Uncommon systems of linear equations

Talk by Dingding Dong

Notes by Sanjana Das

September 19, 2024

This is joint work with Anqi Li and Yufei Zhao.

§1 Introduction

§1.1 Motivation

Overall, our problem is about counting monochromatic solutions to systems of linear equations (in vector spaces over finite fields). To motivate the actual problem, we'll start with a simple theorem.

Theorem 1.1 (Cameron–Cilleruelo–Serra 2007)
Let G be any finite abelian group, and let $f\colon G \to \{0,1\}$ be any 2-coloring of G. Then the number of monochromatic Schur triples — triples $(x,y,z) \in G^3$ with $x+y=z$ — is at least $\frac{1}{4}\left|G\right|^2$.

Proof. First, we can write down the number of monochromatic Schur triples as

$$\sum_{x,y}(f(x)f(y)f(x+y) + (1-f(x))(1-f(y))(1-f(x+y)))$$

(where the first term corresponds to the color 1, and the second to the color 0). The two degree-3 terms cancel out, and we're left with

$$\sum_{x,y}(1 - f(x) - f(y) - f(x+y) + f(x)f(x+y) + f(y)f(x+y) + f(x)f(y)).$$

But for each of the degree-1 terms, we're really just summing over an arbitrary $x \in G$, and for each of the degree-2 terms, we're really summing over arbitrary pairs $(x,y) \in G^2$ — there's no relation between the two terms (e.g., as (x,y) ranges over all possible pairs, so does $(x, x+y)$). So this sum simplifies to

$$\left|G\right|^2 - 3\left|G\right|\sum f + 3\left(\sum f\right)^2.$$

And one can check that no matter what $\sum f$ is, this is always at least $\frac{1}{4}\left|G\right|^2$. □

If we colored the elements of G *randomly*, then we'd expect a $\frac{1}{4}$-fraction of Schur triples to be monochromatic; there are $\left|G\right|^2$ Schur triples in total, so we'd expect $(\frac{1}{4}+o(1))\left|G\right|^2$ monochromatic ones. (The $o(1)$ is because of the few triples where x, y, and $x+y$ are not distinct.)

So Theorem 1.1 means that the number of monochromatic Schur triples in G is minimized by a random coloring; for this reason, we say that Schur triples are *common*.

§1.2 The question

The question we'll look at today is the following:

Question 1.2. What linear patterns are common?

First note that every linear pattern is characterized by a system of linear equations (and vice versa).

Example 1.3

- Schur triples correspond to the single equation $x + y = z$.
- 3-APs correspond to the equation $x - 2y + z = 0$.
- 4-APs correspond to the system of two equations stating that the first three and last three terms both form 3-APs, i.e.,

$$x_1 - 2x_2 + x_3 = 0,$$
$$x_2 - 2x_3 + x_4 = 0.$$

We'll use matrices to denote such systems of equations, because it's kind of unwieldy to write the variables x_i all the time. For example, we'll denote the system for 4-APs by

$$\begin{bmatrix} 1 & -2 & 1 & 0 \\ 0 & 1 & -2 & 1 \end{bmatrix}.$$

We'll call this a 2×4 *linear system*, since there's 2 equations on 4 variables. Also, for a single equation like $x - 2y + z = 0$, we call this a *length*-3 *equation*, since it has 3 variables.

We'll also only consider the case $G = \mathbb{F}_p^n$. The reason for this is that in the proof of commonness for Schur triples, we could simply write $z = x + y$. But if we had an equation $ax + by = cz$ and wanted to write one variable in terms of the others, then we might have division; this would be an issue if the order of G is not coprime to the coefficients, and weird things could happen. One can definitely study this, but it would need extra work that's kind of technical and annoying; so in this work, we study the more canonical case where G is a $\mathbb{F}_p$-vector space.

Here's a definition of commonness (generalizing the notion we mentioned earlier, but working with arbitrary linear systems over finite fields).

