

INDUCED ARITHMETIC REMOVAL: COMPLEXITY 1 PATTERNS OVER FINITE FIELDS

JACOB FOX, JONATHAN TIDOR, AND YUFEI ZHAO

ABSTRACT. We prove an arithmetic analog of the induced graph removal lemma for complexity 1 patterns over finite fields. Informally speaking, we show that given a fixed collection of r -colored complexity 1 arithmetic patterns over $\mathbb{F}_q$, every coloring $\phi: \mathbb{F}_q^n \setminus \{0\} \rightarrow [r]$ with $o(1)$ density of every such pattern can be recolored on an $o(1)$ -fraction of the space so that no such pattern remains.

1. INTRODUCTION

Removal lemmas play a central role in extremal and additive combinatorics (see the survey [7]). The *triangle removal lemma*, first proved by Ruzsa and Szemerédi [23], states that for every $\epsilon > 0$ there exists $\delta > 0$ such that every n -vertex graph with fewer than δn^3 triangles can be made triangle-free by deleting at most ϵn^2 edges.

In computer science, there is much interest in removal lemmas due to the connection between removal lemmas and constant query complexity algorithms for property testing [22, 10]. For example, the triangle removal lemma implies that to distinguish a triangle-free graph from one that is ϵ -far from triangle free, one simply needs to sample a constant number of triples of vertices (this constant depends on ϵ , but not the size of the graph) and check them for triangles. All our discussions and results below can be interpreted using the language of property testing.

An arithmetic analog of the triangle removal lemma was first established by Green [14]. It states that for every $\epsilon > 0$ there exists $\delta > 0$ such that for every finite abelian group $(G, +)$ and every $A \subset G$ with at most $\delta |G|^2$ triples $(x, y, z) \in A^3$ satisfying $x + y = z$, there exists $A' \subseteq A$ with $|A \setminus A'| \leq \epsilon |G|$ such that A' contains no solution to $x + y = z$.

Green's arithmetic triangle removal lemma also has a corresponding property testing algorithm. This algorithm distinguishes between *triangle-free* maps $f: G \rightarrow \{0, 1\}$ (i.e., those with no $x, y \in G$ satisfying $f(x) = f(y) = f(x + y) = 1$) and maps that are ϵ -far from triangle free by sampling a constant number of pairs $(x, y) \in G^2$ and checking if any satisfy $f(x) = f(y) = f(x + y) = 1$.

Generalizing the triangle removal lemma, there is a *graph removal lemma* which states that for every fixed graph H (or a finite collection of graphs), every n -vertex graph with H -density $o(1)$ can be made H -free by removing $o(n^2)$ edges [1, 9]. Furthermore, there is an *induced graph removal lemma* [2] which states that every n -vertex graph with induced H -density $o(1)$ can be made induced H -free by adding or removing $o(n^2)$ edges. Both of these removal lemmas give constant query complexity algorithms for property testing of the corresponding properties.

The graph removal lemma was proved using Szemerédi's graph regularity lemma, an important tool in extremal combinatorics. The arithmetic removal lemma was initially proved by Green via an arithmetic variant of the graph regularity lemma; an alternate proof was later given by Král'–Serra–Vena [18], deducing the arithmetic removal lemma from the graph removal lemma. This technique was extended to deduce a removal lemma for general linear systems from the hypergraph removal lemma [19, 24].

Fox was supported by a Packard Fellowship and by NSF grant DMS-1855635. Tidor was supported by NSF Graduate Research Fellowship Program DGE-1122374. Zhao was supported by NSF Award DMS-1764176, the MIT Solomon Buchsbaum Fund, and a Sloan Research Fellowship.

A natural common extension of the induced graph (or hypergraph) removal lemma and the arithmetic removal lemma would be an induced arithmetic removal lemma, which is not currently known (e.g., [24, Conjecture 5.3]). In particular, the technique for deducing arithmetic removal from graph removal (from [18]) does not appear to extend to the induced setting.

The goal of this paper is to establish an induced arithmetic removal lemma for arithmetic patterns of complexity 1 over finite fields. (See Section 1.2 for the definition of complexity.) Previous results [5] established this claim over $\mathbb{F}_2$, though the proof requires specific Ramsey-theoretic results that fail for all other finite fields. We overcome this difficulty by introducing a novel recoloring technique that works over all base fields.

It remains an open problem to extend the induced arithmetic removal lemma to more general settings, both to higher complexity patterns and to general groups (even $\mathbb{Z}$ is open). For affine-invariant patterns (which are always translation-invariant), a special case of the general problem that avoids an important difficulty (no regularization near the origin; see discussion in Section 2), the problem for bounded complexity patterns over finite fields has been solved [4].

We begin with a discussion motivating the precise setup of our main theorem. In Section 2 we review the background on the regularity proof of the graph removal lemma and its relationship to our proof, and we describe our proof strategy. In Section 3 we prove our main theorem, modulo two key tools, proved in Sections 4 and 5. In Section 6 we give two extensions of our main theorem, to an infinite set of patterns and to inhomogeneous patterns.

Acknowledgements. The authors are grateful to Noga Alon, Freddie Manners, and Tom Sanders for helpful discussions.

1.1. Colored patterns. Induced subgraphs can be equivalently rephrased as colored patterns in a 2-edge-colored clique, where the two colors correspond to edges and non-edges. We consider the natural generalization allowing an arbitrary fixed number of edge colors.

The proof of the induced graph removal lemma gives the following result about colored patterns. Given r -edge-colored cliques G and H , the H -density in G is defined to be the fraction of $|H|$ -vertex subgraphs of G that have edge-coloring isomorphic to H . We say that G is H -free if no such subgraph exists.

Theorem 1.1 (Induced graph removal). *Fix a finite collection $\mathcal{H}$ of r -edge-colored cliques. For every $\epsilon > 0$ there exists $\delta > 0$ so that if an r -edge-coloring of K_n has H -density at most δ for every $H \in \mathcal{H}$, then one can recolor at most ϵn^2 edges of K_n (using the original r colors) so that the resulting coloring is H -free for every $H \in \mathcal{H}$.*

An analogous statement for hypergraphs is also known to hold [21]. In addition, there is even a result allowing an infinite collection of subgraphs $\mathcal{H}$ (with a slightly modified statement); see Section 6.1.

We would like to obtain a similar statement for colored arithmetic patterns. However, there must be some caveats, as illustrated by the following two non-examples.

Non-example 1.2. Consider the 2-coloring of $\mathbb{Z}/N\mathbb{Z}$ with $1, \dots, \lfloor N/2 \rfloor$ colored red and the remaining elements colored blue, and the pattern $x = y + 1$ where x is blue and y is red. There is only one instance of this pattern in the given coloring, but it is impossible to recolor fewer than $\lfloor N/2 \rfloor$ elements to remove this pattern. For an example with a homogeneous equation, a similar situation occurs for $x = 2y$ as long as 2 has large multiplicative order in $(\mathbb{Z}/N\mathbb{Z})^\times$.

The above non-example does not worry us too much since it has infinite complexity and is not over a fixed finite field. Nevertheless, it reminds us that we need to be cautious when formulating a more general conjecture about induced arithmetic removal.

The second non-example is more relevant to our discussion as it is a complexity 1 pattern over a fixed finite field.

Non-example 1.3. Consider the 4-coloring $\phi: \mathbb{F}_5^n \rightarrow [4]$ defined as follows: $\phi((0, \dots, 0)) = 1$ and $\phi((0, \dots, 0, x_k, \dots, x_n)) = x_k$ where $x_k \neq 0$ is the first nonzero coordinate. We consider monochromatic solutions to $x + y + z = 0$. The only such solution is $x = y = z = 0$, but it always remains a monochromatic solution no matter how ϕ is modified. Thus there is no removal lemma here either.

This non-example shows that we need to treat zero with care. While not all patterns exhibit the above behavior, a simple way to remove this issue in general seems to be to ignore 0 entirely. One might attempt to consider a less drastic alternative by only ignoring the all-zero solution, but this does not always work, e.g., we can add some additional unconstrained dummy variables to the system. It seems plausible that for every fixed system of patterns one can identify a “minimal” notion of a trivial solution that should be ignored, but we take the simpler route here by uniformly ignoring zero. An alternative and weaker possible formulation, which is implied by our main theorem, is to only remove *generic* solutions, i.e., solutions whose span has maximum possible dimension.

As an example of our main result, Theorem 1.14, applied to the above system, we prove:

If $\phi: \mathbb{F}_5^n \setminus \{0\} \rightarrow [4]$ has $o(5^{2n})$ monochromatic solutions to $x + y + z = 0$, then one can remove all such solutions by recoloring an $o(1)$ -fraction of the space.

Definition 1.4. An r -colored pattern $H = (A, \psi)$ is represented by an $\ell \times k$ matrix A with entries in $\mathbb{F}_q$ along with a coloring $\psi: [k] \rightarrow [r]$ of the columns of A . Let V be a finite dimensional $\mathbb{F}_q$ -vector space. An H -instance in a coloring $\phi: V \rightarrow [r]$ is some $\mathbf{x} = (x_1, \dots, x_k) \in V^k$ satisfying $A\mathbf{x} = \mathbf{0}$ and $\phi(x_i) = \psi(i)$ for all $i \in [k]$. The H -density is defined to be the number of H -instances divided by $|V|^{k - \text{rank } A}$, i.e., the probability that a uniform random $\mathbf{x} \in V^k$ satisfying $A\mathbf{x} = \mathbf{0}$ also satisfies $\phi(x_i) = \psi(i)$ for all $i \in [k]$. We say that ϕ is H -free if there are no H -instances $\mathbf{x} \in (V \setminus \{0\})^k$.

We generally need to consider a collection of colored patterns.

Example 1.5. Monochromatic solutions to $x + y + z = 0$ can be encoded by a collection of r separate r -colored patterns, $\{(A, \psi_i)\}_{i \in [r]}$, all having the same $A = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}$ where $\psi_i \equiv i$ is the constant function.

Example 1.6. A rainbow solution to $x + 2y + 3z = 0$ is one where x , y , and z have distinct colors, and it can be encoded by a collection of $r(r-1)(r-2)$ separate H 's, all having the same $A = \begin{pmatrix} 1 & 2 & 3 \end{pmatrix}$, and ψ ranging over all injective maps $[3] \rightarrow [r]$.

Example 1.7. The property of having either a blue 3-term arithmetic progression or a red 4-term arithmetic progression can be encoded by two patterns:

$$(A = \begin{pmatrix} 1 & -2 & 1 \end{pmatrix}, \psi \equiv \text{blue}) \quad \text{and} \quad \left(A = \begin{pmatrix} 1 & -2 & 1 & 0 \\ 0 & 1 & -2 & 1 \end{pmatrix}, \psi \equiv \text{red} \right).$$

Note that the second pattern has complexity 2.

Example 1.8. Suppose that p is prime. We consider maps $\phi: \mathbb{F}_p^n \setminus \{0\} \rightarrow \mathbb{F}_p$ as p -colorings of $\mathbb{F}_p^n$. Consider the set of patterns

$$\mathcal{H} = \{(A, \psi_{a,b,c})\}_{\substack{a,b,c \in \mathbb{F}_p \\ a+b \neq c}}$$

where $A = \begin{pmatrix} 1 & 1 & -1 \end{pmatrix}$ and $\psi_{a,b,c}$ is defined by $\psi_{a,b,c}(1) = a, \psi_{a,b,c}(2) = b, \psi_{a,b,c}(3) = c$. Then $\phi: \mathbb{F}_p^n \setminus \{0\} \rightarrow \mathbb{F}_p$ is linear if and only if it is $\mathcal{H}$ -free.

1.2. Complexity 1. Fourier analysis, also known in this context as the Hardy–Littlewood circle method, allows us to analyze certain patterns such as $x + y = z$ or 3-term arithmetic progressions, but not others such as 4-term arithmetic progressions. This distinction is captured by a notion known as the *complexity* of a linear system, introduced by Green and Tao in their work on linear patterns in the primes [16]. Informally, a system has *complexity* s if it is “controlled” by the U^{s+1}

Gowers uniformity norm (the *true complexity* [11] of the system is the smallest such s), which we will not define here in general as in this paper we only address complexity 1 systems.

We consider weighted, or functional, versions of densities of linear patterns. For a finite dimensional $\mathbb{F}_q$ -vector space V , an $\ell \times k$ matrix A over $\mathbb{F}_q$, and functions $f_1, \dots, f_k: V \rightarrow \mathbb{C}$, we write

$$\Lambda_A(f_1, \dots, f_k) := \mathbb{E}_{\mathbf{x}=(x_1, \dots, x_k) \in V^k: A\mathbf{x}=\mathbf{0}} [f_1(x_1) \cdots f_k(x_k)].$$

Then the (A, ψ) -density in $\phi: V \rightarrow [r]$ is given by $\Lambda_A(1_{\phi^{-1}(\psi(1))}, \dots, 1_{\phi^{-1}(\psi(k))})$.

