Electronic Colloquium on Computational Complexity, Report No. 147 (2026)

Improved Soundness for the Line–versus–Point Test

Scott Duke Kominers Justin Thaler Kai Zhe Zheng

July 2026

Abstract

The line–versus–point test asks the following local-to-global question. Suppose a function $f\colon \mathbb{F}_q^m \to \mathbb{F}_q$ is, on average over a random affine line L, correlated with some degree-d polynomial on L. Must f then be globally correlated with a single multivariate polynomial of degree at most d? Beyond being a natural combinatorial question, this test lies at the heart of low-degree testing and has played a central role in the construction of PCPs. Low-degree tests and their higher-dimensional variants have also contributed to results such as the 2-to-2 Games Theorem and two-query PCPs with nearly optimal alphabet–soundness tradeoffs.

Of particular interest is the soundness of the test: the smallest expected local agreement from which one can still deduce nontrivial global agreement. We revisit the line–versus–point test and, building on the algebraic approach of [HKSS24], prove the following cubic threshold. If

$$\Pr_{L,\, x\in L}\big[P_L(x) = f(x)\big] \geq C\left(\frac{d}{q}\right)^{1/3},$$

then there is a polynomial $Q \in \mathbb{F}_q[X_1,\dots,X_m]$ of total degree at most d such that

$$\Pr_{x\in\mathbb{F}_q^m}[Q(x) = f(x)] \geq c \cdot \Pr_{L,\, x\in L}\big[P_L(x) = f(x)\big],$$

for absolute constants $C, c > 0$ and over every finite field $\mathbb{F}_q$. Our proof is comparatively direct and isolates the combinatorial and algebraic mechanisms responsible for the improved soundness.

1 Introduction

The Reed–Muller code is one of the basic algebraic codes in theoretical computer science. Its codewords are the evaluation tables of low-degree polynomials:

$$\mathrm{RM}_q(m,d) := \big\{(Q(x))_{x\in\mathbb{F}_q^m} : Q \in \mathbb{F}_q[X_1,\dots,X_m],\ \deg Q \leq d\big\}.$$

Much of the usefulness of Reed–Muller codes comes from their strong locality. The restriction of a degree-d polynomial to any affine line is a univariate polynomial of degree at most d, and hence a Reed–Solomon codeword. Thus a global Reed–Muller codeword induces a highly structured local view on every affine line. This observation underlies local testing and correction procedures that inspect only a small part of a word of length q^m.

One particularly important test arising from this structure is the line–versus–point test, a two-query consistency test. We are given a point table $f\colon \mathbb{F}_q^m \to \mathbb{F}_q$ and, for every affine line L, a purported degree-d restriction P_L. The test samples a uniformly random incident pair (x, L), queries $f(x)$ and $P_L(x)$, and accepts when

$$P_L(x) = f(x).$$

ISSN 1433-8092

If the point and line tables are induced by a single degree-d polynomial, the test always accepts. Its soundness asks for a converse: how small can the acceptance probability be while still forcing one global degree-d polynomial to agree with f on a noticeable fraction of points?

Low-degree tests, such as the line–versus–point test and its variants, have played a central role in PCP constructions and hardness of approximation since the earliest algebraic PCPs [PS94, RS96, AS98, ALM+98, RS97, AS03]. Their two-query form is particularly useful for constructing projection PCPs and Label Cover instances, which serve as starting points for many hardness-of-approximation reductions [MR10]. Higher-dimensional subspace agreement tests have also been central to the series of works leading to the 2-to-2 Games Theorem [KMS25, DKK+25, DKK+21, KMS23] and to recent two-query and k-query PCPs with nearly optimal alphabet–soundness tradeoffs [MZ24, MZ26]. More recently, the line–versus–point test has been used to construct query-efficient interactive oracle proofs of proximity for NP [MZ25].

Beyond these applications, the line–versus–point test is a natural local-to-global agreement problem in its own right. It asks whether a collection of low-degree polynomials on affine lines that is locally consistent with a point table must contain evidence of a single global low-degree polynomial. This question was introduced by Rubinfeld and Sudan and has been studied through a sequence of increasingly strong analyses [RS96, PS94, FS95, AS03, HKSS24]. It is also closely related to the property-testing problem for Reed–Muller codes, in which one seeks to distinguish low-degree polynomials from functions that are far from every such polynomial using only a small number of queries [AKK+05, KR06, BKS+10, HSS13]. The models and degree regimes in these works are not always identical, but they share the same basic theme of inferring global polynomial structure from local algebraic constraints.

A central quantitative question is how far into the low-agreement, or list-decoding, regime the line–versus–point test remains sound.

Polishchuk and Spielman, and subsequently Friedl and Sudan, initiated the study of the line–versus–point test in the low-error, or unique-decoding, regime [PS94, FS95]. Arora and Sudan subsequently obtained a high-error analysis when the field is polynomially larger than the degree [AS03]. Harsha, Kumar, Saptharishi, and Sudan gave the first high-error analysis when the field size is linear in the degree [HKSS24]. Their work showed that acceptance $(d/q)^{\Omega(1)}$ already forces global structure and introduced the key algebraic mechanism of constructing a low weighted-degree trivariate relation and extracting a bivariate polynomial through Newton–Hensel lifting. They state their all-dimensional result with an unspecified absolute exponent. Tracing the losses in their proof gives the threshold $\varepsilon \gtrsim (d/q)^{1/48}$, while their direct bivariate argument requires $\varepsilon \gtrsim (d/q)^{1/7}$ and initially recovers agreement $\Omega(\varepsilon^4)$.

Building on the analysis of [HKSS24], we introduce new combinatorial and algebraic ideas that improve the soundness threshold to

$$\varepsilon \gtrsim \left(\frac{d}{q}\right)^{1/3}$$

and the recovered agreement to $\Omega(\varepsilon)$. We first establish these bounds in the bivariate case and then preserve them in arbitrary dimension.

1.1 Our results

Throughout, fix a finite field $\mathbb{F}_q$ of order q, where q is an arbitrary prime power; in particular, $\mathbb{F}_q$ need not be a prime field. Let $\mathcal{L}_{q,m}$ be the set of affine lines in $\mathbb{F}_q^m$, where each affine line is given by a translate $a \in \mathbb{F}_q^m$ and a direction $u \in \mathbb{F}_q^m \setminus \{0\}$, so that $L = \{a + tu \mid t \in \mathbb{F}_q\}$. When we work in two dimensions and the field is clear from context, we instead write $\mathcal{L}$ for $\mathcal{L}_{q,2}$. For every affine line L, fix once and for all an affine parameterization, consisting of translate $a_L \in L$ and direction

u_L, and write $L(t) := a_L + tu_L$ to denote the t-th point of L. We slightly abuse notation by using L both for the affine line and, through $L(\cdot)$, for its fixed parameterization.

