Electronic Colloquium on Computational Complexity, Report No. 74 (2025)

Approximate Polymorphisms of Predicates

Yaroslav Alekseev[1] and Yuval Filmus[1,2]

[1]Taub Faculty of Computer Science, Technion Israel Institute of Technology, Haifa, Israel
[2]Faculty of Mathematics, Technion Israel Institute of Technology, Haifa, Israel

June 13, 2025

Abstract

A generalized polymorphism of a predicate $P \subseteq \{0,1\}^m$ is a tuple of functions $f_1, \ldots, f_m \colon \{0,1\}^n \to \{0,1\}$ satisfying the following property: *If $x^{(1)}, \ldots, x^{(m)} \in \{0,1\}^n$ are such that $(x_i^{(1)}, \ldots, x_i^{(m)}) \in P$ for all i, then also $(f_1(x^{(1)}), \ldots, f_m(x^{(m)})) \in P$.*

We show that if $f_1, \ldots, f_m$ satisfy this property for *most* $x^{(1)}, \ldots, x^{(m)}$ (as measured with respect to an arbitrary full support distribution μ on P), then $f_1, \ldots, f_m$ are close to a generalized polymorphism of P (with respect to the marginals of μ).

Our main result generalizes several results in the literature:

- Linearity testing (Blum, Luby, and Rubinfeld): $P = \{(0,0,0), (0,1,1), (1,0,1), (1,1,0)\}$.
- Quantitative Arrow theorems (Kalai; Keller; Mossel): $P = \{x \in \{0,1\}^3 : x \neq (0,0,0), (1,1,1)\}$.
- Approximate intersecting families (Friedgut and Regev): $P = \{(0,0), (0,1), (1,0)\}$.
- AND testing (Filmus, Lifshitz, Minzer, and Mossel): $P = \{(0,0,0), (0,1,0), (1,0,0), (1,1,1)\}$.
- f-testing (Chase, Filmus, Minzer, Mossel, and Saurabh): $P = \{(x, f(x)) : x \in \{0,1\}^m\}$.

In particular, we extend linearity testing to arbitrary distributions.

We use our techniques to significantly improve the parameter dependence in the work of Friedgut and Regev on approximately intersecting families, from tower type to exponential.

We also extend our results to predicates on arbitrary finite alphabets in which all coordinates are "flexible" (for each coordinate j there exists $w \in P$ such that $w^{j \leftarrow \sigma} \in P$ for all σ).

ISSN 1433-8092

Contents

1 Introduction

The classical BLR linearity test [BLR90] states that if $f\colon \{0,1\}^n \to \{0,1\}$ satisfies $\Pr_{x,y\sim\mu_{1/2}}[f(x) \oplus f(y) = f(x \oplus y)] \geq 1-\epsilon$ then there exists a linear function $g\colon \{0,1\}^n \to \{0,1\}$ satisfying $\Pr_{x\sim\mu_{1/2}}[g(x) \neq f(x)] \leq \epsilon$; here $\mu_{1/2}$ is the uniform distribution over $\{0,1\}^n$. Linearity testing is a special case of a more general problem, *approximate polymorphisms*, introduced in [CFM+22], which unifies several other related results in the literature, including Kalai's quantitative Arrow theorem [Kal02] and AND testing [FLMM20].

Definition 1.1 (Polymorphism)**.** A function $f\colon \{0,1\}^n \to \{0,1\}$ is a *polymorphism* of a predicate $P \subseteq \{0,1\}^m$ if the following holds: *Given any vectors $x^{(1)},\dots,x^{(m)} \in \{0,1\}^n$ such that $(x_i^{(1)},\dots,x_i^{(m)}) \in P$ for all $i \in [n]$, also $(f(x^{(1)}),\dots,f(x^{(m)})) \in P$.*

Given a probability distribution μ on P, the function f is a (μ,ϵ)-*approximate polymorphism of P* if

$$\Pr_{\substack{(x_1^{(1)},\dots,x_1^{(m)})\sim P\\ \cdots \\ (x_n^{(1)},\dots,x_n^{(m)})\sim P}}[(f(x^{(1)}),\dots,f(x^{(m)})) \in P] \geq 1-\epsilon.$$

Using this terminology, we can reformulate the BLR linearity test as follows:

Theorem 1.2 (Linearity testing)**.** *Let $P_\oplus = \{(a,b,a\oplus b) : a,b \in \{0,1\}\}$, and let $\mu_\oplus$ be the uniform distribution over $P_\oplus$.*

If $f\colon \{0,1\}^n \to \{0,1\}$ is a $(\mu_\oplus,\epsilon)$-approximate polymorphism of $P_\oplus$ then there exists a polymorphism $g\colon \{0,1\}^n \to \{0,1\}$ of $P_\oplus$ such that $\Pr_{x\sim\mu_{1/2}}[g(x) \neq f(x)] \leq \epsilon$.

In this paper, we extend Theorem 1.2 to arbitrary predicates $P \subseteq \{0,1\}^m$ (for all m) and to arbitrary distributions μ on P with full support.

One might hope for a result of the following form: *Any approximate polymorphism of P is close to a polymorphism of P.* While this holds for some predicates P, it fails for others, as the following counterexample from [CFM+22] demonstrates.

Example 1.3. Let $P_{\mathsf{NAND}} = \{(a,b,\overline{a\wedge b}) : a,b \in \{0,1\}\}$ and let μ_{NAND} be the uniform distribution over P_{NAND}. One can show that the only polymorphisms of P_{NAND} are dictators: $f(x) = x_i$.

For large n, let $f\colon \{0,1\}^n \to \{0,1\}$ be the following function:

$$f(x) = \begin{cases} x_1 \wedge x_2 & \text{if } x_1 + \cdots + x_n \leq 0.6n,\\ x_1 \vee x_2 & \text{otherwise.}\end{cases}$$

If $x,y \sim \mu_{1/2}(\{0,1\}^n)$ then $x_1+\cdots+x_n \approx n/2$ and $y_1+\cdots+y_n \approx n/2$, while $\overline{x_1\wedge y_1}+\cdots+\overline{x_n\wedge y_n} \approx (3/4)n$. Since

$$\overline{(x_1 \wedge x_2) \wedge (y_1 \wedge y_2)} = \overline{x_1\wedge y_1} \vee \overline{x_2 \wedge y_2},$$

the function f is a $(\mu_{\mathsf{NAND}}, o(1))$-approximate polymorphism of P_{NAND}. However, f is not close to any exact polymorphism of P_{NAND}.

While we cannot guarantee that an approximate polymorphism of P is close to a polymorphism of P, we are able to guarantee that it is close to a *generalized polymorphism* of P.

Definition 1.4 (Generalized polymorphism)**.** A tuple of functions $f_1,\dots,f_m\colon \{0,1\}^n \to \{0,1\}$ is a *generalized polymorphism* of a predicate $P \subseteq \{0,1\}^m$ if the following holds: *Given any vectors $x^{(1)},\dots,x^{(m)} \in \{0,1\}^n$ such that $(x_i^{(1)},\dots,x_i^{(m)}) \in P$ for all $i \in [n]$, also $(f_1(x^{(1)}),\dots,f_m(x^{(m)})) \in P$.*

Given a probability distribution μ on P, the tuple $f_1,\dots,f_m$ is a (μ,ϵ)-*approximate generalized polymorphism of P* if

$$\Pr_{(x_i^{(1)},\dots,x_i^{(m)})\sim\mu}[(f_1(x^{(1)}),\dots,f_m(x^{(m)})) \in P] \geq 1-\epsilon,$$

where each tuple $(x_i^{(1)},\dots,x_i^{(m)})$ is sampled independently according to μ.

Using this concept, we can state our main theorem:

Theorem 1.5 (Main). *Let $P \subseteq \{0,1\}^m$ be a non-empty predicate, and let μ be a distribution on P with full support. For every $\epsilon > 0$ there exists $\delta > 0$ such that the following holds for all n.*

If $f_1, \ldots, f_m \colon \{0,1\}^n \to \{0,1\}$ is a (μ, δ)-approximate generalized polymorphism of P then there exists a generalized polymorphism $g_1, \ldots, g_m \colon \{0,1\}^n \to \{0,1\}$ of P such that

$$\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon \text{ for all } j \in [m],$$

where $\mu|_j$ is the marginal distribution of the j'th coordinate.

In the remainder of the introduction, we first state several other results that we prove in the paper (Section 1.1), and then review some of the relevant literature (Section 1.2).

1.1 Other results

Theorem 1.5 has an unspecified dependence between δ and ϵ; the dependence arising from our current proof is of tower type. In two special cases we can improve this dependence.

Theorem 1.6 (Linearity testing for general distributions). *Let $P_{m,b} = \{(a_1, \ldots, a_m) \in \{0,1\}^m : a_1 \oplus \cdots \oplus a_m = b\}$, where $m \geq 3$ and $b \in \{0,1\}$, and let μ be a distribution on $P_{m,b}$ with full support. The following holds for $\delta = \Theta(\epsilon)$.*

If $f_1, \ldots, f_m \colon \{0,1\}^n \to \{0,1\}$ is a (μ, δ)-approximate generalized polymorphism of $P_{m,b}$ then there exists a generalized polymorphism $g_1, \ldots, g_m \colon \{0,1\}^n \to \{0,1\}$ of $P_{m,b}$ such that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon$ for all $j \in [m]$.

Moreover, there exist a set $S \subseteq [m]$ and $b_1, \ldots, b_m \in \{0,1\}$ such that $g_j(x) = \bigoplus_{i \in S} x_i \oplus b_j$.

Furthermore, if $f_i = f_j$ and $\mu|_i = \mu|_j$ then $g_i = g_j$.

Theorem 1.7 (Monotone case). *Let $P \subseteq \{0,1\}^m$ be a non-empty monotone predicate: if $x \in P$ and $y \leq x$ (pointwise) then $y \in P$. Let μ be a distribution on P with full support. There exists a constant $C = C(P, \mu)$ such that the following holds for $\delta = 1/\exp\Theta(1/\epsilon^C)$.*

If $f_1, \ldots, f_m \colon \{0,1\}^n \to \{0,1\}$ is a (μ, δ)-approximate generalized polymorphism of P then there exists a generalized polymorphism $g_1, \ldots, g_m \colon \{0,1\}^n \to \{0,1\}$ of P such that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon$ for all $j \in [m]$. Moreover, $g_j \leq f_j$ pointwise for all $j \in [m]$.

Furthermore, if $f_i = f_j$ and $\mu|_i = \mu|_j$ then $g_i = g_j$.

In both cases we get the additional guarantee that if $f_i = f_j$ and the corresponding marginals of μ coincide, then $g_i = g_j$. This implies that if all marginals of μ are identical then an approximate polymorphism is close to an exact polymorphism.

Intersecting families Friedgut and Regev [FR18] proved the following result on almost intersecting families (see also [DT16]).

Theorem 1.8 (Friedgut–Regev). *Fix $0 < p < 1/2$. For every $\epsilon > 0$ there exist $\delta > 0$ and $j \in \mathbb{N}$ such that the following holds for all n such that pn is an integer.*

If $\mathcal{F} \subseteq \binom{[n]}{pn}$ contains a δ-fraction of the edges of the Kneser graph then there exists an intersecting family $\mathcal{G} \subseteq \binom{[n]}{pn}$ depending on j points such that $|\mathcal{F} \setminus \mathcal{G}| \leq \epsilon\binom{n}{pn}$.

The dependence of δ, j on ϵ is of tower type. We improve the dependence to exponential, at the cost of considering shallow decision trees rather than juntas. (We can convert the decision tree to a junta, losing another exponential in the size of the junta.)

Theorem 1.9 (Improved Friedgut–Regev). *Fix $0 < p < 1/2$. For every $\epsilon > 0$ the following holds for all n such that pn is an integer.*

If $\mathcal{F} \subseteq \binom{[n]}{pn}$ contains a $1/\exp\Theta(1/\epsilon^C)$-fraction of the edges of the Kneser graph then there exists an intersecting family $\mathcal{G} \subseteq \binom{[n]}{pn}$ computed by a decision tree of depth $O(1/\epsilon^C)$ (for some global constant C) such that $|\mathcal{F} \setminus \mathcal{G}| \leq \epsilon\binom{n}{pn}$.

Polymorphisms over larger alphabets So far, we have considered predicates over the binary alphabet. However, the notion of polymorphisms makes sense for every finite alphabet. We conjecture that Theorem 1.5 extends to this setting. While we are unable to prove this conjecture in full generality, we are able to prove the following special case.

Theorem 1.10 (Larger alphabets)**.** *Let Σ be a finite set, let $P \subseteq \Sigma^m$, and let μ be a distribution on P with full support. Suppose that for each $j \in [m]$ there exists $w \in P$ such that w remains in P even if we modify its j'th coordinate arbitrarily. For every $\epsilon > 0$ there exists $\delta > 0$ such that the following holds for all n.*

If $f_1, \ldots, f_m : \Sigma^n \to \Sigma$ is a (μ, δ)-approximate generalized polymorphism of P then there exists a generalized polymorphism $g_1, \ldots, g_m : \Sigma^n \to \Sigma$ of P such that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon$ for all $j \in [m]$.

Input/output predicates The proofs of Theorems 1.7 and 1.10 immediately generalize to a setting which involves two predicates P, Q. Given two predicates $P \subseteq \Sigma^m$ and $Q \subseteq \Delta^m$, a tuple $f_1, \ldots, f_m : \Sigma^n \to \Delta$ is a (P, Q)-generalized polymorphism if the following holds: *Given any vectors $x^{(1)}, \ldots, x^{(m)} \in \Sigma^n$ such that $(x_i^{(1)}, \ldots, x_i^{(m)}) \in P$ for all $i \in [n]$, we have $(f_1(x^{(1)}), \ldots, f_m(x^{(m)})) \in Q$.* This setting arises natural in the study of promise CSPs [AGH17, BG21, BBKO21], where such polymorphisms are often called *weak polymorphisms*.

In the case of Theorem 1.7, we require P to be monotone but Q can be arbitrary, and similarly, in the case of Theorem 1.10, the stated condition need only hold for P. While we do not work out these generalizations explicitly, they follow immediately from the proofs. In contrast, the proof of Theorem 1.5 does rely on the assumption $P = Q$.

1.2 Related work

Arrow's theorem Arrow's celebrated theorem [Arr50] can be expressed in the language of polymorphisms. Let $m \geq 3$. For every permutation $\pi \in S_m$, let $I(\pi) \in \{0,1\}^{\binom{m}{2}}$ be the following vector: $I(\pi)(i,j) = [\pi(i) > \pi(j)]$. Arrow's theorem states that if $f_{i,j}$ is a generalized polymorphism of $P_m := \{I(\pi) : \pi \in S_m\}$ and each $f_{i,j}$ is *unanimous* (satisfies $f_{i,j}(b, \ldots, b) = b$ for $b \in \{0,1\}$) then there exists $k \in [n]$ such that $f_{i,j}(x) = x_k$ for all i, j.

Kalai [Kal02] considered the case $m = 3$. He showed that if $f_{1,2}, f_{2,3}, f_{3,1}$ is a (μ, ϵ)-approximate polymorphism of P_3, where μ is the uniform distribution, and furthermore the $f_{i,j}$ are *balanced* (satisfy $\Pr_{\mu_{1/2}}[f_{i,j} = 1] = 1/2$), then there exists $k \in [n]$ such that either $f_{i,j}(x) = x_k$ for all i, j, or $f_{i,j} = 1 - x_k$ for all i, j. In this case the predicate P_3 consists of all triples (a, b, c) such that a, b, c are not all equal. Keller [Kel10] extended Kalai's result to arbitrary $m \geq 3$ and to arbitrary distributions, under various (severe) restrictions on the $f_{i,j}$.

Mossel [Mos12a] determined all generalized polymorphisms of P_m for all $m \geq 3$ (without assuming unanimity), and proved that approximate generalized polymorphisms of these predicates (with respect to the uniform distribution) are close to exact generalized polymorphisms. His techniques in fact work for arbitrary distributions (the missing piece, reverse hypercontractivity for arbitrary distributions, was proved in [MOS13]). Keller [Kel10] improved on Mossel's result in the case of the uniform distribution by determining the optimal dependence between ϵ and δ.

Linearity testing Blum, Luby and Rubinfeld [BLR90] were the first to propose the BLR test. They analyzed it using self-correction. Bellare et al. [BCH+96] later gave a different argument using Fourier analysis. The test was generalized to arbitrary prime fields by Kiwi [Kiw03].

David et al. [DDG+17] extended the BLR test to the setting of constant weight inputs, which is analogous to the setting of Theorem 1.9. Translated to the setting of Theorem 1.6, they extended the BLR test to a natural distribution μ with $\Pr[\mu|_i = 1] = p$ for all i. Dinur et al. [DFH25] analyzed the related "affine test" $f(x \oplus y \oplus z) = f(x) \oplus f(x \oplus y) \oplus f(x \oplus z)$ for the same distribution. Both of these works used *agreement theorems* to reduce the biased case to the unbiased case: David et al. used the agreement theorem of Dinur and Steurer [DS14], and Dinur et al. proved their own agreement theorem, which we also use to prove Theorem 1.6.

Functional predicates Filmus et al. [FLMM20], prompted by work of Nehama [Neh13] on judgment aggregation, proved Theorem 1.5 for the predicate $P_\wedge = \{(a, b, a \wedge b) : a, b \in \{0,1\}\}$ and various distributions. Their analysis combined Bourgain's tail bound [Bou02, KKO18] with a study of the "one-sided noise operator".

Chase et al. [CFM$^+$22] proved Theorem 1.5 for all predicates of the form $P_f = \{(a_1, \ldots, a_m, f(a_1, \ldots, a_m)) : a_1, \ldots, a_m \in \{0,1\}\}$ for an arbitrary function $f\colon \{0,1\}^m \to \{0,1\}$, with respect to the uniform distribution. They asked whether Theorem 1.5 extends to arbitrary predicates, a question we answer in the affirmative in this paper. Their general approach was similar to the approach taken in this paper, combining Jones' regularity lemma with the It Ain't Over Till It's Over theorem.

Paper organization

We give an outline of our proof techniques in Section 2. We prove Theorem 1.7 in Section 3, Theorem 1.5 in Section 4, Theorem 1.6 in Section 5, Theorem 1.9 in Section 6, and Theorem 1.10 in Section 7. Our proofs require several versions of Jones' regularity lemma [Jon16], proved in Section 8. We close the paper in Section 9 with a few open questions.

Acknowledgments This research was supported by ISF grant no. 507/24.

2 Proof outline

In this section we give a brief overview of the proof of Theorem 1.5. We start with the proof in the monotone case (corresponding to Theorem 1.7), which is more intuitive, and then describe the general case.

2.1 Monotone case

Triangle removal lemma The proof of Theorem 1.7 has the same general outline as the proof of the triangle removal lemma [RS78] using Szémeredi's regularity lemma [Sze78].

Theorem 2.1 (Triangle removal lemma)**.** *For every $\epsilon > 0$ there exists $\delta > 0$ such that the following holds for all n.*

If G is a graph on n vertices with at most $\delta\binom{n}{3}$ triangles, then we can make G triangle-free by removing at most $\epsilon\binom{n}{2}$ edges.

The proof requires Szémeredi's regularity lemma, which we state in a qualitative fashion.

Theorem 2.2 (Regularity lemma for graphs)**.** *For every $\epsilon > 0$ there is $M > 1/\epsilon$ such that the following holds for all graphs G.*

*The vertex set of G can be partitioned into M parts $V_1, \ldots, V_M$ of almost equal size such that all but an ϵ-fraction of pairs (i, j) are ϵ-*regular*: there exists $p_{i,j} \in [0,1]$ such that the edges of G connecting V_i to V_j "behave like" a random bipartite graph with edge density $p_{i,j}$, up to an error ϵ.*

Given the regularity lemma, the proof of the triangle removal lemma is quite simple. Given a graph G, we apply the regularity lemma with an appropriate parameter $\eta > 0$, obtaining a partition $V_1, \ldots, V_M$ in which all but an η-fraction of pairs is η-regular. We remove all edges between V_i and V_j if either (i) $i = j$, or (ii) the pair (i, j) is not η-regular, or (iii) $p_{i,j} \le \eta$. In total, we have removed at most an $(1/M + 2\eta)$-fraction of edges. Choosing η so that $\eta = \epsilon/3$, this fits within our budget.

For an appropriate choice of $\delta > 0$, the resulting graph is triangle-free. Indeed, any remaining triangle must involve three different parts V_i, V_j, V_k. Since the triangle survived the pruning process, we must have $p_{i,j}, p_{i,k}, p_{j,k} \ge \eta$. The *triangle counting lemma* shows that the subgraph of G induced by V_i, V_j, V_k contains $\Omega(\eta^3)$ of the possible triangles between these vertices, and so G contains $\Omega((\eta/M)^3)\binom{n}{3}$ triangles. Choosing δ appropriately, this contradicts the assumption on G.

Approximate polymorphisms of monotone predicates The proof of Theorem 1.7 follows the same plan. Szémeredi's regularity lemma is replaced by Jones' regularity lemma [Jon16].

