaluations $u(\mathbf{1}_B)$ for $|B| \le m$, start with $B = \varnothing$ from which $c_\varnothing$ is implied, and proceed by increasing $|B|$ and retrieving c_B from all previous c_A's and $u(\mathbf{1}_B)$. We get that the prescribed values on H_0 determine every coefficient c_A for $|A| \le m$ uniquely. Complementing all input bits exchanges the two halves and preserves the degree bound, proving the assertion for H_1 as well.

Consequently, every function on the cube has a unique representation $q + hr$, where $q, r \in U_{\le m}$: choose q to interpolate the function on H_0, and r to interpolate the difference on H_1. This gives what we call the *h-basis*, consisting of elements x^A and $x^A h$ with $|A| \le m$. The key additional insight is that, on E, we have $x^A h = x^A P$, whose degree is at most m when $|A| \le m - d$. Thus we can discard these elements, retaining only the ordinary monomials of degree at most m and the elements $x^A h$ with $m - d < |A| \le m$. Counting the retained elements gives $|E| \le 2^{\ell-1} + O(d2^\ell/\sqrt{\ell})$. For a Boolean valued P, applying the same argument to $1 - P$ bounds the complement of E. This is the correlation bound $O(d/\sqrt{\ell})$ of Smolensky [Smo87]; our contribution is to make a similar basis argument work across k blocks.

We note that the above is a slight reformulation of the usual presentation of Smolensky's argument, where one writes a function on E as $q + hr = q + Pr$, with $q, r \in U_{\le m}$, and bounds

$|E|$ by the dimension of the space of polynomials of degree at most $m+d$. In the h-basis, we instead identify explicitly which basis elements are redundant: if $|A| \le m-d$, then on E, $x^A h = x^A P$, and $x^A P$ has degree at most m, so it lies in $U_{\le m}$. Thus only the *high* elements $x^A h$ with $|A| > m-d$ need to be retained. This gives the same dimension bound as the usual degree-counting argument, but the new high/low formulation is what extends naturally to several blocks.

The argument for several blocks. It is natural to consider the products of the single-block bases. Write $\mathbf{A} = (A_1, \dots, A_k)$, where A_i is a set of coordinates from block i, and let $x^{\mathbf{A}} := \prod_{i=1}^k x^{A_i}$. The resulting basis, which we call the h-basis, consists of $x^{\mathbf{A}} h^S$, where $|A_i| \le m$ for every i, $S \subseteq [k]$, and $h^S = \prod_{i \in S} h_i$, where h_i is a copy of h applied to the i-th block. Thus the h-basis is obtained by applying Smolensky's interpolation separately to each block, and then taking products across the blocks.

The single-block argument suggests that when $x^{\mathbf{A}}$ has low degree in some block j, we should try to eliminate the factor h_j from h-basis elements $x^{\mathbf{A}} h^S$ with $j \in S$. In the single-block case, this was done simply by replacing h_j with P, since they agree on E. In the general case, we rely on the fact that on E, for every j,

$$h_j = P \oplus \left(\bigoplus_{i \neq j} h_i \right).$$

Here, $\oplus$ is the binary XOR regardless of the field. In characteristic two, this is simply the substitution

$$h_j = P + \sum_{i \neq j} h_i,$$

which rewrites a basis element containing h_j as a combination of terms in which h_j has disappeared. For characteristic different from two, passing from $\{0,1\}$ to $\{+1,-1\}$ turns XOR into multiplication. Expanding the resulting identity and isolating h_j gives a rewriting rule of the form

$$h_j = P + \sum_{\varnothing \neq U \subseteq [k] \setminus \{j\}} c_U \widetilde{P}\, h^U,$$

where $\widetilde{P} = 1 - 2P$ has degree at most d and the coefficients $c_U = (-2)^{|U|-1}$ lie in $\mathbb{F}$. The important point is that none of the terms on the right-hand side contain h_j. Thus, substituting for h_j removes the selected h_j factor, at the cost of multiplying some terms by a degree-d polynomial.

To make this rewriting strategy work, we use a slightly more conservative notion of low degree than in the single-block argument. We call block i *low* for $\mathbf{A}$ if $|A_i| \le m - 2d$, and *high* otherwise. If every block is high for $\mathbf{A}$, we retain all 2^k basis elements $x^{\mathbf{A}} h^S$. Otherwise, let $j = j(\mathbf{A})$ be the first low block and retain precisely those $x^{\mathbf{A}} h^S$ with $j \notin S$. Let $\mathcal{R}$ denote the span of the retained elements. We will prove that $\mathcal{R}$ still interpolates every function on E, by showing via induction that every basis element agrees on E with a function in $\mathcal{R}$.

Consider a "discarded" basis element $x^{\mathbf{A}} h^S$, so that $\mathbf{A}$ has a low block and $j = j(\mathbf{A}) \in S$. Then, substitute for h_j using the rewriting rule above. Each resulting term has the form $x^{\mathbf{A}} h^T Q$, where $j \notin T$ and $Q \in \{P, \widetilde{P}\}$ has degree at most d. After multiplying by Q and

expanding, the resulting terms may have some blocks with degree larger than m. Whenever this happens, we use interpolation in that block to rewrite the corresponding monomial as $q_i + h_i r_i$, with $\deg(q_i), \deg(r_i) \le m$, simplify using $h_i^2 = h_i$, and continue until the expression is written in the h-basis.

Crucially, this rewriting is never needed in the selected block j. Since $|A_j| \le m - 2d$, multiplication by Q raises the ordinary degree in block j to at most $m - 2d + d \le m$. Hence no h_j factor can be introduced in block j. In particular, if a resulting basis element is still not retained, its selected low block must differ from j, and therefore some block must have changed its low/high status.

It remains to show that this rewriting terminates. Let

$$\mathsf{High}(\mathbf{A}) \coloneqq \{i : |A_i| > m - 2d\},$$

and assign to each basis element the weight

$$w(x^{\mathbf{A}} h^S) \coloneqq 2|\mathbf{A}| - (2d+1)|\mathsf{High}(\mathbf{A})|,$$

where $|\mathbf{A}| = \sum_i |A_i|$. Note that the weight depends only on $\mathbf{A}$. Multiplying by a degree-d polynomial raises the total x-degree by at most d, and hence the weight by at most $2d$, unless some block changes its low/high status. The weight is designed so that any such change saves an additional $2d+1$. The low-to-high direction is immediate. In the other direction, a high-to-low change is possible only if multiplication first pushes the degree in that block above m, after which basis reduction brings it down to at most $m - 2d$. Thus the degree drops from at least $m+1$ to at most $m - 2d$, a saving of at least $2(2d+1)$ in weight, while losing the high-block discount costs only $2d+1$.

Since every non-retained term produced by the rewriting has a changed set of high blocks, its weight is at most

$$w(x^{\mathbf{A}} h^S) + 2d - (2d+1) < w(x^{\mathbf{A}} h^S).$$

Induction on the weight therefore shows that every basis element agrees on E with a function in $\mathcal{R}$, and so $\mathcal{R}$ interpolates every function on E.

It remains to count the retained elements. For every $\mathbf{A}$ having at least one low block, we retain exactly half of the 2^k choices of S; we retain all choices only when every block is high. The fraction of one-block monomials with degree greater than $m - 2d$ is at most $4d/\sqrt{\ell}$, so the fraction of all-high $\mathbf{A}$ is at most

$$\left(\frac{4d}{\sqrt{\ell}}\right)^k.$$

It follows that

$$\frac{|E|}{2^n} \le \frac{\dim \mathcal{R}}{2^n} \le \frac{1}{2}\left(1 + \left(\frac{4d}{\sqrt{\ell}}\right)^k\right).$$

This agreement bound holds for an arbitrary polynomial P over $\mathbb{F}$ of degree at most d.

Correlation over arbitrary fields. If P is Boolean-valued, then $1 - P$ agrees with $H = h_1 \oplus \cdots \oplus h_k$ exactly on the complement of E. Applying the agreement bound to both P and $1 - P$ therefore gives $\mathsf{corr}(H, P) \le (4d/\sqrt{\ell})^k$.

Correlation with a complex phase. For a prime field $\mathbb{F}_p$, we also consider correlation with the complex phase $\omega^{P(x)}$, where $\omega = e^{2\pi i/p}$. We show that

$$\left|\mathop{\mathbb{E}}_{x}[(-1)^{H(x)}\omega^{P(x)}]\right| \leq p \cdot (4d(p-1)/\sqrt{\ell})^k.$$

3 Proof of Main Result

We prove a general agreement bound for XORs of versatile functions over arbitrary fields, from which Theorem 1 follows. We first introduce some notation and conventions used throughout this section.

For a finite set X, we write $x \sim X$ for sampling a uniformly random element from X. Let $\mathbb{F}$ be an arbitrary field, and let p denote its characteristic, which is either zero or a prime. If $\mathbb{F}$ is finite, then p is prime and $|\mathbb{F}| = p^r$ for some integer $r \geq 1$. We work throughout this section with multivariate polynomials $P \in \mathbb{F}[x_1, \ldots, x_n]$ evaluated on $\{0,1\}^n$. Since we are interested in their behavior as functions, such polynomials are without loss of generality multilinear using the identity $x_i^2 = x_i$ that holds over $\{0,1\}$. We identify such polynomials also as functions $P\colon \{0,1\}^n \to \mathbb{F}$. We identify $\{0,1\}$ as a subset of $\mathbb{F}$ in the natural way. We say that a polynomial $P\colon \{0,1\}^n \to \mathbb{F}$ is Boolean-valued if $P(x) \in \{0,1\}$ for every $x \in \{0,1\}^n$. We also may write it as $P\colon \{0,1\}^n \to \{0,1\}$ when the context of the field we're working over is clear. We refer to the degree of a multivariate polynomial as its total degree.