Example 1.9. For the matrix $A = \begin{pmatrix} 1 & -2 & 1 \end{pmatrix}$ corresponding to the pattern $x - 2y + z = 0$, one has

$$\Lambda_A(f_1, f_2, f_3) = \mathbb{E}_{x, y \in V} [f_1(x) f_2(y) f_3(2y - x)].$$

The above expression illustrates the equivalence between expressing a linear pattern as a linear system of equations (e.g., $x - 2y + z = 0$) and as a collection of linear forms of unconstrained variables (e.g., $(x, y, 2y - x)$). The complexity of a system is often defined in the literature in the latter formulation. It is sometimes useful to keep both versions in mind.

The Fourier transform of a function $f: V \rightarrow \mathbb{C}$ is $\hat{f}: V^* \rightarrow \mathbb{C}$ given by

$$\hat{f}(\gamma) := \langle f, \gamma \rangle = \mathbb{E}_{x \in V} f(x) \overline{\gamma(x)}$$

where $\gamma \in V^*$ is a character, i.e., $\gamma: V \rightarrow \mathbb{C}^\times$ satisfying $\gamma(x + y) = \gamma(x)\gamma(y)$ for all $x, y \in V$. For a finite dimensional $\mathbb{F}_q$ -vector space V of characteristic p , one has $V^* \cong V$ by associating $z \in V$ with the character γ_z given by $\gamma_z(x) := \exp(2\pi i \operatorname{tr}(x \cdot z)/p)$ where $x \cdot z := x_1 z_1 + \cdots + x_n z_n$ is the dot product and $\operatorname{tr}: \mathbb{F}_q \rightarrow \mathbb{F}_p$ is the trace map. It is fine to think only about the slightly simpler case of prime $q = p$ so that we can write $\hat{f}(\gamma_z) = \mathbb{E}_{x \in V} f(x) e^{-2\pi i(x \cdot z)/p}$ for $z \in V$.

Definition 1.10. We say that an $\ell \times k$ matrix A over $\mathbb{F}_q$ has *complexity 1* if for every $\epsilon > 0$ there exists $\delta > 0$ so that for every finite dimensional $\mathbb{F}_q$ -vector space V and $f_1, \dots, f_k: V \rightarrow [-1, 1]$, one has

$$|\Lambda_A(f_1, \dots, f_k)| \leq \epsilon \quad \text{whenever} \quad \min_{1 \leq i \leq k} \|\hat{f}_i\|_\infty \leq \delta.$$

It is known that a system given by linear forms $L_1, \dots, L_k$ (i.e., generating the null space of A) has complexity 1 if and only if the quadratic forms $L_1^2, \dots, L_k^2$ are linearly independent (a technical modification is required if q is even) [12, 11, 17]. The general dependence of δ on ϵ is a wide open problem, though conjecturally $\delta = \epsilon^{O_A(1)}$ works. See [20] for some discussion.

Example 1.11. Any single equation $\sum a_i x_i = 0$ with at least three non-zero coefficients is a complexity 1 system. In matrix form, this says that any matrix A of dimensions $1 \times k$ has complexity 1 as long as it has at least three non-zero entries.

Example 1.12. Let $G = (V, E)$ be a directed graph. Suppose A is a matrix with $|E|$ columns satisfying the property that $\mathbf{x} = (x_e)_{e \in E}$ satisfies $A\mathbf{x} = \mathbf{0}$ if and only if the signed sum of x_e around each cycle of G is 0. Such an A is complexity 1. This example can be generalized by scaling each column of A by a non-zero constant. In this form Example 1.11 is a special case corresponding to the cycle.

Example 1.13. There are complexity 1 systems that do not arise from the above constructions. One example is given by

$$A = \begin{pmatrix} 2 & 1 & 1 & -1 & 0 & 0 \\ 1 & 2 & 1 & 0 & -1 & 0 \\ 1 & 1 & 2 & 0 & 0 & -1 \end{pmatrix} \quad \text{over } \mathbb{F}_7.$$

For those familiar with the terminology, this system has Cauchy-Schwarz complexity 2, but true complexity 1. All systems arising from the two previous examples have true complexity and Cauchy-Schwarz complexity 1.

Now we are ready to state our main result.

Theorem 1.14 (Induced arithmetic removal). *Fix a finite set $\mathcal{H}$ of r -colored complexity 1 patterns over $\mathbb{F}_q$. For every $\epsilon > 0$ there exists $\delta = \delta(\epsilon, \mathcal{H}) > 0$ such that the following holds. Given a finite dimensional $\mathbb{F}_q$ -vector space V , if a coloring $\phi: V \rightarrow [r]$ has H -density at most δ for every $H \in \mathcal{H}$, then there exists a recoloring $\phi': V \setminus \{0\} \rightarrow [r]$ that is $\mathcal{H}$ -free and differs from ϕ on at most $\epsilon|V|$ elements of V .*

Remark 1.15. The dependence of $1/\delta$ on $1/\epsilon$ is tower-type. Though we write $\delta = \delta(\epsilon, \mathcal{H})$ for simplicity, the dependence on $\mathcal{H}$ is not strong. We could instead write $\delta = \delta(\epsilon, q, r, k)$ where q is the size of the finite field, r is the number of colors used, and k is the maximum number of columns in any of the patterns in $\mathcal{H}$.

For those familiar with the terminology, our main theorem states the following in the language of property testing: every linear-invariant, linear-subspace hereditary property of complexity 1 is testable with constant query complexity and one-sided error. (To be precise one needs the generalization to an infinite number of patterns; see Theorem 6.2.)

We conjecture that a similar statement holds for arbitrary patterns (with coefficients in $\mathbb{Z}$) in arbitrary abelian groups. However there must be some caveats related to infinite complexity patterns and trivial solutions. As such, we do not state a precise conjecture here.

2. PROOF STRATEGY

We begin by recalling the basic strategy for proving the graph removal lemma. We are given an n -vertex graph G with H -density $o(1)$ and we wish to make the graph H -free by removing $o(n^2)$ edges.

Strategy for graph removal:

- (1) **Regularize.** Apply Szemerédi's regularity lemma to G , partitioning the vertex set into a small number of roughly equal-sized parts, $V(G) = V_1 \cup \dots \cup V_k$, in such a way that G is regular (i.e., quasirandom) between almost all pairs of parts.
- (2) **Clean up.** Modify G by deleting all edges that lie between a pair of parts (V_i, V_j) that either is irregular or has low edge-density.

The resulting graph, after the edge deletions, should be H -free since if there were an H -subgraph it must sit among high-density regular pairs. By a counting lemma this would imply the existence of many copies of H in the original graph G . This would contradict the assumption that G has H -density $o(1)$.

For the induced graph removal lemma, the above strategy needs to be further extended, since removing all the edges between an irregular pair (V_i, V_j) may actually create induced H -subgraphs where there were none. The new tool here is to use a strong regularity lemma to find a “regular model” for G in which every pair of parts is regular.

Strategy for induced graph removal:

- (1) **Find a regular model.** Apply a strong regularity lemma to G to produce a partition of the vertex set into a small number of roughly equal-sized parts, $V(G) = V_1 \cup \dots \cup V_k$, and subsets $U_i \subseteq V_i$ each a positive fraction. This is done in such a way that $U := U_1 \cup \dots \cup U_k$ forms a “regular model” for G in the sense that (U_i, U_j) is a regular pair for *every* i, j and that the edge densities between pairs of V_i 's and the corresponding pairs of U_i 's are close on average.
- (2) **Clean up.** Modify G by deleting all edges that lie between a pair (V_i, V_j) if the edge-density between (U_i, U_j) is close to 0 and by adding all edges that lie between a pair (V_i, V_j) if the edge-density between (U_i, U_j) is close to 1.

The resulting graph, after the edge modifications, should be induced H -free for the following reason. Suppose that there were an induced H -subgraph sitting among some of the V_i 's. Then between the corresponding U_i 's, all the pairs are regular and all must have either high edge-density or high non-edge-density as required. By an appropriate counting lemma, this would produce many more induced copies of H in G , contradicting the assumption that G has induced H -density $o(1)$.

The arithmetic removal lemma is proved analogously to the graph removal lemma. Here a subset $S \subseteq V$ of a finite field vector space V plays the role of a graph and a regular partition of $S \subseteq V$ is given by a subspace $V_1 \leq V$ which has bounded codimension. This induces the partition of V into cosets of V_1 . For this partition to be regular we require $x + V_1$ to be regular (i.e., $S \cap (x + V_1)$ sits inside $x + V_1$ quasirandomly) for almost all $x \in V$.

Suppose H is a complexity 1 system of linear forms over $\mathbb{F}_q$ that has $o(1)$ density in $S \subseteq V$.

Strategy for arithmetic removal:

- (1) **Regularize.** Apply Green's arithmetic regularity lemma to $S \subseteq V$, producing a regular partition determined by a subspace $V_1 \leq V$ of bounded codimension.
- (2) **Clean up.** Modify S by deleting all elements that lie in a coset $x + V_1$ that either is irregular or has low density of S .

The resulting set, after the deletions, should be H -free since if there were an H -instance it must sit in high-density regular cosets. By a counting lemma, this would imply the existence of many copies of H in the original set S . This would contradict the assumption that S has H -density $o(1)$.

A natural attempt to prove the induced arithmetic removal lemma would be to invoke a strong arithmetic regularity lemma to find a regular model. In the arithmetic setting a regular model is a subspace W such that $W \cap (x + V_1)$ is regular for each x .

Suppose H is a 2-colored pattern of complexity 1 that has $o(1)$ density in $S \subseteq V$. (Here we view S and $V \setminus S$ as the two color classes.)

(Incorrect) strategy for induced arithmetic removal:

- (1) **Find a regular model.** Apply a strong arithmetic regularity lemma to $S \subseteq V$ to produce a partition of V into cosets of $V_1 \leq V$ as well as a bounded codimension subspace $W \leq V$ that forms a "regular model" for V in the following sense. We require that $W \cap (x + V_1)$ is a regular coset for *every* x and the densities of S in the $x + V_1$'s and the corresponding $W \cap (x + V_1)$'s are close on average.
- (2) **Clean up.** Modify S by deleting all elements that lie in a coset $x + V_1$ if the density of S in $W \cap (x + V_1)$ is close to 0 and by adding all elements that lie in $x + V_1$ if the density of S in $W \cap (x + V_1)$ is close 1.

If it were possible to carry out this strategy, the resulting set would be induced H -free. However, an important difficulty arises that was not present in the proof of the induced graph removal lemma. It turns out that it is impossible to perform the "find a regular model" step as the subspace $W \cap V_1$ cannot always be made regular. Green and Sanders [15] give an example with the following properties.

For all n , there exists a subset $S \subset \mathbb{F}_3^n$ such that for every positive-dimension subspace $V \leq \mathbb{F}_3^n$ the set $S \cap V$ is *not* an ϵ -regular subset of V for any $\epsilon < \sqrt{3}/6$.

Our strategy for circumventing this difficulty is two-fold. First we get as much out of arithmetic regularity as possible: we say that W is a regular model if $W \cap (x + V_1)$ is regular for every $x \in V \setminus V_1$. Next we develop a new tool to deal with V_1 . Our new tool is a Ramsey-type dichotomy which roughly says the following:

For every finite set $\mathcal{H}$ of 2-colored patterns over $\mathbb{F}_q$, there exist constants n_0 and $\epsilon_0 > 0$ such that either:

- (a) for every finite dimensional $\mathbb{F}_q$ -vector space V satisfying $\dim V \geq n_0$, every $S \subseteq V$ has H -density at least ϵ_0 for some $H \in \mathcal{H}$; or

(b) for every finite dimensional $\mathbb{F}_q$ -vector space V there exists $S \subseteq V$ that is $\mathcal{H}$ -free.

In the above dichotomy we view S and $V \setminus S$ as the two color classes.

Suppose H is a 2-colored pattern of complexity 1 that has $o(1)$ density in $S \subseteq V$.

(Revised) strategy for induced arithmetic removal:

- (1) **Find a regular model.** Apply a strong arithmetic regularity lemma to $S \subseteq V$ to produce a partition of V into cosets of $V_1 \leq V$ and a bounded codimension subspace $W \leq V$ that forms a “regular model” for V in the following sense. We require that $W \cap (x + V_1)$ is a regular coset for *every* $x \notin V_1$ and the densities of S in the $x + V_1$ ’s and the corresponding $W \cap (x + V_1)$ ’s are close on average.
- (2) **Clean up.** Modify S by deleting all elements that lie in a coset $x + V_1$ if the density of S in $W \cap (x + V_1)$ is close to 0 and by adding all elements that lie in $x + V_1$ if the density of S in $W \cap (x + V_1)$ is close to 1.
- (3) **Patch around the origin.** Since $S \subseteq V$ is a subset with H -density $o(1)$, we cannot be in case (a) of the dichotomy above. Therefore there exists a subset $T \subseteq V_1$ that is H -free. Replace $S \cap V_1$ with T . (One actually needs a slightly more involved application of our Ramsey-type dichotomy.)

Suppose that there is an H -instance sitting among some of the $x + V_1$ ’s. If none of these cosets contains the origin, then all the $W \cap (x + V_1)$ are regular and all either have high S -density or high $(V \setminus S)$ -density as required. This is impossible since it would imply the existence of many H -instances by an appropriate counting lemma. It is also impossible for an H -instance to sit entirely in V_1 since V_1 was patched to be H -free. With a little more work we can also rule out the existence of H -instances which are contained partially in V_1 and partially in $V \setminus V_1$.

