Random multilinear maps and the Erdős box problem

David Conlon[*] Cosmin Pohoata[†] Dmitriy Zakharov[‡]

Abstract

By using random multilinear maps, we provide new lower bounds for the Erdős box problem, the problem of estimating the extremal number of the complete d-partite d-uniform hypergraph with two vertices in each part, thereby improving on work of Gunderson, Rödl and Sidorenko.

1 Introduction

Writing $K^{(d)}_{s_1,\dots,s_d}$ for the complete d-partite d-uniform hypergraph with parts of orders $s_1,\dots,s_d$, the extremal number $\mathrm{ex}_d(n, K^{(d)}_{s_1,\dots,s_d})$ is the maximum number of edges in a d-uniform hypergraph on n vertices containing no copy of $K^{(d)}_{s_1,\dots,s_d}$. Already for $d=2$, the problem of determining these extremal numbers is one of the most famous in combinatorics, known as the Zarankiewicz problem. The classic result on this problem, due to Kővári, Sós and Turán [12], says that

$$\mathrm{ex}_2(n, K_{s_1,s_2}) = O\left(n^{2-1/s_1}\right)$$

for all $s_1 \le s_2$. However, this upper bound has only been matched by a construction with $\Omega(n^{2-1/s_1})$ edges when $s_2 > (s_1-1)!$, a result which, in this concise form, is due to Alon, Kollár, Rónyai and Szabó [1, 11], but builds on a long history of earlier work on special cases (see, for example, the comprehensive survey [8]).

Generalizing the Kővári–Sós–Turán bound, Erdős [6] showed that

$$\mathrm{ex}_d(n, K^{(d)}_{s_1,\dots,s_d}) = O\left(n^{d-\frac{1}{s_1\cdots s_{d-1}}}\right) \tag{1}$$

for all $s_1 \le s_2 \le \ldots \le s_d$. An analogue of the Alon–Kollár–Rónyai–Szabó result, due to Ma, Yuan and Zhang [14], is also known in this context and says that (1) is tight up to the constant provided that s_d is sufficiently large in terms of $s_1,\dots,s_{d-1}$. The proof of this result is based on an application of the random algebraic method, introduced by Bukh [2] and further developed in [3] and [4].

Our concern then will be with determining the value of $\mathrm{ex}_d(n, K^{(d)}_{s_1,\dots,s_d})$ in the particular case when $s_1 = \cdots = s_d = 2$. In the literature, this problem, originating in the work of Erdős [6], is sometimes referred to as the *box problem*, owing to a simple reformulation in terms of finding the largest subset of the grid $\{1,2,\dots,n\}^d$ which does not contain the vertices of a d-dimensional box (see also [10] for a connection to a problem in analysis). By (1), we have

$$\mathrm{ex}_d(n, K^{(d)}_{2,\dots,2}) = O\left(n^{d-\frac{1}{2^{d-1}}}\right). \tag{2}$$

[*]Department of Mathematics, California Institute of Technology, USA. Email: dconlon@caltech.edu.
[†]Department of Mathematics, Yale University, USA. Email: andrei.pohoata@yale.edu.
[‡]Laboratory of Combinatorial and Geometric Structures, MIPT, Russia. Email: zakharov2k@gmail.com.

While in the case $d = 2$, it has long been known that $\mathrm{ex}_2(n, K_{2,2}) = \Theta(n^{3/2})$, with a matching construction due to Klein [5] even predating the Kővári–Sós–Turán bound, there has been very little success in finding constructions matching (2) for $d \geq 3$. Indeed, it is unclear whether they should even exist. For $d = 3$, the best available construction is due to Katz, Krop and Maggioni [10], who showed that $\mathrm{ex}_3(n, K^{(3)}_{2,2,2}) = \Omega(n^{8/3})$. For general d, there is a simple, but longstanding, lower bound

$$\mathrm{ex}_d(n, K^{(d)}_{2,\ldots,2}) = \Omega\left(n^{d-\frac{d}{2^d-1}}\right) \tag{3}$$

coming from an application of the probabilistic deletion method. Besides the Katz–Krop–Maggioni construction, the only improvement to this bound is an elegant construction of Gunderson, Rödl and Sidorenko [9] using random hyperplanes which applies for infinitely many values of d.