3.1 Versatile Functions

Our main theorem establishes exponentially decaying correlation bounds between the XOR of k copies of any versatile function on ℓ bits (e.g., the Majority function) and any Boolean-valued polynomial $P\colon \{0,1\}^n \to \mathbb{F}$ of degree at most d.

First, we define the concept of a versatile function. Let $\ell = 2m+1$ be an odd integer. Consider the space U of multilinear polynomials on ℓ variables over $\mathbb{F}$. Its dimension is 2^ℓ. Let $U_{\leq m} \subseteq U$ be the subspace consisting of polynomials of degree at most m. Note that $\dim(U_{\leq m}) = \sum_{i=0}^{m} \binom{\ell}{i} = 2^{\ell-1}$, exactly half the dimension of U.

Definition 1 (Versatile Function [Kop11])**.** We say that a function $h\colon \{0,1\}^\ell \to \mathbb{F}$ is versatile over $\mathbb{F}$ if any function $f\colon \{0,1\}^\ell \to \mathbb{F}$ can be written (uniquely) as

$$f = q + hr, \qquad \text{for } q, r \in U_{\leq m}. \tag{1}$$

We remark that uniqueness comes for free from dimension counting, as $2\dim(U_{\leq m}) = \dim(U)$. The versatility of a *Boolean-valued* h is equivalent to an interpolating property with respect to $h^{-1}(0)$ and $h^{-1}(1)$. In particular, it implies that h is balanced, i.e., $|h^{-1}(0)| = |h^{-1}(1)| = 2^{\ell-1}$.

We record the following useful result and supply its proof for completeness.

Lemma 3.1 (Versatility in terms of Interpolation)**.** *A function $h\colon \{0,1\}^\ell \to \{0,1\}$ is versatile over $\mathbb{F}$ if and only if both $H_0 = h^{-1}(0)$ and $H_1 = h^{-1}(1)$ are interpolating sets for $U_{\leq m}$.*

Proof. In the forward direction, let $g\colon H_0 \to \mathbb{F}$. Extend g to a function g' from $\{0,1\}^\ell \to \mathbb{F}$ by mapping $x \notin H_0$ arbitrarily to values in $\mathbb{F}$. Then, by versatility of h, $g' = q + hr$ on $\{0,1\}^\ell$ for $q, r \in U_{\leq m}$. Thus, on $x \in H_0$ we have $g(x) = g'(x) = q(x)$ and so q interpolates g on H_0.

Similarly, for any $g\colon H_1 \to \mathbb{F}$, extend g to a function g' from $\{0,1\}^\ell \to \mathbb{F}$ by mapping $x \notin H_1$ arbitrarily to values in $\mathbb{F}$. Then, by versatility of h, $g' = q + hr$ on $\{0,1\}^\ell$ for $q, r \in U_{\leq m}$. Thus, on $x \in H_1$ we have $g(x) = g'(x) = q(x) + r(x)$ and so $q + r \in U_{\leq m}$ interpolates g on H_1. Thus, $|H_0|$ and $|H_1|$ are bounded by the dimension of $U_{\leq m}$ which is $2^{\ell-1}$. Combined with $|H_0| + |H_1| = 2^\ell$, this implies $|H_0| = |H_1| = 2^{\ell-1}$, which guarantees uniqueness of the interpolant.

In the other direction, assume H_0, H_1 are interpolating sets for $U_{\leq m}$. Then, both sets have size equal to $\dim(U_{\leq m}) = 2^{\ell-1}$, as otherwise, one of them would have size larger than $\dim(U_{\leq m})$ which contradicts being an interpolating set.

Interpolate f on H_0 and H_1 respectively to get polynomials $q_0, q_1 \in U_{\leq m}$ with $f|_{H_0} \equiv q_0$ and $f|_{H_1} \equiv q_1$. We claim that

$$f = q_0 + h(q_1 - q_0).$$

This is immediate from the cases $h(x) = 1$ and $h(x) = 0$ on H_1 and H_0, respectively. Uniqueness follows from dimension counting. □

For our applications, Majority is versatile over any field $\mathbb{F}$ by the interpolation argument in Section 2. If $p \neq 2$, then Parity_ℓ (parity of ℓ bits) is versatile over $\mathbb{F}$ as well.

Lemma 3.2. *Let $\mathbb{F}$ be a field of characteristic $p \neq 2$. Then Parity_ℓ is versatile over $\mathbb{F}$.*

Proof. As Parity_ℓ is Boolean-valued, it suffices to prove the interpolation property. Write $y_i = 1 - 2x_i = (-1)^{x_i}$. Since $p \neq 2$, the two functions 1 and y_i form a basis for the functions of the i-th input bit. Taking tensor products over the ℓ input bits shows that the characters $y^S = \prod_{i \in S} y_i$, for $S \subseteq [\ell]$, form a basis for U.

Next, the key is the identity

$$y^A = y^{[\ell]} \cdot y^{[\ell] \setminus A},$$

where $y^{[\ell]} = (-1)^{\sum_{i=1}^{\ell} x_i}$. On $\mathsf{Parity}_\ell^{-1}(b)$, the factor $y^{[\ell]}$ is simply the constant $(-1)^b$. Since $\ell = 2m + 1$, one of A and $[\ell] \setminus A$ has size at most m, and hence on $\mathsf{Parity}_\ell^{-1}(b)$ every character is a fixed scalar multiple of a character of degree at most m. There are exactly $2^{\ell-1} = |\mathsf{Parity}_\ell^{-1}(b)|$ such characters, so the interpolation is unique. □

3.2 Statement of the Main Result and Its Corollaries

Theorem 3.3. *Let $\mathbb{F}$ be a field, and let $k \geq 1$ and $m \geq 0$. Set $\ell = 2m + 1$ and $n = k\ell$. Let $h\colon \{0,1\}^\ell \to \{0,1\}$ be a Boolean-valued versatile function over $\mathbb{F}$. Let $H\colon \{0,1\}^n \to \{0,1\}$ be*

$$H(x) = \bigoplus_{i=1}^{k} h_i(x),$$

where $x = (x^{(1)}, \ldots, x^{(k)})$, each block of x, denoted $x^{(i)}$, consists of ℓ bits, and $h_i(x) = h(x^{(i)})$.

Then, for every integer $d \geq 0$ and polynomial $P\colon \{0,1\}^n \to \mathbb{F}$ of total degree at most d,

$$\Pr_{x\sim\{0,1\}^n}[P(x) = H(x)] \leq \frac{1}{2} + \frac{1}{2}\left(\sqrt{\frac{2}{\pi}} \cdot \frac{4d}{\sqrt{\ell}}\right)^k \leq \frac{1}{2} + \frac{1}{2}\left(\frac{4d}{\sqrt{\ell}}\right)^k.$$

Corollary 3.4. *Under the same conditions as Theorem 3.3,*

1. *For any Boolean-valued polynomial $P\colon \{0,1\}^n \to \{0,1\}$ over $\mathbb{F}$ of total degree at most d,*

$$\mathsf{corr}(H, P) = \left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}[(-1)^{H(x)+P(x)}]\right| \leq \left(\sqrt{\frac{2}{\pi}} \cdot \frac{4d}{\sqrt{\ell}}\right)^k.$$

2. *Suppose that $\mathbb{F}$ is a finite field. For any function $g\colon \mathbb{F} \to \{0,1\}$, and any polynomial $P\colon \{0,1\}^n \to \mathbb{F}$ of total degree at most d,*

$$\mathsf{corr}(H, g\circ P) = \left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}[(-1)^{H(x)+g(P(x))}]\right| \leq \left(\sqrt{\frac{2}{\pi}} \cdot \frac{4d(|\mathbb{F}|-1)}{\sqrt{\ell}}\right)^k.$$

3. *Suppose that $\mathbb{F} = \mathbb{F}_p$ is a prime field, and let $\omega = e^{2\pi i/p}$. For any polynomial $P\colon \{0,1\}^n \to \mathbb{F}_p$ of total degree at most d,*

$$\left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}[\omega^{P(x)} \cdot (-1)^{H(x)}]\right| \leq \frac{p}{2}\left(\sqrt{\frac{2}{\pi}} \cdot \frac{4(p-1)d}{\sqrt{\ell}}\right)^k.$$

We now prove the corollary assuming the correctness of Theorem 3.3.

Proof. For an integer $t \geq 0$, write $\delta_t \coloneqq \left(\sqrt{\frac{2}{\pi}} \cdot \frac{4t}{\sqrt{\ell}}\right)^k$.

