Exercises

The following exercises accompany the tutorial, "Erdős–Ko–Rado: Structure and Sparsification," from the 8^{th} Polish Combinatorial Conference, and are intended to help you fill in the gaps from the lectures, practice the techniques, and see how the results can be applied to research problems. If you would like to discuss your solutions, or have any questions, concerns, comments, or funny jokes, please contact Shagnik at `shagnik@mi.fu-berlin.de`.

Spectral graph theory

Exercise 1 Let G be a d-regular graph on n vertices, and let $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$ be its eigenvalues. Show that the following hold.

(a) $|\lambda_i| \leq d$ for every $i \in [n]$.

(b) $\lambda_n = -d$ if and only if G is bipartite.

Exercise 2 In this exercise we shall prove an asymmetric form of the expander-mixing lemma. Let G be a d-regular graph on n vertices with eigenvalues $d = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$. Given vertex sets $U, W \subseteq V(G)$, we define

$$e(U,W) = \left|\{(u,w) \in U \times W : \{u,w\} \in E(G)\}\right|.$$

Show that

$$\left| e(U,W) - \frac{d}{n} |U| \, |W| \right| \leq \lambda \sqrt{|U| \, |W| \left(1 - \frac{|U|}{n}\right)\left(1 - \frac{|W|}{n}\right)},$$

where $\lambda = \max\{|\lambda_2|, |\lambda_n|\}$.

The Kneser Graph and intersecting families

Exercise 3 Let $n > 2k$. We say families $\mathcal{F}, \mathcal{G} \subseteq \binom{[n]}{k}$ are *cross-intersecting* if $F \cap G \neq \emptyset$ for every $F \in \mathcal{F}$ and $G \in \mathcal{G}$.

(a) How large can the product $|\mathcal{F}| \, |\mathcal{G}|$ be if $\mathcal{F}$ and $\mathcal{G}$ are cross-intersecting?

(b) Characterise all pairs of cross-intersecting families that achieve equality in (a).

Bonus What if we want to maximise the sum $|\mathcal{F}| + |\mathcal{G}|$ instead?

Exercise 4 In this exercise we shall derive the eigenvalues of the Kneser Graph.[1] In order to do so, we define two families of matrices. Given $0 \le t, k \le n$, let $C_{t,k}$ and $D_{t,k}$ be the containment and disjointness matrices, whose rows are indexed by the t-sets in $\binom{[n]}{t}$ and columns are indexed by the k-sets in $\binom{[n]}{k}$, and whose entries are defined by

$$C_{t,k}(T,K) = \begin{cases} 1 & \text{if } T \subseteq K, \\ 0 & \text{otherwise,} \end{cases} \quad \text{and} \quad D_{t,k}(T,K) = \begin{cases} 1 & \text{if } T \cap K = \emptyset, \\ 0 & \text{otherwise.} \end{cases}$$

Note that, in particular, $C_{k,k}$ is the $\binom{n}{k} \times \binom{n}{k}$ identity matrix, and $D_{k,k}$ is the adjacency matrix of the Kneser Graph $KG(n,k)$.

(a) Prove the following identities

$$C_{i,j}C_{j,t} = \binom{t-i}{j-i} C_{i,t} \quad \textit{and} \quad C_{i,j}D_{j,t} = \binom{n-t-i}{j-i} D_{i,t}.$$

(b) Prove a couple more identities[2]

$$C_{t,k} = \sum_{i=0}^{t} (-1)^i C_{i,t}^T D_{i,k} \quad \textit{and} \quad D_{t,k} = \sum_{i=0}^{t} (-1)^i C_{i,t}^T C_{i,k}.$$

(c) Given a matrix M, let $\mathrm{row}(M)$ denote its row space (over $\mathbb{R}$, say). Show that the row spaces of our matrices satisfy the following relations:

  (i) $\mathrm{row}(C_{i,k}) \le \mathrm{row}(C_{j,k})$ for $i \le j \le k$.

  (ii) $\mathrm{row}(D_{i,j}) \le \mathrm{row}(D_{j,k})$ for $i \le j \le n-k$.

  (iii) $\mathrm{row}(C_{j,k}) = \mathrm{row}(D_{j,k})$ for $j \le k \le n-j$.