Theorem 1 (Gunderson–Rödl–Sidorenko) *For any $d \geq 2$, let $s = s(d)$ be the smallest positive integer s (if it exists) such that $(sd-1)/(2^d-1)$ is an integer. Then*

$$\mathrm{ex}_d(n, K^{(d)}_{2,\ldots,2}) = \Omega\left(n^{d-\frac{d-1/s}{2^d-1}}\right).$$

It is easy to see that the number $s = s(d)$ exists precisely when d and $2^d - 1$ are relatively prime, which holds, for instance, when d is a prime number or a power of 2, but does not hold for many other numbers, such as $d = 6$, 12, 18, 20, 21. In fact, their result fails to apply for a positive proportion of the positive integers, as may be seen by noting that if the condition $(d, 2^d - 1) = 1$ fails for a given d, then it also fails for all multiples of d.

In this paper, we improve on the lower bound from Theorem 1 by establishing the following result, whose proof relies on a new random algebraic method using multilinear maps rather than high-degree polynomials.

Theorem 2 *For any $d \geq 2$, let r and s be positive integers such that $d(s-1) < (2^d-1)r$. Then*

$$\mathrm{ex}_d(n, K^{(d)}_{2,\ldots,2}) = \Omega\left(n^{d-\frac{r}{s}}\right).$$

This not only improves the lower bound for the box problem provided by Theorem 1 for any d which is not a power of 2, but it also yields a gain over the probabilistic deletion bound (3) for *all* uniformities d. To see this, note that if $d \geq 2$, then d never divides $2^d - 1$, so we may set $r = 1$ and $s = \lceil \frac{2^d-1}{d} \rceil > \frac{2^d-1}{d}$.

Corollary 1 *For any $d \geq 2$,*

$$\mathrm{ex}_d(n, K^{(d)}_{2,\ldots,2}) = \Omega\left(n^{d-\lceil \frac{2^d-1}{d} \rceil^{-1}}\right).$$

By a result of Ferber, McKinley and Samotij [7, Theorem 9], any polynomial gain over the deletion lower bound for the extremal number of a uniform hypergraph $\mathcal{H}$ implies an optimal counting result for the number of $\mathcal{H}$-free graphs on n vertices. In combination with Corollary 1, this implies the following result, generalizing a celebrated theorem of Kleitman and Winston [13] on the $d = 2$ case.

Corollary 2 *For any $d \geq 2$, let $\mathcal{F}_n\left(K^{(d)}_{2,\ldots,2}\right)$ be the set of all (labeled) $K^{(d)}_{2,\ldots,2}$-free graphs with vertex set $\{1,\ldots,n\}$. Then there exists a positive constant C depending only on d and an infinite sequence of positive integers n for which*

$$\left|\mathcal{F}_n\left(K^{(d)}_{2,\ldots,2}\right)\right| \leq 2^{C\cdot \mathrm{ex}_d(n,K^{(d)}_{2,\ldots,2})}.$$

For the reader's convenience, we include below a table comparing the bounds provided by the deletion bound (3), by Gunderson, Rödl and Sidorenko's Theorem 1 and by our Corollary 1. A number α on the d^{th} row of the table means that the corresponding method gives the lower bound

$$\mathrm{ex}_d(n, K_{2,\ldots,2}) = \Omega\left(n^{d-1/\alpha}\right),$$

while an empty cell in the GRS column means that the method does not apply for that value of d. In particular, we note that our method recovers both the fact that $\mathrm{ex}(n, K_{2,2}) = \Theta(n^{3/2})$ and the lower bound $\mathrm{ex}_3(n, K^{(3)}_{2,2,2}) = \Omega(n^{8/3})$ of Katz, Krop and Maggioni.

d	Deletion	GRS	Corollary 1
2	1.50	2.00	2.00
3	2.33	2.50	3.00
4	3.75	4.00	4.00
5	6.20	6.25	7.00
6	10.50		11.00
7	18.14	18.16	19.00
8	31.87	32.00	32.00
9	56.77	56.80	57.00
10	102.30	102.33	103.00
11	186.09	186.10	187.00
12	341.25		342.00
13	630.07	630.08	631.00
14	1170.21	1170.22	1171.00
15	2184.46	2184.50	2185.00
16	4095.93	4096.00	4096.00
17	7710.05	7710.06	7711.00
18	14563.50		14564.00
19	27594.05	27594.05	27595.00
20	52428.75		52429.00
21	99864.33		99865.00

2 New lower bounds for the Erdős box problem

2.1 Linear algebra preliminaries

Let $V_1,\ldots,V_d$ be finite-dimensional spaces over the field $\mathbb{F}_q$. Following standard convention, we call a function $T : V_1 \times \cdots \times V_d \to \mathbb{F}_q$ *multilinear* if, for every $i \in \{1,\ldots,d\}$ and every fixed choice of $x_j \in V_j$ for each $j \neq i$, the function $T(x_1,\ldots,x_{i-1},x,x_{i+1},\ldots,x_d)$, considered as a function on V_i, is linear over $\mathbb{F}_q$.

