Low-Norm Nullstellensatz Hypothesis for the AND Code is False

Zhengzhong Jin
Northeastern University

Abstract

The recent work [Devadas–Hopkins–Kalai–Kothari–Lombardi–Mathialagan, STOC 2026] proposed a low-norm Nullstellensatz hypothesis for the "AND code": every polynomial f vanishing on the "AND-code ideal" should admit a Nullstellensatz decomposition over the local AND constraints whose total coefficient ℓ_1-norm is only polynomially larger than the ℓ_1-norm of f.

We give a counterexample to this conjecture by proving an exponential lower bound on the total coefficient ℓ_1-norm. The core idea of the proof was discovered by ChatGPT 5.5 Pro, and we verified and reorganized the proof to improve its exposition. The proof constructs a dual linear functional, whose analysis leverages the rank of the quadratic forms to bound Fourier correlations.

The counterexample can also be extended to give the first ℓ_1-norm lower bound for Nullstellensatz refutations over the $\{\pm 1\}$-basis. Previously, ℓ_1-norm lower bounds for Nullstellensatz refutations were known only over the $\{0,1\}$-basis, due to Potechin and Zhang [ICALP 2024]. We believe this is of independent interest to proof complexity.

1 Introduction

Constructing succinct non-interactive arguments (SNARGs) for NP from simple, standard cryptographic assumptions is one of the major open problems in cryptography. Recent work of [DHK+26] proposed a new route toward this goal: they conjectured that the "AND-code ideal" admits low-ℓ_1-norm Nullstellensatz decompositions, and showed that this conjecture would imply SNARGs for NP from standard assumptions such as Learning with Errors or k-Lin assumption on bilinear maps. We next recall the conjecture.

Notation. For an integer n, we work over the Boolean quotient ring

$$\mathcal{R}_n := \mathbb{R}[X_1, \ldots, X_n, (Y_{ij})_{1 \le i < j \le n}] / \left\langle X_i^2 - 1,\ Y_{ij}^2 - 1 \right\rangle .$$

Every element of $\mathcal{R}_n$ has a unique multilinear expansion in the monomial basis

$$X^\alpha Y_S := \prod_{i=1}^{n} X_i^{\alpha_i} \prod_{(i,j) \in S} Y_{ij}, \qquad \alpha \in \mathbb{F}_2^n,\ S \subseteq \{(i,j) : 1 \le i < j \le n\}.$$

The coefficient ℓ_1-norm $\|f\|_1$ is the sum of the absolute values of the coefficients in this basis. This ℓ_1-norm complexity measure is relatively new and was first studied in [PZ24], which proved an exponential ℓ_1-norm lower bound in the "$\{0,1\}$-basis", i.e., over the Boolean quotient by the relations $X_i(X_i - 1) = 0$. As observed in [DHK+26], such a lower bound does not appear to directly imply an analogous lower bound in the $\{\pm 1\}$-basis, which is the basis in our setting.

For $X_i, X_j \in \{\pm 1\}$, define the $\{\pm 1\}$-valued AND function by

$$X_i \wedge X_j := \frac{1 + X_i + X_j - X_i X_j}{2}.$$

The AND code is

$$C_{\mathrm{AND},n} := \left\{ (X, Y) \in \{\pm 1\}^{n + \binom{n}{2}} : Y_{ij} = X_i \wedge X_j \text{ for all } i < j \right\}.$$

The generators used in the conjecture are

$$p_{ij}(X, Y) := \frac{1 - Y_{ij}(X_i \wedge X_j)}{2}, \qquad 1 \le i < j \le n.$$

They vanish exactly when $Y_{ij} = X_i \wedge X_j$.

Low-norm Nullstellensatz Conjecture for the AND Code [DHK+26]. There is a fixed polynomial $\mathrm{poly}(n)$ such that every $f \in \mathcal{R}_n$ vanishing on $C_{\mathrm{AND},n}$ has a representation

$$f = \sum_{1 \le i < j \le n} q_{ij} p_{ij}$$

with

$$\sum_{1 \le i < j \le n} \|q_{ij}\|_1 \le \mathrm{poly}(n) \, \|f\|_1 .$$

Conjecture 1.3 and 5.2 of [DHK+26] assert precisely this property for the AND code. Section 2.1 of [DHK+26] explains the SNARG motivation for non-signaling PCPs; the formal low-norm Nullstellensatz conjecture appears in Sections 1.2 and 5.2.

We prove the following theorem:

Theorem 1.1 (Main theorem)**.** *There is a constant $c > 0$ and a family of polynomials $\{f_n\}_{n \in \mathbb{N}}$ such that for all sufficiently large $n \in \mathbb{N}$, the following holds.*

1. *f_n vanishes on the AND code $C_{\mathrm{AND},n}$;*
2. *$\|f_n\|_1 = 1$;*
3. *every representation*

$$f_n = \sum_{1 \le i < j \le n} q_{ij} p_{ij}$$

in the Boolean quotient ring satisfies

$$\sum_{i < j} \|q_{ij}\|_1 \ge 2^{cn}.$$

Consequently, the AND-code low-norm Nullstellensatz hypothesis is false.