Theorem 2.3 (Jones' regularity lemma). *For every $\epsilon, \tau > 0$, $d \in \mathbb{N}$, and $p \in (0,1)$ there exists $M \in \mathbb{N}$ such that the following holds for all n and all functions $f\colon \{0,1\}^n \to \{0,1\}$.*

There exists a set $J \subseteq [n]$ of size at most M such that

$$\Pr_{x \sim \mu_p}[f|_{J \leftarrow x} \text{ is not } (d,\tau)\text{-regular with respect to } \mu_p] \leq \epsilon,$$

where μ_p is the product distribution on $\{0,1\}^J$ with $\Pr[x_j = 1] = p$, and a function $g\colon \{0,1\}^{J^c} \to \{0,1\}$ is (d,τ)-regular if $\mathrm{Inf}_i[g^{\leq d}] \leq \tau$ for all $i \in J^c$, where the influence is computed with respect to μ_p.

Jones' lemma is proved by a simple potential function argument. In applications we need a variant of this lemma for several functions, and also allowing an initial set as a starting point. We also need two extensions of the lemma, one to functions $f\colon \{0,1\}^n \to [0,1]$, for proving Theorem 1.9, and another to functions $f\colon \Sigma^n \to \Sigma$, for proving Theorem 1.10. Moreover, in order to obtain better parameters, in the full proof of Theorem 1.7 we use a version of Jones' lemma which approximates f by a decision tree rather than by a junta. We prove all of these different versions of Jones' regularity lemma in Section 8.

The notion of regularity promised by Jones' regularity lemma is geared toward a result known as *It Ain't Over Till It's Over* [MOO10], which is key to our counting lemma.

Theorem 2.4 (It Ain't Over Till It's Over). *For every $p, q \in (0,1)$ and $\epsilon > 0$ there exist parameters $d \in \mathbb{N}$ and $\tau, \delta > 0$ such that the following holds for all n and all functions $f\colon \{0,1\}^n \to \{0,1\}$ which are (d,τ)-regular with respect to μ_p.*

Let ρ be a random restriction obtained as follows: for each coordinate independently, leave it free with probability q, and otherwise sample it according to μ_p. If $\mathbb{E}_{\mu_p}[f] \geq \epsilon$ then

$$\Pr_{\rho}[\mathop{\mathbb{E}}_{\mu_p}[f|_\rho] \geq \delta] \geq 1 - \epsilon.$$

Applying the theorem to both f and $1-f$, it implies that if f is regular and its expectation is in $[\epsilon, 1-\epsilon]$, then even if we sample all but a q-fraction of inputs, its expectation still lies in $[\delta, 1-\delta]$ (where δ could be much smaller than ϵ), with probability $1-\epsilon$. This explains its moniker.

We outline the proof of Theorem 1.7 in the special case of the predicate $P_{\mathsf{NAND}} = \{(0,0),(1,0),(0,1)\}$; the general case involves no further complications. Let us recall what we would like to prove.

Theorem 2.5 (NAND testing). *Let μ be a distribution on P_{NAND} with full support. For every $\epsilon > 0$ there exists $\delta > 0$ such that the following holds for all n.*

If $f_1, f_2\colon \{0,1\}^n \to \{0,1\}$ is a (μ,δ)-approximate generalized polymorphism of P_{NAND} then there exists a generalized polymorphism $g_1, g_2\colon \{0,1\}^n \to \{0,1\}$ of P_{NAND} such that

$$\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon \text{ for } j \in \{1,2\}.$$

Furthermore, if $f_1 = f_2$ and $\mu|_1 = \mu|_2$ then $g_1 = g_2$.

Suppose we are given f_1, f_2. Apply Jones' regularity lemma with $p := \mu|_1$, $\epsilon := \epsilon/4$, and appropriate d, τ to f_1 and with $p := \mu|_2$ to f_2, obtaining a set J of size at most M, for a parameter M depending on μ, ϵ, d, τ. Then

$$\Pr_{(x^{(1)},x^{(2)}) \sim \mu^J}[f_1|_{J \leftarrow x^{(1)}} \text{ and } f_2|_{J \leftarrow x^{(2)}} \text{ are } (d,\tau)\text{-regular}] \geq 1 - \epsilon/2.$$

We define g_1, g_2 by zeroing outputs of f_1, f_2 as follows. We set $g_j|_{J \leftarrow x^{(j)}} \equiv 0$ if either (i) $f_j|_{J \leftarrow x^{(j)}}$ is not regular, or (ii) $\mathbb{E}[f_j|_{J \leftarrow x^{(j)}}] \leq \epsilon/2$. By construction, $\Pr[g_j \neq f_j] \leq \epsilon$. Also, if $f_1 = f_2$ and $\mu|_1 = \mu|_2$ then $g_1 = g_2$.

For an appropriate choice of δ, the resulting pair (g_1, g_2) is a generalized polymorphism of P_{NAND}. Indeed, if this is not the case, then there exist some $x^{(1)}, x^{(2)} \in \{0,1\}^J$ such that $g_1|_{J \leftarrow x^{(1)}}, g_2|_{J \leftarrow x^{(2)}}$ are both non-zero. If this happens then $g_1|_{J \leftarrow x^{(1)}}, g_2|_{J \leftarrow x^{(2)}}$ are both (d, τ)-regular and their expectations are at least $\epsilon/2$. We will prove a counting lemma which states that for an appropriate choice of d, τ,

$$\Pr_{(y^{(1)}, y^{(2)}) \sim \mu^{J^c}} [g_1|_{J \leftarrow x^{(1)}}(y^{(1)}) = g_2|_{J \leftarrow x^{(2)}}(y^{(2)}) = 1] \geq \gamma,$$

where γ depends on μ and ϵ. Since g_1, g_2 agree with f_1, f_2 on these inputs, this implies that

$$\Pr_{(z^{(1)}, z^{(2)}) \sim \mu^n} [f_1(z^{(1)}) = f_2(z^{(2)}) = 1] \geq \min(\mu)^M \gamma, \text{ where } \min(\mu) = \min_{w \in P_{\mathsf{NAND}}} \mu(w).$$

Choosing δ to be smaller than the right-hand side, we obtain a contradiction to the assumption that (f_1, f_2) is a (μ, δ)-approximate generalized polymorphism.

It remains to prove the counting lemma.

Lemma 2.6 (Counting lemma for NAND)**.** *Let μ be a distribution on P_{NAND} with full support. For every $\epsilon > 0$ there exist $d \in \mathbb{N}$ and $\tau, \gamma > 0$ such that the following holds for all n.*

If $\phi_1, \phi_2 \colon \{0,1\}^n \to \{0,1\}$ are (d, τ)-regular (with respect to $\mu|_j$) and have expectation at least $\epsilon/2$ (with respect to $\mu|_j$) then

$$\Pr_{(y^{(1)}, y^{(2)}) \sim \mu^n} [\phi_1(y^{(1)}) = \phi_2(y^{(2)}) = 1] \geq \gamma.$$

In order to prove this lemma, we will sample μ in two steps. In the first step, we sample a restriction ρ:

$$\rho = \begin{cases} (0,0) & \text{w.p. } (1-q)\mu(0,0) - q, \\ (1,0) & \text{w.p. } (1-q)\mu(1,0), \\ (0,1) & \text{w.p. } (1-q)\mu(0,1), \\ (*,0) & \text{w.p. } q, \\ (0,*) & \text{w.p. } q. \end{cases}$$

We choose the parameter $q > 0$ to be such that all probabilities are positive.

In the second step, if we sampled $(*, 0)$, then we sample the first coordinate using $\mu|_1$. Similarly, if we sampled $(0, *)$, then we sample the second coordinate using $\mu|_2$. The end result has the same distribution as μ by design. Also, the distribution of $\rho|_1$ given that $\rho|_1 \neq *$ is the same as $\mu|_1$, and similarly for $\rho|_2$.

We apply It Ain't Over Till It's Over to the function ϕ_1 with $\epsilon := \epsilon/2$ and $q := q$ to obtain d_1, τ_1, δ_1, and to the function ϕ_2 to obtain d_2, τ_2, δ_2. Choosing $d = \max(d_1, d_2)$ and $\tau = (\tau_1, \tau_2)$, It Ain't Over Till It's Over implies that

$$\Pr_{(u^{(1)}, u^{(2)}) \sim \rho^n} [\mathop{\mathbb{E}}_{\mu|_1} [\phi_1|_{\rho|_1}] \geq \delta_1 \text{ and } \mathop{\mathbb{E}}_{\mu|_2} [\phi_2|_{\rho|_2}] \geq \delta_2] \geq 1 - \epsilon.$$

After apply the restriction ρ, the events $\phi_1|_{\rho|_1} = 1$ and $\phi_2|_{\rho|_2} = 1$ are independent. It follows that

$$\Pr_{(y^{(1)}, y^{(2)}) \sim \mu^n} [\phi_1(y^{(1)}) = \phi_2(y^{(2)}) = 1] \geq (1 - \epsilon)\delta_1 \delta_2.$$

Setting $\gamma := (1 - \epsilon)\delta_1 \delta_2$ completes the proof.

2.2 General case

In the monotone case, we obtained g_j from f_j by zeroing out certain "subfunctions" $f_j|_{J \leftarrow x^{(j)}}$. In the general case, we also need to fix some subfunctions to one. This suggests the following counting lemma, for a given predicate $P \subseteq \{0,1\}^m$.

Lemma 2.7 (Counting lemma for P)**.** *Let μ be a distribution on P with full support. For every $\epsilon > 0$ there exist $d \in \mathbb{N}$ and $\tau, \gamma > 0$ such that the following holds for all n.*

Let $\phi_1, \ldots, \phi_m \colon \{0,1\}^n \to \{0,1\}$ be functions such that ϕ_j is (d,τ)-regular with respect to $\mu|_j$ for all $j \in [m]$. Define a function $\chi_\epsilon \colon [m] \to \{0,1,*\}$ as follows:

$$\chi_\epsilon(j) = \begin{cases} 0 & \text{if } \mathbb{E}_{\mu|_j}[\phi_j] \le \epsilon, \\ 1 & \text{if } \mathbb{E}_{\mu|_j}[\phi_j] \ge 1-\epsilon, \\ * & \text{otherwise.} \end{cases}$$

Let $\alpha \colon [m] \to \{0,1\}$ be any assignment consistent with χ_ϵ. Then

$$\Pr_{(y^{(1)},\ldots,y^{(m)})\sim\mu^n}[(\phi_1(y^{(1)}),\ldots,\phi_m(y^{(m)})) = \alpha] \ge \gamma.$$

Given such a counting lemma, we could hope to complete the proof as in the monotone setting, this time obtaining g_j from f_j by setting subfunctions to constants according to χ_ϵ. There are two main issues with this plan:

1. The counting lemma does not hold for all predicates P.

   A simple example is the predicate $P_= = \{(0,0),(1,1)\}$ with respect to the uniform distribution. If $\phi_1 = \phi_2$ then the lemma holds only for $\alpha \in \{(0,0),(1,1)\}$. The same problem occurs for $P_{\neq} = \{(0,1),(1,0)\}$.

   A more complicated example is the predicate $P_\oplus = \{(a,b,a\oplus b) : a,b \in \{0,1\}\}$ with respect to the uniform distribution. If $\phi_1(x) = \phi_2(x) = \phi_3(x) = x_1 \oplus \cdots \oplus x_n$ (which is (d,τ)-regular for any $d < n$) then the lemma holds only for $\alpha \in P_\oplus$.

   These examples are not surprising: the argument in the monotone case shows that every approximate generalized polymorphism is close to a generalized polymorphism in which moreover each function is a *junta*, implying in particular that every polymorphism is close to a junta. However, this is not the case for the predicates $P_=, P_{\neq}, P_\oplus$.

2. It is not clear how to handle subfunctions which are not regular.

   In the monotone case, it was safe to zero them out, but for general predicates, there is no safe direction.

It turns out that the counting lemma does hold (under the additional assumption that $(\phi_1,\ldots,\phi_m)$ is a (μ,γ)-approximate generalized polymorphism) as long as P satisfies no affine relations: there is no non-empty set S such that $P|_S = \{x \in \{0,1\}^S : \bigoplus_{j\in S} x_j = b\}$. If P satisfies the premise of Theorem 1.10 (for each $j \in [m]$ there is $x^{(j)} \in P$ which remains in P after flipping the j'th coordinate) then this can be proved along the lines of the counting lemma for NAND. The general case requires an argument similar to the one used in [CFM+22], and also uses some ideas from [Mos12b].

In order to handle predicates with affine relations, we first prove Theorem 1.6, a generalization of linearity testing for arbitrary distributions. The proof uses the techniques of [DFH25]. Theorem 1.6 shows that if j is a coordinate involved in an affine relation, then f_j is close to a (possibly negated) XOR, and we fix g_j to be this XOR. We then remove coordinates from P until all affine relations disappear, enabling us to use the counting lemma.

We handle irregular subfunctions using an approach similar to the proof of the counting lemma for NAND. We find a way to sample μ in two steps, first sampling a restriction ρ on the coordinates J^c which leaves at most one coordinate free, and then sampling the free coordinate (if any) according to the correct marginal. We show that

(a) With constant probability over the choice of ρ, if we define $g_j|_{J\leftarrow x^{(j)}}$ by "rounding" $f_j|_{J\leftarrow x^{(j)}, J^c\leftarrow\rho}$ according to χ_η (for an appropriate η) then $(g_1,\ldots,g_m)$ is a generalized polymorphism of P.

(b) For every $(x^{(1)},\ldots,x^{(m)}) \in P^J$ such that all $f_j|_{J\leftarrow x^{(j)}}$ are regular, the coloring χ_η (defined according to $f_j|_{J\leftarrow x^{(j)}, J^c\leftarrow\rho}$) is compatible with the coloring χ_ϵ (defined according to $f_j|_{J\leftarrow x^{(j)}}$) with probability $1-\epsilon$.

The second property guarantees that on average $\Pr_{\mu|_j}[g_j \neq f_j] = O(\epsilon)$, allowing us to find a restriction ρ for which both $(g_1, \ldots, g_m)$ is a generalized polymorphism of P and $\Pr_{\mu|_j}[g_j \neq f_j] = O(\epsilon)$, completing the proof.

Flexible coordinates During the proof of Theorem 1.5, we distinguish between two types of coordinates. A coordinate $j \in [m]$ is *flexible* if there exists a partial input, leaving only the j'th coordinate unset, both of whose completions belong to the predicate. If no such input exists, then the coordinate is *inflexible*.

If the predicate is monotone and no coordinate is constant, then all of its coordinates are flexible. For functional predicates, which are predicates of the form $\{(x, f(x)) : x \in \{0,1\}^m\}$ for some $f\colon \{0,1\}^m \to \{0,1\}$, the first m coordinates are flexible, and the final one is not. For the predicates $\{(1,0,0),(0,1,0),(0,0,1)\}$ and $P_{m,b}$ (all m-ary vectors with parity b), all coordinates are inflexible.

The counting lemma implies that if $(\phi_1, \ldots, \phi_m)$ is a regular generalized polymorphism of a predicate without affine relations then ϕ_j is almost constant for every inflexible coordinate j. This is in fact the only part of the counting lemma which is employed in the rest of the proof.

The counting lemma is proved by induction on m and on the number of coordinates j such that ϕ_j is far from constant. It uses the following dichotomy for *almost full predicates*, which are predicates P such that $P|_{[m]\setminus\{j\}} = \{0,1\}^{[m]\setminus\{j\}}$ for all j:

Either one of the coordinates that P depends on is flexible, or P satisfies an affine relation.

The lack of an analogous property for larger alphabets precludes us from extending Theorem 1.5 to that setting in full generality. Instead, Theorem 1.10 assumes that all coordinates are flexible, and its proof is a simplification of the proof of Theorem 1.5.

3 Monotone case

In this section we prove Theorem 1.7.

Theorem 1.7 (Monotone case)**.** *Let $P \subseteq \{0,1\}^m$ be a non-empty monotone predicate: if $x \in P$ and $y \leq x$ (pointwise) then $y \in P$. Let μ be a distribution on P with full support. There exists a constant $C = C(P,\mu)$ such that the following holds for $\delta = 1/\exp\Theta(1/\epsilon^C)$.*

If $f_1, \ldots, f_m\colon \{0,1\}^n \to \{0,1\}$ is a (μ,δ)-approximate generalized polymorphism of P then there exists a generalized polymorphism $g_1, \ldots, g_m\colon \{0,1\}^n \to \{0,1\}$ of P such that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon$ for all $j \in [m]$. Moreover, $g_j \leq f_j$ pointwise for all $j \in [m]$.

Furthermore, if $f_i = f_j$ and $\mu|_i = \mu|_j$ then $g_i = g_j$.

The proof closely follows the outline in Section 2.1: we first prove an appropriate counting lemma using It Ain't Over Till It's Over, and then deduce the result using Jones' regularity lemma. Both of these results use the concept of (d,τ)-regularity.

Definition 3.1 (Regularity)**.** Let $d \in \mathbb{N}$, $\tau > 0$, and $p \in (0,1)$, and recall that μ_p is the product distribution with $\Pr_{x \sim \mu_p}[x_i = 1] = p$.

A function $f\colon \{0,1\}^n \to \{0,1\}$ is (d,τ)-regular with respect to μ_p if $\mathrm{Inf}_i[f^{\leq d}] \leq \tau$ for all $i \in [n]$, where

$$\mathrm{Inf}_i[f^{\leq d}] = \sum_{\substack{|S| \leq d \\ i \in S}} \hat{f}(S)^2,$$

and $\hat{f}(S)$ is the Fourier expansion of f with respect to μ_p, that is

$$f(x) = \sum_{S \subseteq [n]} \hat{f}(S) \prod_{i \in S} \frac{x_i - p}{\sqrt{p(1-p)}}.$$

We already stated It Ain't Over Till It's Over in Section 2.1. Here we restate it with explicit parameters, which can be read from the proof in [MOO10].

Theorem 3.2 (It Ain't Over Till It's Over)**.** *For every $p, q \in (0,1)$ and $\epsilon > 0$ the following holds for some constant $C = C(p,q)$ and $d = \Theta(\log(1/\epsilon))$, $\tau = \Theta(\epsilon^C)$, and $\delta = \Theta(\epsilon^C)$.*

Let ρ be a random restriction obtained by sampling each coordinate $i \in [n]$ independently according to the following law:

$$\rho_i = \begin{cases} 0 & w.p.\ (1-p)(1-q), \\ 1 & w.p.\ p(1-q), \\ * & w.p.\ q. \end{cases}$$

If $f\colon \{0,1\}^n \to \{0,1\}$ is (d,τ)-regular with respect to μ_p and $\mathbb{E}_{\mu_p}[f] \geq \epsilon$ then

$$\Pr_{\rho}[\mathbb{E}_{\mu_p}[f|_\rho] \geq \delta] \geq 1 - \epsilon.$$

In order to obtain the best possible parameters, we use Jones' regularity lemma in the following form, essentially proved in [CFM+22]. For definiteness, we reprove it in Section 8.

Theorem 3.3 (Jones' regularity lemma)**.** *For every $\epsilon, \tau > 0$, $m, d \in \mathbb{N}$ and $p_1, \ldots, p_m \in (0,1)$, the following holds for $M = O(md/\epsilon\tau)$.*

For all functions $f_1, \ldots, f_m\colon \{0,1\}^n \to \{0,1\}$ there exists a decision tree T of depth at most M such that for all j,

$$\Pr_{\rho\sim T}[f_j|_\rho \text{ is } (d,\tau)\text{-regular with respect to } \mu_{p_j}] \geq 1 - \epsilon,$$

where ρ is sampled by following T, sampling each variable encountered according to μ_{p_j}.

3.1 Counting lemma

We start by stating and proving the counting lemma that we use.

Lemma 3.4 (Counting lemma for monotone predicates)**.** *Let $P \subseteq \{0,1\}^m$ be a non-empty monotone predicate in which no coordinate is constant (i.e., for every $j \in [m]$ there is $x \in P$ with $x_j = 1$). Let μ be a distribution on P with full support. For every $\epsilon > 0$ there exists a constant $C = C(P,\mu)$ such that the following holds for $d = \Theta(\log(1/\epsilon))$, $\tau = \Theta(\epsilon^C)$, and $\gamma = \Theta(\epsilon^{Cm})$.*

Let $\phi_1, \ldots, \phi_m\colon \{0,1\}^n \to \{0,1\}$ be functions such that ϕ_j is (d,τ)-regular with respect to $\mu|_j$. Define $\alpha\colon [m] \to \{0,1\}$ as follows:

$$\alpha(j) = \begin{cases} 0 & \text{if } \mathbb{E}_{\mu|_j}[\phi_j] \leq \epsilon, \\ 1 & \text{otherwise.} \end{cases}$$

If $\alpha \notin P$ then

$$\Pr_{(y^{(1)},\ldots,y^{(m)})\sim\mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) \notin P] \geq \gamma.$$

We first prove the lemma in the special case where P is a NAND predicate, and then deduce the general case.

