



Classical and Quantum Polynomial Freiman-Ruzsa Algorithms

Srinivasan Arunachalam  

IBM Quantum, Almaden Research Center, Almaden, CA, USA

Davi Castro-Silva  

Department of Computer Science and Technology, University of Cambridge, UK

Arkopal Dutt  

IBM Quantum, Cambridge, MA, USA

Tom Gur  

Department of Computer Science and Technology, University of Cambridge, UK

Abstract

We prove algorithmic versions of the polynomial Freiman-Ruzsa theorem of Gowers, Green, Manners, and Tao (Annals of Mathematics, 2025) in additive combinatorics. In particular, we give classical and quantum polynomial-time algorithms that, for $A \subseteq \mathbb{F}_2^n$ with doubling constant K , learn an explicit description of a subspace $V \subseteq \mathbb{F}_2^n$ of size $|V| \leq |A|$ such that A can be covered by K^C translates of V , for a universal constant $C > 1$.

2012 ACM Subject Classification Theory of computation \rightarrow Randomness, geometry and discrete structures

Keywords and phrases Additive combinatorics, sublinear algorithms

Digital Object Identifier 10.4230/LIPIcs.ITCS.2026.11

Related Version *Full Version:* <https://arxiv.org/abs/2509.02338> [2]

Funding *Davi Castro-Silva:* Supported by ESPRC Robust and Reliable Quantum Computing Grant. *Tom Gur:* Supported by ERC Starting Grant 101163189 and UKRI Future Leaders Fellowship MR/X023583/1.

Acknowledgements We want to thank Jop Briët for multiple discussions. We also want to thank the anonymous ITCS and QIP reviewers who helped improve an earlier version of the paper.

1 Introduction

The Freiman-Ruzsa theorem [16, 25] is a cornerstone of additive combinatorics with diverse applications to theoretical computer science (cf. [23]). Loosely speaking, the theorem shows that sets exhibiting approximate combinatorial subgroup behaviour must be algebraically structured. To make this precise, recall that a set A has doubling constant K if $|A + A| \leq K|A|$, where $A + A = \{a + a' ; a, a' \in A\}$. Note that A has doubling constant 1 if and only if it is a subgroup or a coset of a subgroup, and in turn, the doubling constant of A can be thought of as a combinatorial measure of the approximate subgroup behaviour of sets. In this paper, we focus on subsets of \mathbb{F}_2^n . In this setting, the Freiman-Ruzsa theorem states that sets $A \subseteq \mathbb{F}_2^n$ with $|A + A| \leq K|A|$ is covered by $\exp(K)$ translates of a subspace $V \subset \mathbb{F}_2^n$ of size $|V| \leq |A|$.

Marton conjectured that the aforementioned dependency in K can be improved to a polynomial, in what became widely known as the Polynomial Freiman-Ruzsa (PFR) conjecture. Over a decade later, Sanders proved a quasipolynomial Bogolyubov-Ruzsa theorem, which implies a version of the Freiman-Ruzsa theorem with quasipolynomial dependency on the doubling constant K . In a recent breakthrough, the PFR conjecture was proved by Gowers, Green, Manners, and Tao.



© Srinivasan Arunachalam, Davi Castro-Silva, Arkopal Dutt, and Tom Gur;
licensed under Creative Commons License CC-BY 4.0

17th Innovations in Theoretical Computer Science Conference (ITCS 2026).

Editor: Shubhangi Saraf; Article No. 11; pp. 11:1–11:8



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

► **Theorem 1** (Combinatorial PFR theorem [17]). *There exists a polynomial $P_0 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that the following holds. For any $n \geq 1$, if $A \subseteq \mathbb{F}_2^n$ satisfies $|A + A| \leq K|A|$, then A is covered by at most $P_0(K)$ translates of a subspace $V \subset \mathbb{F}_2^n$ of size $|V| \leq |A|$.*

A key reason for the importance of the PFR theorem in additive combinatorics is that it provides means to transition from a combinatorial notion of approximate subgroup structure, captured by constant doubling, to an algebraic notion, captured by a bounded subspace-cover, at only a polynomial cost.

