J Algebr Comb (2007) 25:101–110
DOI 10.1007/s10801-006-0025-4

On the VC-dimension of uniform hypergraphs

Dhruv Mubayi* · **Yi Zhao**†

Received: 2 December 2005 / Accepted: 31 May 2006 /
Published online: 11 July 2006

Abstract Let $\mathcal{F}$ be a k-uniform hypergraph on $[n]$ where $k-1$ is a power of some prime p and $n \geq n_0(k)$. Our main result says that if $|\mathcal{F}| > \binom{n}{k-1} - \log_p n + k!k^k$, then there exists $E_0 \in \mathcal{F}$ such that $\{E \cap E_0 : E \in \mathcal{F}\}$ contains all subsets of E_0. This improves a longstanding bound of $\binom{n}{k-1}$ due to Frankl and Pach [7].

Keywords Trace · Hypergraph · VC-dimension · Extremal problems

1. Introduction

Let G be a set system (or hypergraph) on X and S be a subset of X. The *trace* of G on S is defined as $G|_S = \{E \cap S : E \in G\}$. We treat $G|_S$ as a set and therefore omit multiplicity. We say that S is *shattered* by G if $G|_S = 2^S$, the set of all subsets of S. The Vapnik-Chervonenkis dimension (*VC dimension*) of G is the maximum size of a set shattered by G. Extremal problems on traces started from determining the maximum size of a set system on n vertices with VC dimension $k-1$ (equivalently, without a shattered k-set). Sauer [10], Perles and Shelah [11], and Vapnik and Chervonenkis [12] independently proved that this maximum is $\binom{n}{0} + \cdots + \binom{n}{k-1}$. This and other results on traces have found numerous applications in geometry and computational learning theory (see Füredi and Pach [9] and Section 7.4 of Babai and Frankl [3]).

*Research supported in part by NSF grants DMS-0400812 and an Alfred P. Sloan Research Fellowship.

†Research supported in part by NSA grant H98230-05-1-0079. Part of this research was done while working at University of Illinois at Chicago.

D. Mubayi
Department of Mathematics, Statistics, and Computer Science, University of Illinois, Chicago, IL 60607

Y. Zhao (✉)
Department of Mathematics and Statistics, Georgia State University, Atlanta, GA 30303

Given two set systems G and F, if there exists a set S such that $G|_S$ contains a copy of F as a subhypergraph, we say that G contains F as a trace. In this case we write $G \to F$ ($G \not\to F$ otherwise). Let $\binom{X}{r}$ denote the set of all r-subsets of X. We call G an *r-uniform* hypergraph (r-graph) on X if $G \subseteq \binom{X}{r}$ and call the members of G *edges*. We define $\mathrm{Tr}^r(n, F)$ as the maximum number of edges in an r-graph on $[n] = \{1, \dots, n\}$ not containing F as a trace. Frankl and Pach [7] considered the maximum size of uniform hypergraphs with fixed VC dimension. They showed that $\mathrm{Tr}^r(n, 2^{[k]}) \le \binom{n}{k-1}$ for $k \le r \le n$. They conjectured that $\mathrm{Tr}^k(n, 2^{[k]}) = \binom{n-1}{k-1}$ for sufficiently large n. Obviously if a k-graph G contains a shattered edge, then G contains two disjoint edges (since the empty set appears in the trace). Therefore the conjecture of Frankl and Pach, if true, generalizes the well-known Erdős-Ko-Rado Theorem [5]. However, Ahlswede and Khachatrian [1] disproved it by constructing a $G \subseteq \binom{[n]}{k}$ of size $\binom{n-1}{k-1} + \binom{n-4}{k-3}$ that contains no shattered k-set when $k \ge 3$ and $n \ge 2k$. Combining this with the upper bound in [7], for $k \ge 3$ and $n \ge 2k$,

$$\binom{n-1}{k-1} + \binom{n-4}{k-3} \le \mathrm{Tr}^k\big(n, 2^{[k]}\big) \le \binom{n}{k-1}. \tag{1}$$

Our main result improves the upper bound in (1) in the case that $k-1$ is a prime power and n is large.

Theorem 1. *Let p be a prime, t be a positive integer, $k = p^t + 1$, and $n \ge n_0(k)$. If $\mathcal{F}$ is a k-uniform hypergraph on $[n]$ with more than $\binom{n}{k-1} - \log_p n + k!k^k$ edges, then there is a k-set shattered by $\mathcal{F}$. In other words,*

$$\mathrm{Tr}^k\big(n, 2^{[k]}\big) \le \binom{n}{k-1} - \log_p n + k!k^k.$$

In addition, we find exponentially many k-graphs achieving the lower bound in (1).