3. PROOF OF MAIN THEOREM

In this section we prove the main theorem, modulo two key tools which are proved in Sections 4 and 5. As described in Section 2 the proof has three steps: find a regular model; clean up; patch around the origin. We first give the precise definition of regularity. As is typical in this field we work in the weighted, or functional, setting where functions represent weighted colorings.

Definition 3.1. A function $f: V \rightarrow [-1, 1]$ is ϵ -regular if $\|f - \mathbb{E}f\|_\infty \leq \epsilon$. A coloring $\phi: V \rightarrow [r]$ is ϵ -regular if $1_{\phi^{-1}(i)}$ is ϵ -regular for all $1 \leq i \leq r$.

The next proposition accomplishes the “find a regular model” step. As before, V is partitioned into cosets as $V = \bigcup_{x \in U} (x + V_1)$. The regular model will consist of $U \oplus V_2$ where U is a complement for V_1 in the sense that $U \oplus V_1 = V$. Property (1) states that the model consists of a positive fraction of V , property (2) that the colorings of $x + V_1$ and $x + V_2$ have similar densities, and property (3) that $U \oplus V_2$ is regular. This proposition is proved in Section 5 via an application of a strong arithmetic regularity lemma.

Proposition 3.2 (Regular model). *For a finite dimensional $\mathbb{F}_q$ -vector space V , functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, a subspace $V_0 \leq V$, and a parameter $\epsilon > 0$, there is some $n_{\text{reg}} = n_{\text{reg}}(\epsilon, q, k, \text{codim } V_0)$ such that there exist subspaces $V_2 \leq V_1 \leq V_0$ and a choice of complement U satisfying $U \oplus V_1 = V$ such that*

- (1) $\text{codim } V_2 \leq n_{\text{reg}}$;
- (2) for all but an ϵ -fraction of $x \in U$ we have

$$|\mathbb{E}_{y \in V_1} [f_i(x + y)] - \mathbb{E}_{y \in V_2} [f_i(x + y)]| \leq \epsilon$$

for all $1 \leq i \leq k$;

- (3) for each $x \in U \setminus \{0\}$ and $1 \leq i \leq k$, the function $f_i|_{x+V_2}$ is ϵ -regular (meaning that the function $g_i: V_2 \rightarrow [-1, 1]$ defined by $g_i(y) := f_i(x + y)$ is ϵ -regular).

With Proposition 3.2 in hand it is easy to accomplish the “clean up” step.

Proposition 3.3 (Regularity recoloring). *For a finite dimensional $\mathbb{F}_q$ -vector space V , a coloring $\phi: V \rightarrow [r]$, and parameters $0 < \epsilon, \epsilon' \leq 1$, there is some $n_{\text{reg}} = n_{\text{reg}}(\epsilon, \epsilon', q, r)$ such that there exist subspaces $V_2 \leq V_1 \leq V$, a complement U satisfying $U \oplus V_1 = V$, and a recoloring $\phi': V \rightarrow [r]$ that agrees with ϕ on all but an at most $\epsilon|V|$ values satisfying:*

- (1) $1/\epsilon \leq \text{codim } V_1 \leq \text{codim } V_2 \leq n_{\text{reg}}$;
- (2) *if a color appears in some coset $x + V_1$ under ϕ' , then at least an $\epsilon/(2r)$ -fraction of $x + V_2$ is that color under ϕ ;*
- (3) *for each $x \in U \setminus \{0\}$, the original coloring $\phi|_{x+V_2}$ is ϵ' -regular.*

Proof. For $1 \leq i \leq r$, define $f_i: V \rightarrow [-1, 1]$ by $f_i = 1_{\phi^{-1}(i)}$. Let V_0 be an arbitrary subspace of codimension $\lceil 1/\epsilon \rceil$. Applying Proposition 3.2 with parameter $\epsilon'' = \min(\epsilon/(4r), \epsilon')$ gives subspaces $V_2 \leq V_1 \leq V_0$ and a choice of complement U satisfying $U \oplus V_1 = V$ with the following properties:

- (1') $\text{codim } V_2 \leq n_{\text{reg}}(\epsilon'', q, r, \lceil 1/\epsilon \rceil)$;
- (2') for all but an $\epsilon/(4r)$ -fraction of $x \in U$ we have

$$|\mathbb{E}_{y \in V_1}[f_i(x+y)] - \mathbb{E}_{y \in V_2}[f_i(x+y)]| \leq \frac{\epsilon}{4r}$$

for all $1 \leq i \leq r$;

- (3') for each $x \in U \setminus \{0\}$ and $1 \leq i \leq r$, the function $f_i|_{x+V_2}$ is ϵ' -regular.

Since $V_1 \leq V_0$, clearly $\text{codim } V_1 \geq \text{codim } V_0 \geq 1/\epsilon$. This proves part (1). Part (3) follows from (3'). Next we produce a recoloring $\phi': V \rightarrow [r]$ that satisfies part (2). For each $x \in U$, there is some color i_x that appears with density at least $\epsilon/(2r)$ in $x + V_2$. For each color i which appears with density less than $\epsilon/(2r)$ in $x + V_2$, recolor every element of color i in $x + V_1$ to color i_x . This recoloring satisfies part (2).

All that remains is to show that ϕ and ϕ' agree on all but an ϵ -fraction of V . At least a $(1 - \epsilon/(4r))$ -fraction of $x \in U$ satisfy

$$|\mathbb{E}_{y \in V_1}[f_i(x+y)] - \mathbb{E}_{y \in V_2}[f_i(x+y)]| \leq \frac{\epsilon}{4r} \quad \text{for all } 1 \leq i \leq r. \quad (*)$$

For the $x \in U$ failing to satisfy (*), we use the trivial bound: the number of elements of $x + V_1$ that we recolor is at most all $|V_1|$ of them. For the rest of the $x \in U$, for each color i which occurs with density at most $\epsilon/(2r)$ in $x + V_2$, color i has density at most $\epsilon/(2r) + \epsilon/(4r)$ in $x + V_1$. Thus for each $x \in U$ satisfying (*) we recolor at most a $3\epsilon/4$ -fraction of $x + V_1$. Combining these two cases, ϕ' and ϕ agree on all but at most a $3\epsilon/4 + \epsilon/(4r) \leq \epsilon$ -fraction of V . $\square$

To accomplish the “patch around the origin” step we need the following dichotomy which is proved in Section 4.

Proposition 3.4 (Ramsey dichotomy). *Fix a finite set $\mathcal{H}$ of r -colored patterns over $\mathbb{F}_q$. There exist $n_{\text{rado}} = n_{\text{rado}}(\mathcal{H})$ and $\epsilon_{\text{rado}} = \epsilon_{\text{rado}}(\mathcal{H}) > 0$ such that either:*

- (a) *every coloring $\phi: V \rightarrow [r]$ of V , a finite dimensional $\mathbb{F}_q$ -vector space with $\dim V \geq n_{\text{rado}}$, has H -density at least ϵ_{rado} for some $H \in \mathcal{H}$; or*
- (b) *for every finite dimensional $\mathbb{F}_q$ -vector space V there exists a coloring $\phi: V \setminus \{0\} \rightarrow [r]$ that is $\mathcal{H}$ -free.*

As we saw in the previous discussion, Proposition 3.2 and Proposition 3.3 allow us to recolor $V \setminus V_1$ to be $\mathcal{H}$ -free. Furthermore, applying Proposition 3.4 to $\mathcal{H}$ allows us to recolor $V_1 \setminus \{0\}$ to be $\mathcal{H}$ -free. However this does not rule out H -instances which are contained partially in $V_1 \setminus \{0\}$ and partially in $V \setminus V_1$.

To deal with this problem we apply Proposition 3.4 not only to $\mathcal{H}$ but also to some of the *subpatterns* of the patterns $H \in \mathcal{H}$. The subpatterns of an arithmetic pattern H are analogous to the subgraphs of a graph H .

Definition 3.5. Let $H = (A, \psi)$ be an r -colored pattern where A is an $\ell \times k$ matrix with entries in $\mathbb{F}_q$. For $I \subseteq [k]$ say that $H' = (A', \psi|_I)$ is a *subpattern of H restricted to the variables I* if for all V , a finite dimensional $\mathbb{F}_q$ -vector space, it holds that for $\mathbf{x} \in V^I$ we have $A'\mathbf{x} = \mathbf{0}$ if and only if there exists $\mathbf{y} \in V^k$ satisfying $A\mathbf{y} = \mathbf{0}$ and $x_i = y_i$ for all $i \in I$. (Here $\psi|_I: I \rightarrow [r]$ is the restriction of ψ to I .)

Example 3.6. Consider the following pattern which represents a red 4-term arithmetic progression:

$$H = \left(A = \begin{pmatrix} 1 & -2 & 1 & 0 \\ 0 & 1 & -2 & 1 \end{pmatrix}, \psi \equiv \text{red} \right).$$

The subpattern of H corresponding to the first three variables (i.e., $I = \{1, 2, 3\}$) is a red 3-term arithmetic progression. One possible representation is as:

$$H' = (A' = \begin{pmatrix} 1 & -2 & 1 \end{pmatrix}, \psi \equiv \text{red}).$$

Example 3.7. Consider the matrix

$$A = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \end{pmatrix}.$$

The equation $A\mathbf{x} = \mathbf{0}$ imposes the constraints $x_1 + x_2 + x_3 = 0$ and $x_3 + x_4 + x_5 = 0$. Restricting to variables x_1, x_2, x_3 (i.e., $I = \{1, 2, 3\}$), the only constraint is $x_1 + x_2 + x_3 = 0$ which may be represented by the matrix

$$A' = \begin{pmatrix} 1 & 1 & 1 \end{pmatrix}.$$

Similarly, restricting to variables x_1, x_2, x_4, x_5 (i.e., $I = \{1, 2, 4, 5\}$), the only constraint is $x_1 + x_2 - x_4 - x_5 = 0$ which may be represented by the matrix

$$A' = \begin{pmatrix} 1 & 1 & -1 & -1 \end{pmatrix}.$$

Let $H = (A, \psi)$ be an r -colored pattern where A is an $\ell \times k$ matrix with entries in $\mathbb{F}_q$. Suppose $H' = (A', \psi|_{[j]})$ is a subpattern of H restricted to the first j variables. Our definitions are such that for every finite dimensional $\mathbb{F}_q$ -vector space V and $f_1, \dots, f_k: V \rightarrow [-1, 1]$, one has

$$\Lambda_A(f_1, \dots, f_j, 1, \dots, 1) = \Lambda_{A'}(f_1, \dots, f_j).$$

Proof of Theorem 1.14. We are given $\mathcal{H}$, a finite set of r -colored complexity 1 patterns over $\mathbb{F}_q$, and $\phi: V \rightarrow [r]$, an r -coloring of a finite dimensional $\mathbb{F}_q$ -vector space.

Let $\tilde{\mathcal{H}}$ be a finite set containing one representative of each subpattern of each $H \in \mathcal{H}$. More precisely, for each $H = (A, \psi) \in \mathcal{H}$, a pattern where A is an $\ell \times k$ matrix, and each $I \subseteq [k]$, the set $\tilde{\mathcal{H}}$ contains some $H' = (A', \psi|_I)$, a subpattern of H restricted to variables I . Define constants

$$\epsilon_{\text{rado}} = \min_{\mathcal{H}' \subseteq \tilde{\mathcal{H}}} \epsilon_{\text{rado}}(\mathcal{H}') \quad \text{and} \quad n_{\text{rado}} = \max_{\mathcal{H}' \subseteq \tilde{\mathcal{H}}} n_{\text{rado}}(\mathcal{H}').$$

Write k_{max} for the maximum number of columns in any of the patterns $H \in \mathcal{H}$. With foresight, we define

$$\epsilon_{\text{count}} = \frac{1}{2k_{\text{max}}} \left(\frac{\epsilon}{4r} \right)^{k_{\text{max}}} \epsilon_{\text{rado}}.$$

Define ϵ_{reg} such that for all $H = (A, \psi) \in \tilde{\mathcal{H}}$ and for all $f_1, \dots, f_k: V \rightarrow [-1, 1]$ we have

$$|\Lambda_A(f_1, \dots, f_k)| \leq \epsilon_{\text{count}} \quad \text{whenever} \quad \min_{1 \leq i \leq k} \|\hat{f}_i\|_{\infty} \leq \epsilon_{\text{reg}}. \quad (\dagger)$$

We can do this since all $H \in \mathcal{H}$ are complexity 1 (and this implies that all $H \in \tilde{\mathcal{H}}$ are complexity 1). Define $n_0 = n_{\text{rado}} + n_{\text{reg}}(\epsilon/2, \epsilon_{\text{reg}}, q, r)$. Finally define

$$\delta = \min \left(\frac{1}{4} q^{-k_{\text{max}} \cdot n_{\text{reg}}(\epsilon/2, \epsilon_{\text{reg}}, q, r)} \left(\frac{\epsilon}{4r} \right)^{k_{\text{max}}} \epsilon_{\text{rado}}, q^{-n_0 k_{\text{max}}} \right).$$

First note that the desired result is trivially true for $\dim V < n_0$; if a pattern H with at most $k_{\max}$ columns has density at most $\delta \leq q^{-n_0 k_{\max}}$ in $\phi: V \rightarrow [r]$, then ϕ must have fewer than one H -instance, i.e., be H -free.

