Sets with Few Subset Sums

Ruben Carpenter

24 May 2026

These are notes for the LIMDA seminar at UPC, Barcelona (50 min blackboard talk).

I'm presenting joint work with Colin Defant and Noah Kravitz. I met them at the Duluth REU last summer as a student. Colin (Harvard) is a co-director of the program and Noah (current postdoc at Oxford) is an advisor. The paper I'm discussing is on arxiv with the same title https://arxiv.org/abs/2605.05498

Thanks Juanjo for inviting me!

§1 Intro (setup and aditive combinatorics)

Let's jump in! In any abelian group we can make the following definitions.

Definition 1.1. Given a set A, its **sumset** is

$$A + A := \{a_1 + a_2 : a_1, a_2 \in A\}$$

and its set of **subset sums** is

$$\mathrm{FS}(A) := \left\{ \sum_{a \in B} a : B \subseteq A \text{ finite} \right\}.$$

Motivation: these objects appear in many different areas of math.

- NT: Goldbach's conjecture asks $\mathbb{P} + \mathbb{P} \stackrel{?}{=} 2\mathbb{Z}$.
- Geo: Let $A, B \subseteq \mathbb{R}^d$ have volume 1. Brunn-Minkowski lower bounds $\mathrm{vol}(A + B)$.
- Probability: Let $A = \{a_1, \ldots, a_n\}$. Define the random variable

  $$X = \sum_{i=1}^{n} X_i a_i, \quad X_i \sim \{0, 1\} \text{ i.i.d.}$$

  Here $\mathrm{supp}(X) = \mathrm{FS}(A)$. The classic Littlewood-Offord problem asks about the concentration properties of X (will discuss more later).

Let's get some intuition for what they measure.

Example 1.2

Let $A = [n] = \{1, \ldots, n\}$. Then:

- $A + A = \{2, 3, \ldots, 2n\}$ has size $2n - 1$.
- $\mathrm{FS}(A) = \{0, 1, \ldots, n(n+1)/2\}$ has size $\frac{1}{2}n(n+1) + 1$.

Example 1.3
Let $A = \{1, 2, 2^2 \ldots, 2^{n-1}\}$. Then $|A + A| = \binom{n}{2} + n$ and $|\operatorname{FS}(A)| = 2^n$.

Intuitively, the scale of these quantities measures the structure of A.

¶ **Warm-up bounds.** Let $A = \{a_1 < \cdots < a_n\} \subset \mathbb{R}_+$.

Claim 1.4 (Example 1.2 is extremal) — We have:

1. $|A + A| \geq 2|A| - 1$ (in particular, $\Omega(|A|)$).
2. $|\operatorname{FS}(A)| \geq \frac{1}{2}|A|(|A| + 1) + 1$ (in particular, $\Omega(|A|^2)$).

Proof. For both, induct on $|A|$: consider adding $a_{n+1} > \max A$.

1. We get at least two new elements: $a_{n+1} + a_n < a_{n+1} + a_{n+1}$.
2. Key observation: each residue class $\operatorname{FS}(A) \pmod{a_{n+1}}$ gives a new element in $\operatorname{FS}(A \cup \{a_{n+1}\})$. And we have at least $n + 1$ residue classes, namely $0, a_1, \cdots, a_n$.

(The proof of (2) anticipates ideas that will show up in the proof of Theorem 2.2!) □

Finding equality cases is an easy exercise. But here's a more interesting question – what if you're "close" to the minimum? If $|A + A| \leq 10|A|$, must A to "look like" an AP?

¶ **Inverse Theorems.** If $|A + A| \leq C|A|$, then A has **doubling constant** C.

- It can be a subset of an AP of length $O_C(n)$ (looks like an AP).
- A "random" element increases $|A + A|$ by $O_c(n)$, so we can have $O_C(1)$ of them.
- There are "higher dimensional" examples too!

There are "higher dimensional" examples too! Let's make an important definition.

Definition 1.5. A **generalized arithmetic progression (GAP)** is a set of the form

$$Q := \left\{ \sum_{j=1}^{r} \lambda_j x_j : x_j \in [s_j] \right\}.$$

Vocabulary: We say r is the **rank**. We call λ_j the *generators* and s_j the *sides*. It is *proper* if $|Q| = \prod_j |I_j|$ (The size of additive relations between λ_j determine how far it is from being proper – this is a very important idea that will show up later).