Definition 1.4. Let L be a $m \times k$ linear system (i.e., one with m equations on k variables), and let p be prime. We say L is common over $\mathbb{F}_p$ if for all 2-colorings $f\colon \mathbb{F}_p^n \to \{0, 1\}$, the quantity

$$\#(\text{monochromatic } (x_1, \ldots, x_k) \in \mathbb{F}_p^n \text{ with } Lx = 0)$$

is minimized asymptotically by the random coloring.

By *minimized asymptotically*, we mean that the number of monochromatic solutions under any coloring is at least $(1 - o(1))$ of the expected number under a random one. (This definition is kind of in English; we'll see more mathematical definitions later on.)

§1.3 History and results

First, going back to the proof of Theorem 1.1 for Schur triples, we can see that if we replace $x + y = z$ with *any* equation of odd length, the proof doesn't change at all — we have a sum of two products, the odd-degree terms cancel out, and we're left with a bunch of even-degree terms that get reduced trivially. This leads to the following observation.

Fact 1.5 — Every odd-length equation is common over all $\mathbb{F}_p$.

So the next question is what happens for even-length equations; for this, there's the following theorem.

Theorem 1.6 (Fox–Pham–Zhao 2019)
An even-length equation $a_1x_1 + \cdots + a_kx_k = 0$ is common over all $\mathbb{F}_p$ if and only if $a_1, \ldots, a_k$ can be partitioned into cancelling pairs.

A *cancelling pair* is a pair of the form $(a, -a)$ — so for example, $x_1 + 3x_2 - x_3 - 3x_4 = 0$ is common.

This work characterizes all common linear *equations*, so the next question is what happens for systems with more than one equation. For this, there's the following work.

Theorem 1.7 (Kamčev–Liebenau–Morrison 2021)
Any irredundant 2×4 linear system is uncommon over all sufficiently large $\mathbb{F}_p$.

Definition 1.8. A system is irredundant if the following conditions all hold:

- It doesn't imply any equation of the form $x_i = x_j$.
- The rows (i.e., equations) are linearly independent.
- There is no zero column.

The first condition is there because if a system has two identical variables, then you'll have a different expected number of monochromatic solutions in a random coloring (compared to the case with no identical variables) — this is because you only need to require that x_i is the same color as everything else, and then x_j automatically is as well. So you get a phenomenon that differs from most linear systems; that's why we forbid this.

The second condition essentially states that the equations don't have repeated information; and the third states that we don't have any free variables which are present but not related to anything else.

If L is redundant (meaning that one of these three conditions fail), then one can always find an irredundant subset $L' \subseteq L$ such that L' is common if and only if L is. So when studying commonness, it's fine to only consider irredundant systems.

This was all that was known about commonness of linear systems prior to our work. Generalizing Theorem 1.7, there was a conjecture made by the same group of authors.

Conjecture 1.9 (Kamčev–Liebenau–Morrison) — For any even $k \geq 4$, every irredundant $2 \times k$ linear system with girth $k - 1$ is uncommon over all sufficiently large $\mathbb{F}_p$.

Definition 1.10. The girth of a linear system L is defined as the length of the shortest equation that one can deduce from L (by taking linear combinations of rows).

As a remark, we can actually recover Theorem 1.7 from the conjecture. It might not initially seem like it because the conjecture has an extra girth condition. But when $k = 4$, nothing interesting happens when the girth is less than 3 — if the girth is 2, then you have an equation of the form $x = \lambda y$, and that can be handled easily. So girth 3 is the only interesting case when $k = 4$, which means the conjecture is a generalization of Theorem 1.7.

Our first theorem is that this conjecture is true.

Theorem 1.11 (Dong–Li–Zhao 2024+)
Conjecture 1.9 is true.

Our second result is about 2×5 linear systems (the case $k = 5$). In this case, the authors attempted to characterize all 2×5 common linear systems; they almost succeeded, except for two unknown cases.

