A Note on Subspace Evasive Sets

Avraham Ben-Aroya [*]   Igor Shinkar[†]

Abstract

A subspace-evasive set over a field $\mathbb{F}$ is a subset of $\mathbb{F}^n$ that has small intersection with any low-dimensional affine subspace of $\mathbb{F}^n$. Interest in subspace evasive sets began in the work of Pudlák and Rödl (Quaderni di Matematica 2004). More recently, Guruswami (CCC 2011) showed that obtaining such sets over large fields can be used to construct capacity-achieving list-decodable codes with a constant list size.

Our results in this note are as follows:

- We provide a construction of subspace-evasive sets in $\mathbb{F}^n$ of size $|\mathbb{F}|^{(1-\epsilon)n}$ that intersect any k-dimensional affine subspace of $\mathbb{F}^n$ in at most $(2/\epsilon)^k$ points. This slightly improves a recent construction of Dvir and Lovett (STOC 2012), who constructed similar sets, but with a bound of $(k/\epsilon)^k$ on the size of the intersection. Besides having a smaller intersection, our construction is more elementary. The construction is explicit when k and ϵ are constants. This is sufficient in order to explicitly construct the aforementioned list-decodable codes.

- We show, using the Kővári-Sós-Turán Theorem, that for a certain range of the parameters the subspace-evasive sets obtained using the probabilistic method are optimal (up to a multiplicative constant factor).

Keywords: subspace-evasive sets, explicit constructions, Kővári-Sós-Turán Theorem

1 Introduction

A subspace evasive set over a field $\mathbb{F}$ is a subset of $\mathbb{F}^n$ that nearly avoids all affine subspaces of some specified dimension k. More formally, we say that a set $S \subseteq \mathbb{F}^n$ is (k, c)-*subspace evasive* if for every affine subspace $W \subseteq \mathbb{F}^n$ of dimension k it holds that $|S \cap W| \leq c$.

The interest in subspace evasive sets began in the work of Pudlák and Rödl [5], who showed how to use such sets to construct bipartite Ramsey graphs. Roughly speaking, they showed that a set S, which is $(n/2, c)$-subspace evasive over $\mathbb{F}_2$, can be used to construct a bipartite graph with $|S|$ vertices on each side, that does not contain bipartite cliques nor bipartite independent sets with c vertices on each side. In their setting, the underlying field size is small, the evaded subspaces are of large dimension, and the size of the evasive set itself can vary. Recently, Ben-Sasson and Zewi [7] continued this research direction by showing that such sets might by useful in constructing two-source and affine extractors.

The second motivation for studying subspace evasive sets came from the recent work of Guruswami [2] on capacity-achieving list-decodable codes. An error correcting code of length n is said to be (ρ, L) list-decodable if every Hamming ball of radius ρn contains at most L codewords. A

[*]Department of Computer Science, Weizmann Institute of Science, Rehovot, Israel. avraham.ben-aroya@weizmann.ac.il

[†]Department of Computer Science, Weizmann Institute of Science, Rehovot, Israel. igor.shinkar@weizmann.ac.il

probabilistic argument shows that there are codes of rate R, over large enough alphabets, which are $(1 - R - \epsilon, O(1/\epsilon))$ list-decodable. In [2], Guruswami showed that for a field $\mathbb{F}$ of size $|\mathbb{F}| = n^{O(1/\epsilon^2)}$ if $S \subseteq \mathbb{F}^n$ is a $(1/\epsilon, c)$-subspace evasive set of size $|S| \geq \mathbb{F}^{(1-\epsilon)n}$, then it is possible to construct a code over $\mathbb{F}$ which is $(1 - R - 2\epsilon, c)$ list-decodable. The main point here is that c can be made a constant, i.e., independent of $|\mathbb{F}|$ and n. A similar result was obtained by Guruswami and Wang [3].

Observe that the parameters of the subspace evasive sets needed for constructing Ramsey graphs are quite different than the ones required for the list-decoding application. Specifically, in the latter setting, we think of ϵ as a constant while n is a growing parameter. Hence, the field size $|\mathbb{F}|$ is polynomial n, the evaded subspaces are of constant dimension, and the size of the evasive set must be large. The subspace evasive sets we consider in this paper are in this regime of parameters.

Our results in this paper are twofold. First, we present a simple explicit construction of subspace evasive sets, which slightly improves upon a previous construction of Dvir and Lovett [6] in terms of the intersection size. We then attempt to understand the best possible subspace evasive sets, by giving a lower bound on the evasiveness of any large set.

1.1 Explicit subspace evasive sets

Before discussing our construction, we mention that a trivial probabilistic argument shows that, over any finite field $\mathbb{F}$, there are $(k, O(k/\epsilon))$-subspace evasive sets of size $|\mathbb{F}|^{(1-\epsilon)n}$.

In order to discuss explicit constructions, we have to properly define the term "explicit". Motivated by the application for list-decodable codes, our definition requires an efficient algorithm for outputting the i'th element of the set, as well for computing the intersection with a given subspace. The first algorithm is needed for efficient encoding, while the second is needed for efficient list-decoding.