1.1 Algorithmic PFR

The PFR theorem (and the closely-related quasi-polynomial Bogolyubov-Ruzsa theorem) also provide powerful tools that found diverse applications to theoretical computer science, including linearity testing of maps $f: \mathbb{F}_2^n \rightarrow \mathbb{F}_2^m$ [26], constructions of two-source extractors from affine extractors [28], communication complexity lower bounds [9], super-polynomial lower bounds on locally decodable codes [11], constructions of non-malleable codes [1], higher-order Goldreich-Levin theorems [27, 21, 12], sparsification algorithms for 1-in-3-SAT [8], quantum proofs for classical theorems [14], quantum and classical worst-case to average-case reductions [5, 6], tolerant quantum testing of stabilizer states [4], quantum learning of structured stabilizer decompositions [3], and beyond.

However, when considering applications of the PFR theorem and similar tools to theoretical computer science as above, it is often necessary or desirable to have an efficient algorithmic statement, where an explicit description of the subspace can be learned efficiently, as opposed to an existential combinatorial statement. Indeed, the naive algorithm that extracts the subspace runs in time $O(2^n)$. Fortunately, for the quasi-polynomial Bogolyubov-Ruzsa theorem, Ben-Sasson, Ron-Zewi, Tulsiani, and Wolf [10] showed an algorithmic version, which extracts a subspace in time $O(n^3 \log n)$.

The above motivates a natural question that arose after the resolution of the PFR conjecture [17]. Namely, now that we know that a subset A of constant doubling can be covered by at most a polynomial number of translates of a subspace $H \subset \mathbb{F}_2^n$ of size $|H| \leq |A|$, *can we learn the subspace H efficiently?* We refer to this as the “algorithmic PFR question.” Our main contribution answers this question by proving an algorithmic version of the polynomial Freiman-Ruzsa theorem, where the covering subspace can be learned explicitly in $\text{poly}(n)$ -time. In the following, a query to a set $A \subseteq \mathbb{F}_2^n$ is an evaluation of the characteristic function $\mathbf{1}_A(x)$ for a chosen $x \in \mathbb{F}_2^n$, and a random sample from A is a uniformly chosen element $a \in A$.

► **Theorem 2** (Algorithmic PFR). *There exists a polynomial $P_1 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that the following holds. Let $A \subseteq \mathbb{F}_2^n$ satisfy $|A + A| \leq K|A|$ for a doubling constant $K \geq 1$. There is an $\tilde{O}(n^4)$ -time randomized algorithm that uses $O(\log|A|)$ random samples and $\tilde{O}(\log^2|A|)$ queries to A which, with probability at least $2/3$, returns a subspace $V \leq \mathbb{F}_2^n$ of size at most $|A|$ such that A can be covered by $P_1(K)$ translates of V .*

We remark that the probability of success $2/3$ is arbitrarily chosen and can be amplified via standard error-reduction techniques. The algorithm returns the subspace V specified by an explicit basis. As typically viewed in additive combinatorics, the doubling constant K is a constant independent of n (as constant doubling implies structure), and in turn our asymptotic notation suppresses factors of K .

Quantum algorithms. En route to obtaining our algorithmic PFR theorem, we obtain quantum algorithms with query complexity (which we then dequantize; see the techniques section) and time complexity¹ that is a factor- n lesser than the classical algorithms. Namely, we show the following theorem.

► **Theorem 3** (Quantum Algorithmic PFR). *There exists a polynomial $P_1 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that the following holds. Let $A \subseteq \mathbb{F}_2^n$ satisfy $|A + A| \leq K|A|$ for a doubling constant $K \geq 1$. There is an $O(n^3)$ -time quantum algorithm that uses $O(\log|A|)$ random samples and quantum queries to A which, with probability at least $2/3$, returns a subspace $V \leq \mathbb{F}_2^n$ of size at most $|A|$ such that A can be covered by $P_1(K)$ translates of V .*

Optimality. We complement our algorithms by also showing that both of our classical and quantum algorithms are asymptotically optimal in terms of the *query complexity* dependence in n , up to a logarithmic factor. In particular, we show that $\Omega(n^2)$ queries to the set A are necessary for classical algorithms in order to output the subspace V . Similarly, we show that $\Omega(n)$ quantum queries to the set A are necessary. Our lower bounds are proven via a simple information-theoretic argument. Finally, we note that random samples are necessary to hit A in the case it is sparse, and information-theoretically, it is necessary to obtain $\Omega(\log|A|)$ samples from A in order to hit at least a basis for A . Further details and precise statements are included in the full version of the paper [2].