Now assume that $\dim V \geq n_0$. With these constants chosen we move on to the recoloring algorithm. First, we apply Proposition 3.3 to ϕ with parameters $\epsilon/2, \epsilon_{\text{reg}}$. This produces $V_2 \leq V_1 \leq V$ and U satisfying $U \oplus V_1 = V$ and a recoloring $\phi': V \setminus V_1 \rightarrow [r]$ that differs from ϕ on at most $(\epsilon/2)|V|$ elements that satisfies:

- (1) $2/\epsilon \leq \text{codim } V_1 \leq \text{codim } V_2 \leq n_{\text{reg}}(\epsilon/2, \epsilon_{\text{reg}}, q, r)$;
- (2) if a color appears in some coset $x + V_1$ under ϕ' , then at least an $\epsilon/(4r)$ -fraction of $x + V_2$ is that color under ϕ ;
- (3) for each $x \in U \setminus \{0\}$, the original coloring $\phi|_{x+V_2}$ is ϵ_{reg} -regular.

Second, define $\mathcal{H}' \subseteq \tilde{\mathcal{H}}$ to be the set of all patterns $H \in \tilde{\mathcal{H}}$ which appear with density less than ϵ_{rado} in $\phi|_{V_2}$. Note that we have $\dim V_2 \geq n_{\text{rado}}$. We now apply Proposition 3.4; by the definition of $\mathcal{H}'$, we cannot be in case (a), since $\phi|_{V_2}$ is a counterexample. Thus we are in case (b) which implies the existence coloring $\phi': V_1 \setminus \{0\} \rightarrow [r]$ that is $\mathcal{H}'$ -free.

Combining these two results gives a recoloring $\phi': V \setminus \{0\} \rightarrow [r]$. First note that ϕ' differs from ϕ on at most $(\epsilon/2)|V| + q^{-\text{codim } V_1}|V| \leq \epsilon|V|$ elements of V . Now suppose that $\mathbf{x}$ is an H -instance in ϕ' for some $(A, \psi) = H \in \mathcal{H}$. Write $\mathbf{x} = \mathbf{u} + \mathbf{v}$ with $\mathbf{u} \in U^k$ and $\mathbf{v} \in V_1^k$. We will prove that this implies that the H -density is large in ϕ ; we do this by counting H -instances of the form $\mathbf{u} + \mathbf{y}$ with $\mathbf{u}$ fixed as above and $\mathbf{y} \in V_2^k$.

Relabeling variables as needed, assume that $u_1, \dots, u_j = 0$ and $u_{j+1}, \dots, u_k \neq 0$. Let $H' = (A', \psi|_{[j]}) \in \tilde{\mathcal{H}}$ be a subpattern of H restricted to the first j variables. The existence of $\mathbf{x}$ means that for $i = j+1, \dots, k$, the color $\psi(i)$ appears at least once in $u_i + V_1$ under ϕ' and it means that there exists an H' -instance in $\phi'|_{V_1}$, namely $(x_1, \dots, x_j)$. By property (2) and the $\mathcal{H}'$ -freeness of $\phi'|_{V_1}$, this implies that for $i = j+1, \dots, k$, the density of color $\psi(i)$ in $u_i + V_2$ under ϕ is at least $\epsilon/(4r)$ and that the H' -density in $\phi|_{V_2}$ is at least ϵ_{rado} .

Define $f_i: V \rightarrow [0, 1]$ by $f_i(x) = 1_{\phi^{-1}(\psi(i))}(x)$ and $g_i: V_2 \rightarrow [0, 1]$ by $g_i(x) = 1_{\phi^{-1}(\psi(i))}(x + u_i)$. In this notation, the conclusions of the last paragraph are that $\mathbb{E}g_i \geq \epsilon/(4r)$ for $i = j+1, \dots, k$ and $\Lambda_{A'}(g_1, \dots, g_j) \geq \epsilon_{\text{rado}}$. Finally we compute

$$\begin{aligned}
\Lambda_A(f_1, \dots, f_k) &\geq q^{-(k-\text{rank } A) \cdot \text{codim } V_2} \Lambda_A(g_1, \dots, g_k) \\
&\geq q^{-(k-\text{rank } A) \cdot \text{codim } V_2} (\Lambda_A(g_1, \dots, g_j, \mathbb{E}g_{j+1}, \dots, \mathbb{E}g_k) - (k-j)\epsilon_{\text{count}}) \\
&\geq q^{-(k-\text{rank } A) \cdot \text{codim } V_2} \left(\left(\frac{\epsilon}{4r} \right)^{k-j} \Lambda_{A'}(g_1, \dots, g_j) - (k-j)\epsilon_{\text{count}} \right) \\
&\geq q^{-k \cdot \text{codim } V_2} \left(\left(\frac{\epsilon}{4r} \right)^k \epsilon_{\text{rado}} - k\epsilon_{\text{count}} \right) \\
&> \delta.
\end{aligned}$$

The first line represents restricting from arbitrary H -instances to H -instances of the form $\mathbf{u} + \mathbf{y}$ with $\mathbf{y} \in V_2^k$ (note that the f_i are non-negative). The second line follows by iterating ($\dagger$). The last line follows from our choice of ϵ_{count} and δ .

Thus if the above recoloring procedure does not produce an $\mathcal{H}$ -free coloring, then $\Lambda_A(f_1, \dots, f_k) > \delta$, meaning that the original coloring had H -density more than δ for some $H \in \mathcal{H}$. This completes the proof. $\square$

4. PROOF OF PROPOSITION 3.4: A DENSITY RAMSEY DICHOTOMY

In this section we prove Proposition 3.4. The proof is based on the following trivial dichotomy:

For any $\mathcal{H}$, a finite set of r -colored patterns over $\mathbb{F}_q$, there exists a constant n_0 such that either:

- (a) every r -coloring of every V with $\dim V \geq n_0$ contains an H -instance for some $H \in \mathcal{H}$; or
- (b) for every finite dimensional $\mathbb{F}_q$ -vector space V , there exists an r -coloring of V that is $\mathcal{H}$ -free.

To boost this result to the full strength of Proposition 3.4 we use a sampling argument. The idea of the sampling is simple. Suppose we have some $\mathcal{H}$ that is in case (a). Then for an r -coloring of V we know that there is not just one H -instance, but one H -instance in every dimension n_0 subspace of V . Adding up all of these gives many H -instances in V . Unfortunately we have overcounted each H -instance many times. Double counting more carefully gives us the desired result, but only if the original H -instances are *generic* in the following sense.

Definition 4.1. For A an $\ell \times k$ matrix with entries in $\mathbb{F}_q$ and $\mathbf{x} = (x_1, \dots, x_k) \in V^k$ satisfying $A\mathbf{x} = \mathbf{0}$, say that $\mathbf{x}$ is a *generic solution* if $\dim(\text{span}\{x_1, \dots, x_k\}) = k - \text{rank } A$. For $H = (A, \psi)$ a colored pattern, define a *generic H -instance* to be an H -instance $\mathbf{x}$ that is a generic solution to $A\mathbf{x} = \mathbf{0}$.

To prove Proposition 3.4 we first show that we can always turn an H -instance into a generic H -instance, then we perform the sampling argument.

The first step relies on an argument which is inspired by elements of the proof of Rado's theorem (see [8]). We reduce the question of whether $\mathcal{H}$ lies in case (a) of Proposition 3.4 to a finite check. Instead of checking if there exists an $\mathcal{H}$ -instance in every coloring we only need to check the so-called *canonical colorings*. (With some manipulations this check can be made to resemble the “columns condition” of Rado's theorem.)

Definition 4.2. For $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$ and $n > 0$, define the χ -canonical coloring $\Phi_{n,\chi}: \mathbb{F}_q^n \setminus \{0\} \rightarrow [r]$ by $\Phi_{n,\chi}(x) := \chi(\text{fnz}(x))$ where $\text{fnz}: \mathbb{F}_q^n \setminus \{0\} \rightarrow \mathbb{F}_q \setminus \{0\}$ maps a vector to its first non-zero coordinate.

Lemma 4.3. For $\mathcal{H}$, a finite set of r -colored patterns over $\mathbb{F}_q$, write k for the maximum number of columns in any of the patterns $H \in \mathcal{H}$. There exists some $n_{\text{rado}} = n_{\text{rado}}(\mathcal{H})$ such that for any $n_0 \geq n_{\text{rado}}$ the following are equivalent:

- (a) For every finite dimensional $\mathbb{F}_q$ -vector space V satisfying $\dim V \geq n_0$ and every coloring $\phi: V \setminus \{0\} \rightarrow [r]$, there is an H -instance in ϕ for some $H \in \mathcal{H}$.
- (b) For every $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$, there is an H -instance in $\Phi_{k,\chi}$ for some $H \in \mathcal{H}$.
- (c) For every finite dimensional $\mathbb{F}_q$ -vector space V satisfying $\dim V \geq n_{\text{rado}}$ and every coloring $\phi: V \setminus \{0\} \rightarrow [r]$, there is a generic H -instance in ϕ for some $H \in \mathcal{H}$.

Note that in the above lemma statement, an H -instance always refers to $\mathbf{x}$ where none of the vectors x_i are 0 since all of the colorings considered leave 0 uncolored. We will define the function $n_{\text{rado}}(\mathcal{H})$ in the statement of Proposition 3.4 to be the same as the $n_{\text{rado}}(\mathcal{H})$ in this lemma.

We use the following vector space Ramsey theorem of Graham and Rothschild which was proved with primitive-recursive bounds by Shelah.

Theorem 4.4 (Graham-Rothschild [13], Shelah [25]¹). *There exists $n_{\text{GR}} = n_{\text{GR}}(q, r, k)$ such that for a finite dimensional $\mathbb{F}_q$ -vector space V with $\dim V \geq n_{\text{GR}}$ and $\phi: V \setminus \{0\} \rightarrow [r]$, an r -coloring that is constant on one-dimensional subspaces, there exists a k -dimensional subspace $U \leq V$ such that $U \setminus \{0\}$ is monochromatic under ϕ .*

The best-known upper bound on n_{GR} as a function of k is broadly comparable to the third function in the Ackermann hierarchy.

¹Note that both of these references prove a slight variant of this result, the “affine Ramsey theorem”. Spencer [26] gives an easy deduction of this result, referred to as the “vector space Ramsey theorem”, from the affine Ramsey theorem.

Proof of Lemma 4.3. Define $n_{\text{rado}}(\mathcal{H}) = n_{\text{GR}}(q, r^{q-1}, 2k)$ and fix any $n_0 \geq n_{\text{rado}}$. With these choices we show that $(a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (a)$.

$(a) \Rightarrow (b)$: Fix $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$ and take $n \geq n_0$. Assuming (a), there is an H -instance in $\Phi_{n,\chi}$ for some $H \in \mathcal{H}$. Say that $H = (A, \psi)$ where A is an $\ell \times k'$ matrix and call the H -instance $\mathbf{x} \in (\mathbb{F}_q^n \setminus \{0\})^{k'}$. We use this $\mathbf{x}$ to find an H -instance in $\Phi_{k,\chi}$.

For each $1 \leq j \leq k'$, let b_j be the position of the first non-zero coordinate of x_j . Relabeling the variables as necessary, assume that $b_1 \leq b_2 \leq \dots \leq b_{k'}$. Consider the map $\Pi: \mathbb{F}_q^n \rightarrow \mathbb{F}_q^k$ which sends $v = (v_1, \dots, v_n) \in \mathbb{F}_q^n$ to $(v_{b_1}, v_{b_2}, \dots, v_{b_{k'}}, 0, \dots, 0) \in \mathbb{F}_q^k$, where there are $k - k'$ zeros added at the end. One can easily check that each x_j satisfies $\Phi_{n,\chi}(x_j) = \Phi_{k,\chi}(\Pi(x_j))$. Thus $\Pi(\mathbf{x})$ is an H -instance in $\Phi_{k,\chi}$, as desired.

$(b) \Rightarrow (c)$: Take $n \geq n_{\text{rado}}$ and a coloring $\phi: V \setminus \{0\} \rightarrow [r]$ of a finite dimensional $\mathbb{F}_q$ -vector space V satisfying $\dim V = n$. We define an r^{q-1} -coloring $\bar{\phi}: V \setminus \{0\} \rightarrow [r]^{q-1}$ that is constant on one-dimensional subspaces as follows. Pick an isomorphism $V \cong \mathbb{F}_q^n$. For $x \in V \setminus \{0\}$ there exists a unique $a \in \mathbb{F}_q \setminus \{0\}$ such that the first non-zero coordinate of ax is 1. Then define

$$\bar{\phi}(x) := (\phi(ax))_{b \in \mathbb{F}_q \setminus \{0\}} \in [r]^{q-1}.$$

By Theorem 4.4 and our choice of n_{rado} , there exists $U \leq V$, a $2k$ -dimensional subspace such that $\bar{\phi}$ is constant on $U \setminus \{0\}$. We claim that this implies that there is a coloring $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$ and an isomorphism $\iota: U \xrightarrow{\sim} \mathbb{F}_q^{2k}$ such that ϕ agrees with the canonical coloring $\Phi_{2k,\chi}$ on $U \setminus \{0\}$.