Definition 1.1 *We say that a (k, c)-subspace evasive set $S \subseteq \mathbb{F}^n$ is* explicit *if the following conditions hold:*

1. *There is an algorithm that on input $i \in \{1, \ldots, |S|\}$ outputs the i'th element of S in time* $\mathrm{poly}(n, |\mathbb{F}|)$.

2. *There is an algorithm that when given a subspace $W \subseteq \mathbb{F}^n$ of dimension k (represented as a set of basis vectors) outputs the intersection $W \cap S$ in time* $\mathrm{poly}(c, n, |\mathbb{F}|)$.

Observe that we could have used a more stringent definition which requires the algorithms to run in time polynomial in $\log(|\mathbb{F}|)$. A set which is explicit in the above sense is useful when the underlying field is not too large. Also, note that since in this paper we are considering only subspaces of constant dimension, the representation of the input W (to the second algorithm) is immaterial.

Our first result says that there exists an explicit construction of a set $S \subseteq \mathbb{F}^n$ of size $|\mathbb{F}|^{(1-\epsilon)n}$ that is $\left(k, \left(\frac{2}{\epsilon}\right)^k\right)$-subspace evasive.

Theorem 1.2 *Let $\mathbb{F}$ be an arbitrary finite field. Then, for every constant $k \in \mathbb{N}$ and constant $\epsilon \in (0, 0.5)$, there exists an explicit set $S \subseteq \mathbb{F}^n$ of size $|\mathbb{F}|^{(1-\epsilon)n}$ which is $\left(k, \left(\frac{2}{\epsilon}\right)^k\right)$-subspace evasive.*[1]

We stress that in the setting considered by Guruswami [2] both k and ϵ are constants, and the field size is polynomial in n. Thus, using Theorem 1.2, we get explicit codes of length n and rate R which are $(1 - R - 2\epsilon, \left(\frac{2}{\epsilon}\right)^{1/\epsilon})$ list-decodable. Moreover, they can be encoded and list-decoded in time polynomial in n.

---

[1] The expression $\left(\frac{2}{\epsilon}\right)^k$ can be replaced by $\left(\frac{1+\delta}{\epsilon}\right)^k$, for any constant $\delta > 0$.

Comparison with the construction of Dvir and Lovett [6]. Recently, Dvir and Lovett gave an elegant construction of subspace evasive sets [6]. Their construction gives sets that are $\left(k, \left(\frac{k}{\epsilon}\right)^k\right)$-subspace evasive. Again, this gives explicit codes of rate R which are $(1 - R - 2\epsilon, \left(\frac{1}{\epsilon}\right)^{2/\epsilon})$ list-decodable.

Our construction improves upon [6] in three aspects:

- **Evasiveness.** Our construction improves upon [6] by a multiplicative factor of $(k/2)^k$. This also implies the slightly better list size of $\left(\frac{2}{\epsilon}\right)^{1/\epsilon}$ in the list-decoding application.

- **Elementariness.** Our construction is elementary and does not rely on any heavy machinery. The construction of [6] is algebraic and is based on Bézout's theorem.

- **Underlying field.** Our construction works over any field, while the construction of [6] requires that $|\mathbb{F}| - 1$ does not have many small divisors.

On the other hand, the main advantage of [6] over our construction is its efficiency. Specifically:

- In [6], the running time of the algorithm that computes the i'th element in the set is $\text{poly}(\frac{k}{\epsilon}, \log(|\mathbb{F}|), n)$ while ours runs in time $|\mathbb{F}|^{O(k^2/\epsilon)} + O(n \log(|\mathbb{F}|))$.

- The running time of the algorithm that computes the intersection with a subspace in [6] is $\text{poly}(\left(\frac{k}{\epsilon}\right)^k, \log(|\mathbb{F}|), n)$, while ours runs in time $|\mathbb{F}|^{O(k^2/\epsilon)} + O(n \cdot |\mathbb{F}|^k)$.

The last advantage allows the $(1 - R - 2\epsilon, \left(\frac{1}{\epsilon}\right)^{2/\epsilon})$ list-decodable codes constructed using the subspace evasive sets of [6] to have a quadratic time list-decoding algorithm (again, when ϵ is constant). It should be noted that since the evasiveness parameter in both constructions depends exponentially on k, the list-decoding algorithm for both codes must run in time exponential in $1/\epsilon$.

In terms of techniques, our construction is based on the same high-level approach of the construction of [6]. That is, we first construct a subspace evasive set over a small-sized domain, and then simply take products of this set. In [6] they use an algebraic approach to get a construction for the smaller domain. Our construction, however, is based on a simple derandomization argument through the method of conditional expectations (a similar argument appeared in [5]).