1. Let $E = \{x \in \{0,1\}^n : P(x) = H(x)\}$. By Theorem 3.3, over $x \sim \{0,1\}^n$,

$$\Pr[E] \leq \frac{1}{2} + \frac{\delta_d}{2}.$$

Since P is Boolean-valued, $1 - P$ agrees with H precisely on the complement of E, denoted E^c. Applying Theorem 3.3 to $1 - P$ gives

$$\Pr[E^c] \leq \frac{1}{2} + \frac{\delta_d}{2},$$

thus $\Pr[E] \geq 1/2 - \delta_d/2$. Consequently,

$$\mathsf{corr}(H, P) = |2\Pr[E] - 1| \leq \delta_d.$$

2. Since $\mathbb{F}$ is finite, Lagrange interpolation gives a polynomial $\widetilde{g} \in \mathbb{F}[z]$ of degree at most $|\mathbb{F}| - 1$ that agrees with g on $\mathbb{F}$. Hence $g \circ P$ is represented on $\{0,1\}^n$ by the Boolean-valued polynomial $\widetilde{g} \circ P$, whose degree is at most $(|\mathbb{F}| - 1)d$. The first item therefore gives

$$\mathsf{corr}(H, g\circ P) \leq \delta_{(|\mathbb{F}|-1)d}.$$

3. Versatility implies that h is balanced, and hence so is H, the XOR of its k copies on disjoint blocks. Thus $\mathbb{E}_{x\sim\{0,1\}^n}[(-1)^{H(x)}] = 0$. For each $a \in \mathbb{F}_p$, let $g_a(t) = \mathbf{1}[t = a]$. By the second item,

$$\left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}\left[(-1)^{H(x)+g_a(P(x))}\right]\right| \leq \delta_{(p-1)d}.$$

Using $(-1)^{g_a(P(x))} = 1 - 2\cdot\mathbf{1}[P(x) = a]$ and the balance of H, we obtain

$$\left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}\left[(-1)^{H(x)}\cdot\mathbf{1}[P(x) = a]\right]\right| \leq \frac{\delta_{(p-1)d}}{2}. \tag{2}$$

Finally,

$$\begin{aligned}\left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}\left[\omega^{P(x)}(-1)^{H(x)}\right]\right| &= \left|\sum_{a\in\mathbb{F}_p}\omega^a \mathop{\mathbb{E}}_{x\sim\{0,1\}^n}\left[(-1)^{H(x)}\cdot\mathbf{1}[P(x) = a]\right]\right| \\ &\leq \sum_{a\in\mathbb{F}_p}|\omega^a|\cdot\left|\mathop{\mathbb{E}}_{x\sim\{0,1\}^n}\left[(-1)^{H(x)}\cdot\mathbf{1}[P(x) = a]\right]\right| \\ &\leq p\cdot\frac{\delta_{(p-1)d}}{2}, \qquad \text{(since } |\omega^a| = 1 \text{ and using Equation (2))}\end{aligned}$$

as claimed. □

Comparison with previous bounds. Over a fixed odd prime field $\mathbb{F}_p$, earlier correlation bounds for parity against polynomial phases and Boolean maps of polynomial outputs were obtained by Bourgain [Bou05], Green, Roy, and Straubing [GRS05], and Chattopadhyay [Cha07]. Chattopadhyay's bound has the form $\exp(-\Omega(n/C^d))$, where $C = p2^{p-1}$. For $d \geq \log_2 n$, these estimates do not improve on the classical $O(d/\sqrt{n})$ bound for Boolean-valued polynomials [Smo87, Smo93]. Taking $h = \mathsf{Parity}_\ell$ and $\ell = \Theta(d^2)$ in Corollary 3.4 gives $\exp(-\Omega(n/d^2))$, replacing exponential dependence on d by quadratic dependence. We note that the earlier results worked for composite moduli as well.

Organization. The proof of Theorem 3.3 is an extension of Smolensky's celebrated linear-algebraic encoding argument. We start by introducing a suitable basis for the space of functions from $\{0,1\}^n$ to $\mathbb{F}$ in Section 3.3 based on the versatile function h. We then state our key lemma and prove the main result assuming it in Section 3.4. Finally, in Section 3.5, we prove the lemma using a potential function called *weight* that decreases under repeated applications of the simplification rules.

3.3 The h-basis

We start from a versatile function h on a single block of ℓ bits with the property that every function f on ℓ bits can be written as $f = q + hr$ for polynomials q, r of degree at most m.

Generalizing to $k > 1$ blocks, we consider the space of polynomials whose degree *in each block* is at most m, along with the corresponding functions $h_1, \ldots, h_k$. A monomial *of degree*

at most m in each block is specified by a tuple of subsets $\mathbf{A} = (A_1, \ldots, A_k)$, where $A_i \subseteq [\ell]$ and $|A_i| \leq m$. Let $\mathcal{M}$ be the set of all such tuples, and write

$$x^{\mathbf{A}} \coloneqq \prod_{i=1}^{k} \prod_{j \in A_i} x_j^{(i)}.$$

The degree of this monomial in the i-th block is $|A_i|$. For $S \subseteq [k]$ we write h^S for $\prod_{i \in S} h_i$. It is useful to think of $h_1, \ldots, h_k$ as new formal variables. This brings us to the definition of the h-basis:

Definition 2. The *h-basis* consists of basis functions $\{x^{\mathbf{A}} h^S\}_{\mathbf{A},S}$ where $\mathbf{A} \in \mathcal{M}$ and $S \subseteq [k]$.

Next, we prove that this is indeed a basis for the space of functions from $\{0,1\}^{k\ell}$ to $\mathbb{F}$.

Lemma 3.5. *Every function $f \colon \{0,1\}^{k\ell} \to \mathbb{F}$ can be written uniquely as a linear combination of basis elements in the h-basis.*

Proof. We first show that the h-basis spans every function. As every function from $\{0,1\}^{k\ell}$ to $\mathbb{F}$ can be expressed as a multilinear polynomial over $\mathbb{F}$ in variables $\{x_j^{(i)}\}_{i \in [k], j \in [\ell]}$, it suffices to show that every monomial in these variables lies in the span of the h-basis. To do that, first express each such monomial $x^{\mathbf{B}}$ for $\mathbf{B} = (B_1, \ldots, B_k)$ with $B_i \subseteq [\ell]$ (note that $|B_i|$ can be larger than m) as a product of k monomials, $M_1, \ldots, M_k$, where $M_i = \prod_{j \in B_i} x_j^{(i)}$. Then, use Equation (1) to express M_i as $q_i(x^{(i)}) + h_i r_i(x^{(i)})$, where $\deg(q_i), \deg(r_i) \leq m$ and multiply the results to get that $x^{\mathbf{B}}$ is a linear combination of the basis elements in the h-basis.

Uniqueness follows by counting dimensions, since we have $(2^\ell/2)^k \cdot 2^k = 2^{\ell k}$ basis functions, which is the dimension of the space of functions on $k\ell$ bits. □

3.4 Proof of the Main Theorem by Dimension Counting

Let P be a multilinear polynomial over $\mathbb{F}$ in $k\ell$ variables of total degree at most d. Define the agreement set

$$E \coloneqq \{x \in \{0,1\}^{k\ell} : P(x) = H(x)\}.$$

We need to show that the relative size of E is not much larger than $1/2$. To do so, we will precisely follow Smolensky's idea and show that every function on E is a (restriction of a) function from a space of dimension only slightly larger than half that of the full function space. The new space is defined in the h-basis as follows. For each monomial $x^{\mathbf{A}}$ with a low-degree block, we shall pick the first such block j and keep only the 2^{k-1} basis terms $\{x^{\mathbf{A}} h^S\}_{S \subseteq [k] \setminus \{j\}}$ not involving h_j. For monomials $x^{\mathbf{A}}$ whose degree is high in *every block*, we keep all 2^k terms $\{x^{\mathbf{A}} h^S\}_{S \subseteq [k]}$. A simple calculation bounds the dimension of the space spanned by these basis terms, which then bounds $|E|$.

More precisely, we define the set of *top* monomial indices

$$\mathcal{T} \coloneqq \{\mathbf{A} \in \mathcal{M} : |A_i| > m - 2d \text{ for every } i \in [k]\}.$$

For every non-top $\mathbf{A} \notin \mathcal{T}$ choose a low-degree block[3]

$$j(\mathbf{A}) \coloneqq \min\{i \in [k] : |A_i| \le m - 2d\}.$$

Define $\mathcal{R}$ to be the $\mathbb{F}$-linear span of the following basis functions:

$$\mathsf{Allowed} \coloneqq \{x^{\mathbf{A}} h^S : \mathbf{A} \in \mathcal{T},\ S \subseteq [k]\} \cup \{x^{\mathbf{A}} h^S : \mathbf{A} \notin \mathcal{T},\ S \subseteq [k],\ j(\mathbf{A}) \notin S\}. \tag{3}$$

In words, for a non-top monomial $x^{\mathbf{A}}$ we choose one low-degree block $j(\mathbf{A})$ and forbid all products containing the factor $h_{j(\mathbf{A})}$. For a top monomial we impose no restriction.

Next is the key lemma enabling our dimension count.

Lemma 3.6 (Agreement-Set Interpolation)**.** *For every function $a\colon E \to \mathbb{F}$, there exists a function $r \in \mathcal{R}$ such that $r(x) = a(x)$ for every $x \in E$.*

Assuming this lemma, we can complete the proof of the main theorem by counting the dimension of $\mathcal{R}$.