Concurrent Work. Concurrently and independently, Lombardi [Lom26] gives several counterexamples of the low-norm Nullstellensatz hypothesis for codes other than the AND code. Our work focuses on the low-norm Nullstellensatz for the AND code explicitly conjectured in [DHK+26, Conjecture 1.3 and 5.2].

Implication for [DHK+26]. We emphasize that our result does not preclude the existence of other "Nullstellensatz-friendly codes" as suggested in [DHK+26], nor does it rule out that Hadamard PCP is a non-signaling PCP for NP, nor does it suggest that their SNARG construction has an explicit attack.

The reader may wonder whether, for the SNARG application of [DHK+26], it would suffice to assume the low-norm Nullstellensatz conjecture only for the restricted class of polynomials used in their soundness analysis. We observe that this restriction does not avoid our counterexample. More concretely, we can make our counterexample f_n in the same format as the p^* of Section 7.1 of [DHK+26] by replacing each "high-locality" equation ℓ_i with a constant-local k-Lin "product tree". Then we can recover a sub-exponential lower bound for any "k-Lin induced p^*". The same idea is applied in Section 5, and hence we omit the proof. Whether this lower bound can be strengthened to an exponential one, and whether the k-Lin class defined in [DHK+26] can be further restricted to avoid the counterexample of this work, remains open.

Our Proof Strategy. We design f_n to encode an NP problem: the witness is encoded in $(X_1, \ldots, X_n)$, and the NP verifier checks whether the following quadratic equation system is satisfied:

$$\ell_1(\mathbf{X}) = b_1, \ldots, \ell_{n+1}(\mathbf{X}) = b_{n+1},$$

where each ℓ_i is a quadratic polynomial and $b_i \in \{\pm 1\}$ is a constant bit. Since there are at most 2^n possible witnesses, there must exist $\mathbf{b} = (b_1, \ldots, b_{n+1})$ such that the quadratic equation system is unsatisfiable. Then f_n vanishes on the AND code. We prove Items 1 and 2 in Section 2.

The lower bound of the norm complexity is proved by a dual linear functional $\mathcal{L} : \mathcal{R}_n \to \mathbb{R}$. It satisfies $\mathcal{L}(f_n) = 1$ and $|\mathcal{L}(Mp_{ij})| = 2^{-\Omega(n)}$ for every monomial M and every (i, j). Applying $\mathcal{L}$ to any Nullstellensatz decomposition of f_n therefore forces the total coefficient norm to be exponentially large. We construct $\mathcal{L}$ and prove its properties in Section 3. Then we prove Item 3 in Section 4.

The main technical challenge in the proof is constructing $\mathcal{L}$ and proving that $|\mathcal{L}(Mp_{ij})|$ is small (Lemma 3.3, 3.4, 3.5). The key observation is a "Gauss-sum Lemma" (Lemma 3.4), which is a Fourier-correlation bound for quadratic phases using the rank of the quadratic form; the construction of $\mathcal{L}$ carefully manages ranks so that this estimate applies to every Mp_{ij}.

Our main theorem can be extended to the following lower bound on the ℓ_1-norm complexity of Nullstellensatz refutation.

Theorem 1.2 (ℓ_1-norm Complexity Lower Bound for Nullstellensatz Refutation)**.** *For infinitely many n, there is an unsatisfiable system of $m = \mathrm{poly}(n)$ axioms*

$$p_1, p_2, \ldots, p_m \in \mathbb{R}[x_1, \ldots, x_n] / \left\langle x_i^2 - 1 \right\rangle_{i=1}^n$$

such that each p_t is a 3-junta and $\|p_t\|_1 = O(1)$. Moreover, every Nullstellensatz refutation

$$1 = \sum_{t=1}^m q_t p_t$$

in $\mathbb{R}[x_1, \ldots, x_n] / \left\langle x_i^2 - 1 \right\rangle_{i=1}^n$ satisfies

$$\sum_{t=1}^m \|q_t\|_1 \geq 2^{n^{\Omega(1)}}.$$

We prove this extended theorem in Section 5.

2 Construction of f_n

We present the construction in terms of quadratic equations over $\mathbb{F}_2$ (instead of $\{\pm 1\}$-basis). Write

$$X_i = (-1)^{x_i}, \qquad x_i \in \mathbb{F}_2.$$

Then

$$X_i \wedge X_j = (-1)^{x_i x_j}.$$

Let $\mathcal{Q}_n$ be the $\mathbb{F}_2$-vector space of homogeneous quadratic forms in $x_1, \ldots, x_n$, which contains all polynomials in $\mathbb{F}_2[x_1, \ldots, x_n]$ of the following form

$$a = \sum_{1 \le i < j \le n} a_{ij} x_i x_j \in \mathcal{Q}_n, \qquad a_{ij} \in \mathbb{F}_2.$$