Our results deal with the classic *line–versus–point test*. In this test, we are given a point table $f\colon \mathbb{F}_q^m \to \mathbb{F}_q$ and a degree-d line table $\{P_L\}_{L\in\mathcal{L}_{q,m}}$, where each $P_L : L \to \mathbb{F}_q$ is a univariate polynomial of degree at most d defined over L. Relative to the fixed parameterization of L, let $p_L \in \mathbb{F}_q[t]$ denote the chosen degree-at-most-d polynomial satisfying

$$p_L(t) = P_L(L(t)).$$

We use this notation henceforth. The test samples an incident pair (x, L) uniformly at random, where $x \in \mathbb{F}_q^m$ and $L \in \mathcal{L}_{q,m}$ such that $x \in L$, and accepts if and only if $P_L(x) = f(x)$. Given the point and line tables, we define the acceptance probability of the test as

$$\mathrm{Acc}(f, \{P_L\}) := \Pr_{x\in\mathbb{F}_q^m,\, L\ni x}[P_L(x) = f(x)].$$

Our main theorem shows that one can conclude nontrivial correlation of f with a low-degree polynomial even when the test passes with probability as low as $\Omega\left((d/q)^{1/3}\right)$. This threshold is generally referred to as the *soundness* of the test. The previous state of the art was the work of due to [HKSS24], which required a higher acceptance threshold. While their work does not state it explicitly, we estimate that they require $\Omega\,(d/q)^{1/48}$ pass probability over $\mathbb{F}_q^m$ for $m > 2$ and $\Omega\,(d/q)^{1/7}$ over $\mathbb{F}_q^2$.

Theorem 1.1. *There are absolute constants $C, c > 0$ with the following property. Let $\varepsilon \geq C\left(\frac{d}{q}\right)^{1/3}$. Suppose that $f\colon \mathbb{F}_q^m \to \mathbb{F}_q$ and a degree-d line table $\{P_L\}_{L\in\mathcal{L}_{q,m}}$ satisfy*

$$\mathrm{Acc}(f, \{P_L\}) \geq \varepsilon.$$

Then there is a polynomial $Q \in \mathbb{F}_q[X_1, \ldots, X_m]$ of total degree at most d such that

$$\mathrm{agr}(f, Q) := \Pr_{x\in\mathbb{F}_q^m}[Q(x) = f(x)] \geq c\varepsilon.$$

Mechanization. Theorem 1.1 has been formalized and machine-checked in Lean 4 over the Mathlib library. The bivariate theorem is verified with no assumptions beyond Lean's logical foundations, with certified (deliberately slack) explicit constants $C = 10^{40}$ and $c = 1/8000$; the all-dimensional theorem is verified assuming exactly the two imported results, Theorem 7.2 and Lemma 7.1.

1.2 Proof overview and technical contributions

Our proof proceeds in two stages. We first establish the result in the bivariate case $m = 2$. This bivariate statement is the main lemma underlying the theorem, and its proof contains most of the work. We then bootstrap it to arbitrary dimension using the results of [BR25].

Why the bivariate line case is the core. The difficulty comes from the sparse overlap of line views. In tests based on planes or cubes, two local views can share an entire line or plane, so a consistency event compares many related values at once [RS97, BDN17, MZ23, BR25]. Two distinct lines, by contrast, share at most one point. Thus each accepted line–point incidence supplies only one scalar equality, and most of the proof is devoted to turning this sparse information into a relation that can be propagated globally. Once the bivariate statement is established, the passage to higher dimensions is comparatively direct.

At a high level, we follow the strategy of [HKSS24]. Let G_{acc} be the bipartite graph whose two sides are the affine lines $\mathcal{L}$ and the points $\mathbb{F}_q^2$, with an edge (x, L) whenever $P_L(x) = f(x)$. Thus, the hypothesis that the line–versus–point test accepts with probability at least ε is exactly the statement that the accepted edge set E_{acc} has measure $\mu(E_{\mathrm{acc}}) \geq \varepsilon$ among all incidences (x, L) with $x \in L$.

Starting from these accepted incidences, both proofs construct a trivariate *explainer polynomial* $A(X, Y, Z)$ and arrange that, for many lines L,

$$A(L(t), p_L(t)) \equiv 0.$$

They then find a point b through which many of these lines pass, with $P_L(b) = f(b)$ for every such line. When $f(b)$ is a simple root of the explainer at b, Newton–Hensel lifting produces a global bivariate polynomial $Q(X, Y)$ whose restrictions agree with the corresponding line polynomials.

In the bivariate argument of [HKSS24], the global polynomial Q is obtained by lifting at one suitable point b, and its agreement with f is obtained from the good lines through b. Thus, the final agreement count is tied to one local neighborhood of b in G_{acc}, which is one quantitative limitation of their argument.

To go beyond this limitation, we observe that the polynomial lifted at b can agree with the point table beyond the lines through b. For this purpose, we use the informative part of the explainer A. More precisely, we factor

$$A(X, Y, Z) = H(X, Y)B(X, Y, Z),$$

where H is the greatest common factor of the coefficients of A when it is viewed as a polynomial in Z. The factor H can vanish identically on a line independently of the line polynomial P_L, in which case the identity $A(L(t), p_L(t)) \equiv 0$ gives no information about p_L. Dividing out H removes these trivial vanishing relations and leaves a common explainer $B(X, Y, Z)$ for the line polynomials that remain.

We then arrange that, at every point x in the graph under consideration, the value $f(x)$ is not a repeated root of

$$B(x, Z) \in \mathbb{F}_q[Z].$$

This makes the local lift uniquely determined. Suppose $Q|_L = P_L$ for a line L in the current subgraph, and let (x, L) be an accepted edge. Then

$$Q(x) = P_L(x) = f(x).$$

If M is another accepted line through x, then P_M and $Q|_M$ solve the same restricted explainer equation and agree at the simple root $f(x)$. The algebraic uniqueness of a lift from a simple root therefore gives $Q|_M = P_M$ (later stated in Lemma 2.7). Repeating this step along paths propagates the same polynomial through an entire connected component of the pruned accepted-incidence graph. This allows us to show that Q agrees with f on a linear fraction of the points in that component, rather than only on the agreement sets of the lines through one point.

1. Interpolating a sparse, but global, set of affine lines. We start by finding a subgraph G_0, of G_{acc}, with edge measure $\Omega(\varepsilon)$ and minimum degree $\Omega(\varepsilon q)$ on both sides, and then randomly choose a set of $O(q/\varepsilon)$ from the "line side" of G_0. Denote this set of lines by $\mathcal{R}$. Crucially, $\mathcal{R}$ satisfies the following two properties

- **Small size:** $|\mathcal{R}| = O(q/\varepsilon)$, so that a polynomial interpolating the graphs of $\{P_M\}_{M\in\mathcal{R}}$ can have low weighted degree.
- **Globalness:** for every line L in the dense subgraph, there are $\Omega(\varepsilon q)$ points $x \in L$ such that (x, L) is accepted and x also lies on an accepted line $M \in \mathcal{R}$.

We interpolate a nonzero trivariate polynomial $A(X, Y, Z)$ of weighted degree $D := O\left(\sqrt{\frac{dq}{\varepsilon}}\right)$ such that $A(M(t), p_M(t)) \equiv 0$ for every $M \in \mathcal{R}$. In a sense, A now "encodes" the global agreement of f with a low-degree polynomial, and the later steps extract this agreement out of A.