(d) Since $C_{k,k}$ is the identity, part (i) above gives us a nested sequence of subspaces $\mathrm{row}(C_{0,k}) \le \mathrm{row}(C_{1,k}) \le \ldots \le \mathrm{row}(C_{k,k}) = \mathbb{R}^{\binom{n}{k}}$. Thus, if we let $U_0 = \mathrm{row}(C_{0,k})$ and, for each $1 \le i \le k$, let U_i be the orthogonal complement of $\mathrm{row}(C_{i-1,k})$ in $\mathrm{row}(C_{i,k})$, we obtain an orthogonal decomposition $\mathbb{R}^{\binom{n}{k}} = U_0 \oplus U_1 \oplus \ldots \oplus U_k$.

Complete the proof by showing, for each $0 \le i \le k$, that U_i is an eigenspace of $KG(n,k) = D_{k,k}$ with eigenvalue $\lambda = (-1)^i \binom{n-k-i}{k-i}$

[1] The "proof by exercise" is often a good accompaniment to the "proof by example," I find.

[2] If you think the only reason these identities are separated from those in part (a) is because I couldn't fit all four identities in a single line, you are quite correct.

Supersaturation

One consequence of the Erdős–Ko–Rado Theorem is that, when $n \geq 2k$, any family $\mathcal{F} \subseteq \binom{[n]}{k}$ of size larger than $\binom{n-1}{k-1}$ must contain pairs of disjoint sets. It is an open problem to determine the minimum number of such pairs that must appear in these large families.[3] In the following exercises, you will use the results from the tutorial to provide some initial answers to this question.

Exercise 5 Let $n > 2k$, and let $\mathcal{F} \subseteq \binom{[n]}{k}$ be a k-uniform set family.

(a) Let $\varepsilon > 0$ be such that $\varepsilon = o\left(\frac{n-2k}{n}\right)$. Deduce from the robust stability theorem[4] that if $|\mathcal{F}| \geq (1+\varepsilon)\binom{n-1}{k-1}$, then $\mathcal{F}$ must contain at least $\frac{1}{2}\varepsilon\binom{n-1}{k-1}\binom{n-k-1}{k-1}$ pairs of disjoint sets.

(b) Using the Hilton-Milner Theorem, sharpen the estimate from part (a) in the case when $|\mathcal{F}| = \binom{n-1}{k-1} + 1$ by showing that any such family contains at least $\binom{n-k-1}{k-1}$ pairs of disjoint sets. Give an example to show this bound is best possible.

Exercise 6 Combine the robust stability lemma with some combinatorial arguments to extend the exact result of 5(b) to a wider range. Specifically, prove that if $n \geq n_0(k)$ and $\varepsilon \leq \varepsilon_0(n,k)$, then any family $\mathcal{F} \subseteq \binom{[n]}{k}$ of size at least $(1+\varepsilon)\binom{n-1}{k-1}$ contains at least $\varepsilon\binom{n-1}{k-1}\binom{n-k-1}{k-1}$ pairs of disjoint sets. How small can you make n_0, and how large can you make ε_0?

Sparse Erdős–Ko–Rado

In the Erdős–Ko–Rado Theorem, one is free to choose any set from $\binom{[n]}{k}$ to be in the set family, provided one avoids all pairs of disjoint sets. In line with recent trends in combinatorial research, one might seek to extend the theorem to the sparse random setting. One such direction was proposed by Bollobás, Narayanan and Raigorodskii, who asked when larger families can be constructed if one only wishes to avoid a random subcollection of the pairs of disjoint sets.

[3] A conjecture of Bollobás and Leader [B. Bollobás and I. Leader, Set systems with few disjoint pairs, *Combinatorica* **23** (2003), 559–570.] suggests that the optimal family should always belong to a group of constructions called ℓ-balls. An asymptotic result in this direction was provided by Frankl, Kohayakawa and Rödl [P. Frankl, Y. Kohayakawa and V. Rödl, A note on supersaturated set systems, *European J. Combin.* **51** (2016), 190–199.], and exact results for some ranges of parameters can be found in joint work with Gan and Sudakov [S. Das, W. Gan and B. Sudakov, The minimum number of disjoint pairs in set systems and related problems, *Combinatorica* **36** (2016), 623–660.] and with Balogh, Liu, Sharifzadeh and Tran [J. Balogh, S. Das, H. Liu, M. Sharifzadeh and T. Tran, Structure and supersaturation for intersecting families, *Electron. J. Comb.* **26** (2019), P2.34.]. However, the latter paper also contains a counterexample to the conjecture, suggesting that the truth may be considerably more complex.

[4] That is, the final theorem presented in the tutorial.