For any such $a \in \mathcal{Q}_n$, define the "true quadratic phase":

$$\omega_a(x) := (-1)^{a(x)}.$$

Now we are ready to construct f_n. We will choose an $\mathbb{F}_2$-linear subspace L of $\mathcal{Q}_n$ of dimension $n+1$, and fix a basis of L:

$$\ell_1, \ldots, \ell_{n+1}.$$

Consider the map

$$\Psi : \mathbb{F}_2^n \to \{\pm 1\}^{n+1}, \qquad \Psi(x) = \big(\omega_{\ell_1}(x), \ldots, \omega_{\ell_{n+1}}(x)\big).$$

There are at most 2^n possible inputs x, but there are 2^{n+1} sign vectors in $\{\pm 1\}^{n+1}$. Hence there is a vector

$$\mathbf{b} = (b_1, \ldots, b_{n+1}) \in \{\pm 1\}^{n+1}$$

which is not in $\operatorname{im} \Psi$. This $\mathbf{b}$ is the impossible quadratic-output pattern.

Our f_n is just the arithmetization of the following quadratic equation system

$$\ell_1(x) = b_1 \ \wedge \ \ell_2(x) = b_2 \ \wedge \ \ldots \ \wedge \ \ell_{n+1}(x) = b_{n+1}.$$[1]

However, a direct arithmetization would result in a large ℓ_1-norm, so we first leverage the AND code to "linearize" each ℓ_i.

In general, for any $a = \sum_{i<j} a_{ij} x_i x_j \in \mathcal{Q}_n$, we define its "linearized monomial"

$$Y_a = \prod_{1 \le i < j \le n \,:\, a_{ij} = 1} Y_{ij}.$$

Now we define the polynomial f_n as

$$f_n(Y) := \prod_{i=1}^{n+1} \frac{1 + b_i Y_{\ell_i}}{2}.$$

[1] We slightly abuse the notation in this line by identifying $b_i \in \{\pm 1\}$ with $b_i \in \{0, 1\}$.

Choose L as a "High-Rank" Linear Subspace. Looking ahead, for the dual functional in Section 3, we will apply the Gauss-sum Lemma (Lemma 3.4) to bound the "error" of the functional on the "AND gates". To apply this lemma, we need to choose the linear subspace L such that L only contains "high-rank" quadratic forms (except $0 \in L$). We first formally define rank as follows.

Definition 2.1 (Rank of Quadratic Form)**.** For a quadratic form

$$a = \sum_{1 \le i < j \le n} a_{ij} x_i x_j \in \mathcal{Q}_n,$$

let B_a be the matrix of the associated bilinear form over $\mathbb{F}_2$, defined by

$$u^\top B_a v = a(u+v) + a(u) + a(v) = \sum_{1 \le i < j \le n} a_{ij}(u_i v_j + u_j v_i).$$

Equivalently,

$$B_a = \begin{pmatrix} 0 & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{12} & 0 & a_{23} & \cdots & a_{2n} \\ a_{13} & a_{23} & 0 & \cdots & a_{3n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{1n} & a_{2n} & a_{3n} & \cdots & 0 \end{pmatrix}.$$

We define $\operatorname{rank}(a) := \operatorname{rank}_{\mathbb{F}_2}(B_a)$. In particular, each "AND gate" polynomial $x_i x_j$ has rank at most 2, and rank is subadditive.

We choose the linear subspace L such that *every* (nonzero) quadratic form in L has rank $\Omega(n)$. The existence of such L is given by the following lemma.

Lemma 2.2 (Existence of High-rank Subspaces)**.** *There is an absolute constant $\kappa > 0$ such that, for all sufficiently large n, there is a $\mathbb{F}_2$-linear subspace*

$$L \le \mathcal{Q}_n$$

with $\dim L = n + 1$ *and*

$$\operatorname{rank}(\lambda) \ge \kappa n \qquad \text{for every nonzero } \lambda \in L.$$

Proof. Let $N = \dim \mathcal{Q}_n = \binom{n}{2}$. The number of $n \times n$ matrices over $\mathbb{F}_2$ of rank at most s is at most 2^{2ns}. The map sending a homogeneous quadratic form to the matrix of its associated bilinear form is injective on $\mathcal{Q}_n$, so the number of quadratic forms of rank at most s is at most 2^{2ns}.