Lemma 3.5 (Counting lemma for NAND)**.** *Let $m \geq 2$ and let $P_{\mathsf{NAND}} = \{x \in \{0,1\}^m : x \neq (1,\ldots,1)\}$. Let μ be a distribution on P_{NAND} with full support. For every $\epsilon > 0$ there exists a constant $C = C(P,\mu)$ such that the following holds for $d = \Theta(\log(1/\epsilon))$, $\tau = \Theta(\epsilon^C)$, and $\gamma = \Theta(\epsilon^{Cm})$.*

Let $\phi_1, \ldots, \phi_m\colon \{0,1\}^n \to \{0,1\}$ be functions such that ϕ_j is (d,τ)-regular with respect to $\mu|_j$ and $\mathbb{E}_{\mu|_j}[\phi_j] \geq \epsilon$. Then

$$\Pr_{(y^{(1)},\ldots,y^{(m)})\sim\mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = (1,\ldots,1)] \geq \gamma.$$

Proof. The proof closely follows the proof of the counting lemma in Section 2.1.

Let $q > 0$ be a small enough parameter. We sample a random restriction $\rho \in (\{0,1,*\}^m)^n$ by sampling each coordinate independently according to the following law, where $p_j = \Pr[\mu|_j = 1]$, 0 is the zero vector, $e_1 = (1,0,\ldots,0)$, $s_1 = (*,0,\ldots,0)$, and e_j, s_j are defined analogously by making the special coordinate the j'th one:

- $\rho_i = 0$ with probability $\mu(0) - \sum_j (1 - p_j)q$.
- For each j, $\rho_i = e_j$ with probability $\mu(e_j) - p_j q$.
- For each $w \neq 0, e_1, \ldots, e_m$, $\rho_i = w$ with probability $\mu(w)$.
- For each j, $\rho_i = s_j$ with probability q.

Given ρ, we can obtain a sample of μ by sampling a star in position j according to $\mu|_j$. We denote this distribution by $\mu|\rho$.

Another crucial property of ρ is that the distribution of $\rho|_j$ given that $\rho|_j \neq *$ is the same as the distribution of $\mu|_j$. Indeed,

$$\Pr[\rho|_j = 1 \mid \rho|_j \neq *] = \frac{p_j - p_j q}{1 - q} = p_j.$$

Apply Theorem 3.2 (It Ain't Over Till It's Over) to each ϕ_j with $p := p_j$, $q := q$, $\epsilon := \epsilon$ to obtain d_j, τ_j, δ_j such that if ϕ_j is (d_j, τ_j)-regular and $\mathbb{E}_{\mu|_j}[\phi_j] \geq \epsilon$ then

$$\Pr_\rho[\mathop{\mathbb{E}}_{\mu|_j}[\phi_j|_{\rho|_j}] \geq \delta_j] \geq 1 - \epsilon.$$

Choose $d = \max(d_1, \ldots, d_m)$, $\tau = \min(\tau_1, \ldots, \tau_m)$, so that each ϕ_j is (d_j, τ_j)-regular. Applying the union bound,

$$\Pr_\rho[\mathop{\mathbb{E}}_{\mu|_j}[\phi_j|_{\rho|_j}] \geq \delta_j \text{ for all } j] \geq 1 - m\epsilon,$$

and so

$$\Pr_{(y^{(1)}, \ldots, y^{(m)}) \sim \mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = (1, \ldots, 1)] \geq (1 - m\epsilon)\delta_1 \cdots \delta_m.$$

The proof concludes by taking $\gamma = (1 - m\epsilon)\delta_1 \cdots \delta_m$. □

We prove Lemma 3.4 by applying Lemma 3.5 to each maxterm of P (recall that a maxterm is $x \notin P$ such that $y \in P$ for all $y \lneq x$).

Proof of Lemma 3.4. Let $\mathcal{M}$ be the collection of maxterms of P. Since no coordinate is constant, each maxterm is of the form 1_S for $|S| \geq 2$.

For each $1_S \in \mathcal{M}$, apply Lemma 3.5 with $\mu = \mu|_S$ and $\epsilon := \epsilon$ to obtain d_S, τ_S, γ_S such that the following holds: if ϕ_j is (d_S, τ_S)-regular and $\alpha_j = 1$ for all $j \in S$ then

$$\Pr_{(y^{(1)}, \ldots, y^{(m)}) \sim \mu_n}[\phi_j = 1 \text{ for all } j \in S] \geq \gamma_S.$$

We choose $d = \max(d_S : 1_S \in \mathcal{M})$, $\tau = \min(\tau_S : 1_S \in \mathcal{M})$, and $\gamma = \min(\gamma_S : 1_S \in \mathcal{M})$.

If $\alpha \notin P$ then there exists a maxterm $1_S \in \mathcal{M}$ such that $\alpha_j = 1$ for all $j \in S$. Lemma 3.5 implies that

$$\Pr_{(y^{(1)}, \ldots, y^{(m)}) \sim \mu_n}[\phi_j(y^{(j)}) = 1 \text{ for all } j \in S] \geq \gamma_S \geq \gamma.$$

Since 1_S is a maxterm, this completes the proof. □

3.2 Main result

In this section we complete the proof of Theorem 1.7 using Jones' regularity lemma. We first assume that no coordinate of P is constant, and then show how to get rid of this assumption.

Proof of Theorem 1.7 assuming no coordinate of P is constant. Since no coordinate of P is constant, we can apply Lemma 3.4 (the counting lemma) with $\epsilon := \epsilon/2$ to obtain d, τ, γ for which the lemma holds. Note $d = O(\log(1/epsilon))$, $\tau = \Omega(\epsilon^C)$, and $\delta = \Omega(\epsilon^C)$.

We apply Theorem 3.3 (Jones's regularity lemma) with $\epsilon := \epsilon/2$, $p_1, \dots, p_m$ given by $p_j = \Pr[\mu|_j = 1]$, and the values of d, τ obtained from Lemma 3.4 to obtain a decision tree T of depth $M = O(d/\epsilon\tau) = O((1/\epsilon)^C)$.

We define the functions g_j via the tree T as follows. For each leaf $\rho \in T$ and $x \in \{0,1\}^{(\operatorname{dom}\rho)^c}$,

$$g_j|_\rho(x) = \begin{cases} 0 & \text{if } f_j|_\rho \text{ is not } (d,\tau)\text{-regular or } \mathbb{E}_{\mu|_j}[f_j|_\rho] \leq \epsilon/2, \\ f_j|_\rho(x) & \text{otherwise.} \end{cases}$$

If we sample ρ according to $\mu|_j$ then according to Jones' regularity lemma, $f_j|_\rho$ is (d,τ)-regular (with respect to $\mu|_j$) with probability at least $1 - \epsilon/2$. This shows that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon/2 + \epsilon/2 = \epsilon$.

It remains to show that $(g_1, \dots, g_m)$ is an extended polymorphism of P for small enough δ. If this is not the case, then there exists a partial assignment $\rho \in T$ such that $(g_1|_{\rho|_1}, \dots, g_m|_{\rho|_m})$ is not an extended polymorphism of P.

Apply Lemma 3.4 (the counting lemma) to $\phi_j = f_j|_{\rho|_j}$. Observe that for all $(y^{(1)}, \dots, y^{(m)}) \in \{0,1\}^{(\operatorname{dom}\rho)^c}$, $(g_1|_{\rho_1}(y^{(1)}), \dots, g_m|_{\rho_m}(y^{(m)})) \leq \alpha$, where α is the assignment defined in the lemma. Since $(g_1|_{\rho|_1}, \dots, g_m|_{\rho|_m})$ is not an extended polymorphism of P and P is monotone, necessarily $\alpha \notin P$. Therefore the lemma shows that

$$\Pr_{(y^{(1)},\dots,y^{(m)}) \sim \mu^{(\operatorname{dom}\rho)^c}}[(f_1|_{\rho|_1}(y^{(1)}), \dots, f_m|_{\rho|_m}(y^{(m)})) \notin P] \geq \gamma,$$

implying that $(f_1, \dots, f_m)$ is not a (μ, δ)-approximate generalized polymorphism for $\delta = \min(\mu)^M \gamma/2 = 1/\exp\Theta(1/\epsilon^C)$, where $\min(\mu) = \min_{w \in \mu}(\mu(w))$. Choosing this value of δ completes the proof. □

The proof of Theorem 1.7 in full generality readily follows. Let C_0 be the set of coordinates j such that $x_j = 0$ for all $x \in P$. The premise of the theorem implies that $f_j(0, \dots, 0) = 0$ for all $j \in C_0$.

Let Q be the predicate obtained by removing the coordinates in C_0. We apply the foregoing to the predicate Q, obtaining functions g_j for all $j \notin C_0$. Defining $g_j = f_j$ for all $j \in C_0$ completes the proof.

4 Main theorem

In this section we prove Theorem 1.5.

Theorem 1.5 (Main)**.** *Let $P \subseteq \{0,1\}^m$ be a non-empty predicate, and let μ be a distribution on P with full support. For every $\epsilon > 0$ there exists $\delta > 0$ such that the following holds for all n.*

If $f_1, \dots, f_m \colon \{0,1\}^n \to \{0,1\}$ is a (μ, δ)-approximate generalized polymorphism of P then there exists a generalized polymorphism $g_1, \dots, g_m \colon \{0,1\}^n \to \{0,1\}$ of P such that

$$\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon \text{ for all } j \in [m],$$

where $\mu|_j$ is the marginal distribution of the j'th coordinate.

The proof relies on Theorem 1.6, the special case of Theorem 1.5 for the predicates $P_{m,b} = \{x \in \{0,1\}^m : x_1 \oplus \cdots \oplus x_m = b\}$ for $m \geq 3$.

The proof also relies on a two-sided version of It Ain't Over Till It's Over, which immediately follows from the one-sided version, Theorem 3.2, by applying it to both f and $1 - f$.

Theorem 4.1 (It Ain't Over Till It's Over)**.** *For every $p, q \in (0,1)$ and $\epsilon > 0$ the following holds for some constant $C = C(p,q)$ and $d = \Theta(\log(1/\epsilon))$, $\tau = \Theta(\epsilon^C)$, and $\delta = \Theta(\epsilon^C)$.*

Let ρ be a random restriction obtained by sampling each coordinate $i \in [n]$ independently according to the following law:

$$\rho_i = \begin{cases} 0 & \text{w.p. } (1-p)(1-q), \\ 1 & \text{w.p. } p(1-q), \\ * & \text{w.p. } q. \end{cases}$$

If f is (d, τ)-regular with respect to μ_p and $\epsilon \leq \mathbb{E}_{\mu_p}[f] \leq 1 - \epsilon$ then

$$\Pr_{\rho}[\delta \leq \mathop{\mathbb{E}}_{\mu_p}[f|_\rho] \leq 1 - \delta] \geq 1 - \epsilon.$$

We also use a junta version of Jones' regularity lemma, which we prove in Section 8.

Theorem 4.2 (Jones' regularity lemma)**.** *For every $\epsilon, \tau > 0$, $d, m \in \mathbb{N}$, and $p_1, \ldots, p_m \in (0,1)$ there exists a function $\mathcal{M} \colon \mathbb{N} \to \mathbb{N}$ such that the following holds for all n and all functions $f \colon \{0,1\}^n \to \{0,1\}$.*

For every $J_0 \subseteq [n]$ there exists a set $J \supseteq J_0$ of size at most $\mathcal{M}(|J_0|)$ such that for all j,

$$\Pr_{x \sim \mu_{p_j}}[f_j|_{J \leftarrow x} \text{ is not } (d,\tau)\text{-regular}] \leq \epsilon.$$

The proof consists of several parts:

- We prove a counting lemma for predicates without affine relations in Section 4.1.
- We outline the rounding procedure in Section 4.2.
- We prove Theorem 1.5 under the assumption that no coordinates are constant or duplicate in Section 4.3.
- We prove Theorem 1.5 in full generality in Section 4.4.

4.1 Counting lemma

A predicate $P \subseteq \{0,1\}^m$ has no *affine relations* if there do not exist a non-empty subset S and $b \in \{0,1\}$ such that $\bigoplus_{i \in S} w_i = b$ for all $w \in P$. In particular, this implies that no coordinate of P is constant, and no two coordinates are always equal or always non-equal.

Lemma 4.3 (Counting lemma for predicates without affine relations)**.** *Let P be a predicate without affine relations, and let μ be a distribution on P with full support. For every $\epsilon > 0$ there exist $d \in \mathbb{N}$ and $\tau, \gamma > 0$ such that the following holds for all n.*

Let $\phi_1, \ldots, \phi_m \colon \{0,1\}^n \to \{0,1\}$ be functions such that $(\phi_1, \ldots, \phi_m)$ is a (μ, γ)-approximate generalized polymorphism of P and ϕ_j is (d, τ)-regular with respect to $\mu|_j$ for all $j \in [m]$. Define a function $\chi_\epsilon \colon [m] \to \{0, 1, *\}$ as follows:

$$\chi_\epsilon(j) = \begin{cases} 0 & \text{if } \mathbb{E}_{\mu|_j}[\phi_j] \leq \epsilon, \\ 1 & \text{if } \mathbb{E}_{\mu|_j}[\phi_j] \geq 1 - \epsilon, \\ * & \text{otherwise.} \end{cases}$$

Let $\alpha \colon [m] \to \{0,1\}$ be any assignment consistent with χ_ϵ. Then $\alpha \in P$ and

$$\Pr_{(y^{(1)}, \ldots, y^{(m)}) \sim \mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = \alpha] > \gamma.$$

The proof of the case $\chi_\epsilon \equiv *$ will require the following lemma, whose proof is adapted from [Mos12a].

Lemma 4.4 (Hitting lemma)**.** *Let μ be a distribution over $\{0,1\}^m$ with full support. For every $\epsilon > 0$ there exists $\gamma > 0$ such that the following holds.*

Let $\phi_1, \ldots, \phi_m \colon \{0,1\}^n \to \{0,1\}$ be functions satisfying $\epsilon \leq \mathbb{E}_{\mu|_j}[\phi_j] \leq 1 - \epsilon$ for all $j \in [m]$. Then for all $\alpha \in \{0,1\}^m$,

$$\Pr_{(y^{(1)}, \ldots, y^{(m)}) \sim \mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = \alpha] > \gamma.$$

Proof. The proof is by induction on m. If $m = 1$ then we can take $\gamma = \epsilon/2$, so suppose $m \geq 2$.

Applying the inductive hypothesis with $\mu := \mu|_{\{2,\ldots,m\}}$ and $\epsilon := \epsilon$ gives $\gamma_{m-1} > 0$ such that for all $\beta \in \{0,1\}^{m-1}$,

$$\Pr_{(y^{(2)},\ldots,y^{(m)})\sim\mu|^n_{\{2,\ldots,m\}}}[(\phi_2(y^{(2)}),\ldots,\phi_m(y^{(m)})) = \beta] > \gamma_{m-1}.$$

We now appeal to [MOS13, Lemma 8.3], whose statement is as follows. Suppose that X, Y are finite sets and that ν is a distribution over $X \times Y$ with full support. For every $\eta > 0$ there exists $\zeta > 0$ such that the following holds. If $A \subseteq X^n$ and $B \subseteq Y^n$ have measure at least η (with respect to the corresponding marginals of ν^n) then $\nu^n(A \times B) \geq \zeta$.

Given $\alpha \in \{0,1\}^m$, we apply the lemma with $X = \{0,1\}$, $Y = \{0,1\}^{m-1}$, $\nu = \mu$ and $\eta := \min(\epsilon, \gamma_{m-1})$ to obtain $\zeta > 0$. Taking $A = \{y^{(1)} : \phi_1(y^{(1)}) = \alpha_1\}$, $B = \{(y^{(2)},\ldots,y^{(m)}) : (\phi_2(y^{(2)}),\ldots,\phi_m(y^{(m)})) = (\alpha_2,\ldots,\alpha_m)\}$, the lemma implies that

$$\Pr_{(y^{(1)},\ldots,y^{(m)})\sim\mu^n}[(\phi_1(y^{(1)}),\ldots,\phi_m(y^{(m)})) = \alpha] > \zeta,$$

completing the proof (taking $\gamma := \zeta/2$). □

We can now prove the counting lemma.

Proof of Lemma 4.3. The proof is by double induction: first on m, and then on the set of non-$*$ inputs in χ_ϵ. We also assume, without loss of generality, that $\epsilon < 1/2$.

If $m = 1$ then necessarily $P = \{0,1\}$, since otherwise P has a constant coordinate. Therefore the lemma trivially holds with $\gamma = \epsilon$. From now on, we assume that $m \geq 2$, and induct on the number of non-$*$ inputs in χ_ϵ.

Base case Suppose that $\chi_\epsilon(j) = *$ for all $j \in [m]$. We will find $d' \in \mathbb{N}$ and $\tau', \gamma' > 0$ such that whenever $d \geq d'$, $\tau \leq \tau'$ and $\gamma \leq \gamma'$, the assumptions imply that $P = \{0,1\}^m$. Applying Lemma 4.4 with $\mu := \mu$ and $\epsilon := \epsilon$ gives us $\gamma'' > 0$ such that

$$\Pr_{(y^{(1)},\ldots,y^{(m)})\sim\mu^n}[(\phi_1(y^{(1)}),\ldots,\phi_m(y^{(m)})) = \alpha] > \gamma''$$

for all $\alpha \in \{0,1\}^m$. Taking $\gamma = \min(\gamma', \gamma'')$ will conclude the proof.

Suppose therefore that $P \neq \{0,1\}^m$. We first show that for appropriate d', τ', γ' there exists an index $j_0 \in [m]$ and two inputs $u, v \in P$ such that $u^{\oplus j_0} \in P$ and $v^{\oplus j_0} \notin P$, where the inputs $u^{\oplus j_0}, v^{\oplus j_0}$ are obtained from the inputs u, v by flipping the j_0'th coordinate. ***This is where we use the assumption that P has no affine relations.*** We then show that the probability that $(\phi_1, \ldots, \phi_m)$ evaluates to $v^{\oplus j_0}$ is bounded from below, and obtain a contradiction by setting γ' small enough.

For each $j \in [m]$, we apply the inductive hypothesis to the predicate $P|_{[m]\setminus\{j\}}$, the distribution $\mu|_{[m]\setminus\{j\}}$, and $\epsilon := \epsilon$ to obtain d_j, τ_j, γ_j such that the following holds. If for all $k \neq j$, the function ϕ_k is (d_j, τ_j)-regular with respect to $\mu|_k$ and satisfies $\epsilon \leq \mathbb{E}_{\mu|_k}[\phi_k] \leq 1-\epsilon$, then $\alpha \in P|_{[m]\setminus\{j\}}$ for every $\alpha \in \{0,1\}^{[m]\setminus\{j\}}$, implying that $P|_{[m]\setminus\{j\}} = \{0,1\}^{[m]\setminus\{j\}}$.

Choose $d''' = \max(d_1, \ldots, d_m)$, $\tau''' = \min(\tau_1, \ldots, \tau_m)$, $\gamma''' = \min(\gamma_1, \ldots, \gamma_m)$. If $d \geq d'''$, $\tau''' \leq \tau$, $\gamma''' \leq \gamma$ then the assumptions of the lemma imply that $P|_{[m]\setminus\{j\}} = \{0,1\}^{[m]\setminus\{j\}}$ for all j. This implies that for every $w \in \{0,1\}^m$ and every $j \in [m]$, either $w \in P$ or $w^{\oplus j} \in P$ (or both).

Let J be the set of variables that P depends on: there is $w \in P$ such that $w^{\oplus j} \notin P$. By assumption, $\emptyset \neq P \neq \{0,1\}^m$, and so J is non-empty. If for every $w \in P$ and every $j \in J$ we have $w^{\oplus j} \notin P$ then $P|_J$ consists of all inputs with a given parity, contradicting the assumption that P has no affine relations. Therefore there must exist $j_0 \in J$ and $u \in P$ such that $u^{\oplus j_0} \in P$. Since $j_0 \in J$, there also exists $v \in P$ such that $v^{\oplus j_0} \notin P$.

The remainder of the argument is reminiscent of the corresponding argument in [CFM+22]. As in the proof of the counting lemma for NAND (Lemma 3.5), we will sample μ using a two-step process. Let u_* be obtained from u by setting the j_0'th coordinate to $*$. First, we sample a restriction $\rho \in (\{0,1\}^m \cup \{u_*\})^n$ by sampling each coordinate independently according to the following law, where $q = \min(\mu(u), \mu(u^{\oplus j_0}))$:

- $\rho_i = u$ with probability $\mu(u) - \Pr[\mu|_{j_0} = u_{j_0}]q$.
- $\rho_i = u^{\oplus j_0}$ with probability $\mu(u^{\oplus j_0}) - \Pr[\mu|_{j_0} = u^{\oplus j_0}_{j_0}]q$.
- $\rho_i = w$ with probability $\mu(w)$ for any $w \neq u, u^{\oplus j_0}$.
- $\rho_i = u_*$ with probability q.

Given ρ, we can obtain a sample of μ by sampling the j_0'th coordinate according to $\mu|_{j_0}$ if $\rho_i = u_*$. As in Lemma 3.5, the distribution of $\rho|_j$ given that $\rho|_j \neq *$ is the same as $\mu|_j$.