The color of $U \setminus \{0\}$ under $\bar{\phi}$ is some vector in $[r]^{q-1}$, in other words, a function $\mathbb{F}_q \setminus \{0\} \rightarrow [r]$. This is our χ . Now pick a basis for $U \leq \mathbb{F}_q^n$. Performing Gaussian elimination on this basis transforms it into a basis $u_1, \dots, u_{2k} \in U$ such that the first non-zero coordinate of each of these vectors is 1 and the location of the first non-zero coordinates is strictly increasing from u_1 to u_{2k} . Then define $\iota: U \xrightarrow{\sim} \mathbb{F}_q^{2k}$ to send $c_1 u_1 + \dots + c_{2k} u_{2k} \mapsto (c_1, \dots, c_{2k}) \in \mathbb{F}_q^{2k}$. These χ, ι have the desired properties since $\text{fnz}(c_1 u_1 + \dots + c_{2k} u_{2k}) = \text{fnz}((c_1, \dots, c_{2k}))$.

Consider $\mathbb{F}_q^k \oplus \{0\} \subset \mathbb{F}_q^{2k}$, the subspace of vectors of the form $(*, \dots, *, 0, \dots, 0)$. Assuming (b), there exists an H -instance $\mathbf{x} \in ((\mathbb{F}_q^k \setminus \{0\}) \oplus \{0\})^{k'}$ in $\Phi_{2k,\chi}$. Now for any $\mathbf{y} \in (\{0\} \oplus \mathbb{F}_q^k)^{k'}$ we have that x_j and $x_j + y_j$ are the same color under $\Phi_{2k,\chi}$ for all j . Pick $\mathbf{y}$ any generic solution to $A\mathbf{y} = \mathbf{0}$ in $\{0\} \oplus \mathbb{F}_q^k$. Then $\mathbf{x} + \mathbf{y}$ is a generic H -instance in $\Phi_{2k,\chi}$. Since $\phi|_U$ agrees with $\Phi_{2k,\chi}$, we have found a generic H -instance in ϕ , as desired.

$(c) \Rightarrow (a)$: Obvious. □

Proof of Proposition 3.4. Fix a finite set $\mathcal{H}$ of r -colored patterns over $\mathbb{F}_q$. Write k_{max} for the maximum number of columns in any of the patterns $H \in \mathcal{H}$. Define $n_{\text{rado}} = n_{\text{rado}}(\mathcal{H})$ to be the same as in Lemma 4.3 and, with foresight, define

$$\epsilon_{\text{rado}} = \frac{1}{10|\mathcal{H}|} q^{-n_{\text{rado}} k_{\text{max}}}.$$

Suppose that we are not in case (b) of Proposition 3.4; namely, that there is some finite dimensional $\mathbb{F}_q$ -vector space V_0 such that every coloring $\phi: V_0 \setminus \{0\} \rightarrow [r]$ contains an H -instance for some $H \in \mathcal{H}$. We now apply Lemma 4.3 to $\mathcal{H}$ with parameter $n_0 := \max\{\dim V_0, n_{\text{rado}}\}$. By assumption, statement (a) of Lemma 4.3 is true, so we conclude that statement (c) is also true: for every finite dimensional $\mathbb{F}_q$ -vector space V satisfying $\dim V \geq n_{\text{rado}}$ and every coloring $\phi: V \setminus \{0\} \rightarrow [r]$, there is a generic H -instance in ϕ for some $H \in \mathcal{H}$.

Now we prove we are in case (a) of Proposition 3.4. Fix V , a finite dimensional $\mathbb{F}_q$ -vector space with $\dim V = n \geq n_{\text{rado}}$, and a coloring $\phi: V \rightarrow [r]$. We count the total number of generic $\mathcal{H}$ -instances in ϕ . For ease of notation, given $H = (A, \psi) \in \mathcal{H}$, write k_H for the number of columns in A .

Consider the following random process: pick a uniform random $H = (A, \psi) \in \mathcal{H}$ and then pick a uniform random generic solution $A\mathbf{x} = \mathbf{0}$ with $\mathbf{x} \in V^{k_H}$. Write $N_H(\phi)$ for the number of generic

H -instances in ϕ . Then the probability that this process produces a pair $(H, \mathbf{x})$ where $\mathbf{x}$ is a generic H -instance is

$$\frac{1}{|\mathcal{H}|} \sum_{H=(A,\psi) \in \mathcal{H}} \frac{N_H(\phi)}{(q^n - 1)(q^n - q) \cdots (q^n - q^{k_H - \text{rank } A - 1})} \leq \frac{10}{|\mathcal{H}|} \sum_{H=(A,\psi) \in \mathcal{H}} N_H(\phi) q^{-n(k_H - \text{rank } A)}.$$

Now consider the following random process: pick $U \leq V$, a uniform random subspace of dimension n_{rado} , pick a uniform random $H = (A, \psi) \in \mathcal{H}$, and then pick a uniform random generic solution $A\mathbf{x} = 0$ with $\mathbf{x} \in U^{k_H}$. Since, by assumption, there is at least one generic $\mathcal{H}$ -instance in $\phi|_U$, the probability that this process finds a pair $(H, \mathbf{x})$ where $\mathbf{x}$ is a generic H -instance is at least

$$\frac{1}{|\mathcal{H}|} \frac{1}{(q^{n_{\text{rado}}} - 1)(q^{n_{\text{rado}}} - q) \cdots (q^{n_{\text{rado}}} - q^{k_{\text{max}} - 1})} \geq \frac{1}{|\mathcal{H}|} q^{-n_{\text{rado}} k_{\text{max}}}.$$

Note that these two random processes actually produce the same distribution on pairs $(H, \mathbf{x})$. (This is the key place where we use the assumption that the $\mathbf{x}$ are generic solutions.) Combining these two bounds, we conclude that

$$\sum_{H=(A,\psi) \in \mathcal{H}} N_H(\phi) q^{-n(k_H - \text{rank } A)} \geq \frac{1}{10} q^{-n_{\text{rado}} k_{\text{max}}}.$$

Thus by the pigeonhole principle, there exists some $H = (A, \psi) \in \mathcal{H}$ such that $N_H(\phi) q^{-n(k_H - \text{rank } A)}$, i.e., the H -density in ϕ , is at least ϵ_{rado} , as desired. $\square$

5. PROOF OF PROPOSITION 3.2: ARITHMETIC REGULARITY LEMMAS

We give two proofs of Proposition 3.2, giving wowzer- and tower-type bounds respectively. The first proof can be considered the arithmetic analog of the strong regularity argument in [2] and the second bears some resemblance to an arithmetic analog of the cylinder regularity argument in [6, Theorem 1.3].

Our main technique for proving arithmetic regularity lemmas is the notion of energy.

Definition 5.1. Fix a finite dimensional $\mathbb{F}_q$ -vector space V and functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$. Given $\mathcal{P}$ a partition of $S \subseteq V$, define $(f_i)_{\mathcal{P}}: S \rightarrow [-1, 1]$ to be the projection of $f_i|_S$ onto the σ -algebra generated by $\mathcal{P}$. In other words, $(f_i)_{\mathcal{P}}(x)$ is defined to be the average of f_i over the part of $\mathcal{P}$ containing x . The *energy* of the partition is defined to be

$$\mathcal{E}(\mathcal{P}) := \sum_{i=1}^k \|(f_i)_{\mathcal{P}}\|_{L^2(S)}^2.$$

One case will be of particular interest to us. Let $V_1 \leq V$ be a linear subspace and $S \subseteq V$ any set which can be written as the union of cosets of V_1 . Then write $\mathcal{P}(V_1|S)$ for the partition of S into cosets of V_1 . When S is not given, it is assumed to be all of V .

We record the following three properties of energy.

Proposition 5.2. *For a finite dimensional $\mathbb{F}_q$ -vector space V and functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, we have:*

(1) *for any partition $\mathcal{P}$ of a set $S \subseteq V$,*

$$0 \leq \mathcal{E}(\mathcal{P}) \leq k;$$

(2) *for any partitions $\mathcal{P}, \mathcal{Q}$ of a set $S \subseteq V$ with $\mathcal{Q}$ refining $\mathcal{P}$ (written $\mathcal{Q} \succeq \mathcal{P}$),*

$$\mathcal{E}(\mathcal{Q}) - \mathcal{E}(\mathcal{P}) = \sum_{i=1}^k \|(f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}}\|_{L^2(S)}^2 \geq 0.$$

(3) for $V_1 \leq V$, if there exists $1 \leq i \leq k$ and $x \in V$ such that $f_i|_{x+V_1}$ is not ϵ -regular, there exists $V_2 \leq V_1$ of codimension 1 such that

$$\mathcal{E}(\mathcal{P}(V_2|x+V_1)) - \mathcal{E}(\mathcal{P}(V_1|x+V_1)) > \epsilon^2.$$

Proof. The first result follows since trivially $0 \leq \|(f_i)_{\mathcal{P}}\|_{L^2(S)}^2 \leq 1$.

The second result follows from the Pythagorean theorem: $\|(f_i)_{\mathcal{P}}\|_{L^2(S)}^2 + \|(f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}}\|_{L^2(S)}^2 = \|(f_i)_{\mathcal{Q}}\|_{L^2(S)}^2$ as long as $(f_i)_{\mathcal{P}}$ and $(f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}}$ are orthogonal. We compute

$$\langle (f_i)_{\mathcal{P}}, (f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}} \rangle_{L^2(S)} = \langle f_i, ((f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}})_{\mathcal{P}} \rangle_{L^2(S)} = \langle f_i, 0 \rangle_{L^2(S)} = 0.$$

To prove the third result, suppose that $|\widehat{f_i|_{x+V_1}}(\gamma_z)| > \epsilon$ for some $\gamma_z \in V_1^* \setminus \{0\}$. Recall that γ_z is the character $\gamma_z(x) := \exp(2\pi i \operatorname{tr}(x \cdot z)/p)$. Define $V_2 \leq V_1$ to be $V_2 = \{y \in V_1 : y \cdot z = 0\}$. For ease of notation, write $\mathcal{P}$ for $\mathcal{P}(V_1|x+V_1)$ and $\mathcal{Q}$ for $\mathcal{P}(V_2|x+V_1)$ for the rest of the proof. Note that γ_z is constant on parts of $\mathcal{Q}$. Thus we can write

$$\begin{aligned} \epsilon^2 &< \left| \widehat{f_i|_{x+V_1}}(\gamma_z) \right|^2 \\ &= \left| \langle f_i, \gamma_z \rangle_{L^2(x+V_1)} \right|^2 \\ &= \left| \langle (f_i)_{\mathcal{Q}}, \gamma_z \rangle_{L^2(x+V_1)} \right|^2 \\ &= \left| \langle (f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}}, \gamma_z \rangle_{L^2(x+V_1)} \right|^2 && (\text{since } \mathbb{E}\gamma_z = 0 \text{ and } (f_i)_{\mathcal{P}} \text{ is constant}) \\ &\leq \|(f_i)_{\mathcal{Q}} - (f_i)_{\mathcal{P}}\|_{L^2(x+V_1)}^2 && (\text{Cauchy-Schwarz inequality}) \\ &\leq \mathcal{E}(\mathcal{Q}) - \mathcal{E}(\mathcal{P}). \end{aligned}$$

□

5.1. Proof 1: Strong arithmetic regularity. We first prove Green's arithmetic regularity lemma in finite field vector spaces by the standard energy increment method. We then iterate this arithmetic regularity lemma to prove a strong arithmetic regularity lemma. Given functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, a subspace $V_0 \leq V$, and $\epsilon > 0$, this strong arithmetic regularity lemma gives subspaces $V_2 \leq V_1 \leq V_0$ with several desirable properties. To complete the proof of Proposition 3.2 we choose a random complement U satisfying $U \oplus V_1 = V$ and show that U has the desired properties with positive probability.

For completeness we include the standard proof of Green's arithmetic regularity lemma.

Theorem 5.3 (Green's arithmetic regularity). *For a finite dimensional $\mathbb{F}_q$ -vector space V , functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, a subspace $V_0 \leq V$, and a parameter $\epsilon > 0$, there exists $n_{\text{green-reg}} = n_{\text{green-reg}}(\epsilon, q, k, \operatorname{codim} V_0)$ such that there exists a subspace $V_1 \leq V_0$ such that*

- (1) $\operatorname{codim} V_1 \leq n_{\text{green-reg}}$;
- (2) for each $1 \leq i \leq k$, for all but an ϵ -fraction of $x \in V$ the function $f_i|_{x+V_1}$ is ϵ -regular.

Proof. We iterate to produce a sequence of subspaces $V_0 \geq V_1 \geq \dots \geq V_M$ such that V_M has the desired properties. This sequence will satisfy $\operatorname{codim} V_{m+1} \leq \operatorname{codim} V_m + q^{\operatorname{codim} V_m}$ and $\mathcal{E}(\mathcal{P}(V_{m+1})) > \mathcal{E}(\mathcal{P}(V_m)) + \epsilon^3$.