Choose a uniformly random $(n+1)$-dimensional subspace $L \le \mathcal{Q}_n$. For any fixed nonzero $q \in \mathcal{Q}_n$,

$$\Pr[q \in L] \le 2^{n+1-N}.$$

Therefore the probability that L contains a nonzero quadratic form of rank at most κn is at most

$$2^{2\kappa n^2} \cdot 2^{n+1-N}.$$

For any sufficiently small fixed κ, for instance $\kappa = 1/16$, this tends to 0 as $n \to \infty$. □

Lemma 2.3. *The polynomial f_n vanishes on $C_{\mathrm{AND},n}$, and $\|f_n\|_1 = 1$.*

Proof. We first prove that f_n vanishes on the AND code. Let $(X, Y) \in C_{\mathrm{AND},n}$, and write $X_i = (-1)^{x_i}$ with $x_i \in \mathbb{F}_2$. For every quadratic form

$$a = \sum_{1 \le i < j \le n} a_{ij} x_i x_j \in \mathcal{Q}_n,$$

the definition of the AND code gives

$$Y_a = \prod_{i<j : a_{ij}=1} Y_{ij} = \prod_{i<j : a_{ij}=1} (X_i \wedge X_j) = \prod_{i<j : a_{ij}=1} (-1)^{x_i x_j} = (-1)^{a(x)} = \omega_a(x).$$

In particular, $Y_{\ell_i} = \omega_{\ell_i}(x)$ for every i. Hence

$$f_n(X, Y) = \prod_{i=1}^{n+1} \frac{1 + b_i \omega_{\ell_i}(x)}{2}.$$

Because $\mathbf{b} \notin \operatorname{im} \Psi$, the vector

$$\big(\omega_{\ell_1}(x), \ldots, \omega_{\ell_{n+1}}(x)\big)$$

is not equal to $\mathbf{b}$. Thus for some i, $\omega_{\ell_i}(x) \neq b_i$, and the corresponding factor $(1 + b_i \omega_{\ell_i}(x))/2$ is zero. Therefore $f_n(X, Y) = 0$ on every point of $C_{\mathrm{AND},n}$.

It remains to compute the coefficient ℓ_1-norm. Expanding the definition of f_n,

$$f_n = 2^{-(n+1)} \sum_{S \subseteq \{1,\ldots,n+1\}} \left(\prod_{i \in S} b_i \right) \prod_{i \in S} Y_{\ell_i}.$$

In the Boolean quotient ring,

$$\prod_{i \in S} Y_{\ell_i} = Y_{\sum_{i \in S} \ell_i},$$

where the sum is over $\mathbb{F}_2$. Since $\ell_1, \ldots, \ell_{n+1}$ is a basis of L, the sums $\sum_{i \in S} \ell_i$ are distinct for different subsets S. Thus the expansion above has 2^{n+1} distinct monomials, each with coefficient of absolute value $2^{-(n+1)}$. Therefore

$$\|f_n\|_1 = 2^{n+1} \cdot 2^{-(n+1)} = 1.$$

□

3 Construction of the Linear Functional $\mathcal{L}$

In this section, we build the dual object used for the lower bound. We prove the following lemma.

Lemma 3.1. *There exists a linear functional $\mathcal{L} : \mathcal{R}_n \to \mathbb{R}$ that satisfies the following properties:*

1. $\mathcal{L}(f_n) = \mathcal{L}(1) = 1$
2. *There is a constant $c > 0$ such that for all sufficiently large n, every polynomial $h \in \mathcal{R}_n$, and every $1 \le i < j \le n$,*

$$|\mathcal{L}(h p_{ij})| \le 2^{-cn} \|h\|_1 .$$

Construction of $\mathcal{L}$. We first describe the intuition of the construction. To ensure $\mathcal{L}(f_n) = 1$, observe that if we can find an $\mathcal{L}$ such that $\mathcal{L}$ vanishes on the ideal $\langle Y_{\ell_i} - b_i \rangle_{i=1}^{n+1} \subset \mathcal{R}_n$ then

$$\mathcal{L}(f_n) = \mathcal{L}\left[\prod_{i=1}^{n+1} \frac{1 + b_i Y_{\ell_i}}{2}\right] = \mathcal{L}(1) = 1.$$

The intuition is that such $\mathcal{L}$ "pseudo-satisfies" the constraint $Y_{\ell_i} = b_i$, so we can prove $\mathcal{L}(f_n) = 1$, since f_n just checks whether $Y_{\ell_i} = b$ for all i.

Next, we construct such $\mathcal{L}$ that vanishes on $\langle Y_{\ell_i} - b_i \rangle$, for every $i \in [n+1]$. Due to linearity, to specify $\mathcal{L}$, it suffices to specify the value of $\mathcal{L}$ on each monomial

$$X^\alpha Y_a, \qquad \alpha \in \mathbb{F}_2^n, a \in \mathcal{Q}_n.$$

Clearly, $\mathcal{L}$ vanishes on $\langle Y_{\ell_i} - b_i \rangle$ if and only if for every *monomial* M in $\mathcal{R}_n$,

$$\mathcal{L}(Y_{\ell_i} M) = b_i \mathcal{L}(M).$$

If $M = X^\alpha Y_a$ and note that $Y_{\ell_i} Y_a = Y_{\ell_i + a}$, then this is equivalent to

$$\mathcal{L}(X^\alpha Y_{\ell_i + a}) = b_i \mathcal{L}(X^\alpha Y_a).$$

In other words, for any fixed α, the behavior of the map

$$a \mapsto \mathcal{L}(X^\alpha Y_a)$$

is fully determined as long as we pick one representative in each coset $\mathcal{Q}_n / L$ and specify the value of $\mathcal{L}$ on them. Next, we are going to do this step by step formally.