Proposition 2. *Let $P(n, r)$ denote the number of non-isomorphic r-graphs on $[n]$. Then for $k \ge 3$, there are at least $P(n-4, k-1)/2$ non-isomorphic k-graphs $\mathcal{F}$ on $[n]$ such that $|\mathcal{F}| = \binom{n-1}{k-1} + \binom{n-4}{k-3}$ and $\mathcal{F} \not\to 2^{[k]}$.*

Note that the gap between the upper and lower bounds in (1) is $\binom{n-1}{k-2} - \binom{n-4}{k-3}$. Theorem 1 reduces this gap by essentially $\log n$ for certain values of k. Though this improvement is small, the value of Theorem 1 is perhaps mainly in its proof—a mixture of algebraic and combinatorial arguments. The main tool in proving $\mathrm{Tr}^k(n, 2^{[k]}) \le \binom{n}{k-1}$ in [7] is the so-called *higher-order inclusion matrix*, whose rows are labeled by edges of a hypergraph $\mathcal{F} \subseteq \binom{[n]}{k}$. It was shown that if $\mathcal{F}$ contains no shattered k-sets, then the rows of this matrix are linearly independent. Consequently $|\mathcal{F}|$, the number of the rows, equals to the rank of the matrix, which is at most $\binom{n}{k-1}$. The main idea in proving Theorem 1 is to enlarge the inclusion matrix of $\mathcal{F}$ by adding more rows such that the rows in the enlarged matrix are still linearly independent. The method of adding independent vectors (or functions) to a space has been used before, e.g., on the two-distance problem by Blokhuis [4] and a proof of the Ray-Chaudhuri–Wilson Theorem by Alon, Babai and Suzuki [2].

In order to prove Theorem 1, we also need more combinatorial tools. In particular, the sunflower lemma of Erdős and Rado [6], which is used to prove Lemma 3 below. Note that Lemma 3 and Theorem 4 together prove Theorem 1. Let $2^{[k]-} = 2^{[k]} \backslash \emptyset$.

Lemma 3. *For any $k \leq n$,*

$$\mathrm{Tr}^k\left(n, 2^{[k]}\right) \leq \mathrm{Tr}^k\left(n, 2^{[k]-}\right) + k!k^k.$$

Theorem 4. *Let p be a prime, t be a positive integer, and $k = p^t + 1$. Then $\mathrm{Tr}^k(n, 2^{[k]-}) \leq \binom{n}{k-1} - \log_p n$ for $n \geq n_0(k)$.*

In the next section we prove Proposition 2 and Lemma 3. We prove Theorem 4 in Section 3 and give concluding remarks in the last section.

2. Proofs of Proposition 2 and Lemma 3

Proof of Proposition 2: We construct $\mathcal{F} = \mathcal{F}_0 \cup \mathcal{F}_1 \cup \mathcal{F}_2$ such that $\mathcal{F}_0$ is the set of all k-sets containing 1 and 2, edges in $\mathcal{F}_1$ contain 1 but avoid 2, and edges in $\mathcal{F}_2$ contain 2 but avoid 1. If we let $G_i = \{E \backslash \{i\} : E \in \mathcal{F}_i\}$ denote the link graph of i in $\mathcal{F}_i$, then G_1 and G_2 are $(k-1)$-graphs on $V' = \{3, 4, \ldots, n\}$. Let G_1 and G_2 further satisfy the following conditions:

1. $G_1 \cup G_2 = \binom{V'}{k-1}$
2. $G_1 \cap G_2 = \{E \in \binom{V'}{k-1} : E \supseteq \{3, 4\}\}$
3. $G_1 \supseteq \{E \in \binom{V'}{k-1} : E \ni 3, E \not\ni 4\}$, $G_2 \supseteq \{E \in \binom{V'}{k-1} : E \ni 4, E \not\ni 3\}$.

It is easy to see that $|\mathcal{F}| = \binom{n-1}{k-1} + \binom{n-4}{k-3}$, since $|\mathcal{F}_0| = \binom{n-2}{k-2}$ and

$$|\mathcal{F}_1| + |\mathcal{F}_2| = |G_1| + |G_2| = |G_1 \cup G_2| + |G_1 \cap G_2| = \binom{n-2}{k-1} + \binom{n-4}{k-3}.$$

We claim that $\mathcal{F} \not\to 2^{[k]}$. Suppose to the contrary that some $E \in \binom{[n]}{k}$ is shattered. Then $E \in \mathcal{F}$. Note that every edge in $\mathcal{F}$ contains either 1 or 2. If $\{1, 2\} \subset E$, then $E \backslash \{1, 2\}$ is not contained in $\mathcal{F}|_E$. Without loss of generality, assume that $E \ni 1$ and $E \not\ni 2$. Since $E \backslash \{1\} \in G_1$ is contained in $\mathcal{F}|_E$, we have $(E \backslash \{1\}) \cup \{2\} \in \mathcal{F}$ and consequently $E \backslash \{1\} \in G_1 \cap G_2$. Therefore $E \supseteq \{3, 4\}$. In order to have $E \backslash \{1, 4\} \in \mathcal{F}|_E$, there must be one edge of G_2 containing 3 and not containing 4. But this is impossible because of the third condition on G_1 and G_2.