1.2 A lower bound on evasiveness

Our second result says that for $k \geq \left(\frac{1}{\epsilon}\right)^{\Omega(1)}$ there is no set $S \subseteq \mathbb{F}^n$ of size $|\mathbb{F}|^{(1-\epsilon)n}$ that is $\left(k, o\left(\frac{k}{\epsilon}\right)\right)$-subspace evasive. This means that a set $S \subseteq \mathbb{F}^n$ obtained using the probabilistic method is essentially optimal. Specifically we prove the following:

Theorem 1.3 *Let $\mathbb{F}$ be a finite field. Let $k, c \in \mathbb{N}$ and $\epsilon \in (0, 0.6)$ be such that $\left(\frac{5k}{6c}\right)^c \geq \frac{1}{\epsilon}$, and $k \geq 8c$, and let $n \in \mathbb{N}$ be such that $|\mathbb{F}|^n > 2^{k/\epsilon}$. Then, for any set $S \subseteq \mathbb{F}^n$ of size $|\mathbb{F}|^{(1-\epsilon)n}$, there exists a subspace $W \subseteq \mathbb{F}^n$ of dimension k that intersects S on at least $\left(\frac{1}{75}\right)^c \cdot \frac{k}{\epsilon}$ points.*

Roughly speaking, the proof of Theorem 1.3 uses the Kővári-Sós-Turán Theorem repeatedly, in order to find a large subset of S which has a "product-like" structure. Given such a subset, it is easy to find a subspace that intersects it on many points.

2 The Construction

In this section we prove Theorem 1.2. The proof relies on the following key lemmas. The first lemma states that a random subset of $\mathbb{F}^r$ is a subspace evasive set with high probability. (A similar lemma appeared in [2].)

Lemma 2.1 *Let $\mathbb{F}$ be a finite field, and let $r, k \in \mathbb{N}$ and $\epsilon > 0$ be such that $r \geq 4k/\epsilon$. Let T be a uniformly random subset of $\mathbb{F}^r$ of size $|\mathbb{F}|^{(1-\epsilon)r}$. Then, with probability at least $1 - k \cdot |\mathbb{F}|^{-2k/\epsilon}$ the set T is a $\left(t, \frac{2t}{\epsilon}\right)$-subspace evasive set for every $t \in \{1, \ldots, k\}$.*

The next lemma says that a there exists an efficient algorithm that computes a large subspace evasive set in $\mathbb{F}^r$ in time $|\mathbb{F}|^{O(rk)}$. This means that when the parameters r and k are constant, the running time is $\mathrm{poly}(|\mathbb{F}|)$.

Lemma 2.2 *Let $\mathbb{F}$ be a finite field, and let $r, k \in \mathbb{N}$ and $\epsilon > 0$ be such that $r \geq 4k/\epsilon$. Then, there exists an algorithm that runs in time $|\mathbb{F}|^{O(rk)}$ and outputs all elements of a set $T \subseteq \mathbb{F}^r$ of size $|\mathbb{F}|^{(1-\epsilon)r}$ which is $\left(t, \frac{2t}{\epsilon}\right)$-subspace evasive for every $t \in \{1, \ldots, k\}$.*

Given the set T from Lemma 2.2 we show in Lemma 2.3 an explicit construction (in the sense of Definition 1.1) of a large subspace evasive set $S \subseteq \mathbb{F}^n$. We do this by setting the parameter r (in the two foregoing lemmas) to be a constant that depends on k/ϵ, but is independent of n and $|\mathbb{F}|$.

Lemma 2.3 *Let $\mathbb{F}$ be a finite field, and let $n, r, k \in \mathbb{N}$ and $\epsilon \in (0, 0.5)$ be such that $r = 4k/\epsilon$ and r divides n. Let $T \subseteq \mathbb{F}^r$ be the subspace evasive set from Lemma 2.2. Define a set $S \subseteq \mathbb{F}^n$ by*

$$S = \underbrace{T \times T \times \cdots \times T}_{n/r \text{ times}} \ ,$$

i.e., the set of all concatenations of the form $v_1 \ldots v_{n/r} \in \mathbb{F}^n$, with $v_i \in T$ for each $i \in [n/r]$. Then, the set $S \subseteq \mathbb{F}^n$ has size $|\mathbb{F}|^{(1-\epsilon)n}$, and it is $\left(k, \left(\frac{2}{\epsilon}\right)^k\right)$-subspace evasive. Moreover, for constant k and ϵ, the set S is explicit. Specifically,

1. *There exists an algorithm that on input $i \in \{1, \ldots, |S|\}$ outputs the i'th element of S in time $|\mathbb{F}|^{O(rk)} + O(n \log(|\mathbb{F}|))$.*

2. *There exists an algorithm that when given a subspace $W \subseteq \mathbb{F}^n$ of dimension k, outputs the intersection $W \cap S$ in time $|\mathbb{F}|^{O(rk)} + O(n|\mathbb{F}|^k)$.*

Remark 2.1 *If we are allowed to preform a preprocessing phase prior to running the algorithms that compute the intersection and the i'th element, we can reduce the running time of both algorithms. Specifically, using a preprocessing phase that runs in time $|\mathbb{F}|^{O(rk)}$ and keeps $|\mathbb{F}|^{(1-\epsilon)r}$ vectors from $\mathbb{F}^r$, we can remove the $|\mathbb{F}|^{O(rk)}$ term in the running time of the above algorithms.*

Clearly, Theorem 1.2 follows immediately from Lemma 2.3. We now turn to the proofs of the Lemmas 2.1-2.3.