Step 1: Extend "$\mathcal{L}(Y_{\ell_i}) = b_i$" Linearly in L. Define the character[2] $\beta : L \to \{\pm 1\}$ by

$$\beta(\ell_i) = b_i$$

and extend multiplicatively over the $\mathbb{F}_2$-basis, so that if $\lambda = \sum_{i \in S} \ell_i$, then

$$\beta(\lambda) = \prod_{i \in S} b_i.$$

Step 2: Pick the Representitives of $\mathcal{Q}_n / L$ Carefully. Choose a constant $0 < \tau < \kappa/4$. For each coset $c \in \mathcal{Q}_n / L$, choose a representative $\sigma(c) \in c$ as follows:

- If c contains a representative of rank at most τn, choose that representative. (In particular, $\sigma(0) = 0$.)
- Otherwise choose an arbitrary representative.

The low-rank representative, when it exists, is unique. Indeed, suppose $q, q' \in c$ both have rank at most τn. Then $q + q' \in L$, and

$$\operatorname{rank}(q + q') \le \operatorname{rank}(q) + \operatorname{rank}(q') \le 2\tau n < \kappa n.$$

The minimum-rank property of L forces $q + q' = 0$, and hence $q = q'$.

[2] "Character" means that it satisfies $\beta(a + b) = \beta(a) \cdot \beta(b)$ for any $a, b \in L$.

Step 3: Define $\mathcal{L}$ "Locally" on Representatives. For a monomial $X^\alpha Y_a$, with $\alpha \in \mathbb{F}_2^n$ and $a \in \mathcal{Q}_n$, we can decompose a as

$$a = (a - \bar{a}) + \bar{a},$$

where $\bar{a} = \sigma(a \bmod L)$ is[3] the representitive we picked for the coset that a lies in, and hence $a - \bar{a} \in L$. Define

$$\mathcal{L}(X^\alpha Y_a) := \mathbb{E}_{x \in \mathbb{F}_2^n} \left[(-1)^{\alpha \cdot x} \cdot \beta\big(a - \bar{a}\big) \cdot \omega_{\bar{a}} \right].$$

Then extend $\mathcal{L}$ linearly to all of $\mathcal{R}_n$. Equivalently, the above $\mathcal{L}$ can be viewed as first defining

$$\mathcal{L}(X^\alpha Y_{\bar{a}}) = \mathbb{E}_{x \in \mathbb{F}_2^n} \left[(-1)^{\alpha \cdot x} \omega_{\bar{a}}(x) \right] = \mathbb{E}_{x \in \mathbb{F}_2^n} \left[(-1)^{\alpha \cdot x + \bar{a}(x)} \right]$$

for each representative $\bar{a}$ we picked in Step 2, and extending it "linearly" using β. Intuitively, $\mathcal{L}$ can be viewed as a "uniform distribution" that is "defined locally" on each of the cosets $\mathcal{Q}/L$. Next, we prove Item 1 of Lemma 3.1.

Lemma 3.2. *The functional satisfies $\mathcal{L}(1) = 1$ and $\mathcal{L}(f_n) = 1$.*

Proof. For $1 = X^0 Y_0$, we have $\bar{0} = 0$, $\beta(0) = 1$, and $\omega_0(x) = 1$. Thus

$$\mathcal{L}(1) = \mathbb{E}_x[1] = 1.$$

Next we prove that $\mathcal{L}$ vanishes on the ideal

$$\langle Y_{\ell_i} - b_i \rangle_{i=1}^{n+1}.$$

It is enough to check this on monomials. Let $M = X^\alpha Y_a$. Since a and $a + \ell_i$ lie in the same coset modulo L, they have the same representative $\bar{a}$. Therefore

$$\begin{aligned}
\mathcal{L}(Y_{\ell_i} M) &= \mathcal{L}(X^\alpha Y_{a+\ell_i}) \\
&= \mathbb{E}_x \left[(-1)^{\alpha \cdot x} \beta\big(a + \ell_i - \bar{a}\big) \omega_{\bar{a}}(x) \right] \\
&= b_i \mathbb{E}_x \left[(-1)^{\alpha \cdot x} \beta\big(a - \bar{a}\big) \omega_{\bar{a}}(x) \right] \\
&= b_i \mathcal{L}(M).
\end{aligned}$$

By linearity, $\mathcal{L}$ is zero on every multiple of $Y_{\ell_i} - b_i$.

Finally, modulo this ideal, each factor of f_n is equal to 1:

$$\frac{1 + b_i Y_{\ell_i}}{2} = \frac{1 + b_i^2}{2} = 1.$$

Hence $f_n - 1 \in \langle Y_{\ell_i} - b_i \rangle_{i=1}^{n+1}$, and so

$$\mathcal{L}(f_n) = \mathcal{L}(1) = 1.$$

□

It remains to prove that $\mathcal{L}$ is almost zero on each local AND constraint Mp_{ij} (Item 2 of Lemma 3.1). Fix $1 \le i < j \le n$, and write

$$e = x_i x_j \in \mathcal{Q}_n.$$

[3] We denote the coset $a + L \in \mathcal{Q}_n / L$ as "$a \bmod L$".