We remark that the interpolation step is more efficient than the corresponding step in [HKSS24], although the gain is not simply that we interpolate asymptotically fewer lines at the final threshold. We interpolate only a family $\mathcal{R}$ of size $O(q/\varepsilon)$ while preserving the original ε-scale incidence density. At first this seems to leave too few lines from which to recover substantial agreement. The key is the global property of $\mathcal{R}$: every line in the dense subgraph shares $\Omega(\varepsilon q)$ accepted points with lines in $\mathcal{R}$. These shared incidences extend the interpolated relation to the entire subgraph, and the later propagation step converts this relation into agreement on a whole connected component.

2. Simple roots throughout the incidence graph. Before extracting a bivariate polynomial from the common relation

$$A(L(t), p_L(t)) \equiv 0,$$

we need to arrange two properties. First, the relation on each remaining line must genuinely constrain its line polynomial, rather than vanish for a reason independent of the Z-variable. Second, at every remaining point x where we hope to get global agreement, the value $f(x)$ must be a simple root of the resulting polynomial in Z. The first property makes the line identities informative, while the second is needed to make the later lifting and propagation steps go through.

To obtain the first property, write

$$A(X, Y, Z) = H(X, Y)B(X, Y, Z),$$

where H is the greatest common factor of the coefficients of A when viewed as a polynomial in Z. If $H(L(t)) \equiv 0$, then $A(L(t), p_L(t)) \equiv 0$ for every choice of p_L, so the identity contains no information about the polynomial attached to L. We remove these lines; there are few of them because each contributes a distinct linear factor of H. On every other line, cancellation gives the informative identity

$$B(L(t), p_L(t)) \equiv 0.$$

We next arrange the simple-root property. Choose $\lambda_X, \lambda_Y, \lambda_Z \in \mathbb{F}_q$ so that B and $\lambda_X B_X + \lambda_Y B_Y + \lambda_Z B_Z$ have no common factor. Here B_X, B_Y, and B_Z are the X, Y, and Z derivatives respectively of the formal polynomial B. After pruning a small number of lines, we ensure that every remaining accepted point x on any of the lines actually lies on at least two distinct remaining lines and satisfies

$$(\lambda_X B_X + \lambda_Y B_Y + \lambda_Z B_Z)(x, f(x)) \neq 0 \quad \text{and} \quad B(x, f(x)) = 0.$$

It remains to show that this root is simple. Let L and M be two distinct remaining lines through x, with direction vectors $u_L, u_M \in \mathbb{F}_q^2$. Reparameterize both lines so that $t = 0$ corresponds to x, and write

$$\nabla_{X,Y} B(x, f(x)) := \big(B_X(x, f(x)), B_Y(x, f(x))\big).$$

Now suppose that $B_Z(x, f(x)) = 0$. Differentiating the two formal line identities at $t = 0$ gives

$$\langle u_L, \nabla_{X,Y} B(x, f(x)) \rangle = 0, \qquad \langle u_M, \nabla_{X,Y} B(x, f(x)) \rangle = 0.$$

Here $\langle \cdot, \cdot \rangle$ is the usual inner product on $\mathbb{F}_q^2$ and $u_L, u_M \in \mathbb{F}_q^2$ are the directions of the lines L, M respectively. Because L and M are distinct lines through the same point, u_L and u_M are linearly independent and therefore form a basis of $\mathbb{F}_q^2$. Hence we must have $B_X(x, f(x)) = B_Y(x, f(x)) = 0$. Together with $B_Z(x, f(x)) = 0$, this would imply

$$(\lambda_X B_X + \lambda_Y B_Y + \lambda_Z B_Z)(x, f(x)) = 0,$$

contradicting the pruning. This shows that $B_Z(x, f(x)) \neq 0$ at every remaining point as desired.

3. Componentwise continuation. We can now choose a point b in the remaining incidence graph and lift a degree-d bivariate polynomial Q. The lifting argument guarantees that Q agrees with the line polynomials on a collection of lines through b. By itself, however, this local agreement is not enough to obtain the desired $\Omega(\varepsilon)$ agreement with f.

This is another point at which we must go beyond the analysis of [HKSS24]. Rather than counting only the agreement contributed by the lines through the lifting point b, as is done in [HKSS24], we use the simple-root property arranged in Step 2 to propagate the same polynomial throughout an entire connected component of the incidence graph. Expansion of the point–line incidence graph shows that each such component is sufficiently large, guaranteeing that the agreement of Q with f is $\Omega(\varepsilon)$.

Suppose $Q|_L = P_L$ for one remaining line L, and let $x \in L$ be an accepted incidence. Then

$$Q(x) = P_L(x) = f(x).$$

For any other remaining line M through x, both $Q|_M$ and P_M satisfy the same restricted relation

$$B(M(t), Z) = 0$$

and pass through the same simple root $f(x)$. Algebraic uniqueness of the lift from a simple root then yields $Q|_M = P_M$. Finally, repeating this argument along paths propagates the same polynomial from the lines through b to every line and point in the connected component as desired.

4. Higher dimensions. Finally, we bootstrap the bivariate result to arbitrary dimension. Here we leverage the recent plane–versus–plane theorem of Bhangale and Richelson [BR25], which has the optimal soundness threshold $\Omega(d/q)$. A simple conversion shows that plane–point agreement β induces plane–versus–plane acceptance roughly β^2; thus their theorem applies once $\beta \gtrsim (d/q)^{1/2}$.

The rest of the argument is relatively direct. Most affine planes inherit line–versus–point acceptance $\Omega(\varepsilon)$, so the bivariate theorem gives a degree-d polynomial on each such plane. We apply the Bhangale–Richelson theorem to glue these plane polynomials into one global polynomial, and point–plane expansion recovers $\Omega(\varepsilon)$ agreement with f.

The bootstrap step in [HKSS24], developed before this optimal plane–versus–plane theorem was available, requires a more involved intermediate correction argument and consequently incurs larger quantitative losses. The result of Bhangale and Richelson allows us to avoid that difficulty and pass directly from the bivariate theorem to a consistent global polynomial.

1.3 Organization

Section 2 records notation and preliminaries. Sections 3 and 4 constructs the trivariate explainer polynomial, Section 5 makes its relevant roots simple, and Section 6 lifts and propagates one bivariate polynomial and completes the bivariate proof. Finally, Section 7 handles higher dimensions.

2 Preliminaries

2.1 Measure and affine-plane incidence

Every finite set carries its uniform probability measure. We use μ for point, line, plane, incidence, and conditional measures; the ambient set will be clear from context. In $\mathbb{F}_q^2$,

$$\mu(X) = \frac{|X|}{q^2}, \qquad \mu(T) = \frac{|T|}{q(q+1)}, \qquad \mu(E) = \frac{|E|}{q^2(q+1)}.$$

We identify polynomial functions with their reduced representatives. Thus $Z^q = Z$ as a function on $\mathbb{F}_q$. Factorization and derivatives refer to the corresponding formal polynomials, and a univariate restriction of degree below q that vanishes on all of $\mathbb{F}_q$ is identically zero.