Proof of Theorem 3.3. Let

$$M_\ell \coloneqq \sum_{r=0}^{m} \binom{\ell}{r} = 2^{\ell-1}, \qquad N_{\ell,2d} \coloneqq \sum_{r=\max\{0,m-2d+1\}}^{m} \binom{\ell}{r}.$$

There are M_ℓ^k ordinary monomials in $\mathcal{M}$, and exactly $N_{\ell,2d}^k$ top monomials in $\mathcal{T}$. For each non-top monomial $x^{\mathbf{A}}$, exactly 2^{k-1} subsets $S \subseteq [k]$ avoid $j(\mathbf{A})$. For each top monomial, all 2^k subsets are allowed. Hence

$$\begin{aligned} \dim(\mathcal{R}) &= 2^{k-1}\left(M_\ell^k - N_{\ell,2d}^k\right) + 2^k N_{\ell,2d}^k \\ &= 2^{k-1} M_\ell^k + 2^{k-1} N_{\ell,2d}^k \\ &= 2^{k\ell-1} + 2^{k-1} N_{\ell,2d}^k. \end{aligned} \tag{4}$$

By Lemma 3.6, the restriction map $\mathcal{R} \to \mathbb{F}^E$ is surjective. Since the space of functions on E has dimension $|E|$, Equation (4) gives

$$|E| \le 2^{k\ell-1} + 2^{k-1} N_{\ell,2d}^k.$$

There are at most $2d$ summands in $N_{\ell,2d}$, each at most $\binom{\ell}{m} \le \sqrt{\frac{2}{\pi}}\frac{2^\ell}{\sqrt{\ell}}$, which gives

$$|E| \le 2^{k\ell-1} + 2^{k-1}\left(2d \cdot \sqrt{\frac{2}{\pi}} \cdot \frac{2^\ell}{\sqrt{\ell}}\right)^k = 2^{k\ell} \cdot \left(\frac{1}{2} + \frac{1}{2} \cdot \left(\sqrt{\frac{2}{\pi}} \cdot \frac{4d}{\sqrt{\ell}}\right)^k\right).$$

Consequently,

$$\Pr_{x \sim \{0,1\}^n}[P(x) = H(x)] = \frac{|E|}{2^{k\ell}} \le \frac{1}{2} + \frac{1}{2} \cdot \left(\sqrt{\frac{2}{\pi}} \cdot \frac{4d}{\sqrt{\ell}}\right)^k. \qquad \square$$

[3]Choosing $j(\mathbf{A})$ as the first low-degree block is only for convenience. Any fixed selection rule depending only on the set of low-degree blocks would work.

3.5 Proof of the Agreement-Set Interpolation Lemma

It remains to prove the key Lemma 3.6.

We first describe the rewriting rule. On E, we have $P = H = h_1 \oplus \cdots \oplus h_k$, where $\oplus$ denotes Boolean XOR. In particular, P is Boolean-valued on E, and for each $j \in [k]$ we can isolate

$$h_j = P \oplus \left(\bigoplus_{i \neq j} h_i \right).$$

If the characteristic p is 2, this simplifies on E to $h_j = P + \sum_{i \neq j} h_i$. If $p \neq 2$, then 2 is invertible in $\mathbb{F}$, and we translate the outputs of $h_1, \ldots, h_k$ and P from $\{0,1\}$ to $\{+1,-1\}$ via the invertible map $b \mapsto 1 - 2b$. Under this map XOR becomes multiplication, so on E,

$$1 - 2h_j = (1 - 2P) \prod_{i \neq j} (1 - 2h_i).$$

Isolating h_j gives

$$\begin{aligned} h_j &= \frac{1 - (1-2P)\prod_{i \neq j}(1-2h_i)}{2} \\ &= \frac{1}{2} - \frac{1-2P}{2} \sum_{U \subseteq [k] \setminus \{j\}} (-2)^{|U|} h^U \\ &= P + \sum_{\varnothing \neq U \subseteq [k] \setminus \{j\}} (-2)^{|U|-1} (1-2P) h^U \qquad \text{on } E. \end{aligned} \tag{5}$$

Crucially, none of the monomials h^U on the right-hand side depends on h_j. We note that the expression in Equation (5) reduces to $h_j = P + \sum_{i \neq j} h_i$ in characteristic 2: when $|U| = 1$ the coefficient of h^U is 1 (using $1 - 2P = 1$), while for $|U| \geq 2$ it vanishes.

Thus, over any field $\mathbb{F}$, we have

$$h_j = P + \sum_{\varnothing \neq U \subseteq [k] \setminus \{j\}} (-2)^{|U|-1} (1-2P) h^U \qquad \text{on } E.$$

For ease of notation, set

$$\widetilde{P} := 1 - 2P, \qquad c_U := (-2)^{|U|-1} \in \mathbb{F}.$$

Then the total degree of $\widetilde{P}$ is at most d, and the preceding identity becomes

$$h_j = P + \sum_{\varnothing \neq U \subseteq [k] \setminus \{j\}} c_U \widetilde{P} h^U \qquad \text{on } E.$$

Thus, for a forbidden term $T = x^{\mathbf{A}} h^S$ with $j = j(\mathbf{A}) \in S$, we can write

$$T = x^{\mathbf{A}} h^{S \setminus \{j\}} \cdot h_j = x^{\mathbf{A}} h^{S \setminus \{j\}} \cdot \left(P + \sum_{\varnothing \neq U \subseteq [k] \setminus \{j\}} c_U \widetilde{P} h^U \right) \qquad \text{on } E, \tag{6}$$

where, crucially, none of the displayed h-monomials on the right-hand side of (6) contains h_j.

We then express the right-hand side of (6) in the h-basis and show that each term in its expansion is either an allowed term (in Allowed) or has a smaller weight, to be defined next.

Definition 3 (x-Degree, High Blocks, Weight)**.** For $\mathbf{A} \in \mathcal{M}$, write

$$|\mathbf{A}| \coloneqq \sum_{i=1}^{k} |A_i|, \qquad \mathsf{High}(\mathbf{A}) \coloneqq \{i \in [k] : |A_i| > m - 2d\}.$$

We call block i *high* if $i \in \mathsf{High}(\mathbf{A})$ and *low* otherwise. For a basis function $x^{\mathbf{A}}h^S$, define its *x-degree* as $|\mathbf{A}|$ and its *block-i x-degree* as $|A_i|$. Define its weight by

$$w(x^{\mathbf{A}}h^S) \coloneqq 2|\mathbf{A}| - (2d+1)|\mathsf{High}(\mathbf{A})|.$$

Crucially, the weight of $x^{\mathbf{A}}h^S$ depends only on $\mathbf{A}$, not on S: it is twice the x-degree, with a discount of $2d+1$ for each high block. Thus, multiplying by a polynomial in $h_1, \ldots, h_k$ and reducing via $h_i^2 = h_i$ produces only basis terms of the same weight, since every resulting basis term has the same x-monomial index $\mathbf{A}$. As we show next, multiplication by an input variable can increase the x-degree by at most one, and hence the weight by at most two. Furthermore, if a block changes its low/high status, the discount or the degree loss from basis reduction yields an additional saving of $2d+1$. The next lemma makes this accounting precise.

Lemma 3.7 (Multiplication by One Variable)**.** *Let $T = x^{\mathbf{A}}h^S$ be a basis function, and let $z = x_e^{(i)}$ be an input variable in block i. Suppose $T' = x^{\mathbf{A}'}h^{S'}$ has nonzero coefficient in the h-basis expansion of zT.*

Then $w(T') \le w(T) + 2$. Furthermore, if block i (containing z) changes its low/high status, then

$$w(T') \le w(T) + 2 - (2d+1).$$

Proof. First note that only the block containing z (i.e., block i) changes. Let a and a' be the block-i x-degrees of T and T', respectively. We consider three cases:

1. If $e \in A_i$, then multiplying by $z = x_e^{(i)}$ does nothing as $z^2 = z$ and thus $a' = a$.
2. If $e \notin A_i$ and $a \le m-1$, then $zx^{\mathbf{A}}h^S$ is a valid basis function and thus $T' = zx^{\mathbf{A}}h^S$ and $a' = a+1$.
3. Finally, if $e \notin A_i$ and $a = m$, then multiplication by z increases the block-i x-degree to $m+1$. We then use the versatility of h in the i-th block to rewrite zx^{A_i} as $q_i + h_i r_i$, where $\deg(q_i), \deg(r_i) \le m$. This is called "basis reduction". Each basis term $T' = x^{\mathbf{A}'}h^{S'}$ in this expansion has block-i x-degree at most m, and the other blocks are unaffected.

In all cases, $a' \le a+1$.

If the low/high status of block i is unchanged, $\mathsf{High}(\mathbf{A}) = \mathsf{High}(\mathbf{A}')$, so the weight increases by at most 2.

If the block changes from low to high, then necessarily $a = m - 2d$, $a' = m - 2d + 1$, and no basis reduction occurs. Its block-i x-degree increases by one and it earns the discount $2d+1$, giving weight change $2 - (2d+1)$.

Finally, a high-to-low change requires basis reduction, since without reduction the block-i x-degree can only increase. Hence $a = m$ and $a' \le m - 2d$ and the weight change satisfies

$$w(T') - w(T) \le 2(a' - a) + (2d+1) \le -4d + (2d+1) = 2 - (2d+1).$$

This proves both assertions. □

Corollary 3.8 (Multiplication by a Low-Degree Polynomial)**.** *Let $t \geq 0$ be an integer, let $T = x^{\mathbf{A}} h^S$ be a basis function, and let q be a multivariate polynomial in the input variables of total degree at most t. Suppose that $T' = x^{\mathbf{A}'} h^{S'}$ has nonzero coefficient in the h-basis expansion of qT. Then $w(T') \leq w(T) + 2t$. If $\mathsf{High}(\mathbf{A}) \neq \mathsf{High}(\mathbf{A}')$, then*

$$w(T') \leq w(T) + 2t - (2d+1).$$

Proof. By linearity, it suffices first to take $q = z_1 \cdots z_s$ to be a monomial, where $s \leq t$ and each z_r is an input variable. Multiply by the variables one at a time, expanding in the h-basis after each multiplication. Any basis term T' in the final expansion arises along a sequence of basis terms

$$T = T_0, T_1, \ldots, T_s = T',$$

where T_r has nonzero coefficient in the h-basis expansion of $z_r T_{r-1}$. Applying Lemma 3.7 at every step gives $w(T') \leq w(T) + 2s \leq w(T) + 2t$.