A proper GAP of rank r has doubling at most 2^r, so small doubling! It turns out that these are really the only examples, as the following celebrated result shows.

Theorem 1.6 (Freiman-Rusza, 1973)
Suppose $|A + A| \leq C|A|$. Then there exists a GAP $Q \supseteq A$ with $|Q| = O_C(|A|)$.

This reflects a central theme in additive combinatorics: sets with small doubling have constrained structure.

What about subset sums? What can a set with $|\operatorname{FS}(A)| \leq C|A|^2$ look like?

Arbitrary GAPs don't work! For instance, using notation from above

$$\operatorname{FS}(Q) = \left\{ \sum_{j=1}^{r} \lambda_j x_j : 0 \leq x_j \leq n \cdot s_j \right\}.$$

If the $\{\lambda_j\}$ are generic (i.e. there are no "small" relations between them), then all the values in the LHS will be different, yielding $|\operatorname{FS}(Q)| = \prod_j n s_j \asymp n^{d+1}$. So "the number of subset sums seems to grow with the dimension of the set". This will be a key theme today, that we'll make more precise later.

In our paper we prove the following characterisation.

Theorem 1.7 (C.-Defant-Kravitz, 2026)
Suppose that $|\operatorname{FS}(A)| \leq Cn^2$. Then it must be the case that $A = r \cdot Z \cup R$, where:

- $Z \subset \mathbb{Z}_+$ has sum $\sum_{z \in Z} z = O_C(n^2)$;
- $R \subset \mathbb{R}_+$ has $O_C(1)$ elements.

This is tight up to a constant: all such sets have $O_C(|A|^2)$ subset sums.

§1.1 Relation to Inverse Littlewood-Offord Theory (could skip?)

To give some more background on related literature... questions related to subset sums have often been studied in the context of Littlewood-Offord theory. Recalling the RV X from the introduction, one way to measure its concentration is via

$$\rho(A) := \sup_x \mathbb{P}[X = x].$$

Clearly $|\operatorname{FS}(A)| \cdot \rho(A) \geq 1$. There exist inverse theorems for when $\rho(A)$ is polynomially large (i.e. $|\operatorname{FS}(A)|$ is polynomially small).

Theorem 1.8 (Tao-Vu, 2011)
Let $e > 0$. If $\rho(A) \geq n^{-e}$, there must be a proper symmetric GAP Q of rank $r = O_e(1)$ that contains $(1 - o(1))$ of A and has size $|Q| = n^{O_e(1)}$.

Proof. I'll just mention they use Fourier methods. □

Compare this to our theorem: we start with a stronger assumption, and accordingly get a stronger result. Indeed, we can see that $\rho(A)$ is stronger by considering the equality cases $A = \{1, \ldots, n\}$. Then $\rho(A) = O(n^{-3/2})$ (here is an extra $1/\sqrt{n}$ because in the equality case the distribution has the shape of a Gaussian). Also, there are sets with $\rho(A) \gg n^{-2}$ that don't have $\operatorname{FS}(A) \ll n^2$. So its not clear how to use conclusions about $\rho(A)$ to solve our inverse problem about $|\operatorname{FS}(A)|$.

§2 Sketch Proof of Theorem 1.7

Through this sketch we will get to our second main result, which is of similar importance.

We begin by finding a piece of local structure. Uses a "pigeonholing argument".

Claim 2.1 — If $|\mathrm{FS}(A)| \le Cn^2$, then there is $B \subseteq A$ with $|B| \gg_C |A|$ and

$$|B + B| \ll_C |B|.$$

Proof Sketch. Adding a_{i+1} to $A(i) = \{a_1, \ldots, a_i\}$ introduces at least $|A(i)\hat{+}A(i)|/2$ new elements. But usually we have $O(i)$ new elements, giving the bound. □