In the above construction, every $E \in \binom{V'}{k-1}$ with $E \not\ni 3$, $E \not\ni 4$ could be in either G_1 or G_2. These *undecided* edges form a complete $(k-1)$-graph K_{n-4}^{k-1} on $\{5, \ldots, n\}$. Recall that $P(n-4, k-1)$ is the number of non-isomorphic $(k-1)$-graphs on $n-4$ vertices, or the number of non-isomorphic 2-edge-colorings of K_{n-4}^{k-1}. We claim that the number of non-isomorphic $\mathcal{F}$ satisfying our construction is $P(n-4, k-1)/2$. To see this, let us consider vertex degrees in $\mathcal{F}$. Let $\deg(x)$ be the number of edges in $\mathcal{F}$ containing a vertex x. It is not hard to see that no matter what the undecided edges are, $\deg(1)$ and $\deg(2)$ are always greater than $\deg(3) = \deg(4)$, which is greater than $\deg(x)$ for all $x > 4$, and $\deg(x)$ is fixed for all $x > 4$. Therefore two constructions $\mathcal{F}$ and $\mathcal{F}'$ are isomorphic if and only if $\mathcal{F}|_{\{5,\ldots,n\}}$ and $\mathcal{F}'|_{\{5,\ldots,n\}}$ are isomorphic or one is the complement of the other (since the vertices 1 and 2 are identical). □

Note that the construction in [1] is isomorphic to the case when all undecided E are in G_1.

A *sunflower* (or *Δ-system*) with r petals and a core C is a collection of distinct sets $S_1, \ldots, S_r$ such that $S_i \cap S_j = C$ for all $i \neq j$. Erdős and Rado [6] proved the following simple but extremely useful and fundamental lemma.

Lemma 5 (Sunflower Lemma). *Let G be a k-graph with $|G| > k\,!(r-1)^k$. Then G contains a sunflower with r petals.* □

We call a set S almost-shattered by $\mathcal{F}$ if $\mathcal{F}|_S$ contains $2^S \backslash \emptyset$.

Proof of Lemma 3: Let $\mathcal{F}$ be a k-graph on $[n]$ with $|\mathcal{F}| > \mathrm{Tr}^k(n, 2^{[k]-}) + k!k^k$. We need to show that $\mathcal{F}$ contains a shattered set. Since $|\mathcal{F}| > \mathrm{Tr}^k(n, 2^{[k]-})$, we may find an almost-shattered k-set $E_1 \in \mathcal{F}$. Since $|\mathcal{F}\backslash\{E_1\}| > \mathrm{Tr}^k(n, 2^{[k]-})$, we may find an almost-shattered k-set $E_2 \in \mathcal{F}\backslash\{E_1\}$. Repeating this process, we find distinct almost-shattered sets $E_1, E_2, \ldots, E_{k!k^k} \in \mathcal{F}$. By the Sunflower Lemma, $\mathcal{F}' = \{E_1, \ldots, E_{k!k^k}\}$ contains a sunflower with $k+1$ petals. Let us simply denote it by $E_1, \ldots, E_{k+1}$ and $C = \cap_{i=1}^{k+1} E_i$. Since E_1 is almost-shattered by $\mathcal{F}$ and $E_1 \backslash C \neq \emptyset$, there is $E_0 \in \mathcal{F}$ such that $E_0 \cap E_1 = E_1 \backslash C$. Now $E_1 \cap E_0, E_2 \cap E_0, \ldots, E_{k+1} \cap E_0$ are pairwise disjoint. Since $|E_0| = k < k+1$, there exists $i \neq 1$ such that $E_i \cap E_0 = \emptyset$. This means that $\emptyset \in \mathcal{F}|_{E_i}$. Consequently E_i is shattered by $\mathcal{F}$. □