1.2 Homomorphism testing and structure-vs-randomness decomposition

A key reason for the power and centrality of the PFR theorem is its applications to deriving strong structural theorems regarding homomorphism testing and structure-vs-randomness decomposition. However, while the PFR theorem is known to be equivalent to the aforementioned structural theorems [18, 19], the equivalences are not trivially algorithmic. Nonetheless, we provide versions of these theorems that admit efficient algorithms.

The first theorem is concerned with local-to-global phenomena. Namely, it shows that if a map $f : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ satisfies a local affine-linear constraint with a significant probability, then it must be globally close, in fractional distance, to an affine-linear map, which can be efficiently learned. This statement is useful because it connects the PFR theorem to property testing and coding theory.

► **Theorem 4** (Homomorphism testing). *There exists a polynomial $P_2 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that the following holds. Let $f : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ satisfy*

$$\Pr_{x_1+x_2=x_3+x_4} [f(x_1) + f(x_2) = f(x_3) + f(x_4)] \geq 1/K.$$

Then, there is an affine-linear function $g : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ such that $f(x) = g(x)$ for at least $2^m/P_2(K)$ values of $x \in \mathbb{F}_2^m$. Furthermore, there is an $\tilde{O}((m+n)^3)$ -time randomized algorithm that, with probability at least $2/3$, learns a concise representation of g .

The second theorem also exhibits local-to-global structure, this time by showing that maps that are locally an approximate homomorphisms admit a structured decomposition, which can be efficiently learned.

¹ In the context of quantum algorithms, by time we mean the total number of single and two-qubit quantum gates used in the quantum algorithm.

► **Theorem 5** (Structured approximate homomorphism). *There exists a polynomial $P_3 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that the following holds. Let $f : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ satisfy*

$$|\{f(x) + f(y) - f(x + y) : x, y \in \mathbb{F}_2^m\}| \leq K.$$

Then f may be written as $g + h$, where $g : \mathbb{F}_2^m \rightarrow \mathbb{F}_2^n$ is linear and $|\text{Im}(h)| \leq P_3(K)$. Furthermore, there is an $\tilde{O}((m + n)^3)$ -time randomized algorithm that, with probability at least $2/3$, learns a concise representation of g .

We refer the reader to [2] for the proofs of these results, below, we give an overview of the technical details.

2 Technical overview

A natural approach for proving an algorithmic PFR theorem is to try to algorithmize each step in the proof of the PFR conjecture given in [17], which would in principle provide a result similar to Theorem 2. Unfortunately, the aforementioned proof relies heavily on entropic methods that are non-algorithmic by nature and it is unclear whether such machinery can be transformed into efficient algorithms. To overcome this barrier, we take a detour through quantum algorithms.

2.1 Stabilizer learning and the Gowers U^3 -norm of quantum states

In a recent breakthrough, Chen, Gong, Ye, and Zhang provided an efficient quantum procedure to learn the closest stabilizer state to a given quantum state [13]. The connection between this task and additive combinatorics was first noted by Arunachalam and Dutt [4], who used the following definition of the Gowers U^3 -norm of an arbitrary n -qubit quantum state $|\psi\rangle = \sum_{x \in \mathbb{F}_2^n} f(x) |x\rangle$ (where $(f(x))_x$ is an ℓ_2 -unit vector)

$$\|\psi\|_{U^3} = 2^{n/2} \left[\mathbb{E}_{x, h_1, h_2, h_3 \in \mathbb{F}_2^n} \prod_{\omega \in \mathbb{F}_2^3} C^{|\omega|} f(x + \omega \cdot h) \right]^{1/2^3}, \tag{1}$$

where $C^{|\omega|} f = \bar{f}$ if $\omega := \sum_{j \in [3]} \omega_j$ is odd and f otherwise. This is proportional to the Gowers U^3 -norm of the function f encoded in the amplitudes of the state $|\psi\rangle$, but normalized so that $\|\psi\|_{U^3} \leq \|f\|_{\ell_2} = 1$.