2.1 Proving the key lemmas

We begin with Lemma 2.1 which relies on a standard application of the probabilistic method.

Proof of Lemma 2.1: Denote $q = |\mathbb{F}|$, and let T be a uniformly chosen set of size $q^{(1-\epsilon)r}$. For a fixed subspace W of dimension $t \leq k$ and a fixed set $A \subseteq W$ of size $\lceil 2t/\epsilon \rceil$, the probability that $A \subseteq T$ is

$$\Pr[A \subseteq W] = \frac{|T|}{q^r} \cdot \frac{|T|-1}{q^r-1} \cdot \ldots \cdot \frac{|T|-|A|}{q^r-|A|} < \left(\frac{|T|}{q^r}\right)^{|A|} \leq q^{-r\epsilon \cdot \frac{2t}{\epsilon}} = q^{-2rt}.$$

Therefore, using the union bound over all choices of $t \leq k$, the subspace W, and the set $A \subseteq W$, we get that the probability T is $(t, \frac{2t}{\epsilon})$-subspace evasive for every $t \leq k$ is at least

$$1 - \sum_{t=1}^{k} q^{rt} \cdot q^{2t^2/\epsilon} \cdot q^{-2rt} \geq 1 - \sum_{t=1}^{k} q^{2t^2/\epsilon} \cdot q^{-4kt/\epsilon} \geq 1 - k \cdot q^{-2k/\epsilon}, \tag{1}$$

where the first inequality uses the assumption that $r \geq 4k/\epsilon$. ∎

In order to prove Lemma 2.2 we derandomize Lemma 2.1 using the method of conditional expectations. (A similar argument appeared in [5].)

Proof of Lemma 2.2: Consider a random set $T \subseteq \mathbb{F}^r$ of size $q^{(1-\epsilon)r}$, where $q = |\mathbb{F}|$. For every subspace $W \subseteq \mathbb{F}^r$ of dimension $t \leq k$ let X_W be the indicator random variable of the event $\{|T \cap W| \geq \frac{2t}{\epsilon}\}$. Looking at the proof of Lemma 2.1, it follows that for every subspace $W \subseteq \mathbb{F}^r$ of dimension $t \leq k$ we have

$$\mathbb{E}[X_W] = \Pr[X_W = 1] \leq q^{2t^2/\epsilon - 2rt}.$$

Hence, when summing over all subspaces $W \subseteq \mathbb{F}^r$ of dimension at most k, analogously to the derivation in Eq. (1) we get

$$\mathbb{E}\left[\sum_W X_W\right] = \sum_W \Pr[X_W = 1] \leq k \cdot q^{-2k/\epsilon} < 1,$$

where the expectation is over the random set $T = \{v_1, \ldots, v_M\}$ of size $M = q^{(1-\epsilon)r}$.

Now, instead of choosing T at random, we use the method of conditional expectations. That is, the elements of the set $T = \{v_1, \ldots, v_M\}$, are chosen one by one, so that in each step the expectation of the sum $\sum_W X_W$ is minimized, conditioned on the previously chosen elements. Specifically, the algorithm works as follows:

1. Let $T = \emptyset$, and let $M = q^{(1-\epsilon)r}$.

2. For $i = 1, \ldots, M$ compute the expectation

$$\mathbb{E}\left[\sum_W X_W \,\middle|\, v_1, \ldots, v_{i-1}, v_i = x\right] \tag{2}$$

for each $x \in \mathbb{F}^r \setminus T$, and let v_i be an element x that minimizes this expectation.

3. Output the set $T = \{v_1, \ldots, v_M\}$.

Clearly, the output of the algorithm is a set T such that $\sum_W X_W < 1$, and since $\{X_W\}$ are indicator variables, it follows that $X_W = 0$ for all subspaces $W \subseteq \mathbb{F}^r$ of dimension at most k.

In order to show that the total running time is $q^{O(rk)}$, it is enough to show how the expectation in Eq. (2) can be computed in time $q^{O(rk)}$. Using linearity of expectation the expression in Eq. (2) can be written as

$$\sum_W \mathbb{E}\left[X_W | v_1, \ldots, v_{i-1}, v_i = x\right],$$

where the sum is over all subspaces W of dimension at most k. Since the number of terms in the sum is upper bounded by $k \cdot q^{rk}$, it is enough to show that each term $\mathbb{E}\left[X_W | v_1, \ldots, v_{i-1}, v_i = x\right]$ can be computed in time $q^{O(r)}$. This is done in the following claim.