Apply Theorem 4.1 (It Ain't Over Till It's Over) to ϕ_{j_0} with $p := \Pr[\mu|_{j_0} = 1]$, $q := q$, $\epsilon := \epsilon$ to obtain d'''', τ'''', δ, and take $d'' = \max(d''', d'''')$ and $\tau'' = \min(\tau''', \tau'''')$. The assumptions of the lemma imply that

$$\Pr_{\rho}\big[\Pr_{\mu|_{j_0}}[\phi_{j_0}|_{\rho|_{j_0}} = v^{\oplus j_0}_{j_0}] \geq \delta\big] \geq 1 - \epsilon.$$

Since $P|_{[m]\setminus\{j_0\}} = \{0,1\}^{[m]\setminus\{j_0\}}$, we can apply Lemma 4.4 to $\mu|_{[m]\setminus\{j_0\}}$ and $\epsilon := \epsilon$ to obtain $\gamma'''' > 0$ such that

$$\Pr_{y \sim (\mu|_{[m]\setminus\{j_0\}})^n}[\phi_j(y^{(j)}) = v_j \text{ for all } j \neq j_0] > \gamma''''.$$

It follows that

$$\Pr_{(y^{(1)},\ldots,y^{(m)}) \sim \mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = v^{\oplus j_0}] > (1-\epsilon)\delta\gamma''''.$$

Setting $\gamma' = \min(\gamma''', (1-\epsilon)\delta\gamma'''')$ contradicts the assumption that $(\phi_1, \ldots, \phi_m)$ is a (μ, γ)-approximate generalized polymorphism of P for some $\gamma \leq \gamma'$, and we conclude that necessarily $P = \{0,1\}^m$, as claimed above. This concludes the proof of the base case.

Inductive case Suppose that $\chi_\epsilon(j_0) \neq *$ for some $j_0 \in [m]$.

Let us try to reduce the statement of the lemma to the same statement for $P|_{[m]\setminus\{j_0\}}$. We apply the inductive hypothesis to the predicate $P|_{[m]\setminus\{j_0\}}$, the distribution $\mu|_{[m]\setminus\{j_0\}}$, and $\epsilon := \epsilon$ to obtain d', τ', γ' such that if $d \geq d'$, $\tau \leq \tau'$, $\gamma \leq \gamma'$ then the assumptions imply that for every α consistent with χ_ϵ, $\alpha|_{[m]\setminus\{j_0\}} \in P|_{[m]\setminus\{j_0\}}$, and moreover

$$\Pr_{y \sim (\mu|_{[m]\setminus\{j_0\}})^n}[\phi_j(y^{(j)}) = \alpha_j \text{ for all } j \neq j_0] > \gamma'.$$

If $\Pr_{\mu|_{j_0}}[\phi_{j_0} = \alpha_{j_0} \oplus 1] \leq \gamma'/2$ then

$$\Pr_{(y^{(1)},\ldots,y^{(m)}) \sim \mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = \alpha] > \gamma'/2,$$

implying that the lemma holds if $\gamma \leq \gamma'/2$; we automatically get that $\alpha \in P$ since $(\phi_1, \ldots, \phi_m)$ is a $(\mu, \gamma'/2)$-approximate generalized polymorphism of P.

If $\Pr_{\mu|_{j_0}}[\phi_{j_0} = \alpha_{j_0} \oplus 1] > \gamma'/2$ (implying that $\gamma'/2 < \epsilon$) then we first observe that the set of $*$-inputs of $\chi_{\gamma'/2}$ strictly contains the set of $*$-inputs of χ_ϵ. Indeed, if $\chi_\epsilon = *$ then $\chi_{\gamma'/2} = *$ since $\gamma'/2 < \epsilon$, and moreover $\chi_\epsilon(j_0) \neq *$ whereas $\chi_{\gamma'/2}(j_0) = *$.

This allows us to apply the inductive hypothesis to $P := P$, $\mu := \mu$, and $\epsilon := \gamma'/2$. We obtain d'', τ'', γ'' such that if $d \geq d''$, $\tau \leq \tau''$, $\gamma \leq \gamma''$ then the assumptions imply that for every α consistent with $\chi_{\gamma'/2}$ (in particular, every α consistent with χ_α) we have $\alpha \in P$ and

$$\Pr_{(y^{(1)},\ldots,y^{(m)}) \sim \mu^n}[(\phi_1(y^{(1)}), \ldots, \phi_m(y^{(m)})) = \alpha] > \gamma''.$$

Choosing $d = \max(d', d'')$, $\tau = \min(\tau', \tau'')$, $\gamma = \min(\gamma'/2, \gamma'')$ completes the proof. □

4.2 Rounding

The proof of Theorem 1.5 (for predicates without affine relations) follows the main plan of the proof of Theorem 1.7: we apply Jones' regularity lemma on $f_1, \ldots, f_m$ to construct a junta J such that on average, $f_j|_{J \leftarrow x^{(j)}}$ is regular. Lemma 4.3 suggests a "rounding" procedure to construct the generalized polymorphism $(g_1, \ldots, g_m)$: we define $g_j|_{J \leftarrow x^{(j)}} \equiv b$ if $\Pr[f_j|_{J \leftarrow x^{(j)}} = b] \geq 1-\epsilon$ for some $b \in \{0,1\}$, and $g_j|_{J \leftarrow x^{(j)}} = f_j|_{J \leftarrow x^{(j)}}$ otherwise. The lemma implies (for an appropriate choice of parameters) that $(g_1|_{J \leftarrow x^{(1)}}, \ldots, g_m|_{J \leftarrow x^{(m)}})$ is a generalized polymorphism whenever all of the involved subfunctions are regular, but gives no information otherwise.

In the monotone setting, we handle this difficulty by setting $g_j|_{J \leftarrow x^{(j)}} \equiv 0$ whenever $f_j|_{J \leftarrow x^{(j)}}$ is not regular. In the general case, there is no such safe direction. Instead, we use a two-step sampling procedure, along the lines of the proof of Lemma 3.5 (the counting lemma for NAND), to define $g_1, \ldots, g_m$, and use Lemma 4.3 to show that g_j is close to f_j.

The two-step sampling procedures in Lemma 3.5 and Lemma 4.3 rely on the existence of inputs $w_{(j)} \in P$ such that $w_{(j)}^{\oplus j} \in P$. If such inputs exist for all j then the proof of Theorem 1.5 can be simplified considerably; we take this route when we prove Theorem 1.10. However, in general this is not the case. As an example, if $P = \{(1,0,0),(0,1,0),(0,0,1)\}$ then no such inputs exist for any j. Our argument will need to distinguish between these two types of coordinates.

Definition 4.5 (Flexible coordinates)**.** Let $P \subseteq \{0,1\}^m$ be a non-empty coordinate. A coordinate $j \in [m]$ is *flexible* if there exists $w_{(j)} \in P$ such that $w_{(j)}^{\oplus j} \in P$, and *inflexible* otherwise.

If j is a flexible coordinate, let $w_{(j,0)} \in P$ be some fixed input such that $(w_{(j,0)})_j = 0$ and $w_{(j,1)} := w_{(j,0)}^{\oplus j} \in P$. Also, let $w_{(j,*)}$ be the restriction obtained by changing the j'th coordinate to $*$.

We can now describe the distribution of the restriction ρ in the first step of the two-step sampling process. Each of its coordinates will be sampled according to the following distribution ν on $P \cup \{w_{(j,*)} : j \text{ flexible}\}$, for a small enough $q > 0$:

- For $w \in P$, sample w with probability
$$\mu(w) - \sum_{j\colon w = w_{(j,0)}} (1-p_j)q - \sum_{j\colon w = w_{(j,1)}} p_j q,$$
where the sum is over all flexible coordinates j, and $p_j = \Pr[\mu|_j = 1]$.
- For every flexible j, sample $w_{(j,*)}$ with probability q.

We can sample μ by first sampling ν, and in case $w_{(j,*)}$ was sampled, sampling the j'th coordinate according to $\mu|_j$. Also, the distribution of $\nu|_j$ conditioned on $\nu|_j \neq *$ coincides with $\mu|_j$.

We start by explaining how to use the two-step sampling process to construct a generalized polymorphism $(g_1, \ldots, g_m)$. In the lemma below, J is the set which will be constructed using Jones' regularity lemma. ***This rounding procedure prevents us from ensuring that*** $g_i = g_j$ ***even if*** $f_i = f_j$ ***and*** $\mu|_i = \mu|_j$***.*** In this lemma and below, we use $\min(\mu) := \min_{w \in P} \mu(w)$.

Lemma 4.6 (Rounding lemma)**.** *Let $P \subseteq \{0,1\}^m$ be a non-empty predicate, and let μ be a distribution on P having full support. For every $\epsilon, \zeta > 0$ there exists $\delta > 0$ such that the following holds.*

Let $f_1, \ldots, f_m \colon \{0,1\}^n \to \{0,1\}$, and let $J \subseteq [n]$. For a restriction $\rho \colon J^c \to \{0,1,*\}^m$, define functions $g_1^{\rho,\epsilon}, \ldots, g_m^{\rho,\epsilon} \colon \{0,1\}^n \to \{0,1\}$ as follows: for each $j \in [m]$ and each $x^{(j)} \in \{0,1\}^J$,

$$g_j^{\rho,\epsilon}|_{J \leftarrow x^{(j)}} = \begin{cases} 0 & \text{if } \mathbb{E}_{\mu|_j}[f_j|_{J \leftarrow x^{(j)}, \rho|_j}] \leq \epsilon, \\ 1 & \text{if } \mathbb{E}_{\mu|_j}[f_j|_{J \leftarrow x^{(j)}, \rho|_j}] \geq 1-\epsilon, \\ f_j|_{J \leftarrow x^{(j)}} & \text{otherwise.} \end{cases}$$

For fixed $x^{(1)}, \ldots, x^{(m)}$, this corresponds to a coloring $\chi^{\rho,\epsilon} \colon [m] \to \{0,1,*\}$ in a natural way.

If $(f_1,\dots,f_m)$ *is a* $(\mu,\min(\mu)^{|J|}\delta)$*-approximate generalized polymorphism of* P *then*

$$\Pr_{\rho\sim\nu^{J^c}}[(g_1^{\rho,\epsilon},\dots,g_m^{\rho,\epsilon}) \text{ is a generalized polymorphism of } P] \geq 1-\zeta.$$

If j is inflexible then $\rho|_j \in \{0,1\}^{J^c}$, and so $f_j|_{J\leftarrow x^{(j)},\rho|_j}$ is a constant. In this case $g_j^{\rho,\epsilon}|_{J\leftarrow x^{(j)}}$ always gets set to a constant. As we show later, Lemma 4.3 explains why this is a sound choice.

Proof. Suppose that $(f_1,\dots,f_m)$ is a $(\mu,\min(\mu)^{|J|}\delta)$-approximate generalized polymorphism of P, where we set δ later on. This means that

$$\mathop{\mathbb{E}}_{\rho\sim\nu^{J^c}} \Pr_{\substack{(x^{(1)},\dots,x^{(m)})\sim\mu^J\\(y^{(1)},\dots,y^{(m)})\sim\mu^{J^c}|\rho}}[(f_1|_{J\leftarrow x^{(1)}}(y^{(1)}),\dots,f_m|_{J\leftarrow x^{(m)}}(y^{(m)}))\notin P] \leq \min(\mu)^{|J|}\delta,$$

and so with probability at least $1-\zeta$ over the choice of ρ,

$$\Pr_{\substack{(x^{(1)},\dots,x^{(m)})\sim\mu^J\\(y^{(1)},\dots,y^{(m)})\sim\mu^{J^c}|\rho}}[(f_1|_{J\leftarrow x^{(1)}}(y^{(1)}),\dots,f_m|_{J\leftarrow x^{(m)}}(y^{(m)}))\notin P] \leq \min(\mu)^{|J|}\delta/\zeta.$$

Therefore for every $(x^{(1)},\dots,x^{(m)})\in P^J$ we have

$$\Pr_{(y^{(1)},\dots,y^{(m)})\sim\mu^{J^c}|\rho}[(f_1|_{J\leftarrow x^{(1)}}(y^{(1)}),\dots,f_m|_{J\leftarrow x^{(m)}}(y^{(m)}))\notin P] \leq \delta/\zeta.$$

We will show that if δ is small enough, then this implies that $(g_1^\rho,\dots,g_m^\rho)$ is a generalized polymorphism.

If $(g_1^{\rho,\epsilon},\dots,g_m^{\rho,\epsilon})$ is not a generalized polymorphism of P then there exists $(x^{(1)},\dots,x^{(m)})\in P^J$ such that $(g_1^{\rho,\epsilon}|_{J\leftarrow x^{(1)}},\dots,g_m^{\rho,\epsilon}|_{J\leftarrow x^{(m)}})$ is not a generalized polymorphism of P. This means that the corresponding coloring $\chi^{\rho,\epsilon}$ can be extended to some $\alpha\notin P$. Observe that

$$\Pr_{(y^{(1)},\dots,y^{(m)})\sim\mu^{J^c}|\rho}[(f_1|_{J\leftarrow x^{(1)}}(y^{(1)}),\dots,f_m|_{J\leftarrow x^{(m)}}(y^{(m)}))=\alpha] \geq \epsilon^m.$$

Choosing $\delta=\epsilon^m\zeta/2$ completes the proof. ☐

The next step is to show that on average (over ρ), the function $g_j^{\rho,\eta}$ is close to f_j, for an appropriate choice of η. For flexible coordinates, this follows immediately from Theorem 4.1. For inflexible coordinates, we will appeal to Lemma 4.3, which implies that for such coordinates, $f_j|_{J\leftarrow x^{(j)}}$ is close to constant. Since we apply Lemma 4.3, this argument will only work for certain subfunctions.

Definition 4.7 (Good subfunctions)**.** Let $P\subseteq\{0,1\}^m$ be a predicate, and let μ be a distribution over P having full support. Let $d\in\mathbb{N}$ and $\tau>0$. Let $f_1,\dots,f_m\colon\{0,1\}^n\to\{0,1\}$, and let $J\subseteq[n]$.

A subfunction $f_j|_{J\leftarrow x}$ is (μ,d,τ)*-good* if there exists $(x^{(1)},\dots,x^{(m)})\in P^J$ with $x^{(j)}=x$ such that $f_k|_{J\leftarrow x^{(k)}}$ is (d,τ)-regular with respect to $\mu|_k$ for all $k\in[m]$.

Lemma 4.8 (Most subfunctions are good)**.** *Let* $P\subseteq\{0,1\}^m$ *be a predicate, and let* μ *be a distribution over* P *having full support. Let* $d\in\mathbb{N}$ *and* $\tau>0$*. Let* $f_1,\dots,f_m\colon\{0,1\}^n\to\{0,1\}$*, and let* $J\subseteq[n]$*.*

Suppose that for all j*,*

$$\Pr_{x\sim\mu|_j^J}[f_j|_{J\leftarrow x} \text{ is } (d,\tau)\text{-regular wrt } \mu|_j] \geq 1-\epsilon.$$

Then for all j*,*

$$\Pr_{x\sim\mu|_j^J}[f_j|_{J\leftarrow x} \text{ is } (\mu,d,\tau)\text{-good}] \geq 1-m\epsilon.$$

Proof. Fix $j\in[m]$. The assumption implies that

$$\Pr_{(x^{(1)},\dots,x^{(m)})\sim\mu^J}[f_k|_{J\leftarrow x^{(k)}} \text{ is } (d,\tau)\text{-regular wrt } \mu|_k \text{ for all } k] \geq 1-m\epsilon.$$

If the event happens then $f_j|_{J\leftarrow x^{(j)}}$ is (μ,d,τ)-good. The lemma immediately follows. ☐

We can now show that on average, $g_j^{\rho,\eta}$ is close to f_j, for an appropriate choice of η. Recall that $\min(\mu) = \min_{w \in P} \mu(w)$.

Lemma 4.9 (Soundness of rounding)**.** *Let $P \subseteq \{0,1\}^m$ be a non-empty predicate without affine relations and let μ be a distribution on P having full support. For every $\epsilon > 0$ there exist $d \in \mathbb{N}$ and $\delta, \eta, \tau > 0$ such that the following holds.*

Let $f_1, \ldots, f_m \colon \{0,1\}^m \to \{0,1\}$, and let $J \subseteq [n]$. Suppose that for all $j \in [m]$,

$$\Pr_{x^{(j)} \sim \mu|_j^J}[f_j|_{J \leftarrow x^{(j)}} \text{ is } (d,\tau)\text{-regular}] \geq 1 - \epsilon.$$

If $(f_1, \ldots, f_m)$ is a $(\mu, \min(\mu)^{|J|}\delta)$-approximate generalized polymorphism of P then for all $j \in [m]$,

$$\mathop{\mathbb{E}}_{\rho \sim \nu^J} \Pr_{\mu|_j}[g_j^{\rho,\eta} \neq f_j] = O(\epsilon).$$

Proof. Apply Lemma 4.3 with $\epsilon := \epsilon$ to obtain d', τ', γ' such that the assumptions imply the following, assuming $d \geq d'$, $\tau \leq \tau'$, $\delta \leq \gamma'$. For all $(x^{(1)}, \ldots, x^{(m)}) \in P^J$, if $f_1|_{J \leftarrow x^{(1)}}, \ldots, f|_{J \leftarrow x^{(m)}}$ are all (d,τ)-regular then all assignments extending χ_ϵ belong to P.

For every flexible j, apply Theorem 3.2 (the one-sided version of Theorem 4.1) with $p := \Pr[\mu|_j = 1]$, $q := q$ (the parameter used to define ν) and $\epsilon := \epsilon$ to obtain d_j, τ_j, δ_j such that for $b \in \{0,1\}$ and $x^{(j)} \in \{0,1\}^J$, if $\chi_\epsilon(j) \neq b$ and $\eta \leq \delta_j$ then

$$\Pr_{\rho \sim \nu^J}[\chi^{\rho,\eta}(j) \neq b] \geq 1 - \epsilon.$$

(Note that $\chi_\epsilon(j) \neq b$ iff $\Pr_{\mu|_j}[f_j|_{J \leftarrow x^{(j)}} = b \oplus 1] > \epsilon$, and $\chi^{\rho,\eta}(j) \neq b$ iff $\Pr_{\mu_j}[f_j|_{J \leftarrow x^{(j)}, \rho|_j} = b \oplus 1] > \eta$.)

Let F be the set of flexible coordinates. We take $d = \max(d', (d_j)_{j \in F})$, $\tau = \min(\tau', (\tau_j)_{j \in F})$, $\delta = \gamma'$, $\eta = \min(\epsilon, (\delta_j)_{j \in F})$.

Given j, observe that

$$\mathop{\mathbb{E}}_{\rho \sim \nu^J}[\Pr_{\mu|_j}[g_j^{\rho,\eta} \neq f_j]] \leq \Pr_{x^{(j)} \sim \mu|_j^J}[f_j|_{J \leftarrow x^{(j)}} \text{ is not } (\mu,d,\tau)\text{-good}] + \sum_{x^{(j)} \colon f_j|_{J \leftarrow x^{(j)}} \text{ is } (\mu,d,\tau)\text{-good}} \mu|_j(x^{(j)}) \mathop{\mathbb{E}}_{\rho \sim \nu^J} \Pr_{\mu|_j}[g_j^{\rho,\eta}|_{J \leftarrow x^{(j)}} \neq f_j|_{J \leftarrow x^{(j)}}].$$

The first summand is at most $m\epsilon$, and so it suffices to show that for all $x^{(j)}$ such that $f_j|_{J \leftarrow x^{(j)}}$ is (μ,d,τ)-good,

$$\mathop{\mathbb{E}}_{\rho \sim \nu^J} \Pr_{\mu|_j}[g_j^{\rho,\eta}|_{J \leftarrow x^{(j)}} \neq f_j|_{J \leftarrow x^{(j)}}] = O(\epsilon).$$

Unpacking the definition of $g_j^{\rho,\eta}$, the left-hand side is

$$\Pr_{\rho \sim \nu|_j^J}[\mathbb{E}[f_j|_{J \leftarrow x^{(j)},\rho}] \leq \eta] \cdot \mathop{\mathbb{E}}_{\mu|_j}[f_j|_{J \leftarrow x^{(j)}}] + \Pr_{\rho \sim \nu|_j^J}[\mathbb{E}[1 - f_j|_{J \leftarrow x^{(j)},\rho}] \leq \eta] \cdot \mathop{\mathbb{E}}_{\mu|_j}[1 - f_j|_{J \leftarrow x^{(j)}}].$$

The two summands are similar, so it suffices to bound the first one.

We consider two cases, according to whether j is flexible or not. If j is flexible then either $\chi_\epsilon(j) = 0$, in which case $\mathbb{E}[f_j|_{J \leftarrow x^{(j)}}] \leq \epsilon$, or $\chi_\epsilon(j) \neq 0$, in which case $\Pr_\rho[\chi^{\rho,\eta} = 0] \leq \epsilon$, that is, $\Pr_{\rho \sim \nu|_j^J}[\mathbb{E}[f_j|_{J \leftarrow x^{(j)},\rho}] \leq \eta] \leq \epsilon$. In both cases, the summand is bounded by ϵ.