3. Proof of Theorem 4

3.1. Inclusion matrices and proof outline

The proof of Theorem 4 needs the concept of higher-order inclusion matrices. Let $\mathcal{F}$ be a set system on X. The *incidence matrix* $M(\mathcal{F}, \leq s)$ of $\mathcal{F}$ over $\binom{X}{\leq s}$ is the matrix whose rows (*incidence vectors*) are labeled by the edges of $\mathcal{F}$, columns are labeled by subsets of $[n]$ of size at most s, and entry (E, S), $E \in \mathcal{F}$, $|S| \leq s$, is 1 if $S \subseteq E$ and 0 otherwise. Throughout this paper, we fix $s = k-1$ and simply write $M(\mathcal{F})$ instead of $M(\mathcal{F}, \leq k-1)$. In particular, let

$$I(k) = M\left(\binom{[n]}{k}\right) = M\left(\binom{[n]}{k}, \leq k-1\right).$$

For each $E \subset [n]$, the incidence vector v_E is a $(0,1)$-vector of length $\binom{n}{0} + \cdots + \binom{n}{k-1}$, whose coordinates are labeled by all subsets of $[n]$ of size at most $k-1$. Note that v_E always has a 1 in the position corresponding to $\emptyset$. Let $e_i = v_{\{i\}}$ for each $i \in [n]$.

Let q be 0 or a prime number. As usual, $\mathbb{F}_q$ denotes a field of q elements when q is a prime. Let us define $\mathbb{F}_0$ to be $\mathbb{Q}$, the field of rational numbers. Given a hypergraph $\mathcal{F}$, a *weight* function of $\mathcal{F}$ over $\mathbb{F}_q$ is a function $\alpha : \mathcal{F} \to \mathbb{F}_q$. If $\alpha(E) = 0$ for all $E \in \mathcal{F}$, then we call α the zero function and write $\alpha \equiv 0$. We define

$$v(\mathcal{F}, \alpha) = \sum_{E \in \mathcal{F}} \alpha(E) v_E$$

and write $v(\mathcal{F}) = \sum_{E \in \mathcal{F}} v_E$. We say that $\mathcal{F}$ is *linearly independent* in characteristic q if the rows of $M(\mathcal{F})$ are linearly independent over $\mathbb{F}_q$, namely, $v(\mathcal{F}, \alpha) = 0$ (mod q) implies that $\alpha \equiv 0$.

Part 1 of Lemma 6 below is the key observation to the proof of the upper bound in (1). It implies that if $\mathcal{F} \subseteq \binom{[n]}{k}$ contains no shattered sets, then it is linearly independent in any characteristic. Our proof of Theorem 4 also needs Part 2. We call a set S *near-shattered* by $\mathcal{F}$ if $\mathcal{F}|_S$ contains $2^S \backslash (\{i\} \cup \emptyset)$ for some $i \in S$.

Lemma 6. *Let q be 0 or a prime number. Suppose that $\mathcal{F} \subseteq \binom{[n]}{k}$ and $\alpha : \mathcal{F} \to \mathbb{F}_q$ is a non-zero weight function. Define $d(S) = \sum_{S \subseteq E \in \mathcal{F}} \alpha(E)$ for every subset $S \subset [n]$. Fix $A \in \mathcal{F}$ with $\alpha(A) \neq 0$.*

1. *If $d(S) = 0$ mod q for all $S \subset A$, then A is shattered by $\mathcal{F}$.*
2. *Let $i \in A$. If $d(S) = 0$ mod q for all $S \subset A$ with $S \neq \emptyset$ and $S \neq \{i\}$, then A is near-shattered.*

Proof: Parts 1 and 2 have almost the same proofs. Since Part 1 was proved in [7] and [3] (Theorem 7.27), we only prove Part 2 here.

Since $\mathcal{F}$ is k-uniform, we have $d(A) = \alpha(A) \neq 0$. For $B \subseteq A$, we define $d(A, B) = \sum_{E \in \mathcal{F}, E \cap A = B} \alpha(E)$. The following equality can be considered as a variant of the Inclusion-Exclusion formula.

$$d(A, B) = \sum_{B \subseteq S \subseteq A} (-1)^{|S-B|} d(S). \tag{2}$$

In fact, because $d(B) = d(A, B) + \sum_{E \in \mathcal{F}, B \subset E \cap A} \alpha(E)$, (2) is equivalent to

$$\sum_{E \in \mathcal{F}, B \subset E \cap A} \alpha(E) + \sum_{B \subset S \subseteq A} (-1)^{|S-B|} d(S) = 0.$$

This holds because on the left side, each $\alpha(E)$ with $r = |E \cap A| - |B| > 0$ has coefficient $1 - \binom{r}{1} + \cdots + (-1)^r \binom{r}{r} = 0$.