Suppose that we have produced some V_m that does not satisfy property (2) above. Pick some complement U satisfying $U \oplus V_m = V$. Then there exist some $1 \leq i \leq k$ and $\epsilon|U|$ values of $x \in U$, say $U' \subset U$, such that $f_i|_{x+V_m}$ is not ϵ -regular for each $x \in U'$.

Fix $x \in U'$. Since $f_i|_{x+V_m}$ is not ϵ -regular, part (3) of Proposition 5.2 gives a subspace $V(x) \leq V_m$ of codimension 1 such that $\mathcal{E}(\mathcal{P}(V(x)|x+V_m)) > \mathcal{E}(\mathcal{P}(V_m|x+V_m)) + \epsilon^2$. Define V_{m+1} to be the intersection of all of the $V(x)$'s. Then $\operatorname{codim} V_{m+1} \leq \operatorname{codim} V_m + |U'| \leq \operatorname{codim} V_m + q^{\operatorname{codim} V_m}$, as desired.

Note that for $W \leq V_m$, we have $\mathcal{E}(\mathcal{P}(W)) = \mathbb{E}_{x \in U}[\mathcal{E}(\mathcal{P}(W|x + V_m))]$. Therefore

$$\begin{aligned} \mathcal{E}(\mathcal{P}(V_{m+1})) - \mathcal{E}(\mathcal{P}(V_m)) &= \mathbb{E}_{x \in U}[\mathcal{E}(\mathcal{P}(V_{m+1}|x + V_m)) - \mathcal{E}(\mathcal{P}(V_m|x + V_m))] \\ &\geq \frac{|U'|}{|U|} \mathbb{E}_{x \in U'}[\mathcal{E}(\mathcal{P}(V(x)|x + V_m)) - \mathcal{E}(\mathcal{P}(V_m|x + V_m))] \\ &> \epsilon^3. \end{aligned}$$

Since $0 \leq \mathcal{E}(\mathcal{P}) \leq k$ for all $\mathcal{P}$, we conclude that $M \leq k\epsilon^{-3}$. Thus V_M has the desired properties where $n_{\text{green-reg}}$ grows at a rate comparable to a tower of q 's of height $k\epsilon^{-3}$. $\square$

Theorem 5.4 (Strong arithmetic regularity). *For a finite dimensional $\mathbb{F}_q$ -vector space V , functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, a subspace $V_0 \leq V$, and parameters $\delta > 0$ and $\epsilon = (\epsilon_0, \epsilon_1, \dots)$ satisfying $\epsilon_0 \geq \epsilon_1 \geq \dots > 0$, there exists $n_{\text{strong-reg}} = n_{\text{strong-reg}}(\delta, \epsilon, q, k, \text{codim } V_0)$ such that there exist subspaces $V_2 \leq V_1 \leq V_0$ such that*

- (1) $\text{codim } V_2 \leq n_{\text{strong-reg}}$;
- (2) $\mathcal{E}(\mathcal{P}(V_2)) \leq \mathcal{E}(\mathcal{P}(V_1)) + \delta$;
- (3) for each $1 \leq i \leq k$, for all but an $\epsilon_{\text{codim } V_1}$ -fraction of $x \in V$ the function $f_i|_{x+V_2}$ is $\epsilon_{\text{codim } V_1}$ -regular.

Proof. We iterate Green's arithmetic regularity lemma to produce a sequence of subspaces $V_0 = V^{(0)} \geq V^{(1)} \geq \dots \geq V^{(M)}$ such that setting $V_1 = V^{(M-1)}$ and $V_2 = V^{(M)}$ proves the theorem. Let $V^{(m+1)}$ be the result of Theorem 5.3 to $V^{(m)}$ with parameter $\epsilon_{\text{codim } V_m}$.

Note that $0 \leq \mathcal{E}(\mathcal{P}(V^{(0)})) \leq \mathcal{E}(\mathcal{P}(V^{(1)})) \leq \dots \leq \mathcal{E}(\mathcal{P}(V^{(M)})) \leq k$. Therefore we can stop the iteration at $M \leq k\delta^{-1}$ such that $\mathcal{E}(\mathcal{P}(V^{(M)})) - \mathcal{E}(\mathcal{P}(V^{(M-1)})) \leq \delta$, as desired. $\square$

Proof of Proposition 3.2. We can assume that $|V_0| \leq (\epsilon/4)|V|$. (If this does not hold, simply replace V_0 with a subspace of codimension at most $\lceil \log_q(4\epsilon^{-1}) \rceil$; if this is impossible, then the desired result is trivially true by setting $V_1 = V_2 = \{0\}$.)

Apply Theorem 5.4 to $f_1, \dots, f_k: V \rightarrow [-1, 1]$ and $V_0 \leq V$ with parameters $\delta = \epsilon^3/4$ and $\epsilon_m = \min(\epsilon, q^{-m}/(2k))$. Write $V_2 \leq V_1 \leq V_0$ for the subspaces produced. We claim that a uniform random choice of $U \leq V$ satisfying $U \oplus V_1 = V$ also satisfies the desired properties with positive probability.

For each i , pick $x \in V$ uniformly at random. Property (3) implies that with probability at least $1 - \epsilon_{\text{codim } V_1}$, the function $f_i|_{x+V_2}$ is ϵ -regular. By the union bound, with probability at least $1 - \epsilon_{\text{codim } V_1} \cdot k \cdot (|U| - 1) > 1/2$, for every $x \in U \setminus \{0\}$ and every $1 \leq i \leq k$, the function $f_i|_{x+V_2}$ is ϵ -regular.

Now by part (2) of Proposition 5.2, we write

$$\begin{aligned} \epsilon^3/4 &\geq \mathcal{E}(\mathcal{P}(V_2)) - \mathcal{E}(\mathcal{P}(V_1)) \\ &= \sum_{i=1}^k \left\| (f_i)_{\mathcal{P}(V_2)} - (f_i)_{\mathcal{P}(V_1)} \right\|_{L^2(V)}^2 \\ &= \mathbb{E}_{x \in V} \left[\sum_{i=1}^k (\mathbb{E}_{y \in V_2}[f_i(x+y)] - \mathbb{E}_{y \in V_1}[f_i(x+y)])^2 \right]. \end{aligned}$$

This implies that for a uniform random choice of $x \in V$, with probability at least $1 - \epsilon/4$, this x satisfies

$$|\mathbb{E}_{y \in V_2}[f_i(x+y)] - \mathbb{E}_{y \in V_1}[f_i(x+y)]| \leq \epsilon \quad \text{for all } 1 \leq i \leq k. \quad (\ddagger)$$

Therefore choosing U randomly, the expected number of $x \in U$ failing to satisfy $(\ddagger)$ is at most $1 + (\epsilon/4)(|U| - 1) < (\epsilon/2)|U|$. By Markov's inequality, with probability at least $1/2$, at most an ϵ -fraction of $x \in U$ fail to satisfy $(\ddagger)$. Thus by the union bound, U has the desired properties with positive probability. $\square$

5.2. Proof 2: An improved bound. We prove a weak arithmetic regularity lemma with exponential bounds and then iterate this lemma to produce a strong arithmetic regularity lemma with tower-type bounds that is just strong enough to prove Proposition 3.2.

Definition 5.5. For a finite dimensional $\mathbb{F}_q$ -vector space V and $U \leq V$, a *decomposition of V with respect to U* is a set of subspaces $\mathcal{D} = \{W_1, \dots, W_\ell\}$ all of the same codimension and all transverse to U (i.e., $W_i \cup U$ spans V for each i) such that $W_1 \setminus U, \dots, W_\ell \setminus U$ are disjoint sets that together partition $V \setminus U$. We write $\text{codim } \mathcal{D} := \text{codim } W_1 = \dots = \text{codim } W_\ell$.

One can easily see that every decomposition $\mathcal{D}$ has size

$$|\mathcal{D}| = \frac{q^{\dim V} - q^{\dim U}}{q^{\dim V - \text{codim } \mathcal{D}} - q^{\dim U - \text{codim } \mathcal{D}}} = q^{\text{codim } \mathcal{D}}.$$

To build intuition, we give several constructions of decompositions of increasing complexity. We start with a very simple example.

Example 5.6. Take $V = \mathbb{F}_q^2$ and $U = \{(0, x) : x \in \mathbb{F}_q\}$. We define $\mathcal{D}$, a decomposition of V with respect to U , to be the set of all 1-dimensional subspaces of V other than U itself. More concretely, $\mathcal{D} = \{W_a\}_{a \in \mathbb{F}_q}$ where $W_a = \{(x, ax) : x \in \mathbb{F}_q\}$.

Next we give a slightly more complicated example.

Example 5.7. Take $V = \mathbb{F}_q^{2n}$ and $U = \{0\} \oplus \mathbb{F}_q^n$. We view $V \cong \mathbb{F}_{q^n} \oplus \mathbb{F}_{q^n}$, an isomorphism of $\mathbb{F}_q$ -vector spaces. Then we take the same construction as the previous example: $\mathcal{D} = \{W_a\}_{a \in \mathbb{F}_{q^n}}$ where $W_a = \{(x, ax) : x \in \mathbb{F}_{q^n}\}$.

Both these examples are very special: they both have the property that $W \cap U = \{0\}$ for each $W \in \mathcal{D}$. We can make our examples slightly more interesting as follows.

Example 5.8. Take any $U' \leq U \leq V$ with $\dim V - \dim U = \dim U - \dim U'$. The previous example gives a decomposition $\mathcal{D}$ of V/U' with respect to U/U' . Then $\mathcal{D}' = \{W \oplus U'\}_{W \in \mathcal{D}}$ is a decomposition of V with respect to U .

The decompositions $\mathcal{D}$ produced by Example 5.8 are slightly more interesting, but they satisfy $W \cap U = U'$ for each $W \in \mathcal{D}$. To prove our decomposition regularity lemma, we will need to consider more complicated decompositions.

For two decompositions $\mathcal{D}, \mathcal{D}'$ of V with respect to U , say that $\mathcal{D}'$ refines $\mathcal{D}$ (written as $\mathcal{D}' \succeq \mathcal{D}$) if the partition $\bigcup_{W \in \mathcal{D}'} (W \setminus U)$ refines the partition $\bigcup_{W \in \mathcal{D}} (W \setminus U)$.

We can create more complicated examples of decompositions of V with respect to U as follows. We construct a sequence of decompositions $\{V\} = \mathcal{D}_0 \preceq \mathcal{D}_1 \preceq \dots$ iteratively. To get from $\mathcal{D}_m$ to $\mathcal{D}_{m+1}$, we refine each part of $\mathcal{D}_m$. Fix $W \in \mathcal{D}_m$ and pick any subspace $U' \leq U \cap W$ such that $\dim W - \dim(U \cap W) = \dim(U \cap W) - \dim U'$. Example 5.8 gives a decomposition of W with respect to U . Replacing each part $W \in \mathcal{D}_m$ with the subspaces produces this way gives a new decomposition $\mathcal{D}_{m+1}$.

Fix $f_1, \dots, f_k: V \rightarrow [-1, 1]$ and a subspace $U \leq V$. For a decomposition $\mathcal{D}$ of V with respect to U , we associate to it the partition $\mathcal{P}(\mathcal{D})$ whose parts are, for each $W \in \mathcal{D}$, the cosets of $W \cap U$ contained in $W \setminus U$. As in the previous section, we will consider the energy

$$\mathcal{E}(\mathcal{P}(\mathcal{D})) = \sum_{i=1}^k \|(f_i)_{\mathcal{P}}\|_{L^2(V \setminus U)}^2.$$

Now we look at an example, depicted in Figure 1, to explain why using decompositions can help us prove a more efficient regularity lemma. Suppose $U \leq \mathbb{F}_2^n$ is a codimension 1 subspace. Then $\mathbb{F}_2^n \setminus U$ is just $(n-1)$ -dimensional affine space. Thus a decomposition of $\mathbb{F}_2^n$ with respect to U is simply a partition of $\mathbb{F}_2^{n-1}$ into affine subspaces of the same codimension. Contrast this to the




FIGURE 1. An example of refining partitions and decompositions.

situation in the previous subsection where we partitioned our space into affine subspaces which were all cosets of the same subspace.

This freedom allows us to prove a more efficient regularity lemma. Suppose after some number of iterations we have produced a partition of $\mathbb{F}_2^{n-1}$ depicted in Figure 1(a). In both the proof of arithmetic regularity and our weak regularity lemma, the next step is to find parts of the partition for which the functions $f_1, \dots, f_k$ are not regular and find a refinement of each of these parts which increases the energy. Suppose that the first three parts have refinements which increase the energy as depicted in Figure 1(b).

If we want our next partition to be made of cosets of the same subspace (as in the proof of arithmetic regularity) we need to take the common refinement, pictured in Figure 1(c). However, if all we want is to produce a decomposition (as will be the case in the proof of our weak arithmetic regularity lemma) we simply have to arbitrarily partition the remaining parts to make every part the same codimension, pictured in Figure 1(d).

Therefore we can prove our weak arithmetic regularity via an iteration where the codimension of our decomposition only increases by a constant each iteration where in proving arithmetic regularity the codimension of the partition increases by an exponential at every step of the iteration.