If j is inflexible then we need to use the fact that every extension of χ_ϵ belongs to P. Since j is inflexible, this implies that $\chi_\epsilon(j) \neq *$. We consider two subcases, according to the value of $\chi_\epsilon(j)$. If $\chi_\epsilon(j) = 0$ then $\mathbb{E}[f_j|_{J \leftarrow x^{(j)}}] \leq \epsilon$. If $\chi_\epsilon(j) = 1$ then $\mathbb{E}[f_j|_{J \leftarrow x^{(j)}}] \geq 1 - \epsilon$. Since j is inflexible, $\nu|_j = \mu|_j$, and so $f_j|_{J \leftarrow x^{(j)},\rho}$ is a constant, which equals 0 with probability at most ϵ. In both cases, the summand is bounded by ϵ. $\square$

4.3 Proof without short affine relations

If P has no affine relations then we can complete the proof as in Section 3 as follows. First, we apply Jones' regularity lemma on $f_1, \dots, f_m$ to obtain J. We then combine Lemma 4.6 and Lemma 4.9 to find a restriction ρ such that $(g_1^{\rho,\eta}, \dots, g_m^{\rho,\eta})$ is a generalized polymorphism and $g_j^{\rho,\eta}$ is close to f_j for all j. In this section, we show how to modify this argument to handle linear relations of size at least 3; handling smaller linear relations is easier and will be done in the next section.

For the remainder of the section, we assume that P has no *small* affine relations, meaning no affine relations of length smaller than 3.

The first step is to peel off all affine relations. Below we use the notation

$$\chi_{S,b}(x) = b \oplus \bigoplus_{i \in S} x_i.$$

Lemma 4.10 (Peeling affine relations)**.** *Let P be a predicate without small affine relations, and let μ be a distribution on P with full support. There exist $\epsilon_0 > 0$ and sets $F \subseteq I \subseteq [m]$ such that $P|_I$ has no affine relations, and the following holds for all $\epsilon \le \epsilon_0$.*

Let $f_1, \dots, f_m \colon \{0,1\}^n \to \{0,1\}$ be a (μ, ϵ)-approximate generalized polymorphism of P.

(a) For each $j \notin F$ there exist $b_j \in \{0,1\}$ and $S_j \subseteq [n]$ such that $\Pr_{\mu|_j}[f_j \neq \chi_{S_j,b_j}] = O(\epsilon)$.

(b) If $(g_j)_{j \in I}$ is a generalized polymorphism of $P|_I$ such that $g_j = \chi_{S_j,b_j}$ for all $j \in I \setminus F$, then we can extend it to a generalized polymorphism of P by taking $g_j = \chi_{S_j,b_j}$ for $j \notin I$.

The proof uses Theorem 1.6, proved in Section 5, which is the special case of Theorem 1.5 for affine relations.

Proof. We construct F, I using an iterative process. During the process, each coordinate can be active or not active; originally all are active. Furthermore, each coordinate has a list of characters $\chi_{S,b}$, initially empty. The process ends once there is no affine relation involving only active coordinates.

Each step of the iteration proceeds as follows. Choose an affine relation involving the coordinates in some set A, all of them active. Since P has no small affine relations, $|A| \ge 3$. Therefore we can apply Theorem 1.6 to obtain, for each $j \in A$, a character χ_{S_j,b_j} such that $\Pr_{\mu|_j}[f_j \neq \chi_{S_j,b_j}] = O(\epsilon)$. For each $j \in A$, we add the character χ_{S_j,b_j} to the list of characters for coordinate j. We also pick a coordinate $j_0 \in A$ arbitrarily, and render it inactive.

After the process ends, some of the lists are empty, and they comprise the set F. The set I contains all active coordinates. Property (a) is automatically satisfied.

The list of characters for each $j \notin F$ could contain more than one character. We can rule this out by taking ϵ_0 small enough. Indeed, if the list for j contains two different characters $\chi_{S',b'}, \chi_{S'',b''}$ then

$$\Pr_{\mu|_j}[\chi_{S',b'} \neq \chi_{S'',b''}] = O(\epsilon).$$

If ϵ_0 is smaller than a constant, this rules out $S' = S''$. Take any $i \in S' \triangle S''$, and sample coordinate i last. Whether $\chi_{S',b'} \neq \chi_{S'',b''}$ or not depends on the value of coordinate i, and so the probability that $\chi_{S',b'} \neq \chi_{S'',b''}$ is at least $\min(\mu|_j(0), \mu|_j(1))$. Therefore, choosing $\epsilon_0 = c \min_{j,b} \mu|_j(b)$ for an appropriate $c > 0$ ensures that every non-empty list contains precisely one character.

It remains to prove Property (b). For this, we observe that if $w \in P|_I$ then there is a unique way to extend it to an element in P. This is precisely how we extend the generalized polymorphism in Property (b). ☐

We would like to apply the argument outlined in the beginning of the section to $P|_I$ in such a way that guarantees that $g_j^{\rho,\eta} = \chi_{S_j,b_j}$ for all $j \in I \setminus F$. We do this in two steps. First, we change f_j to $f'_j = \chi_{S_j,b_j}$ for all $j \in I \setminus F$, which increases the error probability in a controlled way. Second, we ensure somehow that $g_j^{\rho,\eta} = f'_j$ for all $j \in I \setminus F$. This can be done in two ways: either J contains S_j, or $S_j \setminus J$ is so large that $\chi^{\rho,\eta} = *$ is very likely, as given by the following lemma.

Lemma 4.11 (Large characters are balanced)**.** *Let μ be a full support distribution on $\{0,1\}$. For every $q \in (0,1)$ and $\eta, \zeta > 0$ there exists $M \in \mathbb{N}$ such that the following holds.*

Let ν be the distribution on $\{0,1,*\}$ given by $\nu(*) = q$ and $\nu(b) = (1-q)\mu(b)$. If $f = \chi_{S,b}$ for $|S| \geq M$ and $b \in \{0,1\}$ then

$$\Pr_{\rho \sim \nu}[\eta < \mathop{\mathbb{E}}_{\mu}[f|_\rho] < 1 - \eta] \geq 1 - \zeta.$$

Proof. The proof is straightforward, so we only outline it. Since μ has full support, we can find $M' \in \mathbb{N}$ such that if $f' = \chi_{S',b'}$ for $|S'| \geq M'$ then $\eta < \mathbb{E}_\mu[f'] < 1 - \eta$. We can find M for which $f|_\rho$ is of this form with probability $1 - \zeta$. $\square$

We now put everything together.

Proof of Theorem 1.5 for predicates without short affine relations. Apply Lemma 4.10 to obtain $F \subseteq I \subseteq [m]$ and $\epsilon_0 > 0$. Recall that the lemma states that $P|_I$ has no affine relations, and that the following holds for every (μ, δ)-approximate polymorphism of P, whenever $\delta \leq \epsilon_0$.

First, for each $j \notin F$, the function f_j is $O(\delta)$-close to some χ_{S_j,b_j}. Second, any generalized polymorphism $(g_j)_{j \in I}$ of $P|_I$ which satisfies $g_j = \chi_{S_j,b_j}$ for $j \in I \setminus F$ can be extended to a generalized polymorphism of P by taking $g_j = \chi_{S_j,b_j}$ for $j \notin I$.

We can assume without loss of generality that $\epsilon \leq \epsilon_0$ (otherwise replace ϵ with ϵ_0). We will prove the theorem with an error probability of $O(\epsilon)$ rather than ϵ for convenience.

Constructing the junta The junta is constructed by applying various lemmas proved in this section. In order to make the argument more readable, we briefly recall the statement of each lemma. We use the following notation: a *random J-subfunction of f'_j* is obtained from f'_j by restricting the coordinates in J according to $\mu|_j$. Also, ρ is a random restriction of the coordinates in J^c sampled according to the distribution ν defined in Section 4.2.

Apply Lemma 4.9 (Soundness of rounding) with the predicate $P|_I$, the distribution $\mu|_I$, and $\epsilon := \epsilon$ to obtain $d', \delta', \eta', \tau'$ such that the following holds. Given a set J, suppose that for each $j \in I$, a random J-subfunction of f'_j is (d', τ')-regular w.p. $1 - \epsilon$; and that $(f_1, \ldots, f_m)$ is a $(\mu, \min(\mu)^{|J|}\delta')$-approximate polymorphism of $P|_I$. Then for all j, $\mathbb{E}_\rho \Pr[g_j^{\rho,\eta'} \neq f'_j] = O(\epsilon)$.

Apply Theorem 4.2 (Jones' regularity lemma) with $\epsilon := \epsilon$, $d = d'$, $\tau = \tau'$, and $p_j = \Pr[\mu|_j = 1]$ for all $j \in F$ to obtain $\mathcal{M} \colon \mathbb{N} \to \mathbb{N}$ such that the following holds. For each set J_0 there exists a set $J \subseteq J_0$ of size at most $\mathcal{M}(|J_0|)$ such that for all $j \in F$, a random J-subfunction of f'_j is (d', τ')-regular w.p. $1 - \epsilon$.

For every $j \in I \setminus F$, apply Lemma 4.11 with $\mu := \mu|_j$, $\eta := \eta'$, $\zeta := 1/(2m)$, and the q in the definition of ν (Section 4.2) to obtain M_j. The lemma implies that if $|S_j \setminus J| \geq M_j$ then for every $x^{(j)} \in \{0,1\}^J$ we have $\eta' < \mathbb{E}[\chi_{S_j,b_j}|_{J \leftarrow x^{(j)}, \rho|_j}] < 1 - \eta'$ with probability $1 - 1/(2m)$ over the choice of ρ.

The set J_0 that we construct will be based on an application of Lemma 4.10, which will yield us a set S_j for each $j \notin F$. We would like the set J constructed by Jones' regularity lemma to satisfy the following, for each $j \in I \setminus F$: either J contains S_j, or $|S_j \setminus J| \geq M_j$. Later on, we determine a value of M, depending only on the function $\mathcal{M}$ and the parameters M_j for $j \in I \setminus F$, such that such a set J can be constructed of size at most M.

Apply Lemma 4.6 (Rounding lemma) with the predicate $P|_I$, the distribution $\mu|_I$, $\epsilon := \eta'$, and $\zeta := 1/(2m)$ to obtain δ'' such that the following holds. If $(f'_1, \ldots, f'_m)$ is a $(\mu, \min(\mu)^{|J|}\delta'')$-approximate polymorphism of $P|_I$ then $(g_j^{\rho,\eta'})_{j \in I}$ is a generalized polymorphism of $P|_I$ w.p. $1 - 1/(2m)$ over the choice of ρ.

In order to apply Lemmas 4.6 and 4.9, we will need the error probability to be at most $\min(\mu)^{|J|} \min(\delta', \delta'')$. Consequently we choose

$$\delta = c \min(\mu)^M \min(\delta', \delta'', \epsilon)/(m+1),$$

for an appropriate constant $c \in (0,1)$; we will see where the additional $m+1$ factor comes from in a moment. (Recall that δ is the parameter such that $(f_1, \ldots, f_m)$ is a (μ, δ)-approximate polymorphism of P.)

Apply Lemma 4.10 to $f_1, \ldots, f_m$ to obtain, for each $j \notin F$, a function $g_j = \chi_{S_j,b_j}$ such that $\Pr_{\mu|_j}[g_j \neq f_j] \le \min(\mu)^M \min(\delta', \delta'', \epsilon)/(m+1)$ (this is where we get c from). Define

$$f'_j = \begin{cases} g_j & \text{if } j \notin F, \\ f_j & \text{if } j \in F. \end{cases}$$

Observe that $(f'_j)_{j \in I}$ is a $(\mu|_I, \min(\mu)^M \min(\delta', \delta'', \epsilon))$-approximate generalized polymorphism of $P|_I$, summing up the probability that $(f_j)_{j\in I} \notin P|_I$ with the probabilities that f_j and f'_j differ.

We can now construct J. First, we construct J_0. Start with $J_0 := \emptyset$, and while some $j \in I \setminus F$ satisfies $|S_j \setminus J_0| < M_j + \mathcal{M}(|J_0|)$, add S_j to J_0. The process terminates with a set J_0 such that for each $j \in I \setminus F$, either $J_0 \supseteq S_j$, or $|S_j \setminus J_0| \ge M_j + \mathcal{M}(|J_0|)$.

We proceed to bound the size of J_0. Let $M' = \max_{j \in I \setminus F} M_j$, and assume without loss of generality that $\mathcal{M}$ is monotone (otherwise, replace it with $\mathcal{M}'(s) := \max_{t \le s} \mathcal{M}(t)$). Taking $B_0 = 0$ and $B_{t+1} = B_t + M' + \mathcal{M}(B_t)$, a simple induction shows that after adding t many S_j's, we have $|J_0| \le B_t$. Thus the final size of J_0 is at most $B_{|I \setminus F|}$.

Finally, we obtain J by applying Jones' regularity lemma to the functions f'_j for $j \in F$. Observe that $|J| \le M := \mathcal{M}(B_{|I\setminus F|})$. Furthermore, for each $j \in I \setminus F$, either $J_0 \supseteq S_j$, in which case $J \supseteq S_j$, or $|S_j \setminus J_0| \ge M_j + \mathcal{M}(|J_0|) \ge M_j + |J|$, in which case $|S_j \setminus J| \ge M_j$.

Finding a good restriction We would like to find a restriction ρ such that all of the following hold:

(a) $(g_j^{\rho,\eta})_{j \in I}$ is a generalized polymorphism of $P|_I$.

(b) For each $j \in I \setminus F$, $g_j^{\rho,\eta} = f'_j$.

(c) For each $j \in F$, $\Pr_{\mu|_j}[g_j^{\rho,\eta} \neq f'_j] = O(\epsilon)$.

Given such ρ, we define $g_j = g_j^{\rho,\eta}$ for $j \in I$ to complete the proof via Lemma 4.10.

Lemma 4.6 implies that the first property holds with probability $1 - 1/(2m)$. We go on to the second property. Let $j \in I \setminus F$. If $S_j \subseteq J$ then $g_j^{\rho,\eta} = f'_j$ always. Otherwise, $|S_j \setminus J| \ge M_j$, and so Lemma 4.11 implies that $g_j^{\rho,\eta} = f'_j$ with probability at least $1 - 1/(2m)$ (crucially, the property $\chi_j^{\rho,\eta} = *$ doesn't depend on the input $x^{(j)}$ to J). In total, the first two properties hold with probability at least $1/2$.

Lemma 4.9 implies that

$$\mathop{\mathbb{E}}_{\rho} \sum_{j \in F} \Pr_{\mu|_j}[g_j^{\rho,\eta} \neq f'_j] = O(\epsilon).$$

Conditioning on the first two properties, this still holds (with the right-hand side doubled). Hence there exists a restriction ρ satisfying all three properties, completing the proof. □

4.4 Concluding the proof

In this short section, we complete the proof of Theorem 1.5 by handling short affine relations.

Proof of Theorem 1.5. We prove the theorem with an error term of $O(\epsilon)$ rather than ϵ, for simplicity.

By negating coordinates, we can assume that all short affine relations are of the following forms: (i) $w_j = 0$ for all $w \in P$, (ii) $w_j = w_k$ for all $w \in P$.

Let Z be the set of coordinates in P that are always 0. Let I consist of a choice of one coordinate from each set of equivalent non-constant coordinates. For each $j \in I$, let I_j be the coordinates equivalent to j; possibly $I_j = \{j\}$. The sets $Z, (I_j)_{j \in I}$ thus partition $[m]$.

The predicate $P|_I$ has no short affine relations, and so the special case handled in Section 4.3 applies to it, giving us a value of δ. We will prove the theorem for $\delta := \min(\delta, \epsilon, 1)$.

Let $(f_1, \ldots, f_m)$ be a (μ, δ)-approximate generalized polymorphism of P. Applying the special case, we obtain a generalized polymorphism $(g_j)|_{j \in I}$ of $P|_I$. We complete it to a generalized polymorphism of P by taking $g_j = 0$ for $j \in Z$ and $g_k = g_j$ for $k \in I_j$ (where $j \in I$).

It remains to bound the distance between f_j and g_j for $j \notin I$. If $j \in Z$ then $f_j(0, \ldots, 0) = 0$ since $(f_1, \ldots, f_m)$ is an approximate polymorphism, and so $\Pr_{\mu|_j}[g_j \neq f_j] = 0$. If $k \in I_j$ then $\Pr_{\mu|_j}[f_k \neq f_j] \leq \delta$ and so $\Pr_{\mu|_j}[g_k \neq f_k] \leq \epsilon + \delta \leq 2\epsilon$. This completes the proof. $\square$

5 Linearity testing

In this section we prove Theorem 1.6, which extends the BLR test to arbitrary distributions. The proof uses ideas from [DFH25], which extended a related test to a specific family of distributions.

Theorem 1.6 (Linearity testing for general distributions)**.** *Let $P_{m,b} = \{(a_1, \ldots, a_m) \in \{0,1\}^m : a_1 \oplus \cdots \oplus a_m = b\}$, where $m \geq 3$ and $b \in \{0,1\}$, and let μ be a distribution on $P_{m,b}$ with full support. The following holds for $\delta = \Theta(\epsilon)$.*

If $f_1, \ldots, f_m \colon \{0,1\}^n \to \{0,1\}$ is a (μ, δ)-approximate generalized polymorphism of $P_{m,b}$ then there exists a generalized polymorphism $g_1, \ldots, g_m \colon \{0,1\}^n \to \{0,1\}$ of $P_{m,b}$ such that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon$ for all $j \in [m]$.

Moreover, there exist a set $S \subseteq [m]$ and $b_1, \ldots, b_m \in \{0,1\}$ such that $g_j(x) = \bigoplus_{i \in S} x_i \oplus b_j$.

Furthermore, if $f_i = f_j$ and $\mu|_i = \mu|_j$ then $g_i = g_j$.

Before describing the proof, we observe when ϵ is small enough (an assumption we can make without loss of generality), we can guarantee the "furthermore" clause. Indeed, if $f_i = f_j$ and $\mu|_i = \mu|_j$ then

$$\Pr_{\mu|_i}[g_i \neq g_j] \leq 2\epsilon.$$

As shown during the course of the proof of Lemma 4.11, for small enough ϵ (as a function only of $\mu|_i$) this implies that $g_i = g_j$.

The starting point of the proof is the case in which μ is the uniform distribution over $P_{m,b}$, which we denote by π. This case is standard, but for completeness we prove it in Section 5.1.

The general case is handled by reduction to the uniform distribution, using an agreement theorem from [DFH25]. The basic idea is to sample $x \sim \mu$ in two steps. In the first step, we sample $x \sim \nu$ with probability $1 - q$ (for appropriate ν, q), and leave it unsampled with probability q. In the second case, if x wasn't sampled, we sample it according to π.

If we stop after the first step, we are in a position to apply the result of Section 5.1, deducing that the resulting subfunctions are close to characters. We show that these characters agree with each other, and apply the agreement theorem to deduce that on average, they are restrictions of the same character (up to sign).

Our argument will require two agreement theorems: one from [DFH25], and a folklore result about mixing Markov chains which we prove in Section 5.2. The reduction itself is described in Section 5.3, where we also provide a more detailed outline.

5.1 Uniform distribution

In this section we prove Theorem 1.6 in the case of the uniform distribution. The proof follows the classical argument of [BCH+96].

Proof of Theorem 1.6 for the uniform distribution. Since we are going to use Fourier analysis, it will be more convenient to switch from $\{0,1\}$ to $\{-1,1\}$. Accordingly, we assume that $f_1, \ldots, f_m \colon \{-1,1\}^n \to \{-1,1\}$ satisfy

$$\Pr_{x^{(1)}, \ldots, x^{(m)}}[f_1(x^{(1)}) \cdots f_m(x^{(m)}) = B] \geq 1 - \delta,$$

where $B = (-1)^b$ and $x^{(1)}, \dots, x^{(m)}$ are sampled as follows: the first $m-1$ are sampled uniformly and independently, and $x_i^{(m)} = Bx_i^{(1)} \cdots x_i^{(m-1)}$.

The assumption is equivalent to

$$\mathop{\mathbb{E}}_{x^{(1)},\dots,x^{(m)}}[Bf_1(x^{(1)}) \cdots f_m(x^{(m)})] \geq 1 - 2\delta.$$

Substituting the Fourier expansions gives

$$1 - 2\delta \leq \sum_{S_1,\dots,S_m} B\hat{f}_1(S_1) \cdots \hat{f}_m(S_m) \mathop{\mathbb{E}}_{x^{(1)},\dots,x^{(m-1)}}[\chi_{S_1}(x^{(1)}) \cdot \chi_{S_{m-1}}(x^{(m-1)})\chi_{S_m}(Bx^{(1)} \cdots x^{(m-1)})],$$

where $\chi_S(x) = \prod_{i \in S} x_i$. Multiplicativity and orthogonality of characters shows that the expectation vanishes unless all S_j are equal, and so we obtain

$$1 - 2\delta \leq \sum_S B^{|S|+1} \hat{f}_1(S) \cdots \hat{f}_m(S) \leq \max_S |\hat{f}_1(S)| \sum_S |\hat{f}_2(S) \cdots \hat{f}_m(S)| \leq \max_S |\hat{f}_1(S)| \sum_S |\hat{f}_2(S)\hat{f}_3(S)|,$$

where we used $m \geq 3$.

At this point, we apply the Cauchy–Schwarz inequality to bound the second factor on the right:

$$1 - 2\delta \leq \max_S |\hat{f}_1(S)| \sqrt{\sum_S \hat{f}_2(S)^2} \sqrt{\sum_S \hat{f}_3(S)^2} = \max_S |\hat{f}_1(S)|.$$

Therefore there exists S_1 such that $|\hat{f}_1(S_1)| \geq 1 - 2\delta$. Let B_1 be the sign of $\hat{f}_1(S_1)$. Then

$$\Pr[f_1 = B_1\chi_{S_1}] \geq 1 - \delta.$$

Accordingly, we choose $\delta = \epsilon$.