Theorem 1.12 (Dong–Li–Zhao 2024+)
Let L be a 2×5 irredundant system that is not isomorphic to

$$\begin{bmatrix} 1 & 1 & -1 & -1 & 0 \\ 1 & -1 & 3 & 0 & -3 \end{bmatrix} \text{ or } \begin{bmatrix} 1 & 1 & -1 & -1 & 0 \\ 2 & -2 & 3 & 0 & -3 \end{bmatrix}.$$

(1) If L is isomorphic to one of

$$\begin{bmatrix} 1 & -1 & 1 & -1 & 0 \\ 1 & 2 & -1 & 0 & -2 \end{bmatrix} \text{ or } \begin{bmatrix} 1 & 0 & -1 & 2 & -2 \\ 0 & 1 & 2 & -1 & -2 \end{bmatrix} \text{ or } \begin{bmatrix} 1 & 1 & -1 & -1 & 0 \\ 2 & -2 & 1 & 0 & -1 \end{bmatrix},$$

or L is isomorphic to

$$\begin{bmatrix} a & b & 0 & 0 & c \\ 0 & 0 & a & b & c \end{bmatrix}$$

for some nonzero a, b, and c, or

$$\begin{bmatrix} a & b & 0 & 0 & -a-b \\ c & -c & d & -d & 0 \end{bmatrix}$$

where a, b, c, d, and $a+b$ are nonzero and $\{|a/b|, |a/(a+b)|, |b/(a+b)|\} \cap \{|c/d|, |d/c|\} \neq \emptyset$, then L is common over all $\mathbb{F}_p$ (for which its coefficients are nonzero).

(2) Otherwise L is uncommon over all sufficiently large $\mathbb{F}_p$.

§2 Proof ideas for Theorem 1.12

We'll now discuss the ideas behind the proof of Theorem 1.12. We'll focus on the case where the girth is 4 (so we'll ignore the last two classes in (1)) — suppose that L is an irredundant 2×5 system with girth 4.

§2.1 Homomorphism densities

We'll work with functions $f: \mathbb{F}_p^n \to [0, 1]$ (so we actually allow the range of f to be an *interval*, rather than $\{0, 1\}$; later, we may convert such a function to a 2-coloring by sampling using it).

Definition 2.1. We define the homomorphism density of L into f as

$$t_L(f) = \mathbb{E}_{Lx=0}[f(x_1)f(x_2)f(x_3)f(x_4)f(x_5)].$$

If f is a coloring, then this is the fraction of solutions to L that are in color 1.

We previously saw an English definition of commonness; we'll now state a more mathematical one in terms of homomorphism densities.

Definition 2.2. We say L is common over $\mathbb{F}_p$ if for all $f\colon \mathbb{F}_p^n \to [0,1]$, we have

$$t_L(f) + t_L(1-f) \geq 2^{1-k}.$$

Here 2^{1-k} is the fraction of monochromatic solutions that we'd expect from a random coloring (because L is irredundant, so we don't have repeated variables).

We'll first expand out the left-hand side — we have

$$t_L(f) + t_L(1-f) = \mathbb{E}_{Lx=0}[f(x_1)\cdots f(x_5) + (1-f(x_1))\cdots(1-f(x_5))].$$

The degree-5 terms again cancel out, and we're left with a bunch of terms like $f(x_1)$ and $f(x_1)f(x_2)$ and so on. But in all the terms of degrees up to 3, their variables are free (e.g., (x_1, x_2, x_3) ranges over all possible triples, because we assumed L has girth 4 so it can't imply any relation between these three variables).

So we can write $t_L(f) + t_L(1-f)$ as a sum of terms with at most 3 variables, which are easy to deal with (they just give us powers of $\mathbb{E}f$), and then

$$\mathbb{E}_{Lx=0}[f(x_1)f(x_2)f(x_3)f(x_4)]$$

and other similar 4-variable terms. When we look at this 4-variable term, we can no longer choose arbitrary values for x_1, x_2, x_3, and x_4 (as we could for the terms with at most 3 variables). Instead, L implies one linear equation on $x_1, \ldots, x_4$, which we call L_5 (because x_5 is missing); and this term corresponds to the homomorphism density $t_{L_5}(f)$.

We csn do the same thing with the other 4-variable terms (we define L_i as the equation implied by L on the variables other than x_i). Then after some normalization, we get the following statement.