Then $Y_e = Y_{ij}$ and $\omega_e(x) = X_i \wedge X_j$. It is slightly cleaner to first use the signed local constraint

$$A_{ij} := Y_{ij} - (X_i \wedge X_j).$$

This is equivalent to p_{ij} up to the harmless factor $Y_{ij}/2$:

$$p_{ij} = \frac{Y_{ij}}{2} A_{ij}.$$

Let $M = X^\alpha Y_a$ be a monomial. By linearity of the definition of $\mathcal{L}$, for every polynomial $P(X)$ in the X-variables,

$$\mathcal{L}(P(X)Y_a) = \mathbb{E}_x \left[P((-1)^x) \beta(a - \bar{a}) \omega_{\bar{a}}(x) \right].$$

Hence

$$\begin{aligned} \mathcal{L}(MA_{ij}) &= \mathcal{L}(X^\alpha Y_{a+e}) - \mathcal{L}(X^\alpha (X_i \wedge X_j) Y_a) \\ &= \mathbb{E}_x \left[(-1)^{\alpha \cdot x} \left(\beta\big(a + e - \overline{a+e}\big) \omega_{\overline{a+e}}(x) - \beta(a - \bar{a}) \omega_{e+\bar{a}}(x) \right) \right]. \end{aligned}$$

Here $\overline{a+e}$ means $\sigma((a+e) \bmod L)$. The desired cancellation would be exact if

$$\overline{a+e} = e + \bar{a}.$$

The next lemma says that when this equality fails, the two quadratic phases left in the difference must both have high rank. This is the only place where the careful choice of the representatives $\sigma(\cdot)$ is used. Looking ahead, we will leverage this high-rank property to apply Gauss-sum Lemma 3.4 to upper bound both terms of $\mathcal{L}(MA_{ij})$. This is proven in Lemma 3.5.

Lemma 3.3 (Local defect)**.** *Fix $a \in \mathcal{Q}_n$ and $e = x_i x_j$. Define*

$$\Delta := e + \bar{a} + \overline{a+e}.$$

Then $\Delta \in L$. If $\Delta = 0$, then $\overline{a+e} = e + \bar{a}$. If $\Delta \neq 0$, then

$$\operatorname{rank}(\overline{a+e}) > \tau n - 2 \qquad \textit{and} \qquad \operatorname{rank}(e + \bar{a}) > \tau n.$$

Proof. Since $\bar{a} \equiv a \pmod{L}$ and $\overline{a+e} \equiv a + e \pmod{L}$, we have

$$e + \bar{a} + \overline{a+e} \equiv e + a + (a+e) = 0 \pmod{L}.$$

Thus $\Delta \in L$. If $\Delta = 0$, then $\overline{a+e} = e + \bar{a}$.

Assume now that $\Delta \neq 0$. Suppose first that

$$\operatorname{rank}(\overline{a+e}) \leq \tau n - 2.$$

Since $\operatorname{rank}(e) \leq 2$, we have

$$\operatorname{rank}(e + \overline{a+e}) \leq \tau n.$$

The form $e + \overline{a+e}$ lies in the same coset as a. By the rule for choosing representatives and the uniqueness of low-rank representatives, this forces

$$\bar{a} = e + \overline{a+e}.$$

But then $\Delta = 0$, contradiction. Therefore

$$\operatorname{rank}(\overline{a+e}) > \tau n - 2.$$

Similarly, if $\operatorname{rank}(e + \bar{a}) \leq \tau n$, then $e + \bar{a}$ is a rank-$\leq \tau n$ representative of the coset of $a + e$. Again by uniqueness, this forces

$$\overline{a+e} = e + \bar{a},$$

and hence $\Delta = 0$, contradiction. Thus $\operatorname{rank}(e + \bar{a}) > \tau n$. □

Lemma 3.4 (Gauss-sum bound)**.** *For every quadratic form $q \in \mathcal{Q}_n$ and every $\alpha \in \mathbb{F}_2^n$,*

$$\left|\mathbb{E}_{x\in\mathbb{F}_2^n}(-1)^{\alpha\cdot x}\omega_q(x)\right| \le 2^{-\operatorname{rank}(q)/2}.$$

Proof. Let

$$S := \mathbb{E}_{x\in\mathbb{F}_2^n}(-1)^{q(x)+\alpha\cdot x}.$$

We prove $|S|^2 \le 2^{-\operatorname{rank}(q)}$. Since the summand is real,

$$|S|^2 = \mathbb{E}_{x,y\in\mathbb{F}_2^n}(-1)^{q(x)+\alpha\cdot x+q(y)+\alpha\cdot y}.$$

Write $y = x + h$. Then

$$|S|^2 = \mathbb{E}_{x,h\in\mathbb{F}_2^n}(-1)^{q(x)+q(x+h)+\alpha\cdot h}.$$

For fixed h, the expression

$$q(x) + q(x + h) + q(h) = x^T B_q h$$

is a linear function of x (see Definition 2.1). Therefore

$$\mathbb{E}_{x\in\mathbb{F}_2^n}(-1)^{q(x)+q(x+h)}$$

is zero unless this linear function is identically zero. In the exceptional case, the inner average has absolute value 1.