By Freiman-Rusza, we immediately get a proper GAP $Q \supseteq B$ with $|Q| \ll |B| \ll n$.[1] We therefore want to think of B as a higher dimensional set. Formally, consider the projection $\pi_\lambda : \mathbb{R}^d \to \mathbb{R}$ sending $(x_1, \ldots, x_d) \mapsto \lambda_1 x_1 + \cdots + \lambda_d x_d$. Then

$$\begin{aligned} \pi_\lambda : \widetilde{Q} = [s_1] \times \cdots \times [s_d] &\mapsto Q \\ \widetilde{B} &\to B \\ \mathrm{FS}(\widetilde{B}) &\mapsto \mathrm{FS}(B) \end{aligned}$$

We want to use $\mathrm{FS}(\widetilde{B})$ as a proxy for $\mathrm{FS}(B)$. Heuristic: The only way $\mathrm{FS}(\widetilde{B})$ collapses to a very small set (size $O_C(n^2)$) is if there are many additive relations between $\{\lambda_j\}$.

So our plan (leave on board if possible) is:

1. Expansion: Show $\mathrm{FS}(\widetilde{B})$ is large.
2. Rank-Reduction: Quantify "π_λ non-injective means small additive relations". Conclude there is $r \in \mathbb{R}_+$ such that $r^{-1} \cdot (\lambda_1, \ldots, \lambda_d) \in (\mathbb{Z} \setminus \{0\})^r$ with $r^{-1}\lambda_j \ll_C \frac{n}{s_j}$. Hence $B \subseteq r \cdot \{0, 1, \ldots, m\}$ where $m \le \sum_j r^{-1}\lambda_j \cdot s_j = O_c(n)$.
3. Endgame: Extend conclusion to $A \setminus B$.

Conceptually, the first step is our main contribution to understanding subset sums.

¶ Step 1: Expansion (in higher dimensions) Here we discuss a question of independent interest that will end up being a main tool in our proof.

In $\mathbb{R}$ we always have $|\mathrm{FS}(A)| \gg n^2$. How about in $\mathbb{R}^d$? This question is boring if we let A be inside a line $\mathbb{R} \subset \mathbb{R}^d$. Yet, we saw before that a "properly d dimensional" GAP has $\Omega(n^{d+1})$ subset sums. How do we make the question meaningful?

One possibility is saying A isn't concentrated in any proper linear subspace.

Theorem 2.2 (C.-Defant-Kravitz, 2026)
Let $0 < \varepsilon < 1$. Let $A \subseteq \mathbb{R}^d$ be an n-element set. If no proper subspace of $\mathbb{R}^d$ contains more than $(1 - \varepsilon)n$ elements of A, then $|\mathrm{FS}(A)| \gg_{d,\epsilon} |A|^{d+1}$.

A more "stability-flavored" phrasing is: if $A \subseteq \mathbb{R}^d$ is an n-element set with $|\mathrm{FS}(A)| = o(n^{d+1})$, then some codimension-1 subspace contains $(1 - o(1))n$ elements of A.

There are several interesting families of examples saturating this statement:

[1] This is already a stronger conclusion than Tao-Vu ILw-O.

- Let $n_1, \ldots, n_d$ be positive integers with $n_1 \cdots n_d \asymp n$ and consider the $\theta(n)$ element set $A \subseteq [0, n_1] \times \cdots \times [0, n_d]$. Then

$$\mathrm{FS}(A) \subseteq [0, nn_1] \times \cdots \times [0, nn_d]$$

has size at most $(1 + nn_1) \cdots (1 + nn_d) \ll n^d n_1 \cdots n_d \ll n^{d+1}$.

- E.g in $d = 2$ take $[n] \times \{1\}$ (this is just a concrete instance of above).

Moreover, it is tight in a precise sense (almost all on x axis, $o(n)$ points away).

Proof of Theorem 2.2. This is much harder than $d = 1$! Naively, following the footsteps from $d = 2$ imagine inserting a_{n+1}. Before we cared about the number of values of $\mathrm{FS}(A)$ (mod a_{n+1}). The natural analogue is to care about the number of values under a different projection, namely

$$\pi_{a_{n+1}} : \mathbb{R}^d \to a_{n+1}^{\perp} \cong \mathbb{R}^{d-1}.$$

We would like the projection $\pi_{a_{n+1}}(\mathrm{FS}(A))$ to have many elements. Turns out this doesn't quite work.