Arunachalam and Dutt showed that stabilizer states are the extremizers of the Gowers U^3 -norm over quantum states² and that a quantum state has non-negligible U^3 -norm if it correlates with a stabilizer state. Moreover, using the PFR theorem (Theorem 1), they showed the converse was true as well and obtained a *polynomial Gowers inverse theorem* for quantum states: the U^3 -norm of a quantum state and its maximal correlation with a stabilizer state are polynomially related (see also [7, 24]). As such, our high-level strategy is to use the stabilizer learning protocol in [13] to obtain a quantum algorithmic version of the polynomial Gowers inverse theorem, which is known to also be equivalent to the PFR theorem due to work of Green and Tao [20] and of Lovett [22]. We then arrive at an algorithmic result, albeit quantum, of a statement that is combinatorially equivalent to PFR.

² Although an explicit description of these extremizers appeared in the literature on higher-order Fourier analysis in 2012 [15], the link to quantum theory appears to have only been made only recently.

However, as discussed in the previous section, it is non-trivial to make such combinatorial equivalences algorithmic. In this paper, we algorithmize a proof of equivalence between these two results (inspired by the proofs of Green-Tao and Lovett), thus allowing us to employ the stabilizer learning algorithm [13] to obtain efficient *quantum* algorithms for the PFR theorem and the structural theorems stated in Section 1.2.

Classical algorithms via dequantization. After obtaining the quantum algorithms above, the last ingredient needed is a method to dequantize these algorithm so as to obtain efficient *classical* algorithms for PFR, as stated in Theorems 2, 4 and 5. Towards this end, we use in our arguments the machinery developed by Briët and Castro-Silva [12] to replace the quantum learning algorithm in [13] by a classical algorithm that emulates it. This allows us to obtain analogous algorithmic results in the classical setting, at the expense of quadratically worse query complexity than in the quantum setting. This quantum-to-classical blow-up is inherent, and indeed, we prove it is necessary and essentially optimal.

2.2 Proof sketch of main theorem

With the strategy above in mind, we proceed to give a high-level outline of the proofs of Theorems 2 and 3, which build on the combinatorial arguments of Green and Tao [20] and Lovett [22]. Our other algorithmic results follow via similar methods from these two theorems and the combinatorial arguments of Green and Ruzsa [19].

Suppose we have sample and query access to a set $A \subseteq \mathbb{F}_2^n$ such that $|A + A| \leq K|A|$ for a doubling constant $K \geq 1$. We shall first need to localize A inside the space \mathbb{F}_2^n , which in applications can be much larger than A itself (indeed, this is the reason why sample access to A is necessary). We do this by first sampling $O(\log|A|)$ uniformly random elements from A and taking their linear span, which we denote by $U \leq \mathbb{F}_2^n$. While U might not contain all of the original set, using the fact that A has bounded doubling we can show that their intersection $A' := A \cap U$ will likely comprise at least half of the points in A (for a careful choice of parameters). This allows us to shift attention from A , which can be arbitrarily sparse inside \mathbb{F}_2^n , to the localization A' , which occupies a positive fraction (at least a 2^{-2K} -fraction) of the vector space U , by the Freiman-Ruzsa theorem. Note that A' will also have small doubling constant:

$$|A' + A'| \leq |A + A| \leq K|A| \leq 2K|A'|.$$

We next wish to obtain a “dense model” of the localized set A' ; that is, a set $S \subseteq \mathbb{F}_2^m$ that is “additively equivalent” to A' , as captured by the notion of Freiman isomorphisms, but which has density at least $1/K^C$ (for a universal constant $C > 1$) inside its ambient space \mathbb{F}_2^m . In particular, we show that with high probability a uniformly random linear map $\pi : U \rightarrow \mathbb{F}_2^m$, for $m = \log|4A'| + 10$, will be a Freiman isomorphism (i.e., isomorphism of additive quadruples) from A' to $S := \pi(A')$. Informally, this means they have the same additive structure, and hence such a dense model can be efficiently obtained by sampling.

Equipped with the localized dense model, we proceed to learn the covering subspace. Denote by $f : S \rightarrow A'$ the inverse of π when restricted to S . By the definition of Freiman isomorphisms, we have that

$$\forall a, b, c, d \in S : a + b = c + d \implies f(a) + f(b) = f(c) + f(d).$$

From this approximate linearity condition of f on S , we can show that the function

$$g(x, y) = \mathbf{1}_S(x)(-1)^{f(x) \cdot y}$$

is approximately quadratic.