If $\mathsf{High}(\mathbf{A}) \neq \mathsf{High}(\mathbf{A}')$, then at least one step in the sequence changes the low/high status of a block. Applying the stronger conclusion of Lemma 3.7 at that step gives

$$w(T') \leq w(T) + 2s - (2d+1) \leq w(T) + 2t - (2d+1).$$

The result for a general polynomial q follows by expanding it into monomials: a basis term with nonzero final coefficient must occur in the expansion of at least one monomial of q. □

Proof of Lemma 3.6. Every function on E extends to the full cube, where it has an h-basis expansion. It therefore suffices to show that each basis function $T = x^{\mathbf{A}} h^S$ agrees with some function in $\mathcal{R}$ on E. We use induction on $w(T)$, with the basis functions of minimum weight as the base case.

If $T \in \mathsf{Allowed}$ (defined in Equation (3)), there is nothing to prove. Otherwise, $\mathbf{A} \notin \mathcal{T}$ and $j = j(\mathbf{A}) \in S$ is a block with $|A_j| \leq m - 2d$. Apply Equation (6) to get

$$T = x^{\mathbf{A}} h^{S \setminus \{j\}} \cdot h_j = x^{\mathbf{A}} h^{S \setminus \{j\}} \cdot \left(P + \sum_{\varnothing \neq U \subseteq [k] \setminus \{j\}} c_U \widetilde{P} h^U \right).$$

We show that each term on the right-hand side expands in the h-basis as a linear combination of allowed terms and terms T' with $w(T') < w(T)$.

Expanding the right-hand side, and using the identities $h_i^2 = h_i$ for $i \in [k]$, each term in it is of the form $c\, x^{\mathbf{A}} h^{U'} P'$, where $c \in \mathbb{F} \setminus \{0\}$, $U' \subseteq [k] \setminus \{j\}$ and $P' \in \{P, \widetilde{P}\}$ has total degree at most d.

Fix such a term $c\, x^{\mathbf{A}} h^{U'} P'$, where $j \notin U'$, and let $T' = x^{\mathbf{A}'} h^{S'}$ be a basis function occurring in the h-basis expansion of $c\, x^{\mathbf{A}} h^{U'} P'$. We claim that $j \notin S'$. Indeed, after multiplication of $c\, x^{\mathbf{A}} h^{U'}$ by any monomial of P', its block-j x-degree is at most

$$|A_j| + d \leq (m - 2d) + d = m - d \leq m.$$

The jth block therefore needs no basis reduction, so h_j cannot reappear.

If T' is not allowed, then $\mathbf{A}'$ is not top and $j(\mathbf{A}') \in S'$. But, since $j \notin S'$, we must have $j(\mathbf{A}') \neq j(\mathbf{A})$. As the selected block is determined by the set of low blocks, it must be the case that $\mathsf{High}(\mathbf{A}') \neq \mathsf{High}(\mathbf{A})$. By Corollary 3.8,

$$w(T') \leq w(x^{\mathbf{A}} h^{U'}) + 2\deg(P') - (2d+1) \leq w(T) - 1,$$

where we used $\deg(P') \leq d$ and the fact that the weight depends only on $\mathbf{A}$, so $w(x^{\mathbf{A}} h^{U'}) = w(x^{\mathbf{A}} h^{S}) = w(T)$.

Induction applies to every such term, while the allowed terms already belong to $\mathcal{R}$. Thus Equation (6) expresses T, on E, as a linear combination of allowed generators. □

4 Applications

In Section 4.1, we translate correlation bounds with XOR of majorities into Fourier bounds. We use these in Section 4.2 to obtain PRGs for Boolean-valued low-degree polynomials and sparse polynomials over any arbitrary field $\mathbb{F}$. In Section 4.3, we obtain PRGs for $\mathsf{AC}^0[q]$ circuits for every fixed prime power q. Finally, in Section 4.4, we construct pseudorandom functions that fool $\mathsf{AC}^0[2]$ oracle circuits.

4.1 Fourier Bounds from Correlation with XOR of Majorities

We show that our correlation bounds with XOR of majorities imply certain Fourier bounds in a generic way; in particular this resolves a conjectured Fourier bound for polynomials [CGL+21]. This has consequences for constructing PRGs for polynomials that we discuss in Section 4.2.

For any Boolean function $f\colon \{0,1\}^n \to \{0,1\}$, define its sign representation $\widetilde{f}\colon \{-1,1\}^n \to \{-1,1\}$ by $\widetilde{f}(x) = (-1)^{f((1-x_1)/2,\ldots,(1-x_n)/2)}$. We write $\widehat{f}(S)$ for the real Fourier coefficients of $\widetilde{f}$, so that

$$\widetilde{f}(x) = \sum_{S \subseteq [n]} \widehat{f}(S) x_S, \qquad \widehat{f}(S) = \mathop{\mathbb{E}}_{x}[\widetilde{f}(x) x_S],$$

where $x_S = \prod_{i \in S} x_i$ and x is uniform on $\{-1,1\}^n$.

Define

$$M_k(f) = \max_{x \in \{-1,1\}^n} \left| \sum_{|S|=k} \widehat{f}(S) x_S \right|.$$

For any class $\mathcal{F}$ of Boolean functions $f\colon \{0,1\}^n \to \{0,1\}$, define

$$M_k(\mathcal{F}) = \max_{f \in \mathcal{F}} M_k(f).$$

Let $\mathsf{Maj}_\ell^{\oplus k}$ denote the XOR of Majority on k disjoint input blocks of ℓ bits, for odd ℓ. Our main result shows that M_k bounds follow from correlation bounds with $\mathsf{Maj}_\ell^{\oplus k}$. Call a class of Boolean functions *nice* if it is closed under input negations, restrictions, permuting input variables, and adding dummy variables.

Theorem 4.1. *Let $\mathcal{F}$ be a class of Boolean functions that is* nice. *Suppose that for any $k \in \{1, \dots, n\}$, any odd $\ell \geq 1$, and any $f \in \mathcal{F}$ on $k\ell$ inputs, we have $\operatorname{corr}(f, \mathsf{Maj}_\ell^{\oplus k}) \leq (b/\sqrt{\ell})^k$.*

Then, for all $1 \leq k \leq n$,

$$M_k(\mathcal{F}) \leq (4b)^k.$$

Before giving the proof, we observe that combining this result with Theorem 1 immediately yields the following Fourier bounds for polynomials, resolving the conjecture of [CGL+21].

Corollary 4.2. *For any field $\mathbb{F}$, let $\mathcal{B}_{n,d}^{\mathbb{F}}$ be the class of Boolean functions on $\{0,1\}^n$ represented by polynomials over $\mathbb{F}$ of degree at most d. Then, for $1 \leq k \leq n$,*

$$M_k(\mathcal{B}_{n,d}^{\mathbb{F}}) \leq (16d)^k.$$

Proof. Since Majority is versatile over every field $\mathbb{F}$, Corollary 3.4 gives correlation at most $(4d/\sqrt{\ell})^k$ with $\mathsf{Maj}_\ell^{\oplus k}$. Apply Theorem 4.1 with $b = 4d$. □

Proof of Theorem 4.1. Note that negating input coordinates multiplies $\widehat{f}(S)$ by the corresponding product of signs. Thus it suffices to bound $|\sum_{|S|=k} \widehat{f}(S)|$, for an arbitrary $f \in \mathcal{F}$.

Consider a uniformly random partition $B = (B_1, \dots, B_k)$ of $[n]$ (allowing empty parts). We will show that

$$\left| \sum_{S:\, |S \cap B_i| = 1 \ \forall i \in [k]} \widehat{f}(S) \right| \leq \left(\frac{4b}{3} \right)^k \tag{7}$$

for each choice of B. Using the inequality $k^k/k! \leq e^k$, this would complete the proof as

$$\mathop{\mathbb{E}}_{B} \left[\left| \sum_{S:\, |S \cap B_i| = 1 \ \forall i \in [k]} \widehat{f}(S) \right| \right] \geq \left| \mathop{\mathbb{E}}_{B} \left[\sum_{S:\, |S \cap B_i| = 1 \ \forall i \in [k]} \widehat{f}(S) \right] \right| = \left| \sum_{S:|S|=k} \widehat{f}(S) \cdot \frac{k!}{k^k} \right|.$$

Now, fix a partition $B_1, \dots, B_k$ as above. It remains to prove Equation (7). Expanding the product gives

$$\sum_{S:|S \cap B_i| = 1 \ \forall i \in [k]} \widehat{f}(S) = \mathop{\mathbb{E}}_{x} \left[\widetilde{f}(x) \prod_{i=1}^{k} \left(\sum_{j \in B_i} x_j \right) \right].$$

To apply our correlation bound, we will approximate each linear factor by a majority padded with many fresh variables.