The sparsified problem is perhaps more naturally stated in terms of the Kneser Graph $KG(n,k)$.[5] Recall that the Erdős–Ko–Rado Theorem asserts that, when $n \geq 2k$, we have $\alpha\left(KG(n,k)\right) = \binom{n-1}{k-1}$. Now, given $p \in [0,1]$, let $KG(n,k)_p$ denote the random spanning subgraph of $KG(n,k)$, where each edge is retained independently with probability p. Since $KG(n,k)_p \subseteq KG(n,k)$, we clearly have $\alpha\left(KG(n,k)_p\right) \geq \binom{n-1}{k-1}$, and we have equality when $p = 1$. The problem is to determine how small p can be while still maintaining equality.[6]

Exercise 7 Show that[7] there is some absolute constant $c > 0$ such that, if $k, n \to \infty$ with $k \leq cn$, then, if we define

$$p_0 = \frac{\ln\left(n\binom{n-1}{k}\right)}{\binom{n-k-1}{k-1}},$$

for every $\varepsilon > 0$ we have

(a) if $p > (1+\varepsilon)p_0$, then with high probability $\alpha\left(KG(n,k)_p\right) = \binom{n-1}{k-1}$, and

(b) if $p < (1-\varepsilon)p_0$, then with high probability $\alpha\left(KG(n,k)_p\right) \geq \binom{n-1}{k-1} + 1$.

Bonus What if we do not require that the Kneser subgraph is random? How sparse can a spanning subgraph $G \subseteq KG(n,k)$ be if we still require that $\alpha(G) = \binom{n-1}{k-1}$?

[5] Indeed, it is a little harder to phrase this problem in terms of intersecting set families. A different question that makes more sense in that setting was introduced by Balogh, Bohman and Mubayi [J. Balogh, T. Bohman and D. Mubayi, Erdős–Ko–Rado in random hypergraphs, *Comb. Probab. Comput.* **18** (2009), 629–646.]. Let $\binom{[n]}{k}_p$ denote the random k-uniform hypergraph, where every k-set survives independently with probability p. The question is then to determine what the largest intersecting subfamilies $\mathcal{F} \subseteq \binom{[n]}{k}_p$ are. As every subfamily of a star is still intersecting, we expect to find intersecting families of size $p\binom{n-1}{k-1}$, but could there be larger families? Following the initial paper, results for this problem were provided by Gauy, Hàn and Oliviera [M. M. Gauy, H. Hàn and I. C. Oliviera, Erdős–Ko–Rado for random hypergraphs: asymptotics and stability, *Comb. Probab. Comput.* **26** (2017), 406–422.], Hamm and Kahn [A. Hamm and J. Kahn, On Erdős–Ko–Rado for random hypergraphs II, *Comb. Probab. Comput.* **28** (2019), 61–80.], and in joint work with Balogh, Delcourt, Liu and Sharifzadeh [J. Balogh, S. Das, M. Delcourt, H. Liu and M. Sharifzadeh, Intersecting families of discrete structures are typically trivial, *J. Comb. Theory Ser. A* **132** (2015), 224–245.]. It remains open to determine the threshold for when the stars give the right answer for larger k; as hinted at in the tutorials, this problem seems to become difficult when $k = \tilde{\Omega}\left(n^{1/2}\right)$.

[6] It is perhaps worth noting that this problem was the original motivation behind our development of the robust stability theorem. In the original paper [B. Bollobás, B. P. Narayanan and A. M. Raigorodskii, On the stability of the Erdős–Ko–Rado theorem, *J. Combin. Theory Ser. A* **137** (2016), 64–78.], the threshold probability was determined for $k = o\left(n^{1/3}\right)$. A subsequent paper of Balogh, Bollobás and Narayanan [J. Balogh, B. Bollobás and B. P. Narayanan, Transference for the Erdős–Ko–Rado theorem, *Forum of Mathematics, Sigma* **3** (2015), e23.] considered the problem for larger k, showing that the threshold must still be very small in this range. After we [S. Das and T. Tran, Removal and stability for Erdős–Ko–Rado, *SIAM J. Discrete Math.* **30** (2016), 1102–1114.] determined the threshold for k bounded away from $\frac{1}{2}n$, Devlin and Kahn [P. Devlin and J. Kahn, On "stability" in the Erdős–Ko–Rado Theorem, *SIAM J. Discrete Math.* **30** (2016), 1283–1289.] completed the picture, solving the problem when k is close to $\frac{1}{2}n$. It remains to prove a sharp threshold when k is very large, and also to prove a sharp hitting time version of this result.

[7] And, in doing so, notice that this shows there is a tremendous amount of redundancy in the Erdős–Ko–Rado Theorem: one only needs to forbid a tiny fraction of the disjoint pairs to achieve the same result!