But we can fix it by adding two elements a_{n+1}, a_{n+2}, with ideas of the same flavor– the proof is really fun and I encourage you to read the paper if you are curious! $\square$

Back to our setting! If B is "properly d dimensional", then $|\,\mathrm{FS}(\widetilde{B})|$ is large.

Else, we can throw out some elements of B and restrict to the (possibly slanted) subspace where it is concentrated. Keep going until we eventually find its true dimension! Every time we throw out dimensions we need to make a change of coordiantes to our generators. This can be bad, because we want the bounding box to have size $\ll n$ (draw example). To choose good coordinates we need to invoke technical lemmas (e.g. discrete John's ellipsoid) – I won't go into details, but want to mention it is annoying :p

Hence, WLOG B is dense in the proper GAP of rank d, and $|\,\mathrm{FS}(\widetilde{B})| \gg n^{d+1}$.

¶ Step 2: Rank Reduction So, the projection π_λ has to be "very non-injective":

$$\max_c |\pi_\lambda^{-1}(c) \cap \mathrm{FS}(\widetilde{B})| \geq \frac{|\,\mathrm{FS}(\widetilde{B})|}{|\,\mathrm{FS}(B)|} \gg \frac{n^{d+1}}{Cn^2} \asymp_C n^{d-1}.$$

Instead of working inside $\mathrm{FS}(\widetilde{B})$, lets make our life easier by working inside a box that contains it:

$$\mathrm{FS}(\widetilde{B}) \subseteq n \cdot \widetilde{Q} = I_1 \times \cdots \times I_d, \qquad I_j := [0, ns_j].$$

Then, there is some c for which

$$\pi_\lambda^{-1}(c) \cap (I_1 \times \cdots I_d) = \left\{(x_1, \ldots, x_d) \in I_1 \times \cdots \times I_d : \sum \lambda_j x_j = c\right\}$$

is very large. We want to say that if there are many solutions inside a small box, then there are many additive relations between the generators. This is the most involved part of the argument and occupies a large part of our paper.

Example 2.3 ($d = 2$)

This is very easy and concrete in two dimensions!

Here $A \subseteq Q = \{x_1t_1 + x_2t_2 : x_i \in [s_i]\}$ has $\mathrm{FS}(A) \le Cn^2$. We aim to show that there is r. So there is a fiber $x_1t_1 + x_2t_2 = c$ with $\gg n$ solutions in $I_1 \times I_2$.

- If $t_1/t_2 \notin \mathbb{Q}$ the fiber can have at most one solution. Not the case!
- Hence we can scale to $\lambda_1, \lambda_2 \in \mathbb{Z}$ which are coprime. Then the solutions (x_1, x_2) to this equation differ by integer multiples of $(\lambda_2, -\lambda_1)$. So $x_i \in I_i$ gives

$$|\pi^{-1}(c) \cap (I_1 \times I_2)| \le \min\left\{\frac{|I_1|}{|\lambda_2|}, \frac{|I_2|}{|\lambda_1|}\right\} = \frac{|I_1| \cdot |I_2|}{\max\{|\lambda_1| \cdot |I_1|, |\lambda_2| \cdot |I_2|\}}.$$

  Because $|I_1| \cdot |I_2| \ll n$ we get $\max |\lambda_j| \cdot |I_j|$, as we wanted!

In higher dimensions two steps become much harder:

- Proving that $(t_1, \ldots, t_r)$ is a scalar multiple of an integer vector $(\lambda_1, \ldots, \lambda_r)$
- Subsequently obtaining the $\max_j\{|\lambda_j| \cdot |I_j|\}$ savings in the fiber bound.

  We heavily used that we can "spell out" the relations between λ as being precisely $\langle(\lambda_2, -\lambda_1)\rangle$. In general we have up to $d-1$ independent "relations" which span everything, and it is harder to bound what combinations of them stay in the box.

  (Perhaps I be concrete and find the basis for $6x + 10y + 15z = 0$ or so)

Overcoming this constitutes a large technical bulk from our paper, which I won't go into detail here. We prove the following lemma, which does most of the heavy lifting for us.