The vector space of all multilinear functions $T : V_1 \times \cdots \times V_d \to \mathbb{F}_q$ can be naturally identified with the space $V_1^* \otimes \cdots \otimes V_d^*$, where V^* denotes the dual space of V. A *uniformly random* multilinear function $T : V_1 \times \cdots \times V_d \to \mathbb{F}_q$ is then a random element of the space $V_1^* \otimes \cdots \otimes V_d^*$, chosen according to the uniform distribution.

If, for each i, we have a subspace $U_i \subset V_i$, then we can define a restriction map

$$r : V_1^* \otimes \cdots \otimes V_d^* \to U_1^* \otimes \cdots \otimes U_d^*.$$

We have the following simple, but important, claim about these restriction maps.

Claim 1 *The restriction $r(T)$ of a uniformly random multilinear function T is again uniformly random.*

Proof: The map r is linear and surjective and so all $T' \in U_1^* \otimes \cdots \otimes U_d^*$ have the same number of preimages in $V_1^* \otimes \cdots \otimes V_d^*$. $\square$

It will also be useful to note the following simple consequence of multilinearity.

Proposition 1 *Suppose that $T : V_1 \times \cdots \times V_d \to \mathbb{F}_q$ is multilinear and, for every $i = 1, \ldots, d$, there are vectors $v_i^0, v_i^1 \in V_i$ such that*

$$T(v_1^{\varepsilon_1}, \ldots, v_d^{\varepsilon_d}) = 1$$

for all 2^d choices of $\varepsilon_i \in \{0,1\}$. Then, for any u_i which lie in the affine hull of v_i^0, v_i^1 for each $i = 1, \ldots, d$,

$$T(u_1, \ldots, u_d) = 1.$$

Proof: Write $u_i = \alpha_i^0 v_i^0 + \alpha_i^1 v_i^1$ for some $\alpha_i^0 + \alpha_i^1 = 1$. Then, by multilinearity, we have

$$\begin{aligned} T(u_1, \ldots, u_d) &= \sum_{\varepsilon_1, \ldots, \varepsilon_d \in \{0,1\}} \alpha_1^{\varepsilon_1} \cdots \alpha_d^{\varepsilon_d} T(v_1^{\varepsilon_1}, \ldots, v_d^{\varepsilon_d}) \\ &= \sum_{\varepsilon_1, \ldots, \varepsilon_d \in \{0,1\}} \alpha_1^{\varepsilon_1} \cdots \alpha_d^{\varepsilon_d} = (\alpha_1^0 + \alpha_1^1) \cdots (\alpha_d^0 + \alpha_d^1) = 1, \end{aligned}$$

as required. $\square$

2.2 Proof of Theorem 2

Fix positive integers d, r and s and let q be a large prime power. Let $V = \mathbb{F}_q^s$ and let $T_1, \ldots, T_r \in V^{*\otimes d}$ be independent uniformly random multilinear functions. Let $\mathcal{H}$ be the d-partite d-uniform hypergraph between d copies of V whose edge set $\mathcal{E}$ consists of all tuples $(v_1, \ldots, v_d) \in V^d$ such that $T_i(v_1, \ldots, v_d) = 1$ for all $i = 1, \ldots, r$. Let us estimate the expected number of edges in $\mathcal{H}$.

Claim 2 $\mathbb{E}\left[|\mathcal{E}|\right] = (q^s - 1)^d q^{-r} \sim q^{ds-r}$.

Proof: Note that if one of $v_1, \ldots, v_d$ is zero, then $T_i(v_1, \ldots, v_d) = 0$, so we may assume that $(v_1, \ldots, v_d)$ is one of the $(q^s - 1)^d$ remaining sequences of non-zero vectors and calculate the probability that it belongs to $\mathcal{E}$. Let $U_i = \langle v_i \rangle \subset V$, a one-dimensional subspace of V. By Claim 1, the restriction T_i' of T_i to $U_1 \times \cdots \times U_d$ is uniformly distributed in $U_1^* \otimes \cdots \otimes U_d^*$. But the latter space is one-dimensional and so $T_i'(v_1, \ldots, v_d)$ takes the value 1 with probability q^{-1}. Since $T_1, \ldots, T_r$ are independent, the functions $T_1', \ldots, T_r'$ are independent, so they all are equal to one at $(v_1, \ldots, v_d)$ with probability exactly q^{-r}. $\square$

We now estimate the expected number of (appropriately ordered) copies of $K^{(d)}_{2,\ldots,2}$ in $\mathcal{H}$.