Claim 2.3 — The system L is common if and only if $\sum_i t_{L_i}(f - \mathbb{E}[f]) \geq 0$ for all f.

§2.2 A common example

Now we'll talk about why the three common classes in Theorem 1.12 are common, the two exceptions are hard, and everything else is uncommon.

Example 2.4

Suppose that L is the system

$$\begin{bmatrix} 1 & -1 & 1 & -1 & 0 \\ 1 & 2 & -1 & 0 & -2 \end{bmatrix}$$

(the first common class from Theorem 1.12).

In this case, if we write the equations $L_5, L_4, \ldots, L_1$ as rows of a matrix, we get

$$\begin{bmatrix} 1 & -1 & 1 & -1 & 0 \\ 1 & 2 & -1 & 0 & -2 \\ 2 & 1 & 0 & -1 & -2 \\ 3 & 0 & 1 & -2 & 2 \\ 0 & -3 & 2 & -1 & 2 \end{bmatrix}$$

(where each row L_i is the relation that the variables other than x_i must satisfy — for example, the third row L_3 comes from adding the two equations in L to eliminate x_3). The good thing is that three of these equations are themselves common; this means they'll have nonnegative contribution to our sum, which will be good for the sum being nonnegative.

Our goal is to show that for all $f\colon \mathbb{F}_p^n \to \mathbb{R}$ with mean 0, we have

$$X = \sum_{i=1}^{5} t_{L_i}(f) \geq 0.$$

(Here we imagine that we've replaced f with the shifted function $f - \mathbb{E}[f]$.)

To show this, one trick we can do is to rewrite each of these homomorphism densities $t_{L_i}(f)$ in terms of the Fourier coefficients of f (in the same way as in the proof of Roth's theorem); this gives

$$X = \sum_{r\in\mathbb{F}_p^n} \Big(\widehat{f}(r)\widehat{f}(r)\widehat{f}(-r)\widehat{f}(-r) + \widehat{f}(r)\widehat{f}(2r)\widehat{f}(-r)\widehat{f}(-2r) + \widehat{f}(2r)\widehat{f}(r)\widehat{f}(-r)\widehat{f}(2r) + \widehat{f}(3r)\widehat{f}(r)\widehat{f}(-2r)\widehat{f}(2r) + \widehat{f}(-3r)\widehat{f}(2r)\widehat{f}(-r)\widehat{f}(2r)\Big).$$

(The coefficients of r in each term come from the coefficients in the corresponding equation.)

And the nice thing is that three of the equations have cancelling pairs, and we get $|\widehat{f}(r)|^4$ in the first term, and $|\widehat{f}(r)|^2|\widehat{f}(2r)|^2$ in the second and third terms. Then we can use Cauchy–Schwarz to cancel out the last two terms using these and show $X \geq 0$.

The second common class in Theorem 1.12 is similar; the third is also similar, but involves a trig identity.

§2.3 An uncommon example

In Theorem 1.12, we can see that most 2×5 equations are *not* common; that's because most don't have such nice cancellation properties. In particular, if none of the five equations $L_1, \ldots, L_5$ are common, then X will have mean 0, so we can conclude that it's sometimes negative. But when one or two of them are common, you need to use specific constructions to show that X is sometimes negative, and that's hard.

Example 2.5

Suppose that L is the system

$$\begin{bmatrix} 1 & -1 & 1 & -1 & 0 \\ -2 & 4 & 3 & 0 & -9 \end{bmatrix}$$

(which is uncommon).

First, one of the equations is symmetric (meaning it forms cancelling pairs) and the other isn't, so this case is nontrivial. As before, we can write down the five length-4 equations $L_1, \ldots, L_5$, to get

$$\begin{bmatrix} 1 & -1 & 1 & -1 & 0 \\ -2 & 4 & 3 & 0 & -9 \\ 5 & -7 & 0 & -3 & 9 \\ 2 & 0 & 7 & -4 & -9 \\ 0 & 2 & 5 & -2 & -9 \end{bmatrix}.$$

So we end up getting

$$Y = \sum_{r\in\mathbb{F}_p^n} \Big(|\widehat{f}(r)|^4 + \widehat{f}(-2r)\widehat{f}(4r)\widehat{f}(3r)\widehat{f}(-9r) + \widehat{f}(5r)\widehat{f}(-7r)\widehat{f}(-3r)\widehat{f}(9r) + \widehat{f}(2r)\widehat{f}(7r)\widehat{f}(4r)\widehat{f}(-9r) + \widehat{f}(2r)\widehat{f}(5r)\widehat{f}(-2r)\widehat{f}(-9r)\Big).$$

And we want to find some f such that $Y < 0$.