Every point lies on $q+1$ lines, every line contains q points, and two distinct points lie on a unique common line.

Lemma 2.1. *For $X \subseteq \mathbb{F}_q^2$ and $T \subseteq \mathcal{L}$,*

$$\left| \Pr_{x\in\mathbb{F}_q^2,\, L\ni x}[x \in X,\ L \in T] - \mu(X)\mu(T) \right| \le \frac{\sqrt{\mu(X)\mu(T)}}{\sqrt{q+1}}.$$

Proof. Put $n_L = |X \cap L|$. Double counting gives

$$\sum_L n_L = (q+1)|X|, \qquad \sum_L n_L^2 = |X|^2 + q|X|,$$

so $\sum_L (n_L - |X|/q)^2 \le q|X|$. Cauchy–Schwarz over $L \in T$ and normalization give the claim. ☐

2.2 Weighted degree and monomial counts

We fix the weights of X, Y, Z to be $1, 1, d$. Henceforth, "weighted degree" means

$$\mathrm{wdeg}(A) := \max\{i + j + dk : [X^i Y^j Z^k]A \neq 0\}.$$

If $\deg p_L \le d$, then

$$\deg_t A(L(t), p_L(t)) \le \mathrm{wdeg}(A). \tag{1}$$

Let $p := \mathrm{char}(\mathbb{F}_q)$ and define

$$\mathcal{V}_D := \mathrm{span}_{\mathbb{F}_q}\{X^i Y^j Z^k : i + j + dk \le D,\ k = 0 \text{ or } p \nmid k\}.$$

The restriction on k is the characteristic-compatible pruning of HKSS [HKSS24].

Lemma 2.2. *For every $D \ge 0$,*

$$\dim \mathcal{V}_D \ge \frac{D^3}{64d}.$$

Proof. The monomials with $i, j \le \lfloor D/4 \rfloor$ and $k \le \lfloor D/(2d) \rfloor$ already number more than $D^3/(32d)$. Pair each excluded Z-degree $mp > 0$ with the retained level $mp - 1$; the latter is at least as large. Thus at least half of the monomials remain. ☐

2.3 Squarefree interpolants and derivatives

Lemma 2.3. *Let $S \subseteq \mathbb{F}_q^3$ be finite, and suppose there is a nonzero polynomial that vanishes on S and has nonzero Z-derivative. If A has minimum weighted degree among all such polynomials, then no nonconstant irreducible factor occurs in A with multiplicity at least two.*

Proof. Write $A = P^e R$ with $e \geq 2$, P irreducible, and $P \nmid R$. Then PR has the same zero set and smaller weighted degree. If $(PR)_Z = 0$, reducing $P_Z R + P R_Z = 0$ modulo P gives $P \mid P_Z$, hence $P_Z = R_Z = 0$, contrary to $A_Z \neq 0$. Thus PR is another admissible interpolant, contradicting minimality. $\square$

Lemma 2.4. *Let $B \in \mathbb{F}_q[X, Y, Z]$ be nonzero and squarefree with $\deg B < q$. There is a vector $\lambda = (\lambda_X, \lambda_Y, \lambda_Z) \in \mathbb{F}_q^3$ such that, for*

$$\mathcal{D}_\lambda := \lambda_X \partial_X + \lambda_Y \partial_Y + \lambda_Z \partial_Z,$$

one has

$$\gcd(B, \mathcal{D}_\lambda B) = 1 \qquad \text{in } \mathbb{F}_q[X, Y, Z].$$

Proof. Write $B = c \prod_{i=1}^{s} P_i$ with distinct irreducible factors. Since $\mathbb{F}_q$ is perfect, each P_i has a nonzero partial derivative. The vectors λ with $\mathcal{D}_\lambda P_i = 0$ form a proper subspace of $\mathbb{F}_q^3$. There are $s \leq \deg B < q$ such subspaces, so choose λ outside their union. Then no P_i divides $\mathcal{D}_\lambda B$. $\square$

2.4 Some algebraic lemmas

Lemma 2.5. *Let $R \in \mathbb{F}_q[X, Y]$ have total degree at most D. If $R(L(t)) \equiv 0$ for more than D distinct affine lines L, then $R = 0$.*

Proof. The linear equation of each line divides R, so a nonzero R would have more than D distinct linear factors. $\square$

If $D < q$, pointwise vanishing on all q points of a line implies the formal identity above.

Lemma 2.6. *Let $U, V \in \mathbb{F}_q[X, Y, Z]$ be coprime nonzero polynomials of total degrees u and v. Let $\mathcal{T} \subseteq \mathcal{L}$ be a family of affine lines, each carrying a polynomial P_L of degree at most d, such that both U and V vanish formally on the graph of P_L:*

$$U(L(t), p_L(t)) \equiv V(L(t), p_L(t)) \equiv 0 \qquad (L \in \mathcal{T}).$$

Then

$$\mu(\mathcal{T}) \leq \frac{uv}{q(q+1)}, \qquad \text{equivalently} \qquad |\mathcal{T}| \leq uv.$$

Proof. If one polynomial has Z-degree zero, the claim follows from Lemma 2.5. Otherwise let $W = \mathrm{Res}_Z(U, V) \in \mathbb{F}_q[X, Y]$. Coprimality gives $W \neq 0$, and $\deg W \leq uv$. For every $L \in \mathcal{T}$, the specializations $U(L(t), Z)$ and $V(L(t), Z)$ share the root $p_L(t)$, so $W(L(t)) \equiv 0$. Apply Lemma 2.5. $\square$

The following standard statement follows from Newton iteration; see [Bür00]. The form stated below appears as Lemma 2.10 of [HKSS24].

Lemma 2.7. *Let $B \in \mathbb{F}_q[X, Y, Z]$ and $\alpha \in \mathbb{F}_q$ satisfy*

$$B(0, 0, \alpha) = 0, \qquad B_Z(0, 0, \alpha) \neq 0.$$

For every integer $k \geq 0$, there is a unique polynomial $\Phi_k \in \mathbb{F}_q[X, Y]$ of total degree at most k such that

$$\Phi_k(0, 0) = \alpha, \qquad B(X, Y, \Phi_k(X, Y)) \equiv 0 \pmod{(X, Y)^{k+1}}.$$

3 Step 1: Preparing the Accepted Incidences for Interpolation

The goal is to choose a small family $\mathcal{R}$ for interpolation without losing contact with the accepted lines. We first prune the accepted incidences to a subgraph with substantial measure and large degree on both sides. We then choose $\mathcal{R}$ so that every line in this subgraph shares many accepted points with lines in $\mathcal{R}$.

3.1 A dense accepted subgraph

Let G_{acc} be the bipartite graph with point side $\mathbb{F}_q^2$, line side $\mathcal{L}$, and edges

$$E(G_{\mathrm{acc}}) = \{(x, L) : x \in L,\ P_L(x) = f(x)\}.$$

Let

$$X_{\mathrm{h}} := \left\{ x : \Pr_{L \ni x}[P_L(x) = f(x)] \geq \frac{\varepsilon}{2} \right\}.$$

The accepted incidences outside X_{h} have measure at most $\varepsilon/2$, so those incident to X_{h} have measure at least $\varepsilon/2$.