In exactly the same way, we can find $B_i\chi_{S_i}$ approximating $f_2, \dots, f_m$:

$$\Pr[f_i = B_i\chi_{S_i}] \geq 1 - \delta.$$

It follows that for $B' = B_1 \cdots B_m$,

$$\Pr_{x^{(1)},\dots,x^{(m)}}[f_1(x^{(1)}) \cdots f_m(x^{(m)}) = B'\chi_{S_1}(x^{(1)}) \cdots \chi_{S_m}(x^{(m)})] \geq 1 - m\delta,$$

and so

$$\Pr_{x^{(1)},\dots,x^{(m)}}[B'\chi_{S_1}(x^{(1)}) \cdots \chi_{S_m}(x^{(m)}) = B] \geq 1 - (m+1)\delta.$$

At this point we recall the definition of $x^{(m)}$, which implies that

$$B'\chi_{S_1}(x^{(1)}) \cdots \chi_{S_m}(x^{(m)}) = B'B^{|S_m|}\chi_{S_1 \triangle S_m}(x^{(1)}) \cdots \chi_{S_{m-1} \triangle S_m}(x^{(m-1)}).$$

If not all S_i are equal then

$$\Pr_{x^{(1)},\dots,x^{(m)}}[B'\chi_{S_1}(x^{(1)}) \cdots \chi_{S_m}(x^{(m)}) = B] = \frac{1}{2},$$

which we can rule out by ensuring, without loss of generality, that $\epsilon = \delta < 1/(2(m+1))$. Thus all S_i are equal, concluding the proof. $\square$

5.2 Agreement lemma for mixing Markov chains

During the proof of Theorem 1.6 we will encounter the following situation. There is a function $f\colon X \to \Sigma$ and a way to sample $x, y \in X$ in a coupled fashion such that (x, y) and (y, x) have the same distribution. Given $\Pr[f(x) \neq f(y)] \leq \epsilon$, we would like to deduce that $\Pr[f(x) \neq \sigma] = O(\epsilon)$ for some $\sigma \in \Sigma$.

We can describe the coupling as a symmetric bistochastic $X \times X$ matrix M whose entries describe the coupling. We denote by $\lambda(M)$ the second largest eigenvalue of M. If $\lambda(M) < 1$ then M has a unique stationary distribution $\mu(M)$.

Lemma 5.1 (Agreement lemma for Markov chains)**.** *Let X, Σ be finite sets, and let M be a symmetric bistochastic $X \times X$ matrix with $\lambda = \lambda(M) < 1$ and stationary distribution $\mu = \mu(M)$.*

If a function $f\colon X \to \Sigma$ satisfies

$$\Pr_{(x,y)\sim M}[f(x) \neq f(y)] \leq \epsilon$$

then there exists $\sigma \in \Sigma$ such that

$$\Pr_{x\sim\mu}[f(x) \neq \sigma] \leq \frac{\epsilon}{1-\lambda}.$$

Proof. For $\sigma \in \Sigma$, let

$$f_\sigma(x) = \begin{cases} 1 & \text{if } f(x) = \sigma, \\ 0 & \text{if } f(x) \neq \sigma. \end{cases}$$

Let $g_\sigma(x) = f_\sigma(x) - \mathbb{E}_\mu[f_\sigma]$, so that $\mathbb{E}_\mu[g_\sigma] = 0$. The assumptions on M imply that

$$\mathop{\mathbb{E}}_{(x,y)\sim M}[g_\sigma(x)g_\sigma(y)] \leq \lambda \mathop{\mathbb{E}}_{x\sim\mu}[g_\sigma(x)^2].$$

We can relate the left-hand side to agreement for f_σ:

$$\begin{aligned}\Pr_{(x,y)\sim M}[f_\sigma(x) \neq f_\sigma(y)] &= \mathop{\mathbb{E}}_{(x,y)\sim M}[(f_\sigma(x) - f_\sigma(y))^2] \\ &= \mathop{\mathbb{E}}_{(x,y)\sim M}[(g_\sigma(x) - g_\sigma(y))^2] = 2 \mathop{\mathbb{E}}_{x\sim\mu}[g_\sigma(x)^2] - 2 \mathop{\mathbb{E}}_{(x,y)\sim M}[g_\sigma(x)g_\sigma(y)].\end{aligned}$$

It follows that

$$\Pr_{(x,y)\sim M}[f_\sigma(x) \neq f_\sigma(y)] \geq 2 \mathop{\mathbb{E}}_{x\sim\mu}[g_\sigma(x)^2] - 2 \mathop{\mathbb{E}}_{(x,y)\sim M}[g_\sigma(x)g_\sigma(y)] \geq 2(1-\lambda) \mathop{\mathbb{E}}_{x\sim\mu}[g_\sigma(x)^2].$$

In a completely analogous way, we obtain

$$\Pr_{x,y\sim\mu}[f_\sigma(x) \neq f_\sigma(y)] = 2 \mathop{\mathbb{E}}_{x\sim\mu}[g_\sigma(x)^2] - 2 \mathop{\mathbb{E}}_{x,y\sim\mu}[g_\sigma(x)g_\sigma(y)] = 2 \mathop{\mathbb{E}}_{x\sim\mu}[g_\sigma(x)^2] \leq \frac{\Pr_{(x,y)\sim M}[f_\sigma(x) \neq f_\sigma(y)]}{1-\lambda}.$$

We now circle back to f. Summing the above inequality over all σ, we obtain

$$\sum_{\sigma\in\Sigma} \Pr_{x,y\sim\mu}[f_\sigma(x) \neq f_\sigma(y)] \leq \frac{1}{1-\lambda} \sum_{\sigma\in\Sigma} \Pr_{(x,y)\sim M}[f_\sigma(x) \neq f_\sigma(y)] = \frac{2\Pr_{(x,y)\sim M}[f(x) \neq f(y)]}{1-\lambda},$$

since given that $f(x) \neq f(y)$, we obtain $f_\sigma(x) \neq f_\sigma(y)$ for precisely two choices of σ, namely $f(x)$ and $f(y)$. As for the left-hand side, it is equal to

$$2\sum_{\sigma\in\Sigma} \Pr_{x\sim\mu}[f(x) \neq \sigma] \Pr_{y\sim\mu}[f(y) = \sigma] \geq 2 \min_{\sigma\in\Sigma} \Pr_{x\sim\mu}[f(x) \neq \sigma] \sum_{\sigma\in\Sigma} \Pr_{y\sim\mu}[f(y) = \sigma] = 2 \min_{\sigma\in\Sigma} \Pr_{x\sim\mu}[f(x) \neq \sigma].$$

Plugging this in the previous inequality, we conclude that

$$\min_{\sigma\in\Sigma} \Pr_{x\sim\mu}[f(x) \neq \sigma] \leq \frac{\Pr_{(x,y)\sim M}[f(x) \neq f(y)]}{1-\lambda} \leq \frac{\epsilon}{1-\lambda}. \qquad \square$$

In our application, M will be a product chain where each factor has full support (a condition which can be replaced by ergodicity) and there are finitely many types of factors. In this case we can bound $\lambda(M)$ irrespective of the number of factors.

Lemma 5.2 (Agreement lemma for product chains)**.** *Let Y, Σ be finite sets, and let $M_1, \ldots, M_t$ be a finite number of symmetric bistochastic $Y \times Y$ matrices with strictly positive entries. There exists $C > 0$ such that the following holds.*

Given $\psi\colon [n] \to [t]$, define the distribution M_ψ on $Y^n \times Y^n$ by sampling the i'th coordinate independently according to $M_{\psi(i)}$. Let μ_ψ be the corresponding stationary distribution, obtained by sampling the i'th coordinate independently according to $\mu(M_{\psi(i)})$.

If a function $f: Y^n \to \Sigma$ satisfies

$$\Pr_{(x,y)\sim M_\psi}[f(x) \neq f(y)] \leq \epsilon$$

then there exists $\sigma \in \Sigma$ such that

$$\Pr_{x\sim\mu_\psi}[f(x) \neq \sigma] \leq C\epsilon.$$

Proof. Given Lemma 5.1, it suffices to find λ such that $\lambda(M_\psi) \leq \lambda$ for all ψ; we can then take $C = 1/(1-\lambda)$.

Since all entries of $M_1, \ldots, M_t$ are positive, we can find $\lambda < 1$ such that all eigenvalues of $M_1, \ldots, M_t$ other than the top ones are bounded by λ in absolute value. It immediately follows that $\lambda(M_\psi) \leq \lambda$. $\square$

5.3 Arbitrary distributions

In this section we prove Theorem 1.6 in full generality, by reduction to the special case $\mu = \pi$ (recall that π is the uniform distribution over $P_{m,b}$).

For small enough $q > 0$, we can find a distribution ν such that

$$\mu = q\pi + (1-q)\nu.$$

Indeed, such a distribution exists whenever $q \leq \min_{w\in P_{m,b}} \mu(w)/\pi(w)$. We furthermore choose q so that $1/q$ is an integer (this will slightly simplify some future argument).

We can sample $x \sim \mu^n$ using a three-step process:

1. Let $S \sim \mu_q^n$. This means that S is a subset of $[n]$ chosen by including each i with probability q independently.
2. If $i \notin S$, sample x_i according to ν.
3. If $i \in S$, sample x_i according to π.

If we stop after the first two steps, then the remaining step is a sample from π^S, which we can analyze using the special case of Theorem 1.6.

Lemma 5.3 (Reduction to uniform distribution)**.** *For every $S \subseteq [m]$ and $\alpha \in P_{m,b}^S$ there exist $b_1(S,\alpha), \ldots, b_m(S,\alpha) \in \{0,1\}$ and $A(S,\alpha) \subseteq S$ such that for all $j \in [m]$,*

$$\mathop{\mathbb{E}}_{(S,\alpha)\sim(\mu_q^n,\nu^{S^c})} \Pr_{\pi|_j}[f_j|_{S^c\leftarrow\alpha|_j} \neq \chi_{A(S,\alpha),b_j(S,\alpha)}] = O(\epsilon).$$

Here $\chi_{A,b}(x) = b \oplus \bigoplus_{i\in A} x_i$. We also use $\chi_A = \chi_{A,0}$.

Proof. For every S and every $\alpha \in P_{m,b}^S$, define

$$\epsilon(S,\alpha) = \Pr_{\beta\sim\pi^S}[(f_1|_{S^c\leftarrow\alpha|_1}(\beta|_1), \ldots, f_m|_{S^c\leftarrow\alpha|_m}(\beta|_m)) \notin P_{m,b}],$$

observing that

$$\mathop{\mathbb{E}}_{(S,\alpha)\sim(\mu_q^n,\nu^{S^c})}[\epsilon(S,\alpha)] = \Pr_{\mu^S}[(f_1,\dots,f_m) \notin P_{m,b}] \le \epsilon.$$

Apply the special case of the uniform distribution to every (S,α) to obtain $b_1(S,\alpha),\dots,b_m(S,\alpha) \in \{0,1\}$ and $A(S,\alpha) \subseteq S$ such that for all $j \in [m]$,

$$\Pr_{\pi|_j}[f_j|_{S^c \leftarrow \alpha|_j} \neq \chi_{A(S,\alpha),b_j(S,\alpha)}] = O(\epsilon(S,\alpha)).$$

The result immediately follows. $\square$

The rest of the proof comprises the following steps:

1. Using the agreement theorem of [DFH25] we show that for each $\alpha \in P_{m,b}^n$ there exists a consensus set $A(\alpha)$ such that typically $A(S,\alpha|_{S^c}) = A(\alpha) \cap S$.
2. Using Lemma 5.2 we show that there exists a consensus set A such that typically $A(\alpha) = A$.
3. Using Lemma 5.2 we show that the functions $f_j \oplus \chi_{A,b_j}$ are close to constant, completing the proof.

5.3.1 Step 1

The first step constitutes the following lemma.

Lemma 5.4 (Agreement over S)**.** *For every $\alpha \in P_{m,b}^n$ there exists $A(\alpha) \subseteq [n]$ such that*

$$\mathop{\mathbb{E}}_{\alpha\sim\nu^n} \Pr_{S\sim\mu_q}[A(S,\alpha|_{S^c}) \neq A(\alpha) \cap S] = O(\epsilon).$$

The proof will require an agreement theorem essentially proved in [DFH25]. The theorem concerns the distribution $\mu_{q,r}^n$, where $0 < q < r < 1$. This is the distribution on triples (S_1,S_2,T) defined as follows:

- Sample $T \sim \mu_r^n$.
- Sample $S_1 \supseteq T$ so that $S_1 \sim \mu_q^n$.
- Sample $S_2 \supseteq T$ so that $S_2 \sim \mu_q^n$.

The sampling in the second and third steps can be performed as follows. If $i \in T$, then we always put $i \in S_1$, and otherwise, we put it in S_1 with probability $\frac{q-r}{1-r}$.

Theorem 5.5 (Agreement theorem)**.** *Let $0 < r < q < 1$. Suppose that for every $S \subseteq [n]$ we have a function $\phi_S : S \to \Sigma$, where Σ is some finite alphabet. If*

$$\Pr_{(S_1,S_2,T)\sim\mu_{q,r}^n}[\phi_{S_1}|_T \neq \phi_{S_2}|_T] \le \epsilon$$

then there exists $\psi : [n] \to \Sigma$ such that

$$\Pr_{S\sim\mu_q^n}[\phi_S \neq \psi|_S] = O(\epsilon).$$

Proof. The slice version of this result is [DFH25, Theorem 3.1]. This version can be proved using the "going to infinity" argument which is used to prove [DFH25, Theorem 5.4]. $\square$

We can now prove the lemma.

Proof of Lemma 5.4. We would like to get into the setting of Theorem 5.5, for $q := q$, an appropriate r, and $\phi_S := A(S, \alpha|_S)$, for various values of α.

For small enough $c > 0$ we can find a distribution λ such that

$$\pi = c\nu + (1-c)\lambda.$$

We take $r = (1-c)q$. We will show that

$$\mathop{\mathbb{E}}_{\alpha \sim \nu^n} \Pr_{(S_1,S_2,T)\sim\mu_{q,r}^n} [A(S_1, \alpha|_{S_1^c}) \cap T \neq A(S_2, \alpha|_{S_2^c}) \cap T] = O(\epsilon). \tag{1}$$

Applying Theorem 5.5 for each α separately then immediately implies the lemma.

In order to show that Equation (1) holds, we use the fact that $f_1|_{S_1 \leftarrow \alpha|_{S_1}}$ is close to $\chi_{A(S_1,\alpha|_{S_1})}$ (up to negation), and similarly $f_1|_{S_2 \leftarrow \alpha|_{S_2}}$ is close to $\chi_{A(S_2,\alpha|_{S_2})}$ (up to negation). Fixing the elements outside of T, we obtain that the same function is close to both $\chi_{A(S_1,\alpha|_{S_1}) \cap T}$ and $\chi_{A(S_2,\alpha|_{S_2}) \cap T}$ (up to negation), which can only happen if $A(S_1, \alpha|_{S_1}) \cap T = A(S_2, \alpha|_{S_2}) \cap T$ (since different characters are far from each other).

The first step in this plan is to massage the conclusion of Lemma 5.3 using the equation for π:

$$\mathop{\mathbb{E}}_{\substack{(S,\alpha')\sim(\mu_q^n,\nu^{S^c}) \\ (T,\alpha'')\sim(\mu_{1-c}^S,\nu^{S\setminus T})}} \min_{B\in\{0,1\}} \Pr_{\lambda|_1^T}[f_1|_{S^c\leftarrow\alpha'|_1, S\setminus T\leftarrow\alpha''|_1} \neq \chi_{A(S,\alpha')\cap T,B}] = O(\epsilon).$$

(We could determine B explicitly, but this is not necessary.)

An equivalent way to sample $(S,\alpha'), (T,\alpha'')$ is to first sample S, T, then sample $\alpha \sim \nu^n$, and take $\alpha' = \alpha|_{S^c}$ and $\alpha'' = \alpha|_{S\setminus T}$. In particular, for $t \in \{1, 2\}$,

$$\mathop{\mathbb{E}}_{\alpha\sim\nu^n} \mathop{\mathbb{E}}_{(S_1,S_2,T)\sim\mu_{q,(1-c)q}^n} \min_{B_t\in\{0,1\}} \Pr_{\lambda|_1^T}[f_1|_{T^c\leftarrow\alpha|_{T^c},1} \neq \chi_{A(S_t,\alpha|_{S_t^c})\cap T,B_t}] = O(\epsilon).$$

Combining this for $t = 1$ and $t = 2$ gives

$$\mathop{\mathbb{E}}_{\alpha\sim\nu^n} \mathop{\mathbb{E}}_{(S_1,S_2,T)\sim\mu_{q,(1-c)q}^n} \min_{B_1,B_2\in\{0,1\}} \Pr_{\lambda|_1^T}[\chi_{A(S_1,\alpha|_{S_1^c})\cap T,B_1} \neq \chi_{A(S_2,\alpha|_{S_2^c})\cap T,B_2}] = O(\epsilon).$$

At this point we use the fact that different characters are far from each other: if $A_1 \neq A_2$ then for all B_1, B_2,

$$\Pr_{\lambda|_1^T}[\chi_{A_1,B_1} \neq \chi_{A_2,B_2}] \geq \min(\lambda|_1(0), \lambda|_1(1)).$$

Since λ is fixed, this immediately implies Equation (1), completing the proof. $\square$

5.3.2 Step 2

The second step constitutes the following lemma.

Lemma 5.6 (Agreement over α)**.** *There exists $A \subseteq [n]$ such that*

$$\Pr_{\alpha\sim\nu^n}[A(\alpha) \neq A] = O(\epsilon).$$

Proof. The proof relies on the fact that while $A(\alpha)$ is a function of all of α, $A(S, \alpha|_{S^c})$ only depends on part of α. In particular, suppose that we sample two copies of ν^n by sampling $\alpha \sim \nu^n$ and then obtaining α' by only resampling the coordinates in S. Denote this distribution by $\nu(S)$. Since $\alpha|_{S^c} = \alpha'|_{S^c}$, Lemma 5.4 shows that

$$\Pr_{(S,\alpha,\alpha')\sim(\mu_q^n,\nu(S))}[A(\alpha) \cap S \neq A(\alpha') \cap S] = O(\epsilon).$$

In order to proceed, we introduce a conceit:

$$\Pr_{\substack{(S,\alpha,\alpha')\sim(\mu_q^n,\nu(S)) \\ T\sim\mu_{1/2}(S)}}[A(\alpha) \cap T \neq A(\alpha') \cap T] \leq \Pr_{(S,\alpha,\alpha')\sim(\mu_q^n,\nu(S))}[A(\alpha) \cap S \neq A(\alpha') \cap S] = O(\epsilon).$$

If we marginalize over S, the distribution of (α, α'), which we denote $\nu'(T)$, becomes the following. First, sample $\alpha \sim \nu^n$. Then, sample $S \supseteq T$ by including any $i \notin T$ with probability $\frac{q/2}{1-q/2}$. Finally, resample all coordinates in S. We deduce

$$\Pr_{(T,\alpha,\alpha')\sim(\mu_{q/2}^n,\nu'(T))}[A(\alpha) \cap T \neq A(\alpha') \cap T] = O(\epsilon).$$

In contrast to the distribution $\nu(S)$, which only mixes the coordinates in S, the distribution $\nu'(T)$ mixes all coordinates. Moreover, there are only two types of coordinates: those in T (which always get resampled) and the rest (which get resampled with probability $\frac{q/2}{1-q/2}$). This allows us to apply Lemma 5.2, obtaining $A'(T) \subseteq T$ satisfying

$$\mathop{\mathbb{E}}_{T\sim\mu_{q/2}^n} \Pr_{\alpha\sim\nu^n}[A(\alpha) \cap T \neq A'(T)] = O(\epsilon).$$

Now it's time for the final trick. Recall that $1/q$ is an integer. Choose $c\colon [n] \to [2/q]$ uniformly at random, and for $t \in [2/q]$, let $T_t(c) = c^{-1}(t)$. Since $T_t \sim \mu_{q/2}^n$ for each t, we have

$$\mathop{\mathbb{E}}_{c} \Pr_{\alpha\sim\nu^n}[A(\alpha) \cap T_t(c) \neq A'(T_t(c)) \text{ for some } t] = O(\epsilon).$$

We can find c such that the probability above is $O(\epsilon)$. Define A via $A \cap T_t(c) = A'(T_t(c))$; this makes sense since $T_1(c), \ldots, T_{2/q}(c)$ partition $[n]$. The lemma immediately follows. □

5.3.3 Step 3

In the final step, we complete the proof of Theorem 1.6.

Proof of Theorem 1.6. The main idea is to "factor out" the set A found in Lemma 5.6. Accordingly, we define functions $h_1, \ldots, h_m\colon \{0,1\}^n \to \{0,1\}$ by

$$h_j = f_j \oplus \chi_A.$$

We will complete the proof by showing that h_j is $O(\epsilon)$-close to a constant.