Pick any $B \subset A$ with $B \neq \emptyset$ and $B \neq \{i\}$. We now show that there exists $E \in \mathcal{F}$ such that $E \cap A = B$. We use (2) and the assumption that $d(S) = 0$ mod q for all S with $B \subseteq S \subset A$ to derive

$$\sum_{E \in \mathcal{F}, E \cap A = B} \alpha(E) = d(A, B) = \sum_{B \subseteq S \subseteq A} (-1)^{|S-B|} d(S) = (-1)^{|A-B|} d(A) \neq 0 \mod q.$$

Hence the sum on the left side is not empty. □

By Lemma 6 Part 1, if $\mathcal{F}$ contains no shattered sets, then the rows of $M(\mathcal{F})$ are linearly independent (over $\mathbb{Q}$) and consequently $|\mathcal{F}| = \text{rank}(M)$. Clearly $\text{rank}(M) \leq \text{rank}(I(k))$. It is well-known that $\text{rank}_{\mathbb{Q}}(I(k)) = \binom{n}{k-1}$ (e.g., see [3] Section 7.3). This immediately gives $\text{Tr}^k(n, 2^{[k]}) \leq \binom{n}{k-1}$, the result of Frankl and Pach [7].

The proof of Theorem 4 proceeds as follows. Suppose that $\mathcal{F} \subseteq \binom{[n]}{k}$ satisfies $\mathcal{F} \not\to 2^{[k]-}$. Recall that $k = p^t + 1$ for some prime p and positive integer t. We will construct a matrix M' obtained from $M = M(\mathcal{F})$ by adding $\log_p n$ new rows. The new rows have the form $e_S = \sum_{i \in S} e_i$, for some set S of size $m = p^{t+1}$. In other words, a new row has entry 1 at m coordinates corresponding to m singletons and 0 otherwise (the entry at $\emptyset$ is 0 because $m = 0 \bmod p$). The main step is to show that these new rows lie in the row space of $I(k)$, and all the rows of M' are still linearly independent. Consequently,

$$|\mathcal{F}| + \log_p n = \text{rank}_{\mathbb{F}_p}(M') \leq \text{rank}_{\mathbb{F}_p}(I(k)) \leq \text{rank}_{\mathbb{Q}}(I(k)) = \binom{n}{k-1},$$

which implies that $|\mathcal{F}| \leq \binom{n}{k-1} - \log_p n$.

We now divide the main step into three lemmas, which we will prove in the next subsection.

Lemma 7. *Suppose that $k = p^t + 1$ and $m = p^{t+1}$ for prime p and $t > 0$. Then for every $S \in \binom{[n]}{m}$, e_S is in the row space of $I(k)$ over $\mathbb{F}_p$.*

Lemma 8 is the key to our proof. For $a, b \in [n]$, let $e_{a,-b} = e_a - e_b$. Thus $e_{a,-b}$ is the vector with a 1 in position $\{a\}$, a -1 in position $\{b\}$, and 0 everywhere else. Lemma 8 says that $e_{a,-b}$ is outside the row space of M for every $a \neq b$.

Lemma 8. *Let $k \geq 2$ and $n \geq n_0(k)$. Suppose that $\mathcal{F} \subseteq \binom{n}{k}$ contains no almost-shattered set,* i.e., *$\mathcal{F} \not\to 2^{[k]-}$. If $|\mathcal{F}| > \binom{n}{k-1} - \log_p n$, then for every two distinct $a, b \in [n]$, the set $\{v_E : E \in \mathcal{F}\} \cup \{e_{a,-b}\}$ is linearly independent in any characteristic.*

Lemma 9. *Given a prime p and $m \geq 1$, let $n \geq n_0(p, m)$ and $r = \log_p n$. Suppose that for every two distinct $a, b \in [n]$, the set $\{v_E : E \in \mathcal{F}\} \cup \{e_{a,-b}\}$ is linearly independent in characteristic p. Then there exist subsets $S_1, \ldots, S_r \in \binom{[n]}{m}$ such that the set $\{v_E : E \in \mathcal{F}\} \cup \{e_{S_1}, \ldots, e_{S_r}\}$ is linearly independent in characteristic p.*

3.2. Proof of lemmas

Given a hypergraph $\mathcal{F}$ on X and a subset $A \subseteq X$, we define the *degree* $\deg_{\mathcal{F}}(A)$ to be the number of edges in $\mathcal{F}$ containing A.

