Erdős-Ko-Rado for μ_p using Katona's circle method

Yuval Filmus

November 7, 2010

Let $\mathcal{F}$ be a family of subsets of $[n] = \{1, \ldots, n\}$, i.e. $\mathcal{F} \subset 2^{2^{[n]}}$. We say that $\mathcal{F}$ is an *intersecting family* if every two $A, B \in \mathcal{F}$ intersect. For example, the family

$$\mathcal{F}_i = \{S \subset [n] : i \in S\}$$

is intersecting; it is called a dictatorship.

For each $p \in [0, 1]$, we define a measure μ_p on families of subsets of $[n]$. Let X_p be a random subset of $[n]$ obtained by selecting each $i \in [n]$ with probability p. Then

$$\mu_p(\mathcal{F}) = \Pr[X_p \in \mathcal{F}].$$

The well-known Erdős-Ko-Rado theorem (in its weighted version) characterizes the intersecting families of maximal measure for $p \leq 1/2$ (when $p > 1/2$, the maximal measure depends on n).

Theorem 1 (EKR)**.** *Let $p \in [0, 1/2]$. Then every intersecting family $\mathcal{F}$ satisfies $\mu_p(\mathcal{F}) \leq p$. Moreover, if $p < 1/2$ then $\mu_p(\mathcal{F}) = p$ iff $\mathcal{F}$ is a dictatorship.*

Proof. Let C be a circle whose circumference has length 1. Pick n uniformly random points x_i on the circle (independently). Let I be any interval of length p on the circle, without its right endpoint. The set

$$X_I = \{i \in [n] : i \in I\}$$

is distributed according to X_p, and therefore $\Pr[X_I \in \mathcal{F}] = \mu_p(\mathcal{F})$. If we pick the interval I itself uniformly at random on the circle, then we still get $\Pr[X_I \in \mathcal{F}] = \mu_p(\mathcal{F})$, since the distribution of the points x_i is rotationally invariant.

Consider now any arrangement of the points on the circle, and let y_1, y_2 be any two starting points of intervals of length p such that $X_{[y_1, y_1+p)}, X_{[y_2, y_2+p)} \in \mathcal{F}$. Since $\mathcal{F}$ is intersecting, both intervals must intersect, i.e. $|y_1 - y_2| \leq p$ (here we use $p \leq 1/2$). Denoting by Y the set of starting points of intervals y such that $X_{[y, y+p)} \in \mathcal{F}$, we see that the diameter of Y is at most p, and so conditioned on the positions of x_i, $\Pr[X_I \in \mathcal{F}] \leq p$. We conclude that $\mu_p(\mathcal{F}) \leq p$.

Suppose now that $p < 1/2$ and that $\mu_p(\mathcal{F}) = p$. Thus for almost all arrangements of points x_i, Y consists of an interval of length p. Consider one such arrangement, and suppose wlog that $Y = (0, p]$. Since $0 \notin Y$, there must be a point x_s whose exact position is p. Let $X = X_{[p,2p)}$. Denote by E the event that the points $x_1, \ldots, x_n$ are chosen so that $x_s = p$, points in X are located in $[p, 2p)$, points not in X are located outside of $(0, 2p+\epsilon)$, and no points coincide, where ϵ is chosen small enough so that $2p + \epsilon < 1$. Since $\Pr[E] > 0$ conditioned on $x_s = p$, there is some arrangement of points satisfying E and $|Y'| = p$ (this follows from rotation invariance). Notice that $p \in Y'$ but $p + \delta \notin Y'$ for $\delta < \epsilon$, since otherwise $p + \delta \in Y$ (here we're using the fact that if $A \in \mathcal{F}$ and $A \subset B$ then $B \in \mathcal{F}$). We conclude that $Y' = (0, p]$. For small enough $\delta > 0$, $X'_{[\delta,\delta+p)} = \{x_s\}$, and so $\mathcal{F} = \{A \subset [n] : s \in A\}$ is a dictatorship. □

Open question: use this proof method to show that if $\mu_p(\mathcal{F}) \geq p - \epsilon$ then there's a point s contained in $1 - c_{p,\epsilon}$ of the sets in $\mathcal{F}$.