In particular, following the approach of Green [19] and Green-Tao [20], we show a slight strengthening of the homomorphism testing formulation of the PFR theorem, which we then proceed to algorithmize, relying on tools such as the quadratic Goldreich-Levin theorem. In more detail, we first prove that there exists a quadratic function $q : \mathbb{F}_2^{m+n} \rightarrow \mathbb{F}_2$ such that

$$\left| \mathbb{E}_{x \in \mathbb{F}_2^m, y \in \mathbb{F}_2^n} \mathbf{1}_S(x) (-1)^{f(x) \cdot y} (-1)^{q(x,y)} \right| \geq \frac{1}{P(K)}, \quad (2)$$

for some polynomial $P: \mathbb{R}_+ \rightarrow \mathbb{R}_+$.

Crucially, we can efficiently *learn* such a high-correlation quadratic function q . This is done relying on the stabilizer learning algorithm of Chen et al [13] in the quantum setting, or its dequantization by Briët and Castro-Silva [12] in the classical setting. In order to use those theorems, however, we need to be able to efficiently query the function $g(x, y)$. This requires making queries to the set $S = \pi(A')$, and inverting the linear map π restricted to A' . We show how this can be done using a $O(n^3)$ time pre-processing step, and an extra cost of $O(n^2)$ time for each query to g . The total time and query complexities of our algorithms follow from this step.

From Eq. (2) and algebraic manipulations, we conclude that the homogeneous bilinear form

$$B(x, y) = q(x, y) - q(x, 0) - q(0, y) + q(0, 0)$$

correlates well with $g(x, y)$. Since we obtained an explicit description of q , we can compute a matrix $M \in \mathbb{F}_2^{n \times m}$ such that $B(x, y) = y^T M x$. By a simple Fourier analytic argument, we can then conclude there is some $v \in \mathbb{F}_2^n$ such that

$$f(x) = Mx + v \quad \text{for at least } 2^m/P'(K) \text{ values } x \in S,$$

where $P': \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is another polynomial we obtain in our proof. This implies that the subspace

$$V = \{Mx : x \in \mathbb{F}_2^m\} \leq \mathbb{F}_2^n$$

satisfies

$$|A' \cap (v + V)| = |\text{Im}(f) \cap (v + V)| \geq 2^m/P'(K) \geq |A|/P'(K).$$

By an application of Ruzsa's covering lemma, we conclude that $P'(K)$ translates of V can cover A . By our choice of m and the fact that A' has bounded doubling, we deduce that V is covered by $2^m/|A| \leq 2^{12}K^4$ translates of any of its subspaces having size $|A|$, which concludes the high-level overview of the proof.

References

- 1 Divesh Aggarwal, Yevgeniy Dodis, and Shachar Lovett. Non-malleable codes from additive combinatorics. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, pages 774–783. Association for Computing Machinery, 2014. doi:10.1145/2591796.2591804.
- 2 Srinivasan Arunachalam, Davi Castro-Silva, Arkopal Dutt, and Tom Gur. Classical and Quantum Polynomial Freiman-Ruzsa Algorithms. *arXiv preprint arXiv:2509.02338*, 2025.
- 3 Srinivasan Arunachalam and Arkopal Dutt. Learning stabilizer structure of quantum states. *arXiv preprint arXiv:2510.05890*, 2025. doi:10.48550/arXiv.2510.05890.
- 4 Srinivasan Arunachalam and Arkopal Dutt. Polynomial-time tolerant testing stabilizer states. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, pages 1234–1241. Association for Computing Machinery, 2025. doi:10.1145/3717823.3718277.