Proof of Lemma 7: Let $K = \binom{S}{k}$. It suffices to prove that $\sum_{E \in K} v_E = c \cdot e_S$ for some nonzero $c \in \mathbb{F}_p$. Equivalently, we need to show that for $T \subset S$, $\deg_K(T) = 0 \bmod p$ when $|T| \geq 2$ or $|T| = 0$, and $\deg(T) = c \neq 0 \bmod p$ when $|T| = 1$. Since K is a complete k-graph, $\deg_K(T) = \binom{m-|T|}{k-|T|}$. By a well-known result of Kummer, the binomial coefficient $\binom{a}{b}$ is divisible by a prime p if and only if, when writing a and b as two numbers in base p, $a = (a_j \dots a_1 a_0)_p$ and $b = (b_j \dots b_1 b_0)_p$, there exists $i \le j$, such that $b_i > a_i$. Since m is a power of p, for any $1 \le k \le m-1$, p divides $\binom{m}{k}$. Hence $\deg_K(\emptyset) = \binom{m}{k} = 0 \bmod p$. Now consider $|T| = s \ge 2$. Since $k = p^t + 1$, we know $k - s < p^t$ and thus write $k - s = (a_{t-1} \cdots a_0)_p$. Since $m = p^{t+1}$, we have $m - s = p^{t+1} - s = (p-1)p^t + k - s - 1$. We thus have $m - s = (p-1\, a_{t-1} \dots a_0)_p - 1$. Hence there exists $i \le t-1$ such that the value of $m - s$ at bit i is less than a_i and consequently $\binom{m-s}{k-s}$ is divisible by p. When $|T| = 1$, we have $m - 1 = p^{t+1} - 1$ and therefore $\binom{m-1}{k-1}$ is not divisible by p for any $1 \le k \le m-1$. □

Proof of Lemma 8: We prove the contrapositive of the claim: If $\mathcal{F} \not\to 2^{[k]-}$ and there exists a non-zero function $\alpha : \mathcal{F} \to \mathbb{F}_q$ such that $v(\mathcal{F}, \alpha) = e_{a,-b}$ for some $a, b \in [n]$ $(a \ne b)$, then $|\mathcal{F}| \le \binom{n}{k-1} - \log_p n$. We claim that it suffices to show that $\deg_{\mathcal{F}}(\{a\}) = O(n^{k-3})$. In fact, suppose $\deg_{\mathcal{F}}(\{a\}) \le c_k n^{k-3}$ for some constant c_k and $|\mathcal{F}| > \binom{n}{k-1} - \log_p n$. After we remove a and all the edges containing a, we obtain a k-graph $\tilde{\mathcal{F}} \subseteq \mathcal{F}$ with $n-1$ vertices satisfying

$$
\begin{aligned}
|\tilde{\mathcal{F}}| &> \binom{n}{k-1} - \log_p n - c_k n^{k-3} \\
&= \binom{n-1}{k-1} + \binom{n-1}{k-2} - \log_p n - c_k n^{k-3} \\
&\ge \binom{n-1}{k-1}
\end{aligned}
$$

where the last inequality holds because $\binom{n-1}{k-2} \ge \log_p n + c_k n^{k-3}$ for $n \ge n_0(k)$. But we showed that $\mathrm{Tr}^k(n, 2^{[k]}) \le \frac{n}{k-1}$ for any $k \le n$, therefore $\tilde{\mathcal{F}} \to 2^{[k]}$, a contradiction.

Suppose that $\sum_{E \in \mathcal{F}} \alpha(E) v_E = e_{a,-b}$. Let $\mathcal{F}' = \{E \in \mathcal{F} : \alpha(E) \ne 0\}$ and $V' = [n] \backslash \{a, b\}$. For a subset $A \subset [n]$, let $d(A) = \sum_{A \subseteq E \in \mathcal{F}'} \alpha(E) \bmod q$. Our assumption $v(\mathcal{F}, \alpha) = e_{a,-b}$ implies that $d(\{a\}) = 1$, $d(\{b\}) = -1$, and $d(A) = 0$ for every $A \ne \{a\}, \{b\}$ and $|A| \le k-1$. Applying Lemma 6 Part 1, we conclude that no $E \in \mathcal{F}'$ satisfies $E \subseteq V'$. In other words, every edge in $\mathcal{F}'$ contains either a or b. Next observe that if $\mathcal{F}'$ contains an edge E such that $a \in E$ and $b \notin E$, then $\mathcal{F}'$ also contains $(E \backslash \{a\}) \cup \{b\}$. Otherwise E is the only edge in $\mathcal{F}'$ containing $E \backslash \{a\}$ and consequently $d(E \backslash \{a\}) = \alpha(E) \ne 0$, a contradiction.

Let $G_a = \{E \backslash \{a\} : E \in \mathcal{F}', a \in E, b \notin E\}$ and define G_b similarly. By the previous observation, we have $G := G_a = G_b$. We then observe that $G \ne \emptyset$ otherwise every edge (of $\mathcal{F}'$) containing a also contains b, and consequently $1 = d(\{a\}) = d(\{a, b\}) = 0$.