Set

$$k_0 := \left\lfloor \frac{\varepsilon q}{100} \right\rfloor.$$

Pruning the subgraph induced by X_{h} gives the following.

Lemma 3.1. *There is a subgraph $G_0 \subseteq G_{\mathrm{acc}}$, with sides $X_0 \subseteq X_{\mathrm{h}}$, $\mathcal{L}_0 \subseteq \mathcal{L}$, and edge set E_0, satisfying*

$$\mathsf{mindeg}(G_0) \geq k_0, \qquad \mu(E_0) := \frac{|E_0|}{q^2(q+1)} \geq \frac{\varepsilon}{4}, \qquad \mu(\mathcal{L}_0) := \frac{|\mathcal{L}_0|}{q(q+1)} \geq \frac{\varepsilon}{4}.$$

Proof. Start with the accepted incidences incident to X_{h} and repeatedly delete vertices of degree below k_0. Deleting point vertices removes incidence measure less than $k_0/(q+1)$, and deleting line vertices removes less than k_0/q. Their sum is below $\varepsilon/50$, so more than $\varepsilon/4$ of the incidence measure remains. The minimum-degree claim is by construction, and $\mu(\mathcal{L}_0) \geq \mu(E_0)$. $\square$

3.2 A small family reaching every line

We need $\mathcal{R}$ to be small enough for low-degree interpolation and spread well enough to reach every line of G_0 through many accepted points.

Lemma 3.2. *Under the hypotheses of Theorem 1.1, after increasing C if necessary, there exists a set $\mathcal{R} \subseteq \mathcal{L}_0$ satisfying the following:*

(i)

$$\frac{100}{q} \leq \mu(\mathcal{R}) \leq \frac{3200}{\varepsilon q}.$$

(ii) *For every $L \in \mathcal{L}_0$,*

$$\big|\{x \in L : (x, L) \in E_0 \text{ and } \exists L' \in \mathcal{R} \setminus \{L\} \text{ with } (x, L') \in E_0\}\big| > \frac{\deg_{G_0}(L)}{2} \geq \frac{k_0}{2}$$

Proof. Select each line independently with probability $\theta = 8/k_0$. For a fixed $L \in \mathcal{L}_0$ and $x \in N_{G_0}(L)$, at least $k_0 - 1$ other accepted lines pass through x, so x is covered with probability at least $3/4$. For distinct points of L these events depend on disjoint line choices. Thus a Chernoff bound gives failure probability at most $(7/8)^{k_0}$ for L. For sufficiently large C, a union bound over all lines shows that all coverage requirements hold with probability at least $3/4$.

Also $\mathbb{E}|\mathcal{R}| = \theta|\mathcal{L}_0|$ and $\mathrm{Var}(|\mathcal{R}|) \le \mathbb{E}|\mathcal{R}|$. For sufficiently large C, $\mathbb{E}|\mathcal{R}| \ge 16$, so Chebyshev gives

$$\frac{1}{2}\theta|\mathcal{L}_0| \le |\mathcal{R}| \le 2\theta|\mathcal{L}_0|$$

with probability at least $3/4$. Since

$$\frac{200}{q} \le \theta\mu(\mathcal{L}_0) \le \frac{1600}{\varepsilon q},$$

one choice satisfies both claims. ☐

4 Step 2: Interpolation on the Selected Lines

Fix $\mathcal{R}$ from Lemma 3.2 and set

$$D := 16\left(\left\lfloor\sqrt{d\,|\mathcal{R}|}\right\rfloor + 1\right), \qquad k_1 := \left\lfloor\frac{\varepsilon q}{1000}\right\rfloor.$$

The goal is to interpolate a trivariate polynomial of weighted degree at most D on the lines in $\mathcal{R}$. We first record the parameter bounds used later.

Lemma 4.1. *After increasing the absolute constant C in Theorem 1.1, the following hold:*

$$\dim\mathcal{V}_D > q(q+1)\mu(\mathcal{R})(D+1), \qquad D < \min\{q, k_0/2, k_1/2\}, \qquad k_1 \ge \frac{\varepsilon q}{2000},$$
$$\mu(\mathcal{R}) > \frac{D}{q(q+1)}, \qquad \frac{Dq + D(D-1)q + Dq(q+1)}{q^2(q+1)} \le \frac{\varepsilon}{16},$$
$$\frac{1}{\sqrt{q+1}} \le \frac{\varepsilon}{8000},$$

Proof. The bounds on $\mathcal{R}$ give

$$100q \le |\mathcal{R}| \le \frac{6400q}{\varepsilon}, \qquad D \le 1280\left(\sqrt{\frac{dq}{\varepsilon}} + 1\right).$$

Since $q \ge Cd/\varepsilon^3$, taking C sufficiently large gives all the stated degree, incidence-loss, and mixing bounds, as well as $|\mathcal{R}| > D$. Finally,

$$\dim\mathcal{V}_D \ge \frac{D^3}{64d} > |\mathcal{R}|(D+1),$$

using $D^2 > 256d|\mathcal{R}|$. ☐

We now construct a squarefree interpolant of low weighted degree.

Lemma 4.2. *There is a nonzero squarefree polynomial $A \in \mathbb{F}_q[X,Y,Z]$ with* $\mathrm{wdeg}(A) \le D$ *and $A_Z \ne 0$ such that*

$$A(M(t), p_M(t)) \equiv 0 \qquad (M \in \mathcal{R}).$$

Proof. Write $\widetilde{A} \in \mathcal{V}_D$ with unknown coefficients. For each $M \in \mathcal{R}$, the identity $\widetilde{A}(M(t), p_M(t)) \equiv 0$ imposes at most $D+1$ homogeneous linear conditions. Since $\dim \mathcal{V}_D > |\mathcal{R}|(D+1)$, there is a nonzero solution.

It depends on Z, since otherwise Lemma 2.5 would force it to be zero; hence $\widetilde{A}_Z \neq 0$ by the definition of $\mathcal{V}_D$. Choose A of minimum weighted degree among all such interpolants. Then $\mathrm{wdeg}(A) \leq D$, and Lemma 2.3 makes A squarefree. Since each restriction has degree below q and vanishes at all $t \in \mathbb{F}_q$, it vanishes identically. $\square$

By the choice of $\mathcal{R}$, this identity extends to every line of $\mathcal{L}_0$.

Lemma 4.3. *For every $L \in \mathcal{L}_0$,*

$$A(L(t), p_L(t)) \equiv 0.$$

Proof. Let $g_L(t) = A(L(t), p_L(t))$, so $\deg g_L \leq D$. By Lemma 3.2, more than $k_0/2$ points $x \in L$ also lie on an accepted line $M \in \mathcal{R} \setminus \{L\}$. At each such point, $P_L(x) = f(x) = P_M(x)$, so $A(x, P_L(x)) = 0$. Thus g_L has more than $k_0/2 > D$ roots and is identically zero. $\square$

5 Removing Trivial and Degenerate Line Roots

The relation $A(L(t), p_L(t)) \equiv 0$ holds on every line of $\mathcal{L}_0$, but it is not yet sufficient for the lifting argument. To apply Lemma 2.7, we need one relation $B(X, Y, Z) = 0$ such that every remaining $p_L(t)$ solves $B(L(t), Z) = 0$ and every remaining value $f(x) = P_L(x)$ is a simple root of $B(x, Z)$.