Claim 2.4 *Let $W \subseteq \mathbb{F}^r$ be a subspace of dimension $t \leq k$, and suppose that $v_1, \ldots, v_i$ are fixed, for some $i \leq |\mathbb{F}|^{(1-\epsilon)r}$. Then, the expectation $\mathbb{E}\left[X_W | v_1, \ldots, v_i\right]$ can be computed in time $q^{O(r)}$.*

Proof: For given $v_1, \ldots, v_i \in \mathbb{F}^r$ and $W \subseteq \mathbb{F}^r$, let ℓ be the number of elements $\{v_j : j = 1, \ldots, i\}$ that belong to W. Then, the expectation $\mathbb{E}\left[X_W | v_1, \ldots, v_i\right]$ is

$$\mathbb{E}\left[X_W | v_1, \ldots, v_i\right] = \frac{1}{\binom{q^r - i}{M - i}} \cdot \sum_{j = \lceil \frac{2t}{\epsilon} \rceil - \ell}^{q^t - \ell} \binom{q^t - \ell}{j}\binom{(q^r - q^t) - (i - \ell)}{M - i - j}.$$

Since each binomial coefficient $\binom{N}{k}$ can be computed in time $O(N^2)$ using dynamic programming (relying on the recurrence relation $\binom{N}{k} = \binom{N-1}{k} + \binom{N-1}{k-1}$), the required expression can be computed in time $q^{O(r)}$ by computing each term individually, and then summing all the terms. The claim follows. ■

This completes the proof of Lemma 2.2. ■

Before proving Lemma 2.3 we claim that the product of subspace evasive sets is also subspace evasive. The claim has appeared implicitly in [6]. The special case of $k = 1$ has also appeared in [1].

Claim 2.5 *Let $T_1 \subseteq \mathbb{F}^{r_1}$ and $T_2 \subseteq \mathbb{F}^{r_2}$ be two sets that are (t, c_t)-evasive for all $t \leq k$. Then, the set $T_1 \times T_2 \subseteq \mathbb{F}^{r_1 + r_2}$ is (t, c_t') subspace evasive for all $t \leq k$, where $c_t' = \max\{c_j \cdot c_{k-j} : j \leq t\}$. In particular, if $c_t = (2/\epsilon)^t$ for all $t \leq k$, then $c_t' = c_t$ for all $t \leq k$.*

For completeness we include a proof of Claim 2.5 in the appendix. We are now ready to prove Lemma 2.3.

Proof of Lemma 2.3: Let $T \subseteq \mathbb{F}^r$ be the set from Lemma 2.2. Then, T is (t, c_t)-subspace evasive with $c_t = (2/\epsilon)^t$ for every $t \leq k$. Let $S = T^{n/r} \subseteq \mathbb{F}^n$ be the n/r-product of T. By Lemma 2.2 the set T if of size $q^{(1-\epsilon)r}$, and hence the size of S is $(q^{(1-\epsilon)r})^{n/r} = q^{(1-\epsilon)n}$, where $q = |\mathbb{F}|$.

We show below that for every $i \in \mathbb{N}$ the set $T^i \subseteq \mathbb{F}^n$ is $\left(t, (2/\epsilon)^t\right)$-subspace evasive for every $t \leq k$. In particular, for $i = n/r$ this implies that S is $\left(k, (2/\epsilon)^k\right)$-subspace evasive.

The proof is by induction on $i \in \mathbb{N}$. The case $i = 1$ follows from Lemma 2.2. For the inductive step, suppose that the set T^i is (t, c_t)-subspace evasive for every $t \leq k$. We also know that T is (t, c_t)-subspace evasive for every $t \leq k$. Therefore, by Claim 2.5 the set $T^{i+1} = T^i \times T$ is also (t, c_t)-subspace evasive for every $t \leq k$.

6

We now turn to show the explicitness of S. This is done in Algorithms 1 and 2 using the algorithm that outputs the set T from Lemma 2.2.

Algorithm 1 computes the i'th element of S. To achieve this, it first expresses the integer i in base $M = q^{(1-\epsilon)r}$, namely $i = \sum_{j=0}^{n/r-1}(c_j - 1)M^j$ for some coefficients $c_j \in \{1, \ldots, M\}$. The i'th element of S is then obtained by concatenation of the elements of T corresponding to the indices given by the c_j's.

Algorithm 1 Computing the i'th element of S

1: Run the algorithm from Lemma 2.2 to obtain the set $T = \{v_1, \ldots, v_M\}$, where $M = q^{(1-\epsilon)r}$.
2: Write $i = \sum_{j=0}^{n/r-1}(c_j - 1)M^j$ for some coefficients $c_j \in \{1, \ldots, M\}$.
3: Output the concatenation $v_{c_0}v_{c_1}\cdots v_{c_{n/r-1}} \in T^{n/r}$.

Algorithm 2 computes the intersection of S with any subspace of constant dimension. That is, given a subspace $W \subseteq \mathbb{F}^n$ of dimension k the algorithm outputs $W \cap S$. It works by going over all elements $w \in W$, and checking whether each such w belongs to S. This, in turn, is done by decomposing w into n/r strings each of length r, and checking whether each substring belongs to T.

Algorithm 2 Computing $W \cap S$

1: Run the algorithm from Lemma 2.2 to obtain the set $T = \{v_1, \ldots, v_M\}$, where $M = q^{(1-\epsilon)r}$.
2: **for** $w \in W$ **do**
3:     Decompose w into n/r blocks of length r.
4:     **if** all blocks of w are in T **then**
5:         Output w.
6:     **end if**
7: **end for**

As mentioned in Remark 2.1 in both algorithms, computing the set T can be done in a preprocessing phase, which runs in time $q^{O(rk)}$ and requires $O(q^r)$ space. $\blacksquare$

3 A Lower Bound on Evasiveness

In this section we prove Theorem 1.3 which we now restate. We mention that we did not attempt to optimize any of the constants that appear in this theorem.