The exceptional h's are exactly the kernel of the symmetric matrix defining $\operatorname{rank}(q)$. Thus, if $r = \operatorname{rank}(q)$, the fraction of exceptional h's is 2^{-r}. Hence

$$|S|^2 \le 2^{-r}.$$

Taking square roots gives

$$|S| \le 2^{-r/2} = 2^{-\operatorname{rank}(q)/2}.$$

□

Lemma 3.5. *There is a constant $c > 0$ such that for all sufficiently large n, every polynomial $h \in \mathcal{R}_n$, and every $1 \le i < j \le n$,*

$$|\mathcal{L}(hp_{ij})| \le 2^{-cn}\|h\|_1.$$

Proof. First consider a monomial $M = X^\alpha Y_a$, and keep $e = x_i x_j$. Set

$$\Delta = e + \bar{a} + \overline{a + e}.$$

The calculation above can be rewritten as

$$\mathcal{L}(MA_{ij}) = \beta(a - \bar{a})\mathbb{E}_x\left[(-1)^{\alpha\cdot x}\left(\beta(\Delta)\omega_{\overline{a+e}}(x) - \omega_{e+\bar{a}}(x)\right)\right].$$

If $\Delta = 0$, then the two phases are identical and the expression is zero. If $\Delta \neq 0$, Lemma 3.3 and Lemma 3.4 give

$$|\mathcal{L}(MA_{ij})| \le 2^{-(\tau n-2)/2} + 2^{-\tau n/2} \le 2^{-cn}$$

for some absolute constant $c > 0$, after taking n large enough.

For a polynomial C, write $C = \sum_M \widehat{C}(M)M$ in the monomial basis. By linearity and the triangle inequality,

$$|\mathcal{L}(CA_{ij})| \le 2^{-cn}\|C\|_1.$$

Finally, since $p_{ij} = Y_{ij}A_{ij}/2$,

$$|\mathcal{L}(hp_{ij})| = \frac{1}{2}|\mathcal{L}(hY_{ij}A_{ij})| \leq 2^{-cn} \|h\|_1 ,$$

where we used that multiplication by Y_{ij} preserves the coefficient ℓ_1-norm, and absorbed the factor $1/2$ into the constant c. □

Together, Lemma 3.2 and Lemma 3.5 prove the lemma stated at the beginning of this section.

4 The Exponential Norm Lower Bound

We now prove the third item of Theorem 1.1; the first two items were proved in Lemma 2.3.

Proof of Theorem 1.1, item 3. Suppose

$$f_n = \sum_{1 \leq i < j \leq n} q_{ij} p_{ij}.$$

Apply $\mathcal{L}$ to both sides. Lemma 3.2 gives $\mathcal{L}(f_n) = 1$, while Lemma 3.5 gives

$$1 = |\mathcal{L}(f_n)| \leq \sum_{1 \leq i < j \leq n} |\mathcal{L}(q_{ij} p_{ij})| \leq 2^{-cn} \sum_{1 \leq i < j \leq n} \|q_{ij}\|_1 .$$

Therefore

$$\sum_{1 \leq i < j \leq n} \|q_{ij}\|_1 \geq 2^{cn}.$$

This proves the claimed exponential lower bound. Since $C_{\mathrm{AND},n}$ has $n + \binom{n}{2} = \Theta(n^2)$ variables, the lower bound is super-polynomial in the ambient dimension, and hence contradicts any polynomial norm-blowup bound. □

5 Extension to Nullstellensatz Refutation Lower Bound

In this section, we extend the counterexample above to the general Nullstellensatz refutations and prove Theorem 1.2. Recall that a polynomial is a k-junta if it depends on at most k variables.

Proof of Theorem 1.2. We start with the construction of $\{f_r\}_{r \in \mathbb{N}}$ from the previous sections on r variables (take n in the previous section as r). We will build a new unsatisfiable system using the original variables $(X_u)_{u \in [r]}$, the AND variables $(Y_{uv})_{1 \leq i < j \leq r}$, and some auxiliary variables. At the end, we relabel all of these variables as $x_1, \ldots, x_n$.

Axioms $\{p_i\}_i$. The axiom includes the original AND constraints $(p_{uv})_{1 \leq u < v \leq r}$. These are 3-juntas and have coefficient ℓ_1-norm $O(1)$. We will add more axioms into the system.