For the case of general $U \leq V$ the proof is slightly more complicated, but the main idea is the same. Instead of a picture like the one above, we have $q^{\text{codim } U} - 1$ similar copies to work with (the cosets of U). Since U is not codimension 1, we cannot use Example 5.6 at each step to refine, but instead have to use Example 5.7. In addition, each step of the iteration produces a smaller energy increment ($\epsilon^2/q^{\text{codim } U}$ instead of ϵ^2), but this increment still produces exponential-type bounds instead of tower-type bounds.

Now we begin the proof.

Theorem 5.9 (Weak decomposition regularity). *For a finite dimensional $\mathbb{F}_q$ -vector space V , functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, a subspace $U \leq V$, and a parameter $\epsilon > 0$, if $\dim V \geq \text{codim } U \cdot q^{\text{codim } U} k \epsilon^{-3}$, then there exists a decomposition $\mathcal{D}$ such that*

- (1) $\text{codim } \mathcal{D} \leq \text{codim } U \cdot q^{\text{codim } U} k \epsilon^{-3}$;
- (2) for each $1 \leq i \leq k$, for all but an ϵ -fraction of $W \in \mathcal{D}$, for all $x \in W \setminus U$ the function $f_i|_{x+W \cap U}$ is ϵ -regular.

Proof. We iterate to produce a sequence of decompositions $\{V\} = \mathcal{D}_0 \preceq \mathcal{D}_1 \preceq \dots \preceq \mathcal{D}_M$ such that $\mathcal{D}_M$ has the desired properties. The sequence will satisfy $\text{codim } \mathcal{D}_{m+1} = \text{codim } \mathcal{D}_m + \text{codim } U$ and $\mathcal{E}(\mathcal{P}(\mathcal{D}_{m+1})) > \mathcal{E}(\mathcal{P}(\mathcal{D}_m)) + \epsilon^3 q^{-\text{codim } U}$.

Suppose that we have produced some $\mathcal{D}_m$ that does not satisfy property (2) above. Then there is some $1 \leq i \leq k$ and $\epsilon|\mathcal{D}_m|$ choices of $W \in \mathcal{D}_m$, say $\mathcal{D}' \subset \mathcal{D}_m$, such that for each $W \in \mathcal{D}'$, there exists $x \in W \setminus U$ such that $f_i|_{x+W \cap U}$ is not ϵ -regular.

Fix $W \in \mathcal{D}'$ with associated $x \in W \setminus U$. Since $f_i|_{x+W \cap U}$ is not ϵ -regular, by part (3) of Proposition 5.2, there exists some $W' \leq W \cap U$ of codimension 1 such that $\mathcal{E}(\mathcal{P}(W'|x+W \cap U)) > \mathcal{E}(\mathcal{P}(W \cap U|x+W \cap U)) + \epsilon^2$. Now we pick $U' \leq W'$ such that $\dim W - \dim(W \cap U) = \dim(W \cap U) - \dim U'$. Applying Example 5.8 to $U' \leq (W \cap U) \leq W$ gives a decomposition $\mathcal{D}_W$ of W with respect to $W \cap U$ such that the partition $\mathcal{P}(\mathcal{D}_W)$ of $W \setminus U$ refines the partition $\mathcal{P}(W'|W \setminus U)$ (the partition of $W \setminus U$ into cosets of W'). This implies that

$$\mathcal{E}(\mathcal{P}(\mathcal{D}_W)) > \mathcal{E}(\mathcal{P}(W \cap U|W \setminus U)) + \frac{\epsilon^2}{q^{\text{codim } U} - 1}.$$

Let $\mathcal{D}_{m+1}$ be the decomposition of V with respect to U defined as follows. For each $W \in \mathcal{D}'$, replace W with $\mathcal{D}_W$, the decomposition defined in the previous paragraph. For the remaining $W \in \mathcal{D}_m \setminus \mathcal{D}'$, replace W with an arbitrary instance of Example 5.8. This refinement satisfies $\text{codim } \mathcal{D}_{m+1} = \text{codim } \mathcal{D}_m + \text{codim } U$.

Now we compute the energy increment

$$\begin{aligned} \mathcal{E}(\mathcal{P}(\mathcal{D}_{m+1})) - \mathcal{E}(\mathcal{P}(\mathcal{D}_m)) &= \mathbb{E}_{W \in \mathcal{D}_m} [\mathcal{E}(\mathcal{P}(\mathcal{D}_{m+1})|_{W \setminus U}) - \mathcal{E}(\mathcal{P}(\mathcal{D}_m)|_{W \setminus U})] \\ &\geq \frac{|\mathcal{D}'|}{|\mathcal{D}_m|} \mathbb{E}_{W \in \mathcal{D}'} [\mathcal{E}(\mathcal{P}(\mathcal{D}_{m+1})|_{W \setminus U}) - \mathcal{E}(\mathcal{P}(\mathcal{D}_m)|_{W \setminus U})] \\ &\geq \epsilon^3 q^{-\text{codim } U}. \end{aligned}$$

Since $0 \leq \mathcal{E}(\mathcal{P}(\mathcal{D})) \leq k$ for any decomposition $\mathcal{D}$, this process halts after $M \leq q^{\text{codim } U} k \epsilon^{-3}$ iterations, giving the desired result. $\square$

Given a decomposition $\mathcal{D}$ of V with respect to U , define

$$V(\mathcal{D}) := \bigcap_{W \in \mathcal{D}} W.$$

Note that this satisfies $\text{codim } V(\mathcal{D}) \leq |\mathcal{D}| \cdot \text{codim } \mathcal{D} = \text{codim } \mathcal{D} \cdot q^{\text{codim } \mathcal{D}}$. Furthermore, note that if $\mathcal{D}$ is non-trivial, i.e., $\text{codim } \mathcal{D} > 0$, then $V(\mathcal{D}) \leq U$.

Theorem 5.10 (Strong decomposition regularity). *For a finite dimensional $\mathbb{F}_q$ -vector space V , functions $f_1, \dots, f_k: V \rightarrow [-1, 1]$, a subspace $V_0 \leq V$, and a parameter $\epsilon > 0$, there exists $n_{\text{improved-reg}} = n_{\text{improved-reg}}(\epsilon, q, k, \text{codim } V_0)$ such that there exists a subspace $V_1 \leq V_0$ and a decomposition $\mathcal{D}$ of V with respect to V_1 such that*

- (1) $\text{codim } V_1, \text{codim } \mathcal{D} \leq n_{\text{improved-reg}}$;
- (2) $\mathcal{E}(\mathcal{P}(V_1|V \setminus V_1)) \leq \mathcal{E}(\mathcal{P}(\mathcal{D})) + \epsilon$;
- (3) for each $1 \leq i \leq k$, for all but an ϵ -fraction of $W \in \mathcal{D}$, for all $x \in W \setminus V_1$ the function $f_i|_{x+W \cap V_1}$ is ϵ -regular.

Proof. We can assume that $|V_0| \leq (\epsilon/4)|V|$.

Let $\mathcal{D}_0$ be a decomposition of V with respect to V_0 formed by applying Example 5.8. We iterate our weak arithmetic regularity lemma to produce a sequence of decompositions $\mathcal{D}_0, \mathcal{D}_1, \dots, \mathcal{D}_M$ such that $V(\mathcal{D}_{M-1}), \mathcal{D}_M$ has the desired properties. Let $V_m = V(\mathcal{D}_m)$ and let $\mathcal{D}_{m+1}$ be the result of applying Theorem 5.9 to V_m with parameter ϵ . (Note that we must have $\dim V \geq \text{codim } V_m \cdot q^{\text{codim } V_m} \cdot k \cdot \epsilon^{-3}$ to apply Theorem 5.9. However, if this ever fails the desired result is trivially true.)

Note that $V_0 \geq V_1 \geq \dots \geq V_M$ so $0 \leq \mathcal{E}(\mathcal{P}(V_0)) \leq \mathcal{E}(\mathcal{P}(V_1)) \leq \dots \leq \mathcal{E}(\mathcal{P}(V_M)) \leq k$. Therefore we can stop the iteration at $M \leq 2k\epsilon^{-1}$ such that $\mathcal{E}(\mathcal{P}(V_M)) \leq \mathcal{E}(\mathcal{P}(V_{M-1})) + \epsilon/2$.

Now since $|V_{M-1}| \leq (\epsilon/4)|V|$, note that $\mathcal{E}(\mathcal{P}(V_M|V \setminus V_{M-1})) \leq \mathcal{E}(\mathcal{P}(V_{M-1}|V \setminus V_{M-1})) + \epsilon$. Furthermore $\mathcal{P}(V_M|V \setminus V_{M-1}) \supseteq \mathcal{P}(\mathcal{D}_M) \supseteq \mathcal{P}(V_{M-1}|V \setminus V_{M-1})$. This implies property (2), completing the proof. $\square$

Proof of Proposition 3.2. We can assume that $|V_0| \leq (\epsilon/4)|V|$.

Apply Theorem 5.10 to $f_1, \dots, f_k: V \rightarrow [-1, 1]$ and $V_0 \leq V$ with parameter $\min(\epsilon^3/4, 1/(4k))$. Write $V_1 \leq V_0$ for the subspace produced and $\mathcal{D}$ for the decomposition produced. For a uniform random choice of $W \in \mathcal{D}$, write $V_2 = V_1 \cap W$ and pick U satisfying $U \oplus V_2 = W$ arbitrarily. We claim that with positive probability, this choice of U, V_2, V_1 satisfies the desired properties.

First note that with probability at least $1 - k(1/(4k)) = 3/4$ we have that for all $1 \leq i \leq k$ and for all $x \in W \setminus V_1$ we have that $f_i|_{x+V_2}$ is ϵ -regular.

Now by part (2) of Proposition 5.2, we write

$$\begin{aligned} \epsilon^3/4 &\geq \mathcal{E}(\mathcal{P}(\mathcal{D})) - \mathcal{E}(\mathcal{P}(V_1|V \setminus V_1)) \\ &= \sum_{i=1}^k \left\| (f_i)_{\mathcal{P}(V_1|V \setminus V_1)} - (f_i)_{\mathcal{P}(\mathcal{D})} \right\|_{L^2(V \setminus V_1)}^2 \\ &= \mathbb{E}_{x \in V \setminus V_1} \left[\sum_{i=1}^k \left(\mathbb{E}_{y \in W(x) \cap V_1} [f_i(x+y)] - \mathbb{E}_{y \in V_1} [f_i(x+y)] \right)^2 \right]. \end{aligned}$$

In the last line above we write $W(x) \in \mathcal{D}$ for the unique part of the decomposition containing x .

This implies that for a uniform random choice of $x \in V \setminus V_1$, with probability at least $1 - \epsilon/4$, this x satisfies

$$|\mathbb{E}_{y \in W(x) \cap V_1} [f_i(x+y)] - \mathbb{E}_{y \in V_1} [f_i(x+y)]| \leq \epsilon \quad \text{for all } 1 \leq i \leq k. \quad (\S)$$

Therefore choosing $W \in \mathcal{D}$ uniformly at random, the expected number of $x \in W$ failing to satisfy $(\S)$ is at most $|W \cap V_1| + (\epsilon/4)(|W| - |W \cap V_1|) \leq (\epsilon/2)|W|$. By Markov's inequality, with probability at least $1/2$, at most an ϵ -fraction of $x \in W$ fail to satisfy $(\S)$. Thus by the union bound, with positive probability this choice of U, V_2, V_1 satisfies all the desired properties. $\square$

6. EXTENSIONS

6.1. Infinite removal. First let us recall the infinite graph removal lemma [3].

Theorem 6.1. *For every (possibly infinite) set of graphs $\mathcal{H}$ and $\epsilon > 0$, there are h_0 and $\delta > 0$ such that the following holds. If a graph G has induced H -density at most δ for each $H \in \mathcal{H}$ on at most h_0 vertices, then G can be made induced $\mathcal{H}$ -free by changing at most an ϵ -fraction of the edges.*

This result is proved by a short modification of the usual proof of induced graph removal. By performing an analogous modification to the proof of induced arithmetic removal it is possible to deduce the following infinite arithmetic removal lemma.

Theorem 6.2. *Fix $\mathcal{H}$ a (possibly infinite) set of r -colored complexity 1 patterns over $\mathbb{F}_q$. For every $\epsilon > 0$ there are k_0 and $\delta > 0$ such that the following holds. Given a finite dimensional $\mathbb{F}_q$ -vector space V , if a coloring $\phi: V \rightarrow [r]$ has H -density at most δ for each $H \in \mathcal{H}$ with at most k_0 columns, then ϕ can be made $\mathcal{H}$ -free by changing the color of at most an ϵ -fraction of the elements.*

To prove this infinite arithmetic removal lemma we use the same recoloring algorithm as in the proof of Theorem 1.14 with one modification. We replace the regularity recoloring result (Proposition 3.3) with the following strong variant.

