PCMI lecture series: from sunflowers to thresholds
Lecture 2 Problems

Lecturer: Shachar Lovett
TA: Anthony Ostuni

July 2025

1. Recall that a family $\mathcal{F}$ of sets over a universe U with common intersection $K = \cap_{S \in \mathcal{F}} S$ is a (p, ε)-*robust sunflower* (with kernel K) if $K \notin \mathcal{F}$ and

$$\Pr_{W \sim \mathrm{Bin}(U,p)}[\exists S \in \mathcal{F}, S \subset K \cup W] \geq 1 - \varepsilon.$$

Prove that any $(1/2r, 1/2)$-robust sunflower contains an r-sunflower.

2. 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 $|\mathcal{F}'| \geq (\log n)^{(1-o(1))n}$.

3. The quantitatively improved sunflower lemma discussed today operates by finding r disjoint sets in any $(cr \log n)$-spread family of n-sets. We will show that this is tight.

Let $R_1, \ldots, R_n$ be pairwise disjoint sets, each of size k, and let U be their union. Consider the set system

$$\mathcal{F} = \{S \subset U : \exists i, R_i \subset S \wedge \forall j \neq i, |S \cap R_j| = 1\}.$$

Pictorially, one may view this as an $n \times k$ grid, where $\mathcal{F}$ is constructed by choosing one element from each row of the grid, as well as one entire row. Note that $\mathcal{F}$ is a family of $(n + k - 1)$-sets.

A) Prove that $\mathcal{F}$ is an intersecting family. (In particular, we cannot even find two disjoint sets.)

B) Prove that $\mathcal{F}$ is $(c \log n / \log \log n)$-spread for an appropriate choice of k.

The following questions (essentially) come from Sam Spiro's notes: `https://samspiro.xyz/Misc/MethodsInCombo.pdf`. (Sam is at PCMI, so say hello!)

4. Let $\mathcal{F}$ be a random family of n-sets on a fixed universe U, where each n-set is included independently with probability p. Suppose $|U|$ is a multiple of n. Prove that if

$$p \geq C \cdot \frac{\log \binom{|U|}{n}}{\binom{|U|}{n}^{n-1}},$$

then $\mathcal{F}$ contains a perfect matching with probability at least $1/2$. (Recall a perfect matching is a set of edges such that every element in U is in exactly one edge.)

Hint: Analyze the spreadness of the set family whose ground set is all n-sets and whose edges are all perfect matchings.

5. Let $\mathcal{F}$ be a random family as in the previous question, and let $\mathcal{G}$ be a fixed family of n-sets. For any set family $\mathcal{H}$, let $U(\mathcal{H})$ denote the size of the universe of $\mathcal{H}$. Define

$$t(\mathcal{G}) = \max_{\mathcal{H} \subset \mathcal{G}} \left\{ \frac{|\mathcal{H}|}{U(\mathcal{H})} \right\}.$$

Prove that there exists some $C = C(\mathcal{G})$ such that if

$$p \geq C \cdot \frac{\log(|\mathcal{G}|)}{|U|^{1/t(\mathcal{G})}},$$

then $\mathcal{F}$ contains a copy of $\mathcal{G}$ with probability at least $1/2$.