Computing Y_{ℓ_i} locally. The only nonlocal condition we need to implement is $Y_{\ell_i} = b_i$. Write

$$\ell_i = \sum_{(u,v) \in S_i} x_u x_v.$$

Then

$$Y_{\ell_i} = \prod_{(u,v) \in S_i} Y_{uv}.$$

For each $i \in [r+1]$, we build a binary tree whose leaves are the variables Y_{uv} with $(u,v) \in S_i$. For every internal node, introduce a new variable Z. If the two children have values U, V, add the multiplication-gate axiom

$$G = \frac{1 - ZUV}{2}.$$

Over the $\{\pm 1\}$-Boolean cube, $G = 0$ exactly enforces $Z = UV$. This axiom is a 3-junta and has ℓ_1-norm 1.

Let Z_i^{root} denote the value at the root of the tree. Add the final axiom

$$R_i = \frac{1 - b_i Z_i^{\mathrm{root}}}{2}.$$

Thus $R_i = 0$ enforces $Z_i^{\mathrm{root}} = b_i$. This is a 1-junta and has ℓ_1-norm 1.

Size of the system. The number of new axioms is

$$O\left(\sum_{i=1}^{r+1} |S_i|\right) = O(r^3),$$

since each ℓ_i has at most $\binom{r}{2}$ monomials. The number of variables is also $O(r^3)$. After relabeling all variables as $x_1, \ldots, x_n$, we have $n = \mathrm{poly}(r)$, and the number of axioms is $m = \mathrm{poly}(n)$. All axioms are 3-juntas and have coefficient ℓ_1-norm $O(1)$.

Unsatisfiability. The system is unsatisfiable. Indeed, the AND axioms force $Y_{uv} = X_u \wedge X_v$, and the computation trees then force

$$Z_i^{\mathrm{root}} = Y_{\ell_i} = \omega_{\ell_i}(x).$$

The final axioms force $Z_i^{\mathrm{root}} = b_i$ for every i, contradicting $\mathbf{b} \notin \operatorname{im} \Psi$.

Projecting a refutation back. Now suppose

$$1 = \sum_{u<v} q_{uv} p_{uv} + \sum_G q_G\, G + \sum_i q_i R_i$$

is any Nullstellensatz refutation. Define a homomorphism π by fixing the original X, Y-variables and replacing every auxiliary variable by the product of the Y-variables computed by its subtree. Then $\pi(G) = 0$ for every multiplication gate, and

$$\pi(R_i) = \frac{1 - b_i Y_{\ell_i}}{2}.$$

Also, π does not increase coefficient ℓ_1-norm, because it sends every monomial to a monomial.

Applying π to the refutation gives

$$1 = \sum_{u<v} \pi(q_{uv}) p_{uv} + \sum_i \pi(q_i) \frac{1 - b_i Y_{\ell_i}}{2}.$$

Killing the final axioms. Let

$$h_i := \frac{1 + b_i Y_{\ell_i}}{2}, \qquad f_r = \prod_{i=1}^{r+1} h_i.$$

Since

$$h_i \cdot \frac{1 - b_i Y_{\ell_i}}{2} = 0$$

in the Boolean quotient ring, multiplying the displayed identity by f_r kills all final-axiom terms and gives

$$f_r = \sum_{u<v} f_r \pi(q_{uv}) p_{uv}.$$

This is a representation of f_r using only the original AND axioms.

Transferring the lower bound. By Theorem 1.1,

$$\sum_{u<v} \|f_r \pi(q_{uv})\|_1 \geq 2^{cr}$$

for some constant $c > 0$. On the other hand, $\|f_r\|_1 = 1$ by Lemma 2.3. By submultiplicativity and since π is norm non-increasing,

$$\|f_r \pi(q_{uv})\|_1 \leq \|\pi(q_{uv})\|_1 \leq \|q_{uv}\|_1 .$$

Therefore

$$\sum_{t=1}^{m} \|q_t\|_1 \geq \sum_{u<v} \|q_{uv}\|_1 \geq 2^{cr}.$$

Relabeling the parameter. The final number of variables satisfies $n = O(r^3)$, so $2^{cr} \geq 2^{n^{\Omega(1)}}$. This proves the theorem for infinitely many values of n. □

Acknowledgments. We thank Liyan Chen for helping verify the correctness of the proof and Aayush Jain and Surya Mathialagan for helpful comments. Zhengzhong Jin is supported in part by a 2025 Amazon Research Award.

References

[DHK+26] Lalita Devadas, Samuel B. Hopkins, Yael Tauman Kalai, Pravesh K. Kothari, Alex Lombardi, and Surya Mathialagan. SNARGs for NP and non-signaling PCPs, revisited. Cryptology ePrint Archive, Paper 2026/006, 2026. https://eprint.iacr.org/2026/006.

[Lom26] Alex Lombardi. Counterexamples to the low-norm nullstellensatz hypothesis, 2026. Private communication.

[PZ24] Aaron Potechin and Aaron Zhang. Bounds on the total coefficient size of nullstellensatz proofs of the pigeonhole principle. In Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson, editors, *ICALP 2024: 51st International Colloquium on Automata, Languages and Programming*, volume 297 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 117:1–117:20, Tallinn, Estonia, July 8–12, 2024. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.