- 5 Vahid R. Asadi, Alexander Golovnev, Tom Gur, and Igor Shinkar. Worst-case to average-case reductions via additive combinatorics. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, pages 1566–1574. Association for Computing Machinery, 2022. doi:10.1145/3519935.3520041.
- 6 Vahid R Asadi, Alexander Golovnev, Tom Gur, Igor Shinkar, and Sathyawageeswar Subramanian. Quantum worst-case to average-case reductions for all linear problems. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2535–2567. SIAM, 2024. doi:10.1137/1.9781611977912.90.
- 7 Zongbo Bao, Philippe van Dordrecht, and Jonas Helsen. Tolerant testing of stabilizer states with a polynomial gap via a generalized uncertainty relation. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, pages 1254–1262. Association for Computing Machinery, 2025. doi:10.1145/3717823.3718201.
- 8 Benjamin Bedert, Tamio-Vesa Nakajima, Karolina Okrasa, and Stanislav Živný. Strong sparsification for 1-in-3-SAT via Polynomial Freiman-Ruzsa. *arXiv preprint arXiv:2507.17878*, 2025.
- 9 Eli Ben-Sasson, Shachar Lovett, and Noga Ron-Zewi. An additive combinatorics approach relating rank to communication complexity. *Journal of the ACM (JACM)*, 61(4):1–18, 2014. doi:10.1145/2629598.
- 10 Eli Ben-Sasson, Noga Ron-Zewi, Madhur Tulsiani, and Julia Wolf. Sampling-based proofs of almost-periodicity results and algorithmic applications. In *International Colloquium on Automata, Languages, and Programming*, pages 955–966. Springer, 2014. doi:10.1007/978-3-662-43948-7_79.
- 11 Abhishek Bhowmick, Zeev Dvir, and Shachar Lovett. New bounds for matching vector families. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 823–832. Association for Computing Machinery, 2013. doi:10.1145/2488608.2488713.
- 12 Jop Briët and Davi Castro-Silva. A near-optimal Quadratic Goldreich-Levin algorithm. *arXiv preprint arXiv:2505.13134*, 2025. doi:10.48550/arXiv.2505.13134.
- 13 Sitan Chen, Weiyuan Gong, Qi Ye, and Zhihan Zhang. Stabilizer bootstrapping: A recipe for efficient agnostic tomography and magic estimation. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, pages 429–438. Association for Computing Machinery, 2025. doi:10.1145/3717823.3718191.
- 14 Andrew Drucker and Ronald de Wolf. Quantum proofs for classical theorems. *arXiv preprint arXiv:0910.3376*, 2009. arXiv:0910.3376.
- 15 Tanja Eisner and Terence Tao. Large values of the Gowers-Host-Kra seminorms. *Journal d'Analyse Mathématique*, 117:133–186, 2012. doi:10.1007/s11854-011-0033-6.
- 16 Gregory A Freiman. What is the structure of k if $k + k$ is small? *Number Theory*, page 109, 1987.
- 17 William Timothy Gowers, Ben Green, Freddie Manners, and Terence Tao. On a conjecture of Marton. *Annals of Mathematics*, 201(2):515–549, 2025.
- 18 Ben Green. Finite field models in additive combinatorics. *arXiv preprint math/0409420*, 2004.
- 19 Ben Green. Notes on the polynomial Freiman-Ruzsa conjecture. *preprint*, 2005.
- 20 Ben Green and Terence Tao. An equivalence between inverse sumset theorems and inverse conjectures for the u_3 norm. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 149, pages 1–19. Cambridge University Press, 2010.
- 21 Dain Kim, Anqi Li, and Jonathan Tidor. Cubic Goldreich-Levin. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4846–4892. SIAM, 2023. doi:10.1137/1.9781611977554.CH178.
- 22 Shachar Lovett. Equivalence of polynomial conjectures in additive combinatorics. *Combinatorica*, 32(5):607–618, 2012. doi:10.1007/S00493-012-2714-Z.
- 23 Shachar Lovett. An exposition of Sanders' quasi-polynomial Freiman-Ruzsa theorem. *Theory of Computing*, pages 1–14, 2015. doi:10.4086/TOC.GS.2015.006.

- 24 Saeed Mehraban and Mehrdad Tahmasbi. Improved bounds for testing low stabilizer complexity states. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC '25, pages 1222–1233. Association for Computing Machinery, 2025. doi:10.1145/3717823.3718228.
- 25 Imre Ruzsa. An analog of Freiman’s theorem in groups. *Astérisque*, 258(199):323–326, 1999.
- 26 Alex Samorodnitsky. Low-degree tests at large distances. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '07, pages 506–515. Association for Computing Machinery, 2007. doi:10.1145/1250790.1250864.
- 27 Madhur Tulsiani and Julia Wolf. Quadratic Goldreich–Levin theorems. *SIAM Journal on Computing*, 43(2):730–766, 2014. doi:10.1137/12086827X.
- 28 Noga Zewi and Eli Ben-Sasson. From affine to two-source extractors via approximate duality. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC '11, pages 177–186. Association for Computing Machinery, 2011. doi:10.1145/1993636.1993661.