Let $\ell = 2m + 1 > n$ be odd and pad each part with fresh variables to length ℓ, and let f ignore these additional variables. Think of n as fixed and let ℓ grow to infinity. On $\{\pm 1\}$, majority is the function $h_\ell(y) = \operatorname{sgn}(y_1 + \cdots + y_\ell)$ and XOR is product. Our assumed correlation bound implies that

$$\left| \mathop{\mathbb{E}}_{x} \left[\widetilde{f}(x) \prod_{i=1}^{k} h_\ell(x^{(i)}) \right] \right| \leq \left(\frac{b}{\sqrt{\ell}} \right)^k, \tag{8}$$

where $x^{(i)}$ contains the variables of B_i and its padding.

We recover the sum of coefficients that use one variable from each part by letting ℓ grow. We prove this as follows. The Fourier coefficients $\widehat{h_\ell}(T)$ of the majority function vanish on even $|T|$, while for odd $|T| = 2j+1$,

$$\widehat{h}_\ell(T) = (-1)^j \frac{\binom{2m}{m}}{2^{2m}} \cdot \frac{\binom{m}{j}}{\binom{2m}{2j}},$$

see, for example, [O'D14, Section 5.3] for a proof.

Expanding in the Fourier basis and dividing by a_ℓ^k, where $a_\ell = \binom{2m}{m} 2^{-2m}$ is the Fourier coefficient of the majority function on sets of size 1, gives

$$\frac{1}{a_\ell^k} \cdot \mathop{\mathbb{E}}_x \left[\widetilde{f}(x) \prod_{i=1}^{k} h_\ell(x^{(i)}) \right] = \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i=1}^{k} \frac{\widehat{h}_\ell(S \cap B_i)}{a_\ell}.$$

If S misses a part, the corresponding factor is zero because majority has mean zero. Otherwise, each factor is 1 when $|S \cap B_i| = 1$, and tends to zero when $1 < |S \cap B_i| \le n$. Since the sum has only 2^n terms, taking $\ell \to \infty$ leaves exactly those sets containing one variable from each part. By Equation (8), we obtain

$$\left| \sum_{S:\, |S \cap B_i| = 1 \ \forall i \in [k]} \widehat{f}(S) \right| \le \lim_{\ell \to \infty} \left(\frac{b}{a_\ell \sqrt{\ell}} \right)^k = \left(b \sqrt{\frac{\pi}{2}} \right)^k \le \left(\frac{4b}{3} \right)^k,$$

where we used $a_\ell \sqrt{\ell} \to \sqrt{2/\pi}$ as $\ell \to \infty$. □

4.2 PRG for Polynomials

We use the PRG from [CGL+21] with the refinement from [HH24, Theorem 5.10]. If a family of Boolean functions $\mathcal{F}$ on n variables is closed under restrictions and negations of the input variables, and satisfies $M_k(\mathcal{F}) \le b^k$ for some $b \ge 1$ and $k = \Theta(\log(b \log n/\varepsilon))$, then it has a PRG of error ε with seed length $\widetilde{O}(b^2 \log^2 n \log(1/\varepsilon))$. Our Fourier bounds proved above imply that we can set $b = 16d$ for Boolean-valued degree-d polynomials over any field $\mathbb{F}$.

Theorem 4.3. *For every field $\mathbb{F}$, there is an explicit PRG that ε-fools Boolean-valued degree-d polynomials over $\mathbb{F}$ on n Boolean inputs with seed length $\widetilde{O}(d^2 \log^2 n \log(1/\varepsilon))$.*

Remark. As mentioned in the introduction, over $\mathbb{F}_2$ our construction improves the dependence on d from exponential to polynomial. Over an odd prime field $\mathbb{F}_p$, the construction of [LMS13] gives seed length $O\big(\log n \cdot \exp^{(2d+1)}((d/\varepsilon)^C)\big)$, where $\exp^{(t)}$ denotes the t-fold iterated exponential and $C > 0$ depends only on p.

Follow-up work. For degree-d polynomials over $\mathbb{F}_2$, Cohen [Coh26] refines the analysis of the PRG in [CGL+21] by combining our new M_k bounds with known $L_{1,k}$ bounds ([CHHL19], see (9) for its definition). This gives seed lengths $\widetilde{O}(d^3 \log n \log(1/\varepsilon))$ and $O(d^4 \log n \log(1/\varepsilon))$, the latter achieving optimal dependence on n at the cost of a worse dependence on d.

PRG for sparse polynomials. The result above can be used to construct PRGs for Boolean-valued sparse polynomials over any field. A polynomial is *S-sparse* if its multilinear representation has at most S nonzero monomials. Note that such polynomials may have monomials of arbitrary degree. Building on earlier work [LVW93, Vio07, ST18], subsequent constructions over $\mathbb{F}_2$ achieved seed lengths $2^{O(\sqrt{\log S})}\log(1/\varepsilon)$ [ST22] and $2^{O(\sqrt{\log S})} + O(\log^{2.01}(1/\varepsilon))$ [Kum25] for $S \geq n$. We reduce the sparsity dependence to polylogarithmic.

Corollary 4.4. *For every field $\mathbb{F}$, there is an explicit PRG that ε-fools all Boolean-valued S-sparse polynomials over $\mathbb{F}$ on n Boolean inputs, with seed length $\widetilde{O}\left(\log^2(S/\varepsilon)\log^2 n \log(1/\varepsilon)\right)$.*

Proof. Set $t = \lceil \log_{4/3}(6S/\varepsilon) \rceil$. Let Y be the output of the degree-t PRG from Theorem 4.3, with error $\varepsilon/6$. Independently sample Z from a t-wise independent distribution on $\{0,1\}^n$ with uniform marginals, using $O(t\log n)$ random bits, and put $D = Y \oplus Z$.

Conditioning on Z, bit flips are affine substitutions over $\mathbb{F}$, so D still $(\varepsilon/6)$-fools Boolean-valued degree-t polynomials. Conditioning on Y shows that D is t-wise independent. Take three independent copies $D^{(1)}, D^{(2)}, D^{(3)}$ and output a distribution X obtained by taking their coordinatewise majority: $X_i = \mathsf{Maj}_3(D_i^{(1)}, D_i^{(2)}, D_i^{(3)})$,[4] where Maj_3 denotes the majority function on 3 bits.

Fix a Boolean-valued S-sparse polynomial P. We compare X with the coordinatewise majority of three independent uniform strings by replacing the three copies of D one at a time. At each replacement, the other two strings are independent and t-wise uniform. Condition on their values a, b. We claim that $P(\mathsf{Maj}_3(x, a, b))$ is a restriction of P: coordinate i is set to a_i if $a_i = b_i$, and remains x_i otherwise. In particular, this restriction remains Boolean-valued.

A monomial of degree greater than t is set to zero if $(a_i, b_i) = (0, 0)$ for any coordinate in its support. Considering any t such coordinates, the probability that the monomial is not set to zero is at most $(3/4)^t$. A union bound therefore shows that, with probability at least $1 - (3/4)^t \cdot S$ over a, b, the restricted polynomial has degree at most t. Consequently, each replacement changes the acceptance probability by at most $\varepsilon/6 + S(3/4)^t$.

The coordinatewise majority of three independent uniform strings is uniform. Summing the errors of the three replacements gives

$$|\Pr[P(X) = 1] - \Pr[P(U_n) = 1]| \leq 3\left(\varepsilon/6 + S(3/4)^t\right) \leq \varepsilon.$$

Substituting $t = O(\log(S/\varepsilon))$ into Theorem 4.3, adding the $O(t\log n)$ random bits for each copy of Z, and taking three independent copies gives the claimed seed length. □

4.3 PRG for $\mathsf{AC}^0[q]$

For a fixed prime power q, the class $\mathsf{AC}^0[q]$ consists of Boolean functions computed by families of polynomial-size, constant-depth circuits with unbounded-fan-in AND, OR, NOT gates and MOD_q gates. A MOD_q gate outputs 0 if the sum of its input bits is congruent to 0 modulo q, and 1 otherwise.

We derive the following PRGs for $\mathsf{AC}^0[q]$, for every fixed prime power q, from our results on fooling Boolean-valued polynomials.

[4] Majority is convenient because it is balanced and, for any choice of fixing two inputs uniformly, makes it constant 0 with probability 1/4. Other choices also work, such as the selector $g(a, b, c) = ab + (1 - a)c$; its induced restrictions may also negate inputs, which preserves degree and Boolean-valuedness.