How do we do this? The first idea is that we want all the complex numbers $\widehat{f}(r)$ to be of the same absolute value. This is because the first term $|\widehat{f}(r)|^4$ is always positive, and if we set one coefficient $\widehat{f}(r)$ to be very large and the others to be very small, then this term will increase the fastest.

Moreover, we want $\operatorname{supp}(\widehat{f})$ to be closed under multiplication by all the coefficients in this expression (i.e., under the maps $r \mapsto \pm 2r, \pm 3r, \pm 4r, \pm 5r, \pm 7r, \pm 9r$). Why? Suppose there's some r for which $\widehat{f}(r)$ is nonzero but $\widehat{f}(-2r)$ is zero. Then the second term goes away, but the first term is still there — so we get a positive contribution from the first term, but don't get the possibility of a negative contribution from the second. And what we want is that every time we get a positive contribution from the first term, there's a chance of getting some negative contribution from the remaining four terms.

Motivated by this, we construct $\operatorname{supp}(\widehat{f})$ to be a *multiplication grid* of the powers of all the primes that are present here — i.e., we take

$$\operatorname{supp}(\widehat{f}) = \{\pm 2^{\alpha_1} 3^{\alpha_2} 5^{\alpha_3} 7^{\alpha_4} \mid 0 \le \alpha_i \le D\}$$

(where D is something very large), so that it's closed under multiplication by these numbers (except near the boundary). And to construct the actual Fourier coefficients, we take

$$\widehat{f}(2^{\alpha_1} \cdots 7^{\alpha_4}) = e^{i(\theta_1 \alpha_1 + \cdots + \theta_4 \alpha_4)}$$

to be a complex number whose argument is a linear combination of our exponents $\alpha_1, \ldots, \alpha_4$ (for some appropriately chosen coefficients $\theta_1, \ldots, \theta_4$). And we take

$$\widehat{f}(-2^{\alpha_1} \cdots 7^{\alpha_4}) = e^{-i(\theta_1 \alpha_1 + \cdots + \theta_4 \alpha_4)}$$

to be the complex conjugate of this.

In this case, the good thing is that every line of our 5×5 matrix has two positive and two negative elements. And if we consider something like $\widehat{f}(2r)\widehat{f}(-3r)$, most things cancel out. So we end up looking at a very finite sum, and by choosing the θ_i appropriately we can make it negative.

§2.4 Cases with two common equations

Finally, we'll talk about why the first two cases in Theorem 1.12 are hard. First, as long as two of the five equations are common, the above approach won't work — moreover, one has to fluctuate $|\widehat{f}(r)|$, or else it won't be possible to prove anything.

When two equations are common, our 5×5 matrix looks like

$$\begin{bmatrix} 1 & 1 & -1 & -1 & 0 \\ 1 & -1 & a & 0 & -a \\ a+1 & a-1 & 0 & -a & -a \\ 2 & 0 & a-1 & -a & -a \\ 0 & 2 & -1-a & -1 & a \end{bmatrix}.$$

What happens with this? Before (in the previous uncommon example), we made a multiplicative grid using the primes that appeared in our coefficients. It turns out that even with two common equations, if there are enough different prime factors that show up, we can use computer search to create a good multiplicative grid and a good periodic assignment of Fourier coefficients. (We need enough prime factors so that the construction can go into multiple dimensions.)

The thing about the two exceptions (the unknown cases in Theorem 1.12) is that for the first one, you only see two prime factors when you do this (2 and 3); for the second, there will be another prime factor, but it turns out they're not aligned in a way that's useful for us.