Proposition 6.3 (Strong variant of Proposition 3.3). *For a finite dimensional $\mathbb{F}_q$ -vector space V , a coloring $\phi: V \rightarrow [r]$, and parameters $0 < \epsilon \leq 1$ and $\epsilon' = (\epsilon'_0, \epsilon'_1, \dots)$ satisfying $\epsilon'_0 \geq \epsilon'_1 \geq \dots > 0$, there is some $n_{\text{reg}} = n_{\text{reg}}(\epsilon, \epsilon', q, r)$ such that there exist subspaces $V_2 \leq V_1 \leq V$, a complement U satisfying $U \oplus V_1 = V$, and a recoloring $\phi': V \rightarrow [r]$ that agrees with ϕ on all but at most $\epsilon|V|$ values satisfying:*

- (1) $1/\epsilon \leq \text{codim } V_1 \leq \text{codim } V_2 \leq n_{\text{reg}}$;
- (2) *if a color appears in some coset $x + V_1$ under ϕ' , then at least an $\epsilon/(2r)$ -fraction of $x + V_2$ is that color under ϕ ;*
- (3) *for each $x \in U \setminus \{0\}$, the original coloring $\phi|_{x+V_2}$ is $\epsilon'_{\text{codim } V_1}$ -regular.*

The only difference between Proposition 3.3 and Proposition 6.3 is that the degree of regularity in part (3) is now allowed to depend on $\text{codim } V_1$. The proof strategy given in Section 5 also proves this strong variant with an appropriate choice of parameters.

Now suppose we perform our recoloring procedure on an r -coloring $\phi: V \rightarrow [r]$. The output is a recoloring $\phi': V \setminus \{0\} \rightarrow [r]$. Suppose that along the way the application of regularity produces $V_2 \leq V_1 \leq V$ where $\text{codim } V_1 = d$. We now define a “summary” of the recoloring $\phi \rightsquigarrow \phi'$ as follows. Let $\Phi: \mathbb{F}_q^d \setminus \{0\} \rightarrow 2^{[r]} \setminus \{\emptyset\}$ be the map which sends a coset $x + V_1$ to the set of colors which appear under ϕ' in $x + V_1$. The recoloring $\phi'|_{V_1}$ is a canonical coloring and thus is determined by a function $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$. We will call (Φ, χ) the summary of the recoloring procedure $\phi \rightsquigarrow \phi'$. Note that the number of possible summaries is bounded in terms of d (and q, r).

Definition 6.4. Let $H = (A, \psi)$ be an r -colored pattern over $\mathbb{F}_q$ where A has dimensions $\ell \times k$. Given $\Phi: \mathbb{F}_q^d \setminus \{0\} \rightarrow 2^{[r]} \setminus \{\emptyset\}$ and $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$, say that (Φ, χ) *partially induces an H -instance* if there exists $\mathbf{x} \in (\mathbb{F}_q^d)^k$ such that $A\mathbf{x} = \mathbf{0}$ and setting $I \subset [k]$ to be the set of $i \in [k]$ such that $x_i = 0$ the following holds:

- (1) $\psi(i) \in \Phi(x_i)$ for each $i \in [k] \setminus I$;
- (2) for some $H' = (A', \psi|_I)$, a subpattern of H restricted to the variables I , the χ -canonical coloring of $\mathbb{F}_q^n \setminus \{0\}$ contains an H' -instance for all large enough n .

This definition is chosen so that if ϕ' contains an H -instance, then the summary (Φ, χ) partially induces an H -instance. A careful reading of the proof of the main theorem shows that it proves the following:

Proposition 6.5. *Let H be an r -colored complexity 1 pattern and fix $\epsilon > 0$. There exist $\epsilon_{\text{strong}} = \epsilon_{\text{strong}}(\epsilon, H) > 0$ and $\delta = \delta(\epsilon, H) > 0$ such that the following holds. Suppose the recoloring procedure of Proposition 6.3 applied to ϕ with parameters ϵ and $\epsilon' = (\epsilon_{\text{strong}}, \epsilon_{\text{strong}}, \dots)$ produces ϕ' . Let (Φ, χ) be the summary of this recoloring. If (Φ, χ) partially induces an H -instance, then the H -density in ϕ is more than δ .*

Now we give the idea of the proof of Theorem 6.2. Fix $\mathcal{H}$ a possibly infinite set of r -colored complexity 1 patterns over $\mathbb{F}_q$. We define finite subsets $\mathcal{H}_1 \subseteq \mathcal{H}_2 \subseteq \dots \subseteq \mathcal{H}$ as follows. Fix some d . Consider the finite set of all possible pairs $\Phi: \mathbb{F}_q^d \setminus \{0\} \rightarrow 2^{[r]} \setminus \{\emptyset\}$ and $\chi: \mathbb{F}_q \setminus \{0\} \rightarrow [r]$. For each (Φ, χ) , if there exists any $H \in \mathcal{H}$ such that (Φ, χ) contains an H -instance add one such H to $\mathcal{H}_d$. This ensures that $\mathcal{H}_d$ is finite.

Given $\epsilon > 0$, we set

$$\epsilon'_d = \min_{H \in \mathcal{H}_d} \epsilon_{\text{strong}}(\epsilon, H) \quad \text{and} \quad \delta_d = \min_{H \in \mathcal{H}_d} \delta(\epsilon, H).$$

We apply Proposition 6.3 with $\epsilon' = (\epsilon'_0, \epsilon'_1, \dots)$; all other parts of the recoloring produces remain the same. At the end we obtain a coloring that we claim is $\mathcal{H}$ -free. However, by the definitions we

made in the previous paragraph, it suffices to check that it is $\mathcal{H}_{\text{codim } V_1}$ -free. Taking $\delta = \delta_{n_{\text{reg}}(\epsilon, \epsilon', q, r)}$, Proposition 6.5 implies that if the recoloring is not $\mathcal{H}_{\text{codim } V_1}$ -free, then it has H -density more than $\delta_{\text{codim } V_1} \geq \delta$ for some $H \in \mathcal{H}_{\text{codim } V_1} \subseteq \mathcal{H}$. Finally taking k_0 to be the maximum number of columns in any of the patterns of $\mathcal{H}_{n_{\text{reg}}(\epsilon, \epsilon', q, r)}$ completes the argument.

6.2. Inhomogeneous patterns.

Definition 6.6. Let $H = (A, \psi)$ be an r -colored pattern over $\mathbb{F}_q$ where A is an $\ell \times k$ matrix. For a finite dimensional $\mathbb{F}_q$ -vector space V , a coloring $\phi: V \rightarrow [r]$, and an offset $\mathbf{b} \in V^\ell$, say that $\mathbf{x} \in V^k$ is an $(H, \mathbf{b})$ -instance if $A\mathbf{x} = \mathbf{b}$ and $\phi(x_i) = \psi(i)$.

Theorem 6.7. Fix a finite set $\mathcal{H}$ of r -colored complexity 1 patterns over $\mathbb{F}_q$. For every $\epsilon > 0$ there exists $\delta = \delta(\epsilon, \mathcal{H}) > 0$ such that the following holds. Given a finite dimensional $\mathbb{F}_q$ -vector space V , if a coloring $\phi: V \rightarrow [r]$, and offsets $\{\mathbf{b}_H\}_{H \in \mathcal{H}}$ (where for $H = (A, \psi) \in \mathcal{H}$ where A has ℓ_H rows we have $\mathbf{b}_H \in V^{\ell_H}$) satisfy that ϕ has $(H, \mathbf{b}_H)$ -density at most δ for every $H \in \mathcal{H}$, then there exists a recoloring $\phi': V \rightarrow [r]$ that is $(H, \mathbf{b}_H)$ -free for all $H \in \mathcal{H}$ and differs from ϕ on at most $\epsilon|V|$ elements. Here

$$B = \text{span}_{H \in \mathcal{H}, i \in [\ell_H]} (\mathbf{b}_H)_i.$$

We prove this by reduction to the main theorem. We have a subspace $B \leq V$ where B is of bounded size. Pick a complement $\tilde{V}$ satisfying $B \oplus \tilde{V} = V$. Now we convert $\phi: V \rightarrow [r]$ to $\tilde{\phi}: \tilde{V} \rightarrow [r]^B$ in the obvious manner: $\tilde{\phi}(x) := (\phi(x + u))_{u \in B}$. One can check that given a pattern $H = (A, \psi)$ where A is an $\ell \times k$ matrix and an offset $\mathbf{b} \in B^\ell$ we can create a set of $N = |B|^{k - \text{rank } A, k(|B|-1)}$ patterns $H_1, \dots, H_N$ so that every $(H, \mathbf{b})$ -instance in ϕ becomes an H_i -instance in $\tilde{\phi}$ for some $1 \leq i \leq N$. Then applying the main theorem to this finite set of homogeneous patterns, we deduce the desired result.

REFERENCES

- [1] Noga Alon, Richard A. Duke, Hanno Lefmann, Vojtěch Rödl, and Raphael Yuster, *The algorithmic aspects of the regularity lemma*, J. Algorithms **16** (1994), 80–109.
- [2] Noga Alon, Eldar Fischer, Michael Krivelevich, and Mario Szegedy, *Efficient testing of large graphs*, Combinatorica **20** (2000), 451–476.
- [3] Noga Alon and Asaf Shapira, *A characterization of the (natural) graph properties testable with one-sided error*, SIAM J. Comput. **37** (2008), 1703–1727.
- [4] Arnab Bhattacharyya, Eldar Fischer, Hamed Hatami, Pooya Hatami, and Shachar Lovett, *Every locally characterized affine-invariant property is testable*, STOC'13—Proceedings of the 2013 ACM Symposium on Theory of Computing, ACM, New York, 2013, pp. 429–435.
- [5] Arnab Bhattacharyya, Elena Grigorescu, and Asaf Shapira, *A unified framework for testing linear-invariant properties*, Random Structures Algorithms **46** (2015), 232–260.
- [6] David Conlon and Jacob Fox, *Bounds for graph regularity and removal lemmas*, Geom. Funct. Anal. **22** (2012), 1191–1256.
- [7] David Conlon and Jacob Fox, *Graph removal lemmas*, Surveys in combinatorics 2013, London Math. Soc. Lecture Note Ser., vol. 409, Cambridge Univ. Press, Cambridge, 2013, pp. 1–49.
- [8] Walter Deuber, *Partition theorems for abelian groups*, J. Combinatorial Theory Ser. A **19** (1975), 95–108.
- [9] Zoltán Füredi, *Extremal hypergraphs and combinatorial geometry*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), Birkhäuser, Basel, 1995, pp. 1343–1352.
- [10] Oded Goldreich, Shafi Goldwasser, and Dana Ron, *Property testing and its connection to learning and approximation*, J. ACM **45** (1998), 653–750.
- [11] W. Timothy Gowers and Julia Wolf, *Linear forms and higher-degree uniformity for functions on $\mathbb{F}_p^n$* , Geom. Funct. Anal. **21** (2011), 36–69.
- [12] W. Timothy Gowers and Julia Wolf, *Linear forms and quadratic uniformity for functions on $\mathbb{F}_p^n$* , Mathematika **57** (2011), 215–237.
- [13] Ronald L. Graham and Bruce L. Rothschild, *Ramsey's theorem for n -parameter sets*, Trans. Amer. Math. Soc. **159** (1971), 257–292.

- [14] Ben Green, *A Szemerédi-type regularity lemma in abelian groups, with applications*, Geom. Funct. Anal. **15** (2005), 340–376.
- [15] Ben Green and Tom Sanders, *Fourier uniformity on subspaces*, arXiv:1510.08739.
- [16] Ben Green and Terence Tao, *Linear equations in primes*, Ann. of Math. (2) **171** (2010), 1753–1850.
- [17] Hamed Hatami, Pooya Hatami, and Shachar Lovett, *General systems of linear forms: equidistribution and true complexity*, Adv. Math. **292** (2016), 446–477.
- [18] Daniel Král, Oriol Serra, and Lluís Vena, *A combinatorial proof of the removal lemma for groups*, J. Combin. Theory Ser. A **116** (2009), 971–978.
- [19] Daniel Král, Oriol Serra, and Lluís Vena, *A removal lemma for systems of linear equations over finite fields*, Israel J. Math. **187** (2012), 193–207.
- [20] Freddie Manners, *Good bounds in certain systems of true complexity one*, Discrete Anal. (2018), Paper No. 21, 40.
- [21] Vojtěch Rödl and Mathias Schacht, *Property testing in hypergraphs and the removal lemma*, STOC’07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing, ACM, New York, 2007, pp. 488–495.
- [22] Ronitt Rubinfeld and Madhu Sudan, *Robust characterizations of polynomials with applications to program testing*, SIAM J. Comput. **25** (1996), 252–271.
- [23] Imre Z. Ruzsa and Endre Szemerédi, *Triple systems with no six points carrying three triangles*, Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II, Colloq. Math. Soc. János Bolyai, vol. 18, North-Holland, Amsterdam-New York, 1978, pp. 939–945.
- [24] Asaf Shapira, *A proof of Green’s conjecture regarding the removal properties of sets of linear equations*, J. Lond. Math. Soc. (2) **81** (2010), 355–373.
- [25] Saharon Shelah, *Primitive recursive bounds for van der Waerden numbers*, J. Amer. Math. Soc. **1** (1988), 683–697.
- [26] Joel H. Spencer, *Ramsey’s theorem for spaces*, Trans. Amer. Math. Soc. **249** (1979), 363–371.

STANFORD UNIVERSITY, STANFORD, CA 94305, USA
E-mail address: jacobfox@stanford.edu

MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE, MA 02139, USA
E-mail address: {jtidor,yufeiz}@mit.edu