Theorem 4.5. *Fix a prime power $q = p^r$ and a depth $\Delta \geq 1$. There is an explicit PRG that ε-fools size-s, depth-Δ $\mathsf{AC}^0[q]$ circuits on n inputs with seed length $\widetilde{O}(\log^{2\Delta}(s) \cdot \log^3(1/\varepsilon))$.*

Proof. Fix such a circuit C. Each MOD_q gate has an exact polynomial of degree at most $q-1$ over $\mathbb{F}_p$. Indeed, if e_t denotes the t-th elementary symmetric polynomial in the gate's input bits, then

$$\mathrm{MOD}_{p^r}(z) = 1 - \prod_{j=0}^{r-1} \left(1 - e_{p^j}(z)^{p-1}\right).$$

By Lucas's theorem [Luc78], $e_{p^j}(z) = \binom{|z|}{p^j} \bmod p$ is the j-th base-p digit of $|z|$, and the polynomial on the right-hand side has degree at most $(p-1)\sum_{j=0}^{r-1} p^j = q-1$. Thus the polynomial approximation method of Razborov and Smolensky [Raz87, Smo87], with the same error-reduction argument as in [KS18], gives a distribution of polynomials P over $\mathbb{F}_p$ of degree at most $d = O_{q,\Delta}(\log^{\Delta-1}(s)\log(1/\varepsilon))$ such that, for every input x, $\Pr_P[P(x) \neq C(x)] \leq \varepsilon/3$. Replace each P by $Q = P^{p-1}$. Then Q is Boolean-valued, has degree at most $(p-1)d$, and agrees with C whenever P does. By Theorem 4.3, there is a PRG that $(\varepsilon/3)$-fools all such Q with seed length $\widetilde{O}(d^2 \log^2 n \log(1/\varepsilon))$. Averaging over Q and applying the triangle inequality, its error for C is at most $\varepsilon/3$ for each of the two approximation errors and $\varepsilon/3$ for fooling Q, hence at most ε. □

4.4 Pseudorandom Functions

Pseudorandom functions are distributions on functions that can be generated from a short seed yet cannot be distinguished from a truly random function by efficient circuits equipped with oracle gates that provide evaluations of the function. They were introduced by Goldreich, Goldwasser, and Micali [GGM86], who constructed them from cryptographic assumptions. We construct unconditional pseudorandom functions (PRFs) with polylogarithmic key lengths that fool $\mathsf{AC}^0[2]$ oracle circuits. To our knowledge, such a construction was not known even for ordinary AC^0 oracle circuits. The construction critically relies on our PRG for circuits with parity gates.

We equip an $\mathsf{AC}^0[2]$ circuit C with oracle gates that evaluate a function $f\colon \{0,1\}^n \to \{0,1\}$. We write C^f for the output of the circuit. Oracle gates count toward both size and depth; their inputs may depend on earlier oracle answers.

Theorem 4.6. *There is an explicit family of functions $f_k\colon \{0,1\}^n \to \{0,1\}$, indexed by a key $k \in \{0,1\}^r$ with length $r = \log^{O(\Delta)}(ns/\varepsilon)$, such that for every size-s, depth-Δ $\mathsf{AC}^0[2]$ oracle circuit C we have*

$$\left|\Pr_k[C^{f_k} = 1] - \Pr_R[C^R = 1]\right| \leq \varepsilon,$$

where k is uniformly random and $R\colon \{0,1\}^n \to \{0,1\}$ is a uniformly random function. Given a key k and an input $x \in \{0,1\}^n$, the value $f_k(x)$ is computable in time polynomial in $n + r$.

The idea is to use Theorem 4.5 to generate a pseudorandom table of polynomially many bits. The pseudorandom function then hashes each query to a position in this table and returns the corresponding bit. The analysis proceeds as follows. With a uniform table, the answers are those of a random function unless two distinct queries collide. For a fixed hash,

replacing the oracle calls with hashing and table lookups gives an ordinary $\mathsf{AC}^0[2]$ circuit on the table bits, so our PRG fools it. Note that hashing introduces parity gates in this circuit; thus the construction uses the $\mathsf{AC}^0[2]$ PRG even for ordinary AC^0 oracle distinguishers.

Construction. Let A be a sufficiently large absolute constant, chosen below, and let Q be the smallest power of two $\geq Ans^2/\varepsilon$. Identify $\mathbb{F}_Q$ with $\{0,1\}^{\log Q}$ using a fixed $\mathbb{F}_2$-basis of $\mathbb{F}_Q$, and index Q-bit tables by $\mathbb{F}_Q$. Take the generator $G\colon \{0,1\}^u \to \{0,1\}^Q$ from Theorem 4.5, that $\varepsilon/2$-fools circuits of size at most $As(Q+n)\log Q$, and depth at most $A\Delta$, which has seed length $u = \log^{O(\Delta)}(ns/\varepsilon)$. Choose $a \in \mathbb{F}_Q$ and $z \in \{0,1\}^u$ independently and uniformly at random. Set

$$H_a(x) \coloneqq \sum_{i=1}^{n} x_i a^{i-1}, \qquad f_{a,z}(x) \coloneqq G(z)_{H_a(x)}.$$

The key is $k = (a, z)$, so its length is $r = \log Q + u = \log^{O(\Delta)}(ns/\varepsilon)$.

Proof. We write $f_{a,z}$ as $G(z) \circ H_a$. Our goal is to show that any size-s, depth-Δ circuit cannot distinguish R from $G(z) \circ H_a$. To this end, we introduce a hybrid $T \circ H_a$, where T is an independent uniform Q-bit table, and show that no such circuit can distinguish R from $T \circ H_a$ or $T \circ H_a$ from $G(z) \circ H_a$, except with small advantage.

Random function versus hashed uniform table. Fix an execution of an oracle circuit, and consider any transcript τ of m query-answer pairs that records only the first occurrence of each distinct query. We have $m \leq s$ as the circuit can only make at most s oracle queries. Note that this transcript determines the queries themselves, as well as the output of the circuit. The transcript occurs with probability 2^{-m} under R. Under the hashed uniform table, the transcript occurs with probability at least 2^{-m} multiplied by the probability π that no two distinct queries collide under the hash function H_a. To lower bound π, note that for distinct x, y, the polynomial $H_a(x) - H_a(y)$ is nonzero and has degree at most $n-1$, so its collision probability is at most $(n-1)/Q$. By a union bound we have

$$\pi \geq 1 - \binom{s}{2}\frac{n-1}{Q} \geq 1 - \frac{\varepsilon}{A}.$$

Summing over the transcripts that make the circuit output 1, and similarly over the transcripts that make the circuit output 0, shows that the distinguishing advantage between the two distributions on functions is at most ε/A.

Uniform table versus PRG table. Fix a. Since H_a is $\mathbb{F}_2$-linear, every coordinate of $H_a(x)$ with respect to the fixed $\mathbb{F}_2$-basis is a parity of the input bits. Replace each oracle gate by this parity layer followed by a table look-up. Each replacement has constant depth and size $O((Q+n)\log Q)$. The resulting circuit on the table bits belongs to $\mathsf{AC}^0[2]$, has size $O(s(Q+n)\log Q)$, and depth $O(\Delta)$. Choose A large enough so that the circuit is covered by the PRG parameters and the collision error is at most $\varepsilon/2$. By pseudorandomness of the PRG, the distinguishing advantage between $T \circ H_a$ and $G(z) \circ H_a$ is at most $\varepsilon/2$. Averaging over all a and applying the triangle inequality completes the pseudorandomness guarantee.

It remains to show that $f_{a,z}(x)$ can be computed efficiently. To compute $f_{a,z}(x)$, we compute $H_a(x)$ and then one specified output bit of $G(z)$. We use the following implementation of the polarizing-walk generator [CHHL19, CGL+21]. (Alternatively, a direct implementation of Nisan's generator [Nis91] gives the desired efficiency.) Its seed specifies N independent samples $Z^{(1)}, \dots, Z^{(N)} \in \{-1, 1\}^Q$, each drawn from a κ-wise independent distribution. To compute output coordinate i, start with $y_0 = 0$ and apply

$$y_j = y_{j-1} + (1 - |y_{j-1}|) c Z_i^{(j)} \qquad (j = 1, \dots, N),$$

where $c = 2^{-\lambda}$ for a positive integer λ. The output bit is $(1 - \mathrm{sign}(y_N))/2$, with $\mathrm{sign}(0) = 1$. For our parameters, the cited analysis allows choices with $N\kappa \log Q = \log^{O(\Delta)}(ns/\varepsilon)$ and $\lambda = O(\log N)$.[5] Standard bounded-independence constructions compute any specified $Z_i^{(j)}$ from its seed and i in time polynomial in κ and $\log Q$, using finite-field polynomial evaluation (see, e.g., the proof of [HH24, Theorem 2.2]). We can therefore compute $G(z)_i$ by following the above recurrence relation. Every intermediate value has denominator dividing $2^{\lambda N}$, so the computation uses exact arithmetic on $O(\lambda N)$-bit integers and takes time polynomial in the seed length. Together with evaluating $H_a(x)$ using $O(n)$ operations in $\mathbb{F}_Q$, this gives evaluation time polynomial in $n + r$. □

5 Open Problems

Our results leave several natural questions.

Better correlation bounds. Our construction gives correlation at most $2^{-\Omega(n/d^2)}$ with degree-d polynomials over $\mathbb{F}_2$. A first target towards improving this would be to obtain an explicit function with correlation $o(1)$ against every degree-d polynomial for some $d = \omega(\sqrt{n})$. More generally, can we obtain correlation $2^{-\Omega(n/d)}$ for an explicit function when $d = o(n)$? Since our bound for the XOR of majorities is tight up to constants, this would require a different hard function. An even more ambitious bound would imply new general Boolean circuit lower bounds: for some constant $\delta > 0$, an explicit function with correlation at most $2^{-\omega(n/\log\log n)}$ against all degree-n^δ polynomials would require superlinear-size bounded-fan-in circuits of depth $O(\log n)$ [Val77, GKW21].

Improved PRGs. Theorem 4.3 gives seed length $d^2 \cdot \mathsf{polylog}(n/\varepsilon)$ for degree-d polynomials. A natural goal would be to reduce the dependence on d from quadratic to linear. We note that non-constructively it is possible to achieve seed length $O(d \log n + \log(1/\varepsilon))$, which is tight up to constants [BV10, ABEK08].