Claim 3 *Let $\mathcal{F}$ denote the family of all $(v_1^0, v_1^1, \ldots, v_d^0, v_d^1) \in V^{2d}$ where $v_j^0 \neq v_j^1$ for all j and $T_i(v_1^{\varepsilon_1}, \ldots, v_d^{\varepsilon_d}) = 1$ for all $i = 1, \ldots, r$ and all choices of $\varepsilon_1, \ldots, \varepsilon_d \in \{0, 1\}$. Then $\mathbb{E}[|\mathcal{F}|] \sim q^{2ds - 2^d r}$.*

Proof: If, for some $j = 1, \ldots, d$, the vectors v_j^0 and v_j^1 are collinear, say $v_j^1 = \lambda v_j^0$ for some $\lambda \neq 1$ (but allowing $\lambda = 0$), then

$$T(v_1^0, \ldots, v_j^1, \ldots, v_d^0) = \lambda T(v_1^0, \ldots, v_j^0, \ldots, v_d^0),$$

so these two numbers cannot be equal to 1 simultaneously. Therefore, we may restrict attention to only those tuples where v_j^0 and v_j^1 are linearly independent for all $j = 1, \ldots, d$.

Fix one of the $(q^s - 1)^d (q^s - q)^d$ remaining tuples $\bar{v} = (v_1^0, v_1^1, \ldots, v_d^0, v_d^1)$ and let us compute the probability that $\bar{v} \in \mathcal{F}$. Let $U_j = \langle v_j^0, v_j^1 \rangle$ be the two-dimensional vector space spanned by v_j^0 and v_j^1. By Claim 1, the restriction T_i' of T_i to $U_1 \times \cdots \times U_d$ is uniformly distributed in $U_1^* \otimes \cdots \otimes U_d^*$. Moreover, the independence of $T_1, \ldots, T_r$ implies that $T_1', \ldots, T_r'$ are also independent. Now observe that the set of 2^d tensors

$$\{v_1^{\varepsilon_1} \otimes \cdots \otimes v_d^{\varepsilon_d} : \ \varepsilon_j \in \{0, 1\}\}$$

forms a basis for the space $U_1 \otimes \cdots \otimes U_d$. Therefore, there exists a unique $R \in U_1^* \otimes \cdots \otimes U_d^*$ such that $R(v_1^{\varepsilon_1}, \ldots, v_d^{\varepsilon_d}) = 1$ for all $\varepsilon_j \in \{0, 1\}$. Moreover, since there are q^{2^d} different choices for the value of a function in $U_1^* \otimes \cdots \otimes U_d^*$ at the $(v_1^{\varepsilon_1}, \ldots, v_d^{\varepsilon_d})$ and each such choice determines a unique function, the probability that $T_i' = R$ is q^{-2^d}. Since $\bar{v} \in \mathcal{F}$ if and only if $T_i' = R$ for all $i = 1, \ldots, r$, the independence of the T_i' implies that the probability $\bar{v} \in \mathcal{F}$ is $q^{-2^d r}$. Thus,

$$\mathbb{E}[|\mathcal{F}|] = (q^s - 1)^d (q^s - q)^d q^{-2^d r} \sim q^{2ds - 2^d r},$$

as required. $\square$

The next step is crucial.

Lemma 1 *Let $\mathcal{B}$ be the family of all $(v_1, \ldots, v_d) \in \mathcal{E}$ for which there exists $(v_1', \ldots, v_d') \in V^d$ such that $(v_1, v_1', \ldots, v_d, v_d') \in \mathcal{F}$. Then*

$$\mathbb{E}[|\mathcal{B}|] \leq (1 + o(1)) q^{-d} \mathbb{E}[|\mathcal{F}|].$$

Proof: Given a sequence of affine lines $l_1, \ldots, l_d \subset V$, denote by $P(l_1, \ldots, l_d)$ the set of all sequences $(x_1, x_1', \ldots, x_d, x_d') \in V^{2d}$ such that x_j and x_j' are distinct and lie on l_j for all j. Clearly,

$$|P(l_1, \ldots, l_d)| = q^d (q - 1)^d \sim q^{2d}.$$

Note that:

1. If $(l_1, \dots, l_d) \neq (l'_1, \dots, l'_d)$, then $P(l_1, \dots, l_d) \cap P(l'_1, \dots, l'_d) = \emptyset$, since the lines $l_1, \dots, l_d$ are uniquely determined by any member of $P(l_1, \dots, l_d)$.
2. If $P(l_1, \dots, l_d) \cap \mathcal{F} \neq \emptyset$, then $P(l_1, \dots, l_d) \subset \mathcal{F}$ by Proposition 1.
3. Any $\bar{v} \in \mathcal{F}$ is contained in $P(l_1, \dots, l_d)$ for some $l_1, \dots, l_d$.