As our starting point, we combine Lemmas 5.3, 5.4 and 5.6 to obtain

$$\mathop{\mathbb{E}}_{(S,\alpha)\sim(\mu_q^n,\nu^{S^c})} \Pr_{\pi|_j^S}[f_j|_{S^c\leftarrow\alpha|_j} \neq \chi_{A\cap S,b_j(S,\alpha)}] = O(\epsilon).$$

Plugging in h_j, we obtain

$$\mathop{\mathbb{E}}_{(S,\alpha)\sim(\mu_q^n,\nu^{S^c})} \Pr_{\pi|_j^S}[h_j|_{S^c\leftarrow\alpha|_j} \neq \chi_{A\cap S^c,b_j(S,\alpha)}(\alpha|_{S^c,j})] = O(\epsilon).$$

The expression $\chi_{A\cap S^c,b_j(S,\alpha)}(\alpha|_{S^c,j})$ doesn't depend on the coordinates in S, and so

$$\mathop{\mathbb{E}}_{(S,\alpha)\sim(\mu_q^n,\nu^{S^c})} \Pr_{\beta,\beta'\sim\pi|_j^S}[h_j|_{S^c\leftarrow\alpha|_j}(\beta) \neq h_j|_{S^c\leftarrow\alpha|_j}(\beta')] = O(\epsilon).$$

Let $\gamma \in P_{m,b}^n$ be the vector defined by $\gamma|_{S^c} = \alpha$ and $\gamma|_S = \beta$, and let $\gamma' \in P_{m,b}^n$ be defined similarly with $\gamma'|_S = \beta'$. The inputs to h_j in the formula above are thus $\gamma|_j$ and $\gamma'|_j$.

We can sample γ, γ' as follows. First, sample $\gamma \in \mu^n$. Next, sample S given γ; the (non-zero) probability that $i \in S$ depends only on γ_i, and can be calculated using Bayes' law. Finally, γ' is obtained by resampling the coordinates in S according to π. This means that the distribution of (γ, γ') corresponds to some product chain in the sense of Lemma 5.2, with a single type of factor. The same holds for $(\gamma|_j, \gamma'|_j)$, and so Lemma 5.2 implies that there exists $b_j \in \{0,1\}$ such that

$$\Pr_{\mu^n}[h_j \neq b_j] = O(\epsilon).$$

Finally, we take $g_j = \chi_{A,b_j}$. Observe that $\Pr_{\mu|_j^n}[g_j \neq f_j] = O(\epsilon)$, and so

$$\Pr_{\mu^n}[(g_1,\dots,g_m) \notin P_{m,b}] = O(\epsilon).$$

Whether $(g_1,\dots,g_m) \in P_{m,b}$ or not depends only on $b_1,\dots,b_m$ (the dependence on the input cancels out). We can assume without loss of generality that ϵ is small enough to make the right-hand side smaller than 1. This implies that $(g_1,\dots,g_m)$ is a generalized polymorphism of $P_{m,b}$, completing the proof. □

6 Intersecting families

In this section we prove Theorem 1.9, which improves on results of Friedgut and Regev [FR18] (their paper, in turn, improves on [DFR08, DF09]).

Theorem 1.9 (Improved Friedgut–Regev). *Fix $0 < p < 1/2$. For every $\epsilon > 0$ the following holds for all n such that pn is an integer.*

If $\mathcal{F} \subseteq \binom{[n]}{pn}$ contains a $1/\exp\Theta(1/\epsilon^C)$-fraction of the edges of the Kneser graph then there exists an intersecting family $\mathcal{G} \subseteq \binom{[n]}{pn}$ computed by a decision tree of depth $O(1/\epsilon^C)$ (for some global constant C) such that $|\mathcal{F} \setminus \mathcal{G}| \leq \epsilon\binom{n}{pn}$.

The proof relies on the machinery of Friedgut and Regev, which reduces it to the following statement.

Theorem 6.1 (Fractional improved Friedgut–Regev). *Fix $p \in (0,1/2)$. Let $P_{\mathsf{NAND}} = \{(0,0),(1,0),(0,1)\}$, and let μ be the following distribution: $\mu(0,0) = 1-2p$, $\mu(1,0) = \mu(0,1) = p$. For every $\epsilon > 0$ there exist $\delta = 1/\exp\Theta(1/\epsilon^C)$ and $M = O(1/\epsilon^C)$ (for some global constant C) such that the following holds.*

Suppose that $f_1, f_2 \colon \{0,1\}^n \to [0,1]$ satisfy

$$\mathop{\mathbb{E}}_{(x^{(1)},x^{(2)})\sim\mu^n}[f_1(x^{(1)})f_2(x^{(2)})] \leq \delta.$$

Then there exist $g_1, g_2 \colon \{0,1\}^n \to \{0,1\}$, computed by a decision tree of depth at most M, such that (g_1,g_2) is a generalized polymorphism of P_{NAND}, and moreover for $j \in \{1,2\}$,

$$\mathop{\mathbb{E}}_{x^{(j)}\sim\mu|_j^n}[(1-g_j(x^{(j)}))f_j(x^{(j)})] \leq \epsilon.$$

Furthermore, if $f_1 = f_2$ then $g_1 = g_2$.

Notice that in this statement, the functions f_1, f_2 are not necessarily Boolean, but they take values in the interval $[0,1]$.

Let us briefly explain how Theorem 1.9 follows from Theorem 6.1. Given $\mathcal{F}$, which we identify with the corresponding Boolean function, we define

$$f_1(x) = f_2(x) = \begin{cases} \mathop{\mathbb{E}}_{\substack{x' \leq x \\ |x'|=k}}[\mathcal{F}(x')] & \text{if } |x| \geq k, \\ 0 & \text{otherwise.} \end{cases}$$

[FR18, Lemma 7.3] shows that the assumption of Theorem 1.9 implies that of Theorem 6.1. We apply Theorem 6.1, and convert $g_1 = g_2$ to a family $\mathcal{G}$ in the natural way. [FR18, Lemma 7.4] shows that the conclusion of Theorem 6.1 implies the conclusion of Theorem 1.9.

The proof of Theorem 6.1 is very similar to the proof of Theorem 1.7. First, we need versions of Theorem 3.2 (It Ain't Over Till It's Over) and Theorem 3.3 (Jones' regularity lemma) for functions taking values in $[0,1]$.

The original proof of Theorem 3.2 [MOO10] was in fact formulated for functions taking values in $[0,1]$. (Note that our definition of regularity makes sense for arbitrary real-valued functions.)

Theorem 6.2 (Fractional It Ain't Over Till It's Over). *Theorem 3.2 holds for functions $f\colon \{0,1\}^n \to [0,1]$ (with the same parameters).*

A simple modification of the proof of Jones' regularity lemma, which we present in Section 8, extends it to the same setting.

Theorem 6.3 (Fractional Jones' regularity lemma). *Theorem 3.3 holds for functions $f\colon \{0,1\}^n \to [0,1]$ (with the same parameters).*

Upon inspection, the proof of Lemma 3.5 (Counting lemma for NAND) translates to the following statement.

Lemma 6.4 (Fractional counting lemma for NAND). *Fix $p \in (0,1/2)$, and let P_{NAND}, μ be as in Theorem 6.1. For every $\epsilon > 0$ there exists a constant $C = C(P,\mu)$ such that the following holds for $d = \Theta(\log(1/\epsilon))$, $\tau = \Theta(\epsilon^C)$, and $\gamma = \Theta(\epsilon^{2C})$.*

Let $\phi_1,\ldots,\phi_m\colon \{0,1\}^n \to [0,1]$ be functions such that ϕ_j is (d,τ)-regular with respect to $\mu|_j$ and $\mathbb{E}_{\mu|_j}[\phi_j] \geq \epsilon$. Then

$$\mathop{\mathbb{E}}_{(y^{(1)},\ldots,y^{(2)})\sim\mu^n}[\phi_1(y^{(1)})\phi_2(y^{(2)})] \geq \gamma.$$

Proof sketch. In the proof of Lemma 3.5 we define a random restriction $\rho \sim \nu^n$, where ν is supported on inputs in $\{0,1,*\}^m$ with at most one free coordinate. We show that for an appropriate choice of $d = \Theta(\log(1/\epsilon))$ and $\tau, \delta_1, \delta_2 = \Theta(1/\epsilon^C)$,

$$\Pr_{\rho}[\mathop{\mathbb{E}}_{\mu|_j}[\phi_j|_{\rho|_j}] \geq \delta_j] \geq 1-\epsilon \text{ for } j \in \{1,2\}.$$

If we sample $(y^{(1)},y^{(2)})|\rho$ then $\phi_1(y^{(1)}), \phi_2(y^{(2)})$ are independent since they depend on disjoint sets of coordinates. Therefore

$$\mathop{\mathbb{E}}_{(y^{(1)},\ldots,y^{(2)})\sim\mu^n}[\phi_1(y^{(1)})\phi_2(y^{(2)})] \geq (1-\epsilon)\delta_1\delta_2.$$

□

We can now prove Theorem 6.1, closely following the proof of Theorem 1.7.

Proof of Theorem 6.1. Apply Lemma 6.4 with $\epsilon := \epsilon/2$ to obtain $d = \Theta(\log(1/\epsilon))$ and $\tau, \gamma = \Theta(1/\epsilon^C)$.

We apply Theorem 6.3 with $\epsilon := \epsilon/2$ and the parameters d, τ to f_1, f_2 to obtain a decision tree T of depth $M = O(d/\epsilon\tau) = \tilde{O}(1/\epsilon^{2C})$. We define the functions $g_1, g_2\colon \{0,1\}^n \to \{0,1\}$ as follows. For each leaf $\rho \in T$,

$$g_j|_\rho = \begin{cases} 0 & \text{if } f_j|_\rho \text{ is not } (d,\tau)\text{-regular or } \mathbb{E}_{\mu|_j}[f_j|_\rho] \leq \epsilon/2, \\ 1 & \text{otherwise.} \end{cases}$$

If we sample ρ according to $\mu|_j$ then according to Jones' regularity lemma, $f_j|_\rho$ is (d,τ)-regular with probability at least $1-\epsilon/2$. This shows that $\mathbb{E}_{\mu|_j}[(1-g_j)f_j] \leq \epsilon/2 + \epsilon/2 = \epsilon$.

It remains to show that (g_1, g_2) is a generalized polymorphism of P_{NAND} for small enough δ. If this is not the case, then there exists a partial assignment $\rho \in T$ such that $(g_1|_{\rho|_1}, g_2|_{\rho|_2})$ is not a generalized polymorphism of P_{NAND}, which implies that $f_1|_{\rho|_1}, f_2|_{\rho|_2}$ are both regular and have expectation at least $\epsilon/2$. Lemma 6.4 thus implies that $\mathbb{E}_\mu[f_1|_{\rho|_1} f_2|_{\rho|_2}] \geq \gamma$, which is ruled out by defining $\delta = \min(\mu)^M\gamma/2 = 1/\exp\tilde{\Theta}(1/\epsilon^{2C})$. □

The proofs of Theorems 1.9 and 6.1 extend to other predicates. We leave this to future work.

7 General alphabets

In this section we prove Theorem 1.10, an analog of Theorem 1.5 for some predicates over larger alphabets.

Theorem 1.10 (Larger alphabets)**.** *Let Σ be a finite set, let $P \subseteq \Sigma^m$, and let μ be a distribution on P with full support. Suppose that for each $j \in [m]$ there exists $w \in P$ such that w remains in P even if we modify its j'th coordinate arbitrarily. For every $\epsilon > 0$ there exists $\delta > 0$ such that the following holds for all n.*

If $f_1, \ldots, f_m \colon \Sigma^n \to \Sigma$ is a (μ, δ)-approximate generalized polymorphism of P then there exists a generalized polymorphism $g_1, \ldots, g_m \colon \Sigma^n \to \Sigma$ of P such that $\Pr_{\mu|_j}[g_j \neq f_j] \leq \epsilon$ for all $j \in [m]$.

The proof follows the general outline of the proof of Theorem 1.5, but is much simpler. First, the counting lemma is direct rather than inductive. Second, there is no need to accommodate affine relations.

Before starting the proof proper, we need to generalize the notion of regularity to the setting of functions $f \colon \Sigma^n \to \Sigma$. The original proof of It Ain't Over Till It's Over [MOO10] actually works for arbitrary alphabets. Given a function $f \colon \Sigma^n \to \{0, 1\}$ and a distribution μ over Σ with full support, the Efron–Stein decomposition is the unique decomposition

$$f = \sum_S f_S$$

in which f_S depends only on the coordinates in S, the functions f_S are orthogonal, and f_S has zero expectation if we fix the values of coordinates in some set $T \not\supseteq S$. The definition of low-degree influences readily extends:

$$\mathrm{Inf}_i[f^{\leq d}] = \sum_{\substack{|S| \leq d \\ i \in S}} \|f_S\|^2,$$

where the norm is computed according to μ.

This prompts the following definition of regularity.

Definition 7.1 (Regularity for arbitrary alphabets)**.** Let Σ be a finite alphabet of size at least 2, let μ be a distribution on Σ with full support, and let $d \in \mathbb{N}$ and $\tau > 0$.

A function $f \colon \Sigma^n \to \Sigma$ is (d, τ)-regular with respect to μ if $\mathrm{Inf}_i[(f^{=\sigma})^{\leq d}]$ for all $i \in [n]$ and $\sigma = \Sigma$, where

$$f^{=\sigma}(x) = \begin{cases} 1 & \text{if } f(x) = \sigma, \\ 0 & \text{otherwise.} \end{cases}$$

With this definition, we can extend both It Ain't Over Till It's Over and Jones' regularity lemma.

Theorem 7.2 (It Ain't Over Till It's Over for arbitrary alphabets)**.** *For every alphabet Σ, full support distribution μ, $q \in (0, 1)$ and $\epsilon > 0$ the following holds for some constant C and $d = \Theta(\log(1/\epsilon))$, $\tau = \Theta(\epsilon^C)$, and $\delta = \Theta(\epsilon^C)$.*

Let ρ be a random restriction obtained by sampling each coordinate $i \in [n]$ independently according to the following law: with probability $1 - q$, draw a sample from μ, and otherwise, draw $*$.

If $f \colon \Sigma^n \to \Sigma$ is (d, τ)-regular with respect to μ then for every $\sigma \in \Sigma$ such that $\Pr_\mu[f = \sigma] \geq \epsilon$ we have

$$\Pr_\rho[\Pr[f|_\rho = \sigma] \geq \delta] \geq 1 - \epsilon.$$

Proof. Follows immediately from [MOO10] by considering the functions $f^{=\sigma}$. □

Jones' regularity lemma can also be extended, as we show in Section 8.

Theorem 7.3 (Jones' regularity lemma for arbitrary alphabets)**.** *For every alphabet Σ, $m \in \mathbb{N}$, full support distributions $\mu_1, \ldots, \mu_m$, and every $\epsilon, \tau > 0$, $d \in \mathbb{N}$ the following holds for some function $M \in \mathbb{N}$.*

For all functions $f_1, \ldots, f_m \colon \Sigma^n \to \Sigma$ there exists a set J of size at most M such that for all j,

$$\Pr_{x \sim \mu_j}[f|_{J \leftarrow x} \text{ is } (d, \tau)\text{-regular with respect to } \mu_j] \geq 1 - \epsilon.$$

While we can prove a counting lemma in the style of Lemma 3.4 or Lemma 4.3, the counting lemma that is useful here is similar to the one implicitly used to prove Lemma 4.6. We prove this lemma in Section 7.1, where we also prove an analog of Lemma 4.9. Combining these with Jones' regularity lemma, we complete the proof of Theorem 1.10 in Section 7.2.

For the rest of this section, we fix P and μ.

7.1 Counting lemma

The counting lemma that we prove uses a restriction similar to the one appearing in Section 4.2. In that section, we had to distinguish between flexible coordinates and inflexible coordinates. In our case, all coordinates are flexible by assumption.

Definition 7.4 (Restriction)**.** For every $j \in [m]$, let $w_{(j,*)}$ be a partial input, missing only the j'th coordinate, such that all of its completions $w_{(j,\sigma)}$ belong to P.

The distribution ν is a distribution on $Q := P \cup \{w_{(j,*)} : j \in [m]\}$ defined as follows, for a small enough $q > 0$:

- For $w \in P$, sample w with probability
$$\mu(w) - q \sum_{(j,\sigma)\colon w=w_{(j,\sigma)}} \mu|_j(\sigma).$$
- For every $j \in [m]$, sample $w_{(j,*)}$ with probability q.

Concretely, it suffices to take $q = \min(\mu)/m$, where $\min(\mu) = \min_{w \in P} \mu(w)$.

Given $\rho \in Q$, the distribution $\mu|\rho$ is obtained by sampling the missing coordinate in $w_{(j,*)}$ using $\mu|_j$.

By construction, if $\rho \sim \nu$ and $x \sim \mu|\rho$ then $x \sim \mu$. Moreover, the marginal distribution of ρ_j given $\rho_j \neq *$ is $\mu|_j$.

Lemma 7.5 (Counting lemma)**.** *Let $\phi_1, \ldots, \phi_m \colon \Sigma^n \to \Sigma$, and let $\rho \in Q^n$. Let $w \in \Sigma^m$ be such that for each $j \in [m]$,*
$$\Pr_{\mu|_j}[\phi_j|_{\rho|_j} = w_j] \geq \epsilon.$$
Then
$$\Pr_{\mu|\rho}[(\phi_1, \ldots, \phi_m) = w] \geq \epsilon^m.$$

Proof. Since ρ is supported on inputs with at most one free coordinate, the events $\phi_j|_{\rho|_j} = w_j$ are independent. The lemma immediately follows. □

We will eventually choose ρ and δ (a bound on $\Pr_\mu[(f_1, \ldots, f_m) \notin P]$) so that the counting lemma implies that any w satisfying the conditions in the lemma belongs to P. This suggests the following rounding procedure.

Definition 7.6 (Rounding)**.** Let $\phi_1, \ldots, \phi_m \in \Sigma^n \to \Sigma$ and let $\rho \in Q^n$.

For a parameter $\epsilon \in (0, 1)$, we define $\mathrm{round}_j^{\rho,\epsilon}(\phi_j) \colon \Sigma^n \to \Sigma$ as follows. First, let
$$\Sigma_{\geq\epsilon} := \{\sigma \in \Sigma : \Pr_{\mu|_j}[\phi_j|_{\rho_j} = \sigma] \geq \epsilon\},$$
and choose $\sigma_0 \in \Sigma_{\geq\epsilon}$ arbitrarily (say the symbol with the largest probability). Then
$$\mathrm{round}_j^{\rho,\epsilon}(\phi_j)(x) = \begin{cases} \phi_j(x) & \text{if } \phi_j(x) \in \Sigma_{\geq\epsilon}, \\ \sigma_0 & \text{otherwise.} \end{cases}$$

The following lemma bounds the error in rounding, and roughly corresponds to Lemma 4.9.

Lemma 7.7 (Rounding lemma)**.** *For every $\epsilon > 0$ there exist $d \in \mathbb{N}$ and $\tau, \eta > 0$ such that the following holds.*

Let $\phi_1, \ldots, \phi_m \colon \Sigma^n \to \Sigma$. If ϕ_j is (d,τ)-regular with respect to $\mu|_j$ for all j then for all j,

$$\mathop{\mathbb{E}}_{\rho \sim \nu^n} \Big[\Pr_{\mu|_j}[\mathrm{round}_j^{\rho,\eta}(\phi_j) \neq \phi_j]\Big] \leq \epsilon.$$

Proof. Apply Theorem 7.2 (It Ain't Over Till It's Over) with Σ, μ, q and $\epsilon := \epsilon/|\Sigma|$ to obtain d, τ, δ. We take $\eta := \delta$.

Recall the definition of $\Sigma_{\geq \eta}$ in Definition 7.6. By definition,

$$\Pr_{\mu|_j}[\mathrm{round}_j^{\rho,\eta} \neq \phi_j] = \sum_{\sigma \notin \Sigma_{\geq \eta}} \Pr_{\mu|_j}[\phi_j = \sigma].$$

Therefore

$$\mathop{\mathbb{E}}_{\rho \sim \nu^n} \Big[\Pr_{\mu|_j}[\mathrm{round}_j^{\rho,\eta} \neq \phi_j]\Big] = \sum_{\sigma \in \Sigma} \Pr[\sigma \notin \Sigma_{\geq \eta}] \Pr_{\mu|_j}[\phi_j = \sigma].$$

For each $\sigma \in \Sigma$, we consider two cases. If $\Pr_{\mu|_j}[\phi_j = \sigma] < \epsilon/|\Sigma|$ then the summand is clearly at most $\epsilon/|\Sigma|$. Otherwise, $\Pr[\sigma \notin \Sigma_{\geq \eta}] \leq \epsilon/|\Sigma|$ by Theorem 7.2. In both cases, the summand is at most $\epsilon/|\Sigma|$. Summing over all $\sigma \in \Sigma$ yields the result. ☐

7.2 Main result

We are now ready to prove Theorem 1.10.

Proof of Theorem 1.10. We prove the theorem with an error probability of $O(\epsilon)$ rather than ϵ.