Theorem 1.3 *Let $\mathbb{F}$ be a finite field. Let $k, c \in \mathbb{N}$ and $\epsilon \in (0, 0.6)$ be such that $\left(\frac{5k}{6c}\right)^c \geq \frac{1}{\epsilon}$, and $k \geq 8c$, and let $n \in \mathbb{N}$ be such that $|\mathbb{F}|^n > 2^{k/\epsilon}$. Then, for any set $S \subseteq \mathbb{F}^n$ of size $|\mathbb{F}|^{(1-\epsilon)n}$, there exists a subspace $W \subseteq \mathbb{F}^n$ of dimension k that intersects S on at least $\left(\frac{1}{75}\right)^c \cdot \frac{k}{\epsilon}$ points.*

The proof of Theorem 1.3 relies on the following theorem of Kővári, Sós, and Turán [4], stating that any sufficiently dense bipartite graph contains a large bi-clique as a subgraph.

Theorem 3.1 (Kővári-Sós-Turán Theorem) *Let N, M, r, s be natural numbers that satisfy $N \geq r$ and $M \geq s \geq r$. Then, any bipartite graph $G = (U \cup V, E)$ with $|U| = N$, $|V| = M$ and $|E| > \sqrt[r]{s-1} \cdot NM^{1-1/r} + rM$ contains a copy of $K_{r,s}$ as a subgraph with r vertices in U and s vertices in V.*

The following corollary of Theorem 3.1 states that for any set $S \subseteq \mathbb{F}^m$ of cardinality $|\mathbb{F}|^{(1-\delta)m}$ we can find sets X, Y of large cardinality such that $X \times Y \subseteq S$.

Corollary 3.2 *Let $\mathbb{F}$ be a finite field, $\ell \geq 2$ and m be integers, and $\delta, \alpha \in (0,1)$ be such that $(\ell + 1 + 2\alpha)\delta < 1$ and $|\mathbb{F}|^m > 2^{\ell/\alpha\delta}$. Then, for every set $S \subseteq \mathbb{F}^m$ of size $|\mathbb{F}|^{(1-\delta)m}$ there exists $X \subseteq \mathbb{F}^{(1+\alpha)\delta m}$ of size ℓ and $Y \subseteq \mathbb{F}^{(1-(1+\alpha)\delta)m}$ of size $|\mathbb{F}|^{(1-(\ell+1+2\alpha)\delta)m}$ such that $X \times Y \subseteq S$.*

Proof: Define a bipartite graph $G = (\mathbb{F}^{(1+\alpha)\delta m} \cup \mathbb{F}^{(1-(1+\alpha)\delta)m}, E)$, where $(x,y) \in E$ if and only if $xy \in S$. Denoting the size of $\mathbb{F}$ by q, the number of edges in G is $|E| = q^{(1-\delta)m}$. Therefore, by Theorem 3.1 it contains a copy of $K_{r,s}$ as a subgraph with $r = \ell$ and $s = q^{(1-(\ell+1+2\alpha)\delta)m}$. Indeed, setting $N = q^{(1+\alpha)\delta m}$ and $M = q^{(1-(1+\alpha)\delta)m}$ we see that the condition of Theorem 3.1 holds, namely

$$
\begin{aligned}
\left(\frac{s-1}{M}\right)^{1/r} N + r &= \left(\frac{q^{(1-(\ell+1+2\alpha)\delta)m} - 1}{q^{(1-(1+\alpha)\delta)m}}\right)^{1/\ell} q^{(1+\alpha)\delta m} + \ell \\
&< q^{((-\ell-\alpha)\delta/\ell + (1+\alpha)\delta)m} + \ell \\
&= q^{(\alpha - \alpha/\ell)\delta m} + \ell \\
&= q^{\alpha\delta m} \cdot q^{\frac{-\alpha\delta}{\ell}m} + \ell \\
&< q^{\alpha\delta m} = \frac{|E|}{M}.
\end{aligned}
$$

Denoting by $X \subseteq \mathbb{F}^{(1+\alpha)\delta m}$ and $Y \subseteq \mathbb{F}^{(1-(1+\alpha)\delta)m}$ the vertices of this subgraph, we have $|X| = r = \ell$ and $|Y| = s = q^{(1-(\ell+1+2\alpha)\delta)m}$, as required. $\blacksquare$

We are ready to prove Theorem 1.3. The main idea is to repeatedly apply Corollary 3.2 to obtain a large subset of S which has a product structure.