This section removes the two obstructions: lines on which A vanishes independently of p_L, and incidences where the relevant root is degenerate. At the end, the hypotheses needed for lifting and propagation hold throughout the remaining graph.

Write

$$A(X, Y, Z) = \sum_j a_j(X, Y) Z^j, \qquad H(X, Y) = \gcd_j a_j(X, Y), \qquad B = \frac{A}{H}.$$

If $H(L(t)) \equiv 0$, then A vanishes above the whole line L and says nothing about p_L. Dividing by H removes this possibility.

The polynomial B is squarefree and has degree at most $D < q$. By Lemma 2.4, choose

$$\mathcal{D} = \lambda_X \partial_X + \lambda_Y \partial_Y + \lambda_Z \partial_Z$$

with $\gcd(B, \mathcal{D}B) = 1$ and define the following sets of lines:

$$\mathcal{L}_{\mathrm{triv}} := \{L \in \mathcal{L}_0 : H(L(t)) \equiv 0\},$$

$$\mathcal{L}_{\mathrm{deg}} := \{L \in \mathcal{L}_0 \setminus \mathcal{L}_{\mathrm{triv}} : (\mathcal{D}B)(L(t), p_L(t)) \equiv 0\}.$$

Lemma 5.1. *For every $L \in \mathcal{L}_0 \setminus \mathcal{L}_{\mathrm{triv}}$,*

$$B(L(t), p_L(t)) \equiv 0.$$

Moreover,

$$|\mathcal{L}_{\mathrm{triv}}| \leq D, \qquad |\mathcal{L}_{\mathrm{deg}}| \leq D(D-1).$$

Proof. Every line in $\mathcal{L}_{\mathrm{triv}}$ contributes a distinct linear factor to H, so there are at most D of them. If $L \notin \mathcal{L}_{\mathrm{triv}}$, then $H(L(t))B(L(t), p_L(t)) \equiv 0$ and the first factor is nonzero. For $L \in \mathcal{L}_{\mathrm{deg}}$, both B and $\mathcal{D}B$ vanish on the same line graph, so Lemma 2.6 gives $|\mathcal{L}_{\mathrm{deg}}| \leq D(D-1)$. $\square$

Delete every accepted incidence on a line in $\mathcal{L}_{\text{triv}} \cup \mathcal{L}_{\text{deg}}$ and every remaining incidence (x, L) with $(\mathcal{D}B)(x, f(x)) = 0$. Their total measure is at most

$$\frac{D}{q(q+1)} + \frac{D(D-1)}{q(q+1)} + \frac{D}{q} \leq \frac{\varepsilon}{16},$$

because the restriction of $\mathcal{D}B$ to every other line is a nonzero polynomial of degree at most D.

After these deletions, repeatedly remove vertices of degree below k_1.

Lemma 5.2. *The pruning process leaves a nonempty graph G_1, with point side X_1, line side $\mathcal{L}_1$, and edge set E_1. It satisfies*

$$\mathsf{mindeg}(G_1) \geq k_1, \qquad \mu(E_1) \geq \frac{\varepsilon}{8}, \qquad D < k_1/2.$$

Moreover:

(i) for every $x \in X_1$,

$$B(x, f(x)) = 0, \qquad B_Z(x, f(x)) \neq 0;$$

(ii) for every $L \in \mathcal{L}_1$,

$$B(L(t), p_L(t)) \equiv 0.$$

Proof. The first deletion leaves incidence measure at least $3\varepsilon/16$. Degree pruning removes less than

$$\frac{k_1}{q+1} + \frac{k_1}{q} < \frac{\varepsilon}{16},$$

so $\mu(E_1) \geq \varepsilon/8$.

Every remaining line satisfies $B(L(t), p_L(t)) \equiv 0$, and every remaining incidence satisfies $(\mathcal{D}B)(x, f(x)) \neq 0$. Hence $B(x, f(x)) = 0$ at every remaining point. To prove $B_Z(x, f(x)) \neq 0$, choose two distinct incident lines, parameterized so that $L(0) = M(0) = x$, with independent directions u_L, u_M. Differentiating their two identities at x gives

$$\langle u_L, \nabla_{X,Y} B \rangle + p_L'(0) B_Z = 0, \qquad \langle u_M, \nabla_{X,Y} B \rangle + p_M'(0) B_Z = 0.$$

If $B_Z = 0$, independence forces $B_X = B_Y = 0$, contradicting $\mathcal{D}B(x, f(x)) \neq 0$. □

6 Constructing and Propagating the Bivariate Polynomial

The goal of this section is to convert the local line roots retained in G_1 into a single degree-d bivariate polynomial. Starting from any point of a connected component, the simple-root condition first allows us to lift the incident line polynomials to a bivariate root Q of B. Initially, this polynomial agrees only with the line polynomials through the starting point. The same simple-root condition then allows us to propagate Q through the entire connected component, producing a single polynomial that agrees with all line polynomials and point values in that component. Finally, a standard application of the expander mixing lemma to the point–line incidence graph gives the desired $\Omega(\varepsilon)$ global agreement with f and proves the $m = 2$ case of Theorem 1.1.

The next lemma is essentially taken from Lemma 3.1 of [HKSS24], stated in the form needed here. It carries out the initial lifting step at a single point. We require a slightly different form of the lemma for our purposes, so we restate its proof below for completeness.

Lemma 6.1. *Suppose $B \in \mathbb{F}_q[X,Y,Z]$ has weighted degree at most D. Let $b \in \mathbb{F}_q^2$ and $\alpha \in \mathbb{F}_q$ satisfy*

$$B(b,\alpha) = 0, \qquad B_Z(b,\alpha) \neq 0.$$

Suppose $\mathcal{U}$ contains more than D distinct lines through b, and each $L \in \mathcal{U}$ carries a polynomial P_L of degree at most d such that

$$P_L(b) = \alpha, \qquad B(L(t), p_L(t)) \equiv 0.$$

Then there is a polynomial $Q \in \mathbb{F}_q[X,Y]$ of total degree at most d such that

$$B(X,Y,Q(X,Y)) \equiv 0 \qquad \textit{and} \qquad Q|_L = P_L \quad (L \in \mathcal{U}).$$

Proof. Translate so that $b = 0$ and parameterize each line as $L(t) = tu_L$. Apply Lemma 2.7 with $k = d$ to obtain Φ_d. For every $L \in \mathcal{U}$, both $\Phi_d(tu_L)$ and $p_L(t)$ have degree at most d, constant term α, and satisfy the same equation modulo t^{d+1}. The uniqueness property in Lemma 2.7 gives $\Phi_d|_L = P_L$.

Set $Q = \Phi_d$ and $R(X,Y) = B(X,Y,Q(X,Y))$. Then $\deg R \leq D$, and $R(L(t)) \equiv 0$ for more than D lines. Hence Lemma 2.5 gives $R = 0$. $\square$

Next, we show that no connected component of G_1 is too small. This follows from the minimum-degree conditions in G_1 and the point–line mixing lemma.