Lemma 2.4 (Lemma 4.5 in our paper)

Let $\lambda = (\lambda_1, \ldots, \lambda_d) \in (\mathbb{R}\backslash\{0\})^d$, and $I_1, \ldots, I_r \subset \mathbb{Z}$ be intervals. Define $\pi_\lambda : \mathbb{R}^d \to \mathbb{R}$ as above. Let $c \in \mathbb{R}$ be any real number. Then:

1. If $\dim \pi_\lambda^{-1}(c) = d - 1$, and $\lambda_1, \ldots, \lambda_r \in \mathbb{Z}$ with $\gcd\{\lambda_j\} = 1$, then

$$|\pi_\lambda^{-1}(c) \cap (I_1 \times \cdots \times I_d)| \ll_d \frac{|I_1| \cdots |I_d|}{\max_j(|\lambda_j| \cdot |I_j|)}.$$

2. If $\dim \pi_\lambda^{-1}(c) < d - 1$, then

$$|\pi_\lambda^{-1}(c) \cap (I_1 \times \cdots \times I_d)| \le \frac{|I_1| \cdots |I_d|}{\max_j |I_j| \cdot \min_j |I_j|}.$$

Proof. The main tools come from the "geometry of numbers" – the details are spelled out in the paper.

(Maybe I'll say that: morally, the denominator bound should be $\mathrm{vol}(I_1 \times I_d)/\mathrm{covol}(\Lambda) + o(1)$... but the error term is really hard to handle, and saving this requires theory). □

We can now implement the logic from $d = 2$:

- We rule out the case that $\dim \pi^{-1}(c) < d - 1$. Indeed, using case (2) of 2.4 gives

$$n^{d-1} \ll_C \max_c |\pi_\lambda^{-1}(c) \cap (I_1 \times \cdots \times I_d)| \ll_C \frac{|I_1| \cdots |I_d|}{\max_j |I_j| \cdot \min_j |I_j|}$$

$$\ll \frac{n^d \cdot s_1 \cdots s_d}{\underbrace{\max_j |ns_j|}_{\gg n \cdot n^{1/d}} \cdot \underbrace{\min_j |ns_j|}_{\geq n}}$$
$$\ll_C \frac{n^{d+1}}{n \cdot n^{1/d} \cdot n} = n^{d-1-\frac{1}{d}},$$

contradiction.

- So $\dim \pi^{-1}(c) = d-1$. This means there is $r \in \mathbb{R}^+$ with $r^{-1}\lambda_j \in \mathbb{Z}$ and $\gcd\{\lambda_j\} = 1$. Then case (1) of 2.4 gives $\max_j\{|\lambda_j| \cdot |I_j|\} \ll_C n$ and we are done!

Altogether, we deduce that $B = \pi(\widetilde{B})$ is contained in a homogeneous arithmetic progression of length $O_C(n)$.

¶ Step 3: Endgame. Given that B looks like an AP, what does that tell us about A? This step is hands-on combinatorial analysis – the details are in the paper.

(Maybe I can say that we use the fact that $\mathrm{FS}(B)$ is a dense ($\Omega_C(1)$-fraction) subset of an interval of length $\Theta_C(n^2)$ and hence $\mathrm{FS}(A) = \mathrm{FS}(A \setminus B) + \mathrm{FS}(B)$ cannot contain too many disjoint translates of $\mathrm{FS}(B)$).

This concludes the proof of Theorem 1.7!

§3 Further Directions and Related Problems

- One of the main takeaways of the paper is that "higher dimensional" sets have faster subset sum expansion.

  What happens if we enforce a stricter condition that "not concentrated in a subspace"? An alternative leads to an open problem.

  Conjecture 3.1. Let $A \subset \mathbb{R}^d$. If any d elements form a basis, $|\mathrm{FS}(A)| \gg |A|^{2d-1}$.

  The best construction we found is the moment curve $\{(1, i, i^2, \ldots, i^{d-1}) : i \in [n]\}$ taken modulo some prime $p \asymp n$. Can we improve this?

- What inverse theorems can we say about sets A with $|\mathrm{FS}(A)| \leq n^{2.01}$ or $|\mathrm{FS}(A)| \leq n^3$? How do higher-rank generalized arithmetic progressions change the picture?

Thank you for following along! :-)