PCMI lecture series: from sunflowers to thresholds
Lecture 3 Problems

Lecturer: Shachar Lovett
TA: Anthony Ostuni

July 2025

Problem (2) on the original version of yesterday's problem set had an error in it. Here is a corrected (and more computationally friendly) version.

1. Prove that the robust sunflower lemma is sharp by considering the following example. Let $R_1, \ldots, R_n$ be pairwise disjoint sets, each of size $k = \log(n/c)$ for some $1 \leq c \ll n$ to be chosen later, and let U be their union. Consider the set system

$$\mathcal{F} = \{S \subset U : |S \cap R_i| = 1 \ \forall i \in [n]\}.$$

Note that $\mathcal{F}$ is a family of n-sets of size $|\mathcal{F}| = k^n$. Let $W \sim \mathrm{Bin}(U, 1/2)$.

A) Show that $\mathcal{F}$ *does* contain a $(1/2, 1/2)$-robust sunflower.

It turns out that by refining $\mathcal{F}$, we can remove all such robust sunflowers. Set $\varepsilon = 1/c$, and define $\mathcal{F}' \subset \mathcal{F}$ to satisfy

$$|S \cap S'| \leq (1-\varepsilon)n \quad \text{for all } S, S' \in \mathcal{F}', S \neq S'.$$

By a greedy approach, one can find such an $\mathcal{F}'$ of size $|\mathcal{F}'| \geq 2^{-n} k^{(1-\varepsilon)n}$. (You may assume this without proof.) Assume by contradiction $\mathcal{F}'$ does contain a $(1/2, 1/2)$-robust sunflower.

B) Show that $\Pr_W[W \text{ is disjoint from some set } R_i] > 1/2$.

C) Use this to argue that $\mathcal{F}'$ does not contain a $(1/2, 1/2)$-robust sunflower.

Observe that by setting $c = \sqrt{n}$, we have $|\mathcal{F}'| \geq (\log n)^{(1-o(1))n}$.

2. Let $\mathcal{F} = \{S_1, \ldots, S_m\}$ be a family of sets (possibly with repetitions) that is k-spread. Prove the following.

A) Let $\mathcal{F}' = \{S'_1, \ldots, S'_m\}$ be a family of sets with $S'_i \subset S_i$. Then $\mathcal{F}'$ is also k-spread.

B) Let $\mathcal{F}'$ be a sub-family of $\mathcal{F}$ of size $|\mathcal{F}'| \geq (1-\varepsilon)|\mathcal{F}|$. Then $\mathcal{F}'$ is $((1-\varepsilon)k)$-spread.

Recall the spread lemma.

Lemma 0.1. *Let $\mathcal{F}$ be a family of n-sets defined over a universe U. Let $p \in (0,1)$ and let $W \sim Bin(U,p)$. Assume that $\mathcal{F}$ is k-spread, where $k = cp^{-1}\log(n)$ for a large enough absolute constant c. Then*

$$\Pr_W[\exists S \in \mathcal{F}, S \subset W] \geq 1/2.$$

3. Suppose we only wanted a random $W \sim \mathrm{Bin}(U,p)$ to contain a $(1-\delta)$-fraction of an element of $\mathcal{F}$ with probability at least $1/2$. What spreadness would we require?

4. In yesterday's lecture, Shachar mentioned that if $p \in (0,1), W \sim \mathrm{Bin}(U,p)$, and $\mathcal{F}$ is a k-spread family of n-sets, where $k = cp^{-1}\log(n/\varepsilon)$ for a large enough absolute constant c, then

$$\Pr_W[\exists S \in \mathcal{F}, S \subset W] \geq 1-\varepsilon.$$

A) As a warm-up, prove the weaker result with $k = cp^{-1}\log(n)\log(1/\varepsilon)$. The proof only requires Lemma 0.1 as a "blackbox".

B) Prove the stronger result with $k = cp^{-1}\log(n/\varepsilon)$. (You will need to modify the proof of Lemma 0.1.)

At this point, you may either revisit problems (4) and (5) from yesterday if you did not have time to consider them, or work on the following open question.

5. (Open Question) In yesterday's problem session, you analyzed an example of a roughly $(\log n)$-spread family of n-sets where every two sets intersect. What is the largest spreadness a family of n-sets can have where every three sets intersect? It is known there is an example that is roughly $(\log\log n)$-spread. Can you match or beat that bound?