Lemma 6.2. *Let (X',T') be the point and line sides of a connected component of G_1. Then*

$$\mu(X'), \mu(T') \geq \frac{k_1}{4q} \geq \frac{\varepsilon}{8000}.$$

Proof. Let E' be the set of edges in the connected component. Minimum degree gives

$$\mu(E') \geq \frac{k_1}{2q}\mu(X'), \qquad \mu(E') \geq \frac{k_1}{q}\mu(T').$$

By Lemma 2.1,

$$\mu(E') \leq \mu(X')\mu(T') + \frac{\sqrt{\mu(X')\mu(T')}}{\sqrt{q+1}}.$$

Substituting $1/\sqrt{q+1} \leq \varepsilon/8000 \leq k_1/(4q)$ gives both desired bounds. $\square$

It remains to show that a single degree-d polynomial governs all of the line and point labels in such a component. Starting from one point, Lemma 6.1 constructs the polynomial on the incident lines, and simple-root uniqueness propagates it through the component.

Lemma 6.3. *For every connected component (X',T') of G_1, there is a polynomial $Q \in \mathbb{F}_q[X,Y]$ of total degree at most d such that $Q(x) = f(x)$ for all $x \in X'$ and $Q|_L = P_L$ for all $L \in T'$.*

Proof. Choose $b \in X'$ arbitrarily. By Lemmas 5.2 and 6.1, there is a degree-d polynomial Q with $B(X,Y,Q(X,Y)) \equiv 0$ and $Q|_L = P_L$ for every line adjacent to b.

Suppose $Q|_L = P_L$ for a line already reached, and let $(x,L) \in E_1$. Then $Q(x) = f(x)$. If M is another line adjacent to x, parameterize it with $M(0) = x$. The polynomials $p_M(t)$ and $Q(M(t))$ have degree at most d, take the value $f(x)$ at $t = 0$, and satisfy

$$B(M(t), p_M(t)) \equiv 0, \qquad B(M(t), Q(M(t))) \equiv 0.$$

Since $B_Z(x, f(x)) \neq 0$, the uniqueness case in the one-variable version of Lemma 2.7 gives $Q|_M = P_M$. Connectivity propagates this identity to every line of T', and gives $Q(x) = f(x)$ for every $x \in X'$. $\square$

We are now able to prove our main theorem in the dimension $m = 2$ case.

Proof of Theorem 1.1. Choose the absolute constant in Theorem 1.1 larger than all constants required in Lemmas 3.2 and 4.1. Construct G_0 using Lemma 3.1. Choose the line set $\mathcal{R}$ by Lemma 3.2, construct the minimum interpolating polynomial by Lemma 4.2, and use Lemma 4.3 to extend its vanishing identity to every line of G_0. The parameter regime of Lemma 4.1 is used twice, at two different scales: this extension needs $D < k_0/2$, while the lifting and propagation below need the stronger $D < k_1/2$ after the second round of pruning.

Pass from A to B, remove the exceptional lines and incidences described in Section 5, and prune to the subgraph G_1 using Lemma 5.2. Choose any connected component (X', T') of G_1. By Lemma 6.3, one degree-d polynomial Q agrees with f at every point of X'. By Lemma 6.2,

$$\operatorname{agr}(f, Q) \geq \mu(X') \geq \frac{k_1}{4q} \geq \frac{\varepsilon}{8000}.$$

Thus the theorem holds with $c = 1/8000$. □

7 Bootstrapping from Planes to Arbitrary Dimension

Throughout this section let $m \geq 3$, and let $\mathcal{P}_{m,q}$ denote the affine planes in $\mathbb{F}_q^m$. We use the following standard sampler estimates for affine Grassmann graphs; see, for example, [BDN17, BR25].

Lemma 7.1. *For $m \geq 3$ and $q \geq 2$, the following hold.*

(i) *If $a\colon \mathcal{L}_{q,m} \to [0,1]$, $\bar{a} = \mathbb{E}_L[a(L)]$, and $a(P) = \mathbb{E}_{L \subseteq P}[a(L)]$, then*

$$\mathbb{E}_P[(a(P) - \bar{a})^2] \leq \frac{4}{q}.$$

(ii) *If $\mathcal{B} \subseteq \mathcal{P}_{m,q}$ and $X \subseteq \mathbb{F}_q^m$, then*

$$\left| \Pr_{P \in \mathcal{P}_{m,q},\, x \in P}[P \in \mathcal{B},\ x \in X] - \mu(\mathcal{B})\mu(X) \right| \leq \frac{2}{q}\sqrt{\mu(\mathcal{B})\mu(X)}.$$

We use the following theorem of Bhangale and Richelson [BR25].

Theorem 7.2. *There is an absolute constant $C_{\mathrm{BR}} > 0$ such that, if a degree-d plane table T satisfies*

$$\mathrm{Acc}_{\mathrm{PVP}}(T) \geq C_{\mathrm{BR}} \frac{d}{q},$$

then there is a global polynomial $Q \in \mathbb{F}_q[X_1, \ldots, X_m]$ of total degree at most d satisfying

$$\Pr_{P \in \mathcal{P}_{m,q}}[T(P) = Q|_P] \geq \frac{1}{10} \mathrm{Acc}_{\mathrm{PVP}}(T).$$

The consequence we need is the following plane–versus–point statement.

Lemma 7.3. *There are absolute constants $C_{\mathrm{PP}}, c_{\mathrm{PP}} > 0$ such that the following holds for every agreement threshold $\beta \geq C_{\mathrm{PP}}\sqrt{d/q}$. Let $\mathcal{G} \subseteq \mathcal{P}_{m,q}$ have measure at least $1/2$, and suppose that for every $P \in \mathcal{G}$ we are given a degree-d polynomial $T(P)$ satisfying*

$$\Pr_{x \in P}[T(P)(x) = f(x)] \geq \beta.$$

Then there exists a global degree-d polynomial Q such that

$$\Pr_{x \in \mathbb{F}_q^m}[Q(x) = f(x)] \geq c_{\mathrm{PP}}\beta.$$

Proof. Complete the table outside $\mathcal{G}$ by partitioning $\mathcal{P}_{m,q} \setminus \mathcal{G}$ into q classes whose sizes differ by at most one, and assigning a different constant polynomial to each class. Thus every constant is assigned to either

$$\left\lfloor \frac{|\mathcal{P}_{m,q} \setminus \mathcal{G}|}{q} \right\rfloor \quad \text{or} \quad \left\lceil \frac{|\mathcal{P}_{m,q} \setminus \mathcal{G}|}{q} \right\rceil$$

planes. Consequently, if Q is constant, the measure of planes in $\mathcal{P}_{m,q} \setminus \mathcal{G}$ on which $T(P) = Q|_P$ is at most

$$\frac{1}{q} + \frac{1}{|\mathcal{P}_{m,q}|} \le \frac{1}{q} + \frac{1}{q^3} \le \frac{2d}{q}.$$

If Q is nonconstant, then Q can agree with at most a $2d/q$ fraction of the planes in $\mathcal{P}_{m,q} \setminus \mathcal{G}$ by the Schwartz–Zippel lemma. Hence, for every global degree-d polynomial Q, the measure of planes on which it can agree with the table T coming from $\mathcal{P}_{m,q} \setminus \mathcal{G}$ is at most $2d/q$.