Stronger Fourier bounds. We bound $M_k(f)$, which takes the absolute value after summing the signed level-k Fourier coefficients and thus allows cancellation. It remains open whether

[5] For this application, we do not need the randomized-rounding refinement of [HH24, Theorem 5.10] used to obtain the parameters in Theorem 4.3. Since $Q \geq 1/\varepsilon$, we have $\log(Q/\varepsilon) = O(\log Q)$, so the original walk with deterministic sign rounding already achieves the claimed polylogarithmic key length.

the sum of their absolute values is also bounded by $O(d)^k$ for every degree-d polynomial. For a Boolean function $f : \{0,1\}^n \to \{0,1\}$, define

$$L_{1,k}(f) \coloneqq \sum_{|S|=k} \left|\widehat{f}(S)\right|. \tag{9}$$

For any $f \in \mathcal{B}_{n,d}^{\mathbb{F}_2}$, the bound $L_{1,k}(f) \leq (Cd)^k$, for an absolute constant C, was conjectured for $k = 2$ in [CHLT19] and for all k in [BIJ+21]. The conjecture was proved for symmetric polynomials in [BIJ+21]. For general polynomials, the best known bound is $L_{1,k}(f) \leq (k2^{3d})^k$ [CHHL19]. A confirmation of this conjecture would yield a PRG for degree-d polynomials over $\mathbb{F}_2$ with seed length $\widetilde{O}(d^2 \log n \log(1/\varepsilon))$, saving a factor of $\log n$ over the bound in Theorem 4.3.

References

[ABEK08] Noga Alon, Ido Ben-Eliezer, and Michael Krivelevich. Small sample spaces cannot fool low degree polynomials. In *12th Workshop on Randomization and Computation (RANDOM)*, pages 266–275. Springer, 2008. (cit. on page 22.)

[BIJ+21] Jaroslaw Blasiok, Peter Ivanov, Yaonan Jin, Chin Ho Lee, Rocco A. Servedio, and Emanuele Viola. Fourier growth of structured $\mathbb{F}_2$-polynomials and applications. In *Workshop on Randomization and Computation (RANDOM)*, 2021. Invited to special issue. (cit. on page 23.)

[Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *ACM Symp. on the Theory of Computing (STOC)*, pages 21–30, 2005. (cit. on page 3.)

[Bou05] Jean Bourgain. Estimation of certain exponential sums arising in complexity theory. *Comptes Rendus Mathématique. Académie des Sciences. Paris*, 340(9):627–631, 2005. (cit. on page 10.)

[BV10] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM J. on Computing*, 39(6):2464–2486, 2010. (cit. on page 22.)

[BVdW07] Harry Buhrman, Nikolay Vereshchagin, and Ronald de Wolf. On computation and communication with small bias. In *Twenty-Second Annual IEEE Conference on Computational Complexity (CCC'07)*, pages 24–32. IEEE, 2007. (cit. on page 2.)

[CGL+21] Eshan Chattopadhyay, Jason Gaitonde, Chin Ho Lee, Shachar Lovett, and Abhishek Shetty. Fractional pseudorandom generators from any Fourier level. In *36th Computational Complexity Conference (CCC 2021)*, volume 200 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:24. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021. (cit. on pages 3, 16, 17, 18, and 22.)

[Cha07] Arkadev Chattopadhyay. Discrepancy and the power of bottom fan-in in depth-three circuits. In *48th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 449–458. IEEE, 2007. (cit. on page 10.)

[CHH+20] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, Shachar Lovett, and David Zuckerman. XOR lemmas for resilient functions against polynomials. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 234–246, 2020. (cit. on page 2.)

[CHHL19] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. Pseudorandom generators from polarizing random walks. *Theory Comput.*, 15:1–26, 2019. (cit. on pages 3, 18, 22, and 23.)

[CHLT19] Eshan Chattopadhyay, Pooya Hatami, Shachar Lovett, and Avishay Tal. Pseudorandom Generators from the Second Fourier Level and Applications to AC0 with Parity Gates. In *ACM Innovations in Theoretical Computer Science conf. (ITCS)*, pages 22:1–22:15, 2019. (cit. on page 23.)

[Coh26] Gil Cohen. Fooling low-degree polynomials: PRGs with optimal dependence on input length and expander walks. 2026. ECCC TR26-243. (cit. on page 18.)

[FSUV13] Bill Fefferman, Ronen Shaltiel, Christopher Umans, and Emanuele Viola. On beating the hybrid argument. *Theory of Computing*, 9:809–843, 2013. (cit. on page 3.)

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *J. of the ACM*, 33(4):792–807, October 1986. (cit. on page 20.)

[GKW21] Alexander Golovnev, Alexander S. Kulikov, and R. Ryan Williams. Circuit Depth Reductions. In James R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*, volume 185 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. (cit. on page 22.)

[GRS05] Frederic Green, Amitabha Roy, and Howard Straubing. Bounds on an exponential sum arising in Boolean circuit complexity. *Comptes Rendus Mathématique. Académie des Sciences. Paris*, 341(5):279–282, 2005. (cit. on page 10.)

[HH24] Pooya Hatami and William Hoza. Paradigms for unconditional pseudorandom generators. *Foundations and Trends in Theoretical Computer Science*, 16(1–2):1–210, 2024. (cit. on pages 18 and 22.)

[HMP+93] András Hajnal, Wolfgang Maass, Pavel Pudlák, Mário Szegedy, and György Turán. Threshold circuits of bounded depth. *J. of Computer and System Sciences*, 46(2):129–154, 1993. (cit. on page 2.)

[Kop11] Swastik Kopparty. On the complexity of powering in finite fields. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*, pages 489–498, 2011. (cit. on page 7.)

[KS18] Swastik Kopparty and Srikanth Srinivasan. Certifying polynomials for $\mathrm{AC}^0[\oplus]$ circuits, with applications to lower bounds and circuit compression. *Theory of Computing*, 14(1):1–24, 2018. (cit. on pages 2 and 20.)

[Kum25] Vinayak M. Kumar. New pseudorandom generators and correlation bounds using extractors. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, volume 325 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 68:1–68:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. (cit. on pages 3 and 19.)

[LMS13] Shachar Lovett, Partha Mukhopadhyay, and Amir Shpilka. Pseudorandom generators for $CC^0[p]$ and the Fourier spectrum of low-degree polynomials over finite fields. *Computational Complexity*, 22(4):679–725, 2013. (cit. on page 18.)

[Lov09] Shachar Lovett. Unconditional pseudorandom generators for low degree polynomials. *Theory of Computing*, 5(1):69–82, 2009. (cit. on page 3.)

[Luc78] Edouard Lucas. Théorie des fonctions numériques simplement périodiques. *American Journal of Mathematics*, pages 289–321, 1878. (cit. on page 20.)

[LVW93] Michael Luby, Boban Veličković, and Avi Wigderson. Deterministic approximate counting of depth-2 circuits. In *2nd Israeli Symposium on Theoretical Computer Science (ISTCS)*, pages 18–24, 1993. (cit. on pages 3 and 19.)

[Nis91] Noam Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica. A Journal on Combinatorics and the Theory of Computing*, 11(1):63–70, 1991. (cit. on page 22.)

[NW94] Noam Nisan and Avi Wigderson. Hardness vs Randomness. *Journal of Computer and System Sciences*, 49(2):149–167, 1994. (cit. on page 2.)

[O'D14] Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. (cit. on page 18.)

[Raz87] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 1987. English translation in Mathematical Notes of the Academy of Sci. of the USSR, 41(4):333-338, 1987. (cit. on pages 2 and 20.)

[Smo87] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *19th ACM Symp. on the Theory of Computing (STOC)*, pages 77–82. ACM, 1987. (cit. on pages 2, 4, 10, and 20.)

[Smo93] Roman Smolensky. On representations by low-degree polynomials. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 130–138, 1993. (cit. on pages 2, 4, and 10.)

[ST18] Rocco A. Servedio and Li-Yang Tan. Luby-Velickovic-Wigderson revisited: Improved correlation bounds and pseudorandom generators for depth-two circuits. In *Workshop on Randomization and Computation (RANDOM)*, volume 116 of *LIPIcs*, pages 56:1–56:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. (cit. on pages 3 and 19.)

[ST22] Rocco A. Servedio and Li-Yang Tan. Improved pseudorandom generators from pseudorandom multi-switching lemmas. *Theory of Computing*, 18(4):1–46, 2022. (cit. on pages 3 and 19.)

[Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *6th Symposium on Mathematical Foundations of Computer Science*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977. (cit. on page 22.)

[Vio] Emanuele Viola. *Mathematics of the impossible*. (cit. on page 1.)

[Vio06] Emanuele Viola. New correlation bounds for GF(2) polynomials using Gowers uniformity. *Electronic Colloquium on Computational Complexity*, Technical Report TR06-097, 2006. www.eccc.uni-trier.de/. (cit. on page 2.)

[Vio07] Emanuele Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM J. on Computing*, 36(5):1387–1403, 2007. (cit. on pages 3 and 19.)

[Vio09] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. *Computational Complexity*, 18(2):209–217, 2009. (cit. on page 3.)

[Vio21] Emanuele Viola. Fourier conjectures, correlation bounds, and majority. In *Coll. on Automata, Languages and Programming (ICALP)*, 2021. Available at http://www.ccs.neu.edu/home/viola/. (cit. on page 2.)

[Vio22] Emanuele Viola. Correlation bounds against polynomials, a survey. 2022. (cit. on page 1.)

[VW08] Emanuele Viola and Avi Wigderson. Norms, XOR lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4(7):137–168, 2008. (cit. on page 2.)

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il