Fix an edge $E_0 \in \mathcal{F}'$ containing a but not b. Applying Lemma 6 Part 2, we conclude that E_0 is near-shattered, i.e., all subsets of E_0 are in the trace $\mathcal{F}'|_{E_0}$ except for $\{a\}$ and $\emptyset$. If another edge $E \in \mathcal{F}$ satisfies $E \cap E_0 = \{a\}$, then E_0 becomes almost-shattered, contradicting the assumption that $F \not\to 2^{[k]-}$. We may therefore assume that every $E \in \mathcal{F}$ containing a also contains some other element of E_0. Below we show that there exists $H \subseteq G$ with at most $2k$ vertices and transversal number at least 2 (i.e., no element lies in all sets of H). Therefore every $E \in \mathcal{F}$ containing a has at least two vertices in H and consequently $\deg_{\mathcal{F}}(\{a\}) \leq \binom{2k}{2}\binom{n-3}{k-3} = O(n^{k-3})$.

Pick $A \in G_a$ (thus $|A| = k-1$). We claim that for every $S \subset A$, $|S| = k-2$, there exists $B \in G_a$ such that $A \cap B = S$. Suppose instead, that for some $S \in \binom{A}{k-2}$, no such B exists. In this case, $A \cup \{a\}$ and $S \cup \{a, b\}$ are the only possible edges in $\mathcal{F}'$ containing $S \cup \{a\}$. We thus have $S \cup \{a, b\} \in \mathcal{F}'$, otherwise $d(S \cup \{a\}) = \alpha(A \cup \{a\}) \neq 0$. Because $G_a = G_b$, no $B \in G_b$ satisfies $A \cap B = S$. We now have a contradiction since

$$\begin{aligned} d(S) &= \alpha(A \cup \{a\}) + \alpha(A \cup \{b\}) + \alpha(S \cup \{a, b\}) \\ &= d(A) + \alpha(S \cup \{a, b\}) = \alpha(S \cup \{a, b\}) \neq 0. \end{aligned}$$

Now, for every $S \in \binom{A}{k-2}$, we choose exactly one set $B = B(S) \in G_a$ such that $A \cap B = S$. Let $H = \{A\} \cup \{B(S) : S \in \binom{A}{k-2}\}$. Clearly H contains at most $2k$ vertices. It is easy to see that there is no $x \in \cap_{E \in H} E$. In fact, if such $x \in A$, then $B(A \setminus \{x\})$ misses x. If $x \notin A$, then A misses x. Therefore the transversal number of H is at least 2, and the proof is complete. □

Proof of Lemma 9: Let M be the inclusion matrix of $\mathcal{F}$. We sequentially add vectors $e_{S_1}, \ldots, e_{S_i}$ with $S_1, \ldots, S_i \in \binom{[n]}{m}$ to M such that $e_{S_1}, \ldots, e_{S_i}$ and the rows of M are linearly independent. We claim that this can be done as long as $i \leq \log_p n$. Suppose to the contrary, that there exists $i \leq \log_p n - 1$ such that we fail to add a new vector at step $i+1$. In other words, we have chosen $e_{S_1}, \ldots, e_{S_i}$ successfully, but for every $S \in \binom{[n]}{m} \setminus \{S_1, \ldots, S_i\}$, there exist a weight function α and $c_1, \ldots, c_i \in \mathbb{F}_p$ such that

$$e_S = v(\mathcal{F}, \alpha) + \sum_{j=1}^{i} c_j \, e_{S_j}. \tag{3}$$

We observe that for fixed $c_1, \ldots, c_i$, the set of m-sets satisfying (3) forms a partial Steiner system $PS(n, m, m-1)$ (an m-graph on $[n]$ such that each $(m-1)$-subset of $[n]$ is contained in at most one edge). In fact, if two m-sets S, S' with $|S \cap S'| = m-1$ both satisfy (3), with weight functions α_1 and α_2 respectively, then $v(\mathcal{F}, \alpha_1 - \alpha_2) = e_{a,-b}$, where $\{a\} = S \setminus S'$ and $\{b\} = S' \setminus S$. This is a contradiction to our assumption. Consequently for fixed $c_1, \ldots, c_i$, the number of m-sets satisfying (3) is at most $\binom{n}{m-1}/m$. As a result, the number of m-sets that cannot be chosen is at most $p^i \binom{n}{m-1}/m$. We thus obtain

$$\left| \binom{[n]}{m} \setminus \{S_1, \ldots, S_i\} \right| = \binom{n}{m} - i \leq p^i \frac{1}{m} \binom{n}{m-1},$$

which implies that

$$(n - m + 1) - \frac{im}{\binom{n}{m-1}} \leq p^i.$$

Since $i \leq \log_p n - 1$, we have $p^i \leq n/p$, and consequently $n - m + 1 - im/\binom{n}{m-1} \leq n/p$, which is impossible for fixed $p \geq 2$, m and sufficiently large n. □