Proof of Theorem 1.3: We prove the theorem by induction on c. Suppose first that $c = 1$, and let $S \subseteq \mathbb{F}^n$ be a set of size $q^{(1-\epsilon)n}$, where $q = |\mathbb{F}|$. Using Corollary 3.2 on the set S with $\ell = 1/25\epsilon$, $\delta = \epsilon$, and $\alpha = 0.1$ we get sets $X \subseteq \mathbb{F}^{1.1\epsilon n}, Y \subseteq \mathbb{F}^{(1-1.1\epsilon)n}$ with $X \times Y \subseteq S$. Their cardinalities are $|X| = 1/25\epsilon$ and $|Y| \geq q^{\frac{6n}{25}} > \frac{k}{3}$, where the inequality holds for all $k \leq n$. Note that we can apply Corollary 3.2 since $(\ell + 1 + 2\alpha)\delta = 1/25 + 1.2\epsilon < 1$, and by the assumption of the theorem we have $q^n > 2^{k/\epsilon} > 2^{6/5\epsilon^2} > 2^{\ell/\alpha\delta}$.

Let $Y' \subseteq Y$ be any subset of Y of size $\frac{k}{3}$, and let

$$ A = \{x0^{(1-2\epsilon)n} : x \in X\} \cup \{0^{2\epsilon n} y : y \in Y'\} \subseteq \mathbb{F}^n. $$

Using the assumption $k \geq \frac{6}{5\epsilon}$ we get that $|A| = \frac{1}{25\epsilon} + \frac{k}{3} \leq k$. Choosing $W \subseteq \mathbb{F}^n$ to be an arbitrary k dimensional subspace that contains A we get that $X \times Y' \subseteq W$. Hence, $|W \cap S| \geq |X \times Y'| = \frac{1}{25\epsilon} \cdot \frac{k}{3} > \frac{k}{75\epsilon}$, as required.

Next we prove the induction step, i.e., we prove the theorem for c, assuming it hold for all $c' < c$. Let $c_{\min}$ be the minimal c_0 such that $\left(\frac{5k}{6c_0}\right)^{c_0} \geq \frac{1}{\epsilon}$. If $c_{\min} < c$, then we can use the induction hypothesis on $c_{\min}$ to obtain the required subspace W. Therefore, we may assume $c = c_{\min} \geq 2$. Note that this implies that $\frac{5k}{6} < \frac{1}{\epsilon}$.

We use Corollary 3.2 on the set S, setting $\ell = \frac{k}{c} - 2$, $\delta = \epsilon$, and $\alpha = 0.5$ to obtain sets $X \subseteq \mathbb{F}^{1.5\epsilon n}, Y \subseteq \mathbb{F}^{(1-1.5\epsilon)n}$ such that $X \times Y \subseteq S$. Their sizes are $|X| = \frac{k}{c} - 2$ and $|Y| = q^{\left(1-\frac{k}{c}\epsilon\right)n}$. (Observe that this is possible since $(\ell + 1 + 2\alpha) \cdot \delta = \frac{k}{c}\epsilon < 1$, where we use the assumption that $c = c_{\min} \geq 2$ and $\frac{5k}{6} < \frac{1}{\epsilon}$. Also, we have $q^n > 2^{k/\epsilon} > 2^{\ell/\alpha\delta}$, as required by Corollary 3.2.)

Our next goal is to apply the induction hypothesis on Y. Since Y resides in $\mathbb{F}^{(1-1.5\epsilon)n}$ we define n' to be $(1-1.5\epsilon)n$. Set $\epsilon' = \frac{k/c-1.5}{1-1.5\epsilon}\epsilon$ so that $|Y| = q^{(1-\epsilon')n'}$. Next, we define $k' = k - \ell = \frac{c-1}{c}k + 2$ and apply the induction hypothesis on Y with k', ϵ' and $c-1$. We verify that the required conditions hold, namely that (1) $\left(\frac{k'}{2(c-1)}\right)^{c-1} > \frac{1}{\epsilon'}$, (2) $k' > 8(c-1)$, (3) $\epsilon' < 0.6$, and (4) $q^{n'} > 2^{k'/\epsilon'}$. Indeed, the first condition holds since

$$\left(\frac{k'}{2(c-1)}\right)^{c-1} > \left(\frac{k}{2c}\right)^{c-1} \geq \frac{1}{\epsilon} \cdot \frac{2c}{k} = \frac{1}{\epsilon'} \cdot \frac{k/c - 1.5}{1-1.5\epsilon} \cdot \frac{2c}{k} \geq \frac{1}{\epsilon'},$$

where the last inequality follows from the assumption that $k \geq 8c$. The second condition holds as $k' > \frac{c-1}{c} \cdot k \geq \frac{c-1}{c} \cdot 8c > 8(c-1)$. For the third condition we have $\epsilon' = \frac{k/c-1.5}{1-1.5\epsilon}\epsilon < \frac{0.6\epsilon - 1.5}{1-1.5\epsilon}\epsilon = \frac{0.6-1.5\epsilon}{1-1.5\epsilon} < 0.6$, where the first inequality uses the assumptions that $c \geq 2$ and $\frac{5k}{6} < \frac{1}{\epsilon}$. Finally, the last condition is easily verified using the facts that $k' < k$, $n' = (1-1.5\epsilon)n$ and $\epsilon' > \frac{\epsilon}{(1-1.5\epsilon)}$.