Apply Lemma 7.7 with ϵ to obtain d, τ, η. Apply Theorem 7.3 with alphabet Σ, distributions $\mu|_1, \ldots, \mu|_m$, and parameters ϵ, d, τ to obtain M and a set J of size at most M such that for all j,

$$\Pr_{x \sim \mu_j^J}[f|_{J \leftarrow x} \text{ is } (d,\tau)\text{-regular with respect to } \mu_j] \geq 1 - \epsilon.$$

As in the proof of Theorem 1.5, we say that a subfunction $f_j|_{J \leftarrow x}$ is *good* if there exist $(x^{(1)}, \ldots, x^{(m)}) \in P^J$ such that $x^{(j)} = x$ and $f_k|_{J \leftarrow x^{(k)}}$ is (d,τ)-regular for all k. A simple argument (written down explicitly as Lemma 4.8) shows that for all j,

$$\Pr_{x \sim \mu|_j}[f_j|_{J \leftarrow x} \text{ is good}] \geq 1 - m\epsilon.$$

With hindsight, define

$$\delta = \min(\mu)^M \eta^m / 3.$$

We have

$$\mathop{\mathbb{E}}_{\rho \sim \nu^{J^c}} \mathop{\mathbb{E}}_{(x^{(1)}, \ldots, x^{(m)}) \sim \mu^J} \Pr_{\mu^{J^c}|\rho}[(f_1|_{J \leftarrow x^{(1)}}, \ldots, f_m|_{J \leftarrow x^{(m)}}) \notin P] = \Pr_{\mu^n}[(f_1, \ldots, f_m) \notin P] \leq \delta.$$

Therefore with probability at least $1/2$ over the choice of ρ,

$$\mathop{\mathbb{E}}_{(x^{(1)}, \ldots, x^{(m)}) \sim \mu^J} \Pr_{\mu^{J^c}|\rho}[(f_1|_{J \leftarrow x^{(1)}}, \ldots, f_m|_{J \leftarrow x^{(m)}}) \notin P] \leq 2\delta < \min(\mu)^M \eta^m.$$

In this case, we say that ρ is *good*.

If ρ is good then for any $(x^{(1)}, \ldots, x^{(m)}) \in P^J$,

$$\Pr_{\mu|\rho}[(f_1|_{J \leftarrow x^{(1)}}, \ldots, f_m|_{J \leftarrow x^{(m)}}) \notin P] < \eta^M.$$

Applying Lemma 7.5, this shows that $(\mathrm{round}_1^{\rho,\eta}(f_1|_{J\leftarrow x^{(1)}}),\dots,\mathrm{round}_m^{(\rho,\eta}(f_m|_{J\leftarrow x^{(m)}}))$ is a generalized polymorphism of P. Accordingly, we define $g_1^\rho,\dots,g_m^\rho$ by

$$g_j^\rho|_{J\leftarrow x} = \mathrm{round}_j^{(\rho,\eta)}(f_j|_{J\leftarrow x}).$$

For every good ρ, the functions $(g_1,\dots,g_m)$ are a generalized polymorphism of P.

If $f_j|_{J\leftarrow x}$ is good then Lemma 7.7 shows that

$$\mathop{\mathbb{E}}_{\substack{\rho\sim\nu^{J^c}\\ \rho \text{ good}}}\Big[\Pr_{\mu|_j^{J^c}}[g_j^\rho|_{J\leftarrow x}\neq f_j|_{J\leftarrow x}]\Big] \le \epsilon/\Pr[\rho \text{ is good}] \le 2\epsilon.$$

Since the probability that $f_j|_{J\leftarrow x}$ is not good is at most $m\epsilon$, it follows that

$$\sum_{j=1}^m \mathop{\mathbb{E}}_{\substack{\rho\sim\nu^{J^c}\\ \rho \text{ good}}}\Big[\Pr_{\mu_j^n}[g_j^\rho\neq f_j]\Big] \le m(m+2)\epsilon.$$

In particular, we can find ρ for which this sum is at most $m(m+2)\epsilon$. Taking $g_j = g_j^\rho$ for all j completes the proof. $\square$

8 Regularity lemmas

In this section we prove the various versions of Jones' regularity lemma: Theorems 3.3, 4.2, 6.3 and 7.3.

We derive all of them from the following two versions.

Theorem 8.1 (Jones' regularity lemma (decision tree version))**.** *For every alphabet Σ, every $m\in\mathbb{N}$, every full support distributions $\mu_1,\dots,\mu_m$, and every $\epsilon,\tau>0$, $d\in\mathbb{N}$ the following holds for $M=O(md/\epsilon\tau)$.*

For all functions $f_1,\dots,f_m\colon\Sigma^n\to[0,1]$ and every decision tree T_0 there exists a decision tree T extending T_0 with additional depth at most M such that for all j,

$$\Pr_{\rho\sim T}[f_j|_\rho \text{ is } (d,\tau)\text{-regular with respect to } \mu_j] \ge 1-\epsilon,$$

where ρ is sampled by following T, sampling each variable encountered according to μ_j.

Theorem 8.2 (Jones' regularity lemma (junta version))**.** *For every alphabet Σ, every $m\in\mathbb{N}$, every full support distributions $\mu_1,\dots,\mu_m$, and every $\epsilon,\tau>0$, $d\in\mathbb{N}$ the following holds for some function $\mathcal{M}\colon\mathbb{N}\to\mathbb{N}$.*

For all functions $f_1,\dots,f_m\colon\Sigma^n\to[0,1]$ and set $J_0\subseteq[n]$ there exists set $J\supseteq J_0$ of size at most $\mathcal{M}(|J|)$ such that for all j,

$$\Pr_{x\sim\mu_j}[f_j|_{J\leftarrow x} \text{ is } (d,\tau)\text{-regular with respect to } \mu_j] \ge 1-\epsilon.$$

Theorems 3.3 and 6.3 follow from Theorem 8.1, and Theorems 4.2 and 7.3 follow from Theorem 8.2.

The proofs proceed via a different notion of regularity, using noisy influences.

Definition 8.3 (Noisy influences)**.** For $\rho\in(0,1)$, let N_ρ be the distribution of pairs $x,y\in\Sigma^n$ sampled as follows. We sample $x\sim\mu^n$. We sample y by resampling each coordinate of x with probability $1-\rho$. The noise stability of a function $f\colon\Sigma^n\to\mathbb{R}$ is defined as

$$\mathrm{Stab}_\rho[f] = \mathop{\mathbb{E}}_{(x,y)\sim N_\rho}[f(x)f(y)] = \sum_S \rho^{|S|}\|f_S\|^2,$$

where f_S are the components of the Efron–Stein decomposition of f, and the norm is computed according to μ.

For a coordinate i, let $E_i f$ be obtained by averaging over the coordinate i:

$$E_i f(x) = \mathop{\mathbb{E}}_{a \sim \mu}[f(x|_{i \leftarrow a})] = \sum_{i \notin S} f_S.$$

The noisy influences of f are

$$\mathrm{Inf}_i^{\rho}[f] = \mathrm{Stab}_{\rho}[f - E_i f] = \sum_{i \in S} \rho^{|S|} \|f_S\|^2.$$

We say that f is (ρ, τ)-noisy-regular if $\mathrm{Inf}_i^{\rho}[f] \le \tau$ for all i.

We derive Theorem 8.1 from the following version for noisy influences.

Theorem 8.4 (Jones' regularity lemma for noisy influences (decision tree version))**.** *For every alphabet Σ, every $m \in \mathbb{N}$, every full support distributions $\mu_1, \ldots, \mu_m$, and every $\epsilon, \tau > 0$, $d \in \mathbb{N}$ the following holds for $M = m/\epsilon\tau \cdot \rho/(1-\rho)$.*

For all functions $f_1, \ldots, f_m \colon \Sigma^n \to [0,1]$ and every decision tree T_0 there exists a decision tree T extending T_0 with an additional depth at most M such that for all j,

$$\Pr_{\rho \sim T}[f_j|_\rho \text{ is } (\rho, \tau)\text{-noisy-regular with respect to } \mu_j] \ge 1 - \epsilon,$$

where ρ is sampled by following T, sampling each variable encountered according to μ_j.

The proof of Theorem 8.4 uses a potential argument. For a decision tree T, we define

$$\Phi^{\rho}_{f,\mu}(T) = \mathop{\mathbb{E}}_{\ell \in T}[\mathrm{Stab}_{\rho}[f|_\ell]],$$

where ℓ is a random leaf of T sampled using μ. The potential function we use is the sum of these potentials for all f_j:

$$\Phi(T) = \sum_j \Phi^{\rho}_{f_j, \mu_j}(T).$$

All properties of this potential function follow from the following simple calculation, which corresponds to splitting a node in the tree.

Lemma 8.5 (Splitting a node)**.** *For every i we have*

$$\mathop{\mathbb{E}}_{a \sim \mu}[\mathrm{Stab}_{\rho}[f|_{i \leftarrow a}]] = \mathrm{Stab}_{\rho}[f] + \frac{1-\rho}{\rho} \mathrm{Inf}_i^{\rho}[f].$$

Proof. The definition of noise stability shows that

$$\mathrm{Stab}_{\rho}[f] = \rho \mathop{\mathbb{E}}_{a \sim \mu}[\mathrm{Stab}_{\rho}[f|_{i \leftarrow a}]] + (1-\rho) \mathop{\mathbb{E}}_{\substack{a,b \sim \mu \\ (x',y') \sim N_\rho}}[f|_{i \leftarrow a}(x') f|_{i \leftarrow b}(y')].$$

Similarly, the noisy influence is

$$\mathrm{Inf}_i^{\rho}[f] = \mathop{\mathbb{E}}_{\substack{(x,y) \sim N_\rho \\ a,b \sim \mu}}[(f(x) - f(x|_{i \leftarrow a}))(f(y) - f(y|_{i \leftarrow b}))] = \mathrm{Stab}_{\rho}[f] - \mathop{\mathbb{E}}_{\substack{a,b \sim \mu \\ (x',y') \sim N_\rho}}[f|_{i \leftarrow a}(x') f|_{i \leftarrow b}(y')],$$

where the second term is the result of summing three identical terms with different signs. Substituting the earlier formula gives

$$\mathrm{Inf}_i^{\rho}[f] = \rho \mathop{\mathbb{E}}_{a \sim \mu}[\mathrm{Stab}_{\rho}[f|_{i \leftarrow a}]] - \rho \mathop{\mathbb{E}}_{\substack{a,b \sim \mu \\ (x',y') \sim N_\rho}}[f|_{i \leftarrow a}(x') f|_{i \leftarrow b}(y')].$$

The lemma immediately follows. $\square$

We can now prove Theorem 8.4.

Proof of Theorem 8.4. We construct a sequence of trees T_t as follows. The starting point is T_0. If T_t doesn't satisfy the conclusion of the lemma for f_j, then Lemma 8.5 implies that we can add one level to T_t, forming a tree T_{t+1} satisfying

$$\Phi^{\rho}_{f_j,\mu_j}(T_{t+1}) \geq \Phi^{\rho}_{f_j,\mu_j}(T_t) + \frac{1-\rho}{\rho}\epsilon\tau.$$

For $k \neq j$, the lemma gives the guarantee

$$\Phi^{\rho}_{f_k,\mu_k}(T_{t+1}) \geq \Phi^{\rho}_{f_k,\mu_k}(T_t),$$

and so

$$\Phi(T_{t+1}) \geq \Phi(T_t) + \frac{1-\rho}{\rho}\epsilon\tau.$$

Since $\mathrm{Stab}_\rho[f] \leq \|f\|^2$ for every function f, it is easy to check that $0 \leq \Phi(T) \leq m$ for all T. In particular, the process above must terminate after at most $m/\epsilon\tau \cdot \rho/(1-\rho)$ steps. Taking the last tree to be our T completes the proof. $\square$

We deduce Theorem 8.1 by relating low-degree influences and noisy influences.

Proof of Theorem 8.1. If a function $g\colon \Sigma^n \to [0,1]$ is (ρ,τ)-noisy-regular then for every $\rho \in (0,1)$ and every $i \in [n]$ we have

$$\mathrm{Inf}_i[g^{\leq d}] = \sum_{\substack{|S|\leq d\\ i\in S}} \|f_S\|^2 \leq \rho^{-d} \sum_{i\in S} \rho^{|S|}\|f_S\|^2 = \rho^{-d}\,\mathrm{Inf}^{\rho}_i[g].$$

We choose $\rho = 1 - 1/d$ and apply Theorem 8.4 with $\tau := \rho^d\tau$ to complete the proof, observing that $\rho^d = \Theta(1)$. $\square$

The proof of Theorem 8.2 similarly follows from a version for noisy influences.

Theorem 8.6 (Jones' regularity lemma for noisy influences (junta version))**.** *For every alphabet Σ, every $m \in \mathbb{N}$, every full support distributions $\mu_1,\ldots,\mu_m$, and every $\epsilon,\tau > 0$, $d \in \mathbb{N}$ the following holds for some function $\mathcal{M}\colon \mathbb{N} \to \mathbb{N}$.*

For all functions $f_1,\ldots,f_m\colon \Sigma^n \to [0,1]$ and every $J_0 \subseteq [n]$ there exists a set $J \supseteq J_0$ of size at most $\mathcal{M}(|J_0|)$ such that for all j,

$$\Pr_{x\sim\mu_j}[f_j|_{J\leftarrow x} \text{ is } (\rho,\tau)\text{-noisy-regular with respect to } \mu_j \text{ for all } j] \geq 1-\epsilon.$$

The proof is a modification of the proof of Theorem 8.4.

Proof. For a set J, let $C(J)$ be the decision tree querying all coordinates in J.

We construct a sequence of sets J_t as follows. The start point is J_0. If J_t doesn't satisfy the conclusion of the lemma for f_j, by Lemma 8.5 we can add one more level to $C(J_t)$ to obtain a tree T' such that

$$\Phi(T') \geq \Phi(C(J_t)) + \frac{1-\rho}{\rho}\epsilon\tau.$$

Let J_{t+1} consist of all variables appearing in T'. We can obtain a tree with the same set of leaves as $C(J_{t+1})$ by splitting nodes of T', and so Lemma 8.5 implies that

$$\Phi(C(J_{t+1})) \geq \Phi(T') \geq \Phi(C(J_t)) + \frac{1-\rho}{\rho}\epsilon\tau.$$

As in the proof of Theorem 8.4, the process terminates within $1/\epsilon\tau \cdot \rho/(1-\rho)$ steps. Since $|J_{t+1}| \leq |J_t| + |\Sigma|^{|J_t|}$, the final size of the resulting set J can be bounded independently of n. $\square$

The proof makes it clear that the dependence of $\mathcal{M}$ on the parameters is much worse in Theorem 4.2 than in Theorem 3.3. For this reason, it would be interesting to find a proof of Theorem 1.5 which uses the decision tree version rather than the junta version.

9 Open questions

We would like to highlight the following open questions:

1. Extend Theorem 1.10 to cover all predicates $P \subseteq \Sigma^m$.

   The simplest predicate that our techniques cannot handle is $\{(a, b, c) \in \mathbb{Z}_3 : a + b + c \neq 0\}$.

2. Prove Theorem 1.5 with the additional guarantee that if $f_i = f_j$ and $\mu|_i = \mu|_j$ then $g_i = g_j$.

   This guarantee follows from Theorem 1.5 when all generalized polymorphisms $g_1, \dots, g_m$ of P are such that if $\mu|_i = \mu_j$ then either $g_i = g_j$ or $\Pr_{\mu|_i}[g_i \neq g_j] = \Omega(1)$. This is the case for the predicates $P_{m,b}$ considered in Theorem 1.6. We are also able to provide this guarantee in the monotone setting (Theorem 1.7).

   There are other situations in which we can prove this guarantee as a corollary of Theorem 1.5. One example is the predicate $P_\wedge = \{(a, b, a \wedge b) : a, b \in \{0, 1\}\}$. In this case the generalized polymorphisms g_1, g_2, g_3 come in two types: (i) $g_1 = g_2 = g_3 = \bigwedge_{i \in S} x_i$ for some S; (ii) $g_1 = g_3 = 0$ or $g_2 = g_3 = 0$. In the first case, we automatically get the guarantee. In the second case, say $g_1 = g_3 = 0$, the only claim which doesn't automatically hold is that if $f_1 = f_2$ and $\mu|_1 = \mu|_2$ then $g_1 = g_2$. In this case g_2 is close to g_1 (since $f_1 = f_2$ and g_1, g_2 are close to f_1, f_2), and so we can set $g_2 \equiv 0$ to satisfy the guarantee while maintaining the other properties in the theorem.

   Many predicates satisfy a similar but more restricted type of guarantee: every generalized polymorphism $(g_1, \dots, g_m)$ is such that either (i) all functions are of the form x_i or $1 - x_i$ (for the same i), or (ii) some of the functions are constant, and these constants constitute a "certificate" for the predicate. This case can be handled just as the case of $P_\wedge$. A characterizations of this type of behavior can be found in [Fil25].

3. Determine the optimal relation between ϵ and δ in Theorem 1.5. We conjecture that the optimal dependence is polynomial or even linear. This holds in the case of the predicates $P_{m,b}$ of Theorem 1.6, and also for many functional predicates (see [CFM+22, Appendix D] for an illustrative example).

References

[AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2 + \varepsilon)$-Sat is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017.

[Arr50] Kenneth J. Arrow. A difficulty in the concept of social welfare. *Journal of Political Economy*, 58(4):328–346, 1950.

[BBKO21] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):Art. 28, 66, 2021.

[BCH+96] Mihir Bellare, Don Coppersmith, Johan Håstad, Marcos Kiwi, and Madhu Sudan. Linearity testing in characteristic two. *IEEE Trans. Inform. Theory*, 42(6, part 1):1781–1795, 1996. Codes and complexity.

[BG21] Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: algebraic structure and a symmetric Boolean dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021.

[BLR90] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. In *STOC*, pages 73–83, 1990.

[Bou02] J. Bourgain. On the distributions of the Fourier spectrum of Boolean functions. *Israel J. Math.*, 131:269–276, 2002.

[CFM+22] Gilad Chase, Yuval Filmus, Dor Minzer, Elchanan Mossel, and Nitin Saurabh. Approximate polymorphisms. In *STOC'22*, 2022.

[DDG+17] Roee David, Irit Dinur, Elazar Goldenberg, Guy Kindler, and Igor Shinkar. Direct sum testing. *SIAM J. Comput.*, 46(4):1336–1369, 2017.

[DF09] Irit Dinur and Ehud Friedgut. Intersecting families are essentially contained in juntas. *Combin. Probab. Comput.*, 18(1-2):107–122, 2009.

[DFH25] Irit Dinur, Yuval Filmus, and Prahladh Harsha. Agreement tests on graphs and hypergraphs. *SIAM J. Comput.*, 54(2):279–320, 2025.

[DFR08] Irit Dinur, Ehud Friedgut, and Oded Regev. Independent sets in graph powers are almost contained in juntas. *Geom. Funct. Anal.*, 18(1):77–97, 2008.

[DS14] Irit Dinur and David Steurer. Direct product testing. In *IEEE 29th Conference on Computational Complexity—CCC 2014*, pages 188–196. IEEE Computer Soc., Los Alamitos, CA, 2014.

[DT16] Shagnik Das and Tuan Tran. Removal and stability for Erdős–Ko–Rado. *SIAM J. Discrete Math.*, 30(2):1102–1114, 2016.

[Fil25] Yuval Filmus. Aggregation of evaluations without unanimity. Online manuscript, 2025.

[FLMM20] Yuval Filmus, Noam Lifshitz, Dor Minzer, and Elchanan Mossel. AND testing and robust judgement aggregation. In *52nd ACM Symposium on Theory of Computing (STOC'20)*, 2020.

[FR18] Ehud Friedgut and Oded Regev. Kneser graphs are like Swiss cheese. *Discrete Analysis*, jan 23 2018.

[Jon16] Chris Jones. A noisy-influence regularity lemma for boolean functions, 2016.

[Kal02] Gil Kalai. A Fourier-theoretic perspective on the Condorcet paradox and Arrow's theorem. *Adv. in Appl. Math.*, 29(3):412–426, 2002.

[Kel10] Nathan Keller. On the probability of a rational outcome for generalized social welfare functions on three alternatives. *J. Combin. Theory Ser. A*, 117(4):389–410, 2010.

[Kiw03] M. Kiwi. Algebraic testing and weight distributions of codes. *Theoret. Comput. Sci.*, 299(1-3):81–106, 2003.

[KKO18] Guy Kindler, Naomi Kirshner, and Ryan O'Donnell. Gaussian noise sensitivity and Fourier tails. *Israel J. Math.*, 225(1):71–109, 2018.

[MOO10] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, 171(1):295–341, 2010.

[Mos12a] Elchanan Mossel. A quantitative Arrow theorem. *Probab. Theory Related Fields*, 154(1-2):49–88, 2012.

[Mos12b] Elchanan Mossel. A quantitative Arrow theorem. *Probab. Theory Related Fields*, 154(1-2):49–88, 2012.

[MOS13] Elchanan Mossel, Krzysztof Oleszkiewicz, and Arnab Sen. On reverse hypercontractivity. *Geometric and Functional Analysis*, 23(3):1062–1097, 2013.

[Neh13] Ilan Nehama. Approximately classic judgement aggregation. *Ann. Math. Artif. Intell.*, 68(1-3):91–134, 2013.

[RS78] I. Z. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. In *Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II*, volume 18 of *Colloq. Math. Soc. János Bolyai*, pages 939–945. North-Holland, Amsterdam-New York, 1978.

[Sze78] Endre Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il