Denote the family of all tuples $(l_1, \dots, l_d)$ such that $P(l_1, \dots, l_d) \cap \mathcal{F} \neq \emptyset$ by $\mathcal{L}$. By the observations above, we have that

$$|\mathcal{L}| q^d (q-1)^d = |\mathcal{F}|.$$

On the other hand, it is clear that

$$\mathcal{B} = \bigcup_{(l_1, \dots, l_d) \in \mathcal{L}} l_1 \times l_2 \times \dots \times l_d,$$

so that

$$|\mathcal{B}| \leq q^d |\mathcal{L}| = (q-1)^{-d} |\mathcal{F}|.$$

Taking expectations, we obtain the required result. □

By definition, the subgraph $\mathcal{H}'$ of $\mathcal{H}$ with edge set $\mathcal{E} \setminus \mathcal{B}$ is $K_{2,\dots,2}$-free. By Lemma 1 and Claim 3,

$$\mathbb{E}[|\mathcal{B}|] \leq (1 + o(1)) q^{-d} \mathbb{E}[|\mathcal{F}|] = (1 + o(1)) q^{2ds - 2^d r - d}.$$

On the other hand, by Claim 2, $\mathbb{E}[|\mathcal{E}|] \sim q^{ds-r}$. By the assumption on r and s from the statement of Theorem 2, we have

$$2ds - 2^d r - d < ds - r,$$

which immediately implies that $\mathbb{E}[|\mathcal{B}|] = o(\mathbb{E}[|\mathcal{E}|])$. Therefore, there must exist a $K^{(d)}_{2,\dots,2}$-free hypergraph $\mathcal{H}'$ on a ground set of size $n = dq^s$ with edge set $\mathcal{E}'$ satisfying

$$|\mathcal{E}'| = (1 + o(1)) q^{ds-r} = (c + o(1)) n^{d - \frac{r}{s}},$$

where $c = d^{\frac{r}{s} - d}$, completing the proof of Theorem 2.

References

[1] N. Alon, L. Rónyai and T. Szabó, Norm-graphs: variations and applications, *J. Combin. Theory Ser. B* **76** (1999), 280–290.

[2] B. Bukh, Random algebraic construction of extremal graphs, *Bull. London Math. Soc.* **47** (2015), 939–945.

[3] B. Bukh and D. Conlon, Rational exponents in extremal graph theory, *J. Eur. Math. Soc.* **20** (2018), 1747–1757.

[4] D. Conlon, Graphs with few paths of prescribed length between any two vertices, *Bull. Lond. Math. Soc.* **51** (2019), 1015–1021.

[5] P. Erdős, On sequences of integers no one of which divides the product of two others and on some related problems, *Mitt. Forsch.-Inst. Math. Mech. Univ. Tomsk* **2** (1938), 74–82.

[6] P. Erdős, On extremal problems of graphs and generalized hypergraphs, *Israel J. Math.* **2** (1964), 183–190.

[7] A. Ferber, G. McKinley and W. Samotij, Supersaturated sparse graphs and hypergraphs, *Int. Math. Res. Not. IMRN* **2020** (2020), 378–402.

[8] Z. Füredi and M. Simonovits, The history of degenerate (bipartite) extremal graph problems, in Erdős centennial, 169–264, Bolyai Soc. Math. Stud., 25, János Bolyai Math. Soc., Budapest, 2013.

[9] D. S. Gunderson, V. Rödl and A. Sidorenko, Extremal problems for sets forming Boolean algebras and complete partite hypergraphs, *J. Combin. Theory Ser. A* **88** (1999), 342–367.

[10] N. Katz, E. Krop and M. Maggioni, Remarks on the box problem, *Math. Res. Lett.* **9** (2002), 515–519.

[11] J. Kollár, L. Rónyai and T. Szabó, Norm-graphs and bipartite Turán numbers, *Combinatorica* **16** (1996), 399–406.

[12] T. Kővári, V. T. Sós and P. Turán, On a problem of K. Zarankiewicz, *Colloq. Math.* **3** (1954), 50–57.

[13] D. Kleitman and D. Wilson, On the number of graphs without 4-cycles, *Discrete Math.* **41** (1982), 167–172.

[14] J. Ma, X. Yuan and M. Zhang, Some extremal results on complete degenerate hypergraphs, *J. Combin. Theory Ser. A* **154** (2018), 598–609.