The completed table passes the plane–versus–point test with probability at least $\beta/2$. We show that it passes the plane–versus–plane test with probability $\Omega(\beta^2)$. Choose two random distinct planes meeting in a line and then a random point on that line. The probability that both plane polynomials equal f at this point is at least $\beta^2/4 - 1/(q^2+q)$. If their restrictions to the line are distinct, they agree at no more than d points. Hence

$$\mathrm{Acc}_{\mathrm{PVP}}(T) \ge \frac{\beta^2}{4} - \frac{d}{q} - \frac{1}{q^2+q} = \Omega(\beta^2).$$

For sufficiently large C_{PP}, Theorem 7.2 therefore gives a global polynomial Q agreeing with T on a plane set $\mathcal{M}$ of measure $\Omega(\beta^2)$. Since only $O(d/q) = o(\beta^2)$ of these matches can come from the completion, $\mathcal{M} \cap \mathcal{G}$ still has measure $\Omega(\beta^2)$.

Let $X_Q = \{x : Q(x) = f(x)\}$. Every plane in $\mathcal{M} \cap \mathcal{G}$ contains at least a β-fraction of points from X_Q. Applying Lemma 7.1(ii) gives

$$\beta \le \mu(X_Q) + O\left(\frac{1}{q\beta}\right).$$

The error is at most $\beta/2$ once C_{PP} is sufficiently large, which proves the claim. ☐

7.1 Proof in higher dimensions

Proof of Theorem 1.1 for $m \ge 3$. Let $C_0, c_0 > 0$ be the constants established for the bivariate case, and enlarge the constant C in Theorem 1.1 as needed.

1. Decode almost every plane. For an affine plane P, let

$$\alpha_P := \Pr_{x \in P,\, L \subseteq P,\, x \in L}[P_L(x) = f(x)].$$

Then $\mathbb{E}_P[\alpha_P] \ge \varepsilon$. By Lemma 7.1(i), $\Pr_P[\alpha_P < \varepsilon/2] \le \frac{16}{q\varepsilon^2}$. Thus, after increasing C, the set

$$\mathcal{G} := \{P : \alpha_P \ge \varepsilon/2\}$$

has measure at least 1/2. Moreover, every $P \in \mathcal{G}$ satisfies $\alpha_P \ge C_0(d/q)^{1/3}$. Applying the bivariate case of Theorem 1.1 inside P, choose a degree-d polynomial $T(P)$ such that

$$\Pr_{x \in P}[T(P)(x) = f(x)] \ge c_0 \alpha_P \ge \frac{c_0 \varepsilon}{2}.$$

2. Complete the plane table and apply the plane–versus–point theorem. Set $\beta = c_0\varepsilon/2$. Since

$$\varepsilon \geq C\left(\frac{d}{q}\right)^{1/3} \geq C\left(\frac{d}{q}\right)^{1/2},$$

we have $\beta \geq C_{\mathrm{PP}}\sqrt{d/q}$ after increasing C once more. Apply Lemma 7.3 to the plane polynomials constructed on $\mathcal{G}$. It produces a global degree-d polynomial Q satisfying

$$\Pr_{x\in\mathbb{F}_q^m}[Q(x) = f(x)] \geq c_{\mathrm{PP}}\beta = \Omega(\varepsilon),$$

which proves Theorem 1.1 for $m \geq 3$. □

References

[AKK+05] Noga Alon, Tali Kaufman, Michael Krivelevich, Simon Litsyn, and Dana Ron. Testing Reed–Muller codes. *IEEE Transactions on Information Theory*, 51(11):4032–4039, 2005.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[AS03] Sanjeev Arora and Madhu Sudan. Improved low-degree testing and its applications. *Combinatorica*, 23(3):365–426, 2003.

[BDN17] Amey Bhangale, Irit Dinur, and Inbal Livni Navon. Cube versus cube low-degree test. In *Proceedings of the 8th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 67 of *LIPIcs*, pages 40:1–40:31, 2017.

[BKS+10] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed–Muller codes. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 488–497, 2010.

[BR25] Amey Bhangale and Silas Richelson. Plane versus plane low-degree test. Technical Report TR25-126, Electronic Colloquium on Computational Complexity, 2025.

[Bür00] Peter Bürgisser. *Completeness and Reduction in Algebraic Complexity Theory*, volume 7 of *Algorithms and Computation in Mathematics*. Springer, 2000.

[DKK+21] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. *Israel Journal of Mathematics*, 243(1):377–420, 2021.

[DKK+25] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture. *Theory of Computing*, 21(11):1–50, 2025.

[FS95] Katalin Friedl and Madhu Sudan. Some improvements to total degree tests. In *Proceedings of the 3rd Israel Symposium on the Theory of Computing Systems*, pages 190–198, 1995.

[HKSS24] Prahladh Harsha, Mrinal Kumar, Ramprasad Saptharishi, and Madhu Sudan. An improved line-point low-degree test. In *Proceedings of the 65th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1883–1892, 2024.

[HSS13] Elad Haramaty, Amir Shpilka, and Madhu Sudan. Optimal testing of multivariate polynomials over small prime fields. *SIAM Journal on Computing*, 42(2):536–562, 2013.

[KMS23] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in grassmann graph have near-perfect expansion. *Annals of Mathematics*, 198(1):1–92, 2023.

[KMS25] Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and grassmann graphs. *Theory of Computing*, 21(10):1–55, 2025.

[KR06] Tali Kaufman and Dana Ron. Testing polynomials over general fields. *SIAM Journal on Computing*, 36(3):779–802, 2006.

[MR10] Dana Moshkovitz and Ran Raz. Two-query PCP with subconstant error. *Journal of the ACM*, 57(5):29:1–29:29, 2010.

[MZ23] Dor Minzer and Kai Zhe Zheng. Approaching the soundness barrier: A near optimal analysis of the cube versus cube test. In *Proceedings of the 34th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pages 2761–2776, 2023.

[MZ24] Dor Minzer and Kai Zhe Zheng. Near optimal alphabet–soundness tradeoff PCPs. Technical Report TR24-027, Electronic Colloquium on Computational Complexity, 2024. Revised 2026.

[MZ25] Dor Minzer and Kai Zhe Zheng. Improved round-by-round soundness IOPs via Reed–Muller codes. In *Proceedings of the 66th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1286–1294, 2025.

[MZ26] Dor Minzer and Kai Zhe Zheng. Near optimal hardness of approximating k-CSP. In Aditya Bhaskara and Artur Czumaj, editors, *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC)*, pages 72–83. ACM, 2026.

[PS94] Alexander Polishchuk and Daniel A. Spielman. Nearly-linear size holographic proofs. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC)*, pages 194–203, 1994.

[RS96] Ronitt Rubinfeld and Madhu Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.

[RS97] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*, pages 475–484, 1997.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il