The induction hypothesis guarantees a subspace $W' \subseteq \mathbb{F}^{(1-1.5\epsilon)n}$ of dimension k' such that $|Y \cap W'| \geq \left(\frac{1}{75}\right)^{c-1} \cdot \frac{k'}{\epsilon'}$. As in the base case, let

$$A = \{x0^{(1-2\epsilon)n} : x \in X\} \cup \{0^{2\epsilon n}w' : w' \in W'\} \subseteq \mathbb{F}^n,$$

and let $W \subseteq \mathbb{F}^n$ be a subspace of dimension k containing A. (Such a subspace exists since $\dim(W') = k' = k - |X|$.) Then, for every $x \in X$ and $y \in Y \cap W'$ we have $xy \in W \cap S$, and therefore

$$
\begin{aligned}
|W \cap S| \geq |X| \cdot |Y \cap W'| &\geq \left(\frac{k}{c} - 2\right) \cdot \left(\left(\frac{1}{75}\right)^{c-1} \cdot \frac{k'}{\epsilon'}\right) \\
&\geq \left(\frac{k}{c} - 2\right) \cdot \left(\frac{1}{75}\right)^{c-1} \cdot \frac{c}{20k} \cdot \frac{k}{\epsilon} \\
&\geq \left(\frac{1}{75}\right)^c \cdot \frac{k}{\epsilon},
\end{aligned}
$$

where the second inequality uses $\frac{k'}{\epsilon'} \geq \frac{(c-1)(1-1.5\epsilon)}{(k-1.5c)} \cdot \frac{k}{\epsilon}$, which can be (crudely) lower bounded by $\frac{c}{20k} \cdot \frac{k}{\epsilon}$ for $\epsilon < 0.6$. The last inequality $\left(\frac{k}{c} - 2\right) \cdot \frac{c}{20k} \geq \frac{1}{75}$ follows from the assumption that $k \geq 8c$. This completes the induction step. ∎

References

[1] J. Bierbrauer. Large caps. *Journal of Geometry*, 76:16–51, 2003.

[2] V. Guruswami. Linear-algebraic list decoding of folded reed-solomon codes. In *Proceedings of the 26th IEEE Conference on Computational Complexity*, 2011.

[3] V. Guruswami and C. Wang. Optimal rate list decoding via derivative codes. In *RANDOM*, pages 593–604. Springer, 2011.

[4] T. Kővári, V. Sós, and P. Turán. On a problem of Zarankiewicz. *Colloquium Math*, 3:50–57, 1954.

[5] P. Pudlák and V. Rödl. Pseudorandom sets and explicit construction of Ramsey graphs. *Quaderni di Matematica*, 13:327–346, 2004.

[6] Z.Dvir and S. Lovett. Subspace evasive sets. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, 2012.

[7] N. Zewi and E. Ben-Sasson. From affine to two-source extractors via approximate duality. In *Proceedings of the 43rd annual ACM symposium on Theory of computing*, pages 177–186. ACM, 2011.

A   Proof of Claim 2.5

Claim 2.5   *Let $T_1 \subseteq \mathbb{F}^{r_1}$ and $T_2 \subseteq \mathbb{F}^{r_2}$ be two sets that are (t, c_t)-evasive for all $t \leq k$. Then, the set $T_1 \times T_2 \subseteq \mathbb{F}^{r_1 + r_2}$ is (t, c'_t) subspace evasive for all $t \leq k$, where $c'_t = \max\{c_j \cdot c_{t-j} : j \leq t\}$.*

Proof: Let $W \subseteq \mathbb{F}^{r_1 + r_2}$ be an affine subspace of dimension $t \leq k$. We shall show that

$$|W \cap (T_1 \times T_2)| \leq \max\{c_j \cdot c_{t-j} : j \leq t\}.$$

Let $W_1 = Proj_1(W)$ be the projection of W on the first r_1 coordinates, and denote its dimension by $\ell = \dim(W_1)$. For every $x \in W_1$, let $W_2(x) \overset{\text{def}}{=} \{y \in \mathbb{F}^{r_2} : xy \in W\}$.

It is easy to see that $\dim(W_2(x)) = t - \ell$ for every $x \in W_1$. Indeed, for every $x \in W_1$ consider the affine transformation $P_x : W \to \mathbb{F}^{r_1}$ defined as $P_x(w) = Proj_1(w) - x$. Then, P_x satisfies the following equalities: (1) $\dim(\text{Im}(P_x)) = \dim(Proj_1(W))$, and (2) $\dim(W_2(x)) = \dim(\text{Ker}(P_x))$. Using the formula $\dim(\text{Im}(P_x)) + \dim(\text{Ker}(P_x)) = \dim(W)$ we infer that $\dim(W_2(x)) = t - \ell$.

By the assumption that T_1 is a (ℓ, c_ℓ) subspace evasive set, it follows that $|W_1 \cap T_1| \leq c_\ell$. Analogously we have $|W_2(x) \cap T_2| \leq c_{t-\ell}$ for every $x \in W_1$. Therefore, for every $x \in W_1 \cap T_1$, there are at most $c_{t-\ell}$ possible y's in $W_2(x) \cap T_2$. We conclude that

$$|W \cap (T_1 \times T_2)| = |\{xy \in W : x \in T_1, y \in T_2\}| \leq c_\ell \cdot c_{t-\ell}.$$

The claim follows.   ∎