4. Concluding remarks

We believe that the lower bound in (1) is correct, though verifying this for all k may be hard because Proposition 2 gives exponentially many extremal hypergraphs. In order to reduce the bound in Theorem 1, one probably wants to look for a better way to find independent vectors than the greedy algorithm we used in the proof of Lemma 9. It may not be very hard to check this for the $k = 3$ case, namely, to verify that $\mathrm{Tr}^3(n, 2^{[3]}) = \binom{n-1}{2} + 1$. Using more involved combinatorial arguments, instead of the Sunflower Lemma, we can prove that $\mathrm{Tr}^3(n, 2^{[3]}) \leq \binom{n}{2} - \log_2 n$.

Improving the upper bound further for other values of k will most likely need some new ideas. Our approach uses incidence vectors of a family of singletons. The following proposition shows that this approach requires $k - 1$ to be a prime power.

Proposition 10. *Let p be a prime and $k \geq 2$. Suppose that $\mathcal{F} \subseteq \binom{[n]}{k}$ and $\alpha : \mathcal{F} \to \mathbb{F}_p$ is a non-zero weight function. Define $d(S) = \sum_{S \subseteq E \in \mathcal{F}} \alpha(E)$ for every subset $S \subset [n]$. If there exists a vertex $x \in [n]$ such that $d(\{x\}) \neq 0$ and $d(S) = 0$ mod p for every $S \ni x$ with $2 \leq |S| \leq k - 1$, then $k - 1$ is a power of p.*

Proof: Let $2 \leq s \leq k - 1$. When we sum up $d(S)$ for all $S \ni x$ with $|S| = s$, we over-count $d(\{x\})$ by a factor of $\binom{k-1}{s-1}$. In other words,

$$d(\{x\}) = \frac{1}{\binom{k-1}{s-1}} \sum_{x \in S, |S| = s} d(S).$$

Since $d(\{x\}) \neq 0$ but $d(S) = 0$ mod p for all S in the right-hand side, it must be the case that p divides $\binom{k-1}{s-1}$. We thus conclude that p divides $\binom{k-1}{i}$ for *all* $1 \leq i \leq k - 1$. By the result of Kummer on binomial coefficients, this happens only if $k - 1$ is a power of p. □

Acknowledgements We thank the referees for their comments, which improved the presentation of the paper and in particular, shortened the proof of Lemma 9.

References

1. R. Ahlswede and L.H. Khachatrian, "Counterexample to the Frankl-Pach conjecture for uniform, dense families," *Combinatorica* **17**(2) (1997), 299–301.
2. N. Alon, L. Babai, and H. Suzuki, "Multilinear polynomials and Frankl–Ray-Chaudhuri–Wilson type intersection theorems," *J. Combin. Theory Ser. A* **58**(2) (1991), 165–180.
3. L. Babai and P. Frankl, *Linear Algebra Method in Combinatorics*, preliminary version 2, University of Chicago, 1992.
4. A. Blokhuis, "A new upper bound for the cardinality of 2-distance sets in Euclidean space," Convexity and graph theory (Jerusalem, 1981), North-Holland Math. Stud., 87, North-Holland, Amsterdam, 1984, pp. 65–66.

5. P. Erdős, C. Ko, and R. Rado, “Intersection theorems for systems of finite sets,” *Quart. J. Math. Oxford Ser.* **12**(2) (1961), 313–320.
6. P. Erdős and R. Rado, “Intersection theorems for systems of sets,” *J. London Math. Soc.* **35** (1960), 85–90.
7. P. Frankl and J. Pach, “On disjointly representable sets,” *Combinatorica* **4** (1984), 39–45.
8. K. Friedl and L. Rónyai, “Order shattering and Wilson’s theorem,” *Discrete Math.* **270**(1–3) (2003), 127–136.
9. Z. Füredi and J. Pach, “Traces of finite sets: extremal problems and geometric applications,” *Extremal Problems for Finite Sets* (Visegrád, 1991), Bolyai Soc. Math. Stud., 3, János Bolyai Math. Soc., Budapest, 1994, pp. 251–282.
10. N. Sauer, “On the density of families of sets,” *J. Combinatorial Theory Ser. A* **13** (1972), 145–147.
11. S. Shelah, “A combinatorial problem; stability and order for models and theories in infinitary languages,” *Pacifc J. Math.* **41** (1972), 247–261.
12. V.N. Vapnik and A. Ya Chervonenkis, “On the uniform convergence of relative frequencies of events to their probabilities,” *Theory Probab. Appl.* **16** (1971), 264–280.