

V. Szemerédi's regularity lemma

1. Szemerédi's regularity lemma

The ‘regularity lemma’ of Endre Szemerédi [5] roughly asserts that, for each $\varepsilon > 0$, there exists a number k such that the vertex set V of any graph $G = (V, E)$ can be partitioned into at most k *almost* equal-sized classes so that between *almost* any two classes, the edges are distributed *almost* homogeneously. Here *almost* depends on ε . The important issue is that k only depends on ε , and not on the size of the graph.

Let $G = (V, E)$ be a graph. For nonempty $I, J \subseteq V$, let $d(I, J)$ denote the number of adjacent pairs of vertices in $I \times J$ divided by $|I \times J|$. Call the pair (I, J) ε -*regular* if for all $X \subseteq I, Y \subseteq J$:

$$(1) \quad \text{if } |X| > \varepsilon|I| \text{ and } |Y| > \varepsilon|J| \text{ then } |d(X, Y) - d(I, J)| \leq \varepsilon.$$

A partition P of V is called ε -*regular* if

$$(2) \quad \sum_{\substack{I, J \in P \\ (I, J) \text{ } \varepsilon\text{-irregular}}} |I||J| \leq \varepsilon|V|^2.$$

Moreover, P is called ε -*balanced* if P contains a subcollection C such that all sets in C have the same size and such that $|V \setminus \bigcup C| \leq \varepsilon|V|$.

For $I, J \subseteq V$, let $L_{I, J}$ be the linear subspace of $\mathbb{R}^{V \times V}$ consisting of all scalar multiples of the incidence matrix of $I \times J$ in $\mathbb{R}^{V \times V}$. For any $M \in \mathbb{R}^{V \times V}$, let $M_{I, J}$ be the orthogonal projection of M onto $L_{I, J}$ (with respect to the inner product $\text{Tr}(MN^T)$ for matrices $M, N \in \mathbb{R}^{V \times V}$). So the entries of $M_{I, J}$ on $I \times J$ are all equal to the average value of M on $I \times J$.

If P is a partition of V , let L_P be the sum of the spaces $L_{I, J}$ with $I, J \in P$, and let M_P be the orthogonal projection of M onto L_P . So $M_P = \sum_{I, J \in P} M_{I, J}$.

Define $f_\varepsilon(x) := (1 + \varepsilon^{-1})x4^x$ for $x \in \mathbb{R}$.

Lemma 1. *Let $\varepsilon > 0$ and $G = (V, E)$ be a graph, with adjacency matrix A . Then each ε -irregular partition P has an ε -balanced refinement Q with $|Q| \leq f_\varepsilon(|P|)$ and $\|A_Q\|^2 > \|A_P\|^2 + \varepsilon^5|V|^2$.*

Proof. Let $(I_1, J_1), \dots, (I_n, J_n)$ be the ε -irregular pairs in P^2 . For each $i = 1, \dots, n$, we can choose (by definition (1)) subsets $X_i \subseteq I_i$ and $Y_i \subseteq J_i$ with $|X_i| > \varepsilon|I_i|$, $|Y_i| > \varepsilon|J_i|$ and $|d(X_i, Y_i) - d(I_i, J_i)| > \varepsilon$. For any fixed $K \in P$, there exists a partition R_K of K such that each X_i with $I_i = K$ and each Y_i with $J_i = K$ is a union of classes of R_K and such that $|R_K| \leq 2^{2|P|} = 4^{|P|}$. Let $R := \bigcup_{K \in P} R_K$. Then R is a refinement of P such that each X_i and each Y_i is a union of classes of R . Moreover, $|R| \leq |P|4^{|P|}$.

Now note that for each i , since $(A_R)_{X_i, Y_i} = A_{X_i, Y_i}$ (as $L_{X_i, Y_i} \subseteq L_R$) and since A_{X_i, Y_i} and A_P are constant on $X_i \times Y_i$, with values $d(X_i, Y_i)$ and $d(I_i, J_i)$, respectively:

$$(3) \quad \|(A_R - A_P)_{X_i, Y_i}\|^2 = \|A_{X_i, Y_i} - (A_P)_{X_i, Y_i}\|^2 = |X_i||Y_i|(d(X_i, Y_i) - d(I_i, J_i))^2 > \varepsilon^4|I_i||J_i|.$$

Then negating (2) gives with Pythagoras, as A_P is orthogonal to $A_R - A_P$ (as $L_P \subseteq L_R$), and as the spaces L_{X_i, Y_i} are pairwise orthogonal,

$$(4) \quad \|A_R\|^2 - \|A_P\|^2 = \|A_R - A_P\|^2 \geq \sum_{i=1}^n \|(A_R - A_P)_{X_i, Y_i}\|^2 \geq \sum_{i=1}^n \varepsilon^4 |I_i| |J_i| > \varepsilon^5 |V|^2.$$

To obtain an ε -balanced partition Q , define $t := \varepsilon|V|/|R|$. Split each class of R into classes, each of size $\lceil t \rceil$, except for at most one of size less than t . This gives partition Q . Then $|Q| \leq |R| + |V|/t = (1 + \varepsilon^{-1})|R| \leq f_\varepsilon(|P|)$. Moreover, the union of the classes of Q of size less than t has size at most $|R|t = \varepsilon|V|$. So Q is ε -balanced. As $L_R \subseteq L_Q$, we have, using (4), $\|A_Q\|^2 \geq \|A_R\|^2 > \|A_P\|^2 + \varepsilon^5|V|^2$. \blacksquare

For $n \in \mathbb{N}$, f_ε^n denotes the n -th iterate of f_ε .

Theorem 1 (Szemerédi's regularity lemma). *For each $\varepsilon > 0$ and graph $G = (V, E)$, each partition P of V has an ε -balanced ε -regular refinement of size $\leq f_\varepsilon^{\lceil \varepsilon^{-5} \rceil}(|P|)$.*

Proof. Let A be the adjacency matrix of G . Set $P_0 = P$. For $i \geq 0$, if P_i has been set, let P_{i+1} be an ε -balanced refinement of P_i with $|P_{i+1}| \leq f_\varepsilon(|P_i|)$ and with $\|A_{P_{i+1}}\|$ maximal. As $\|A_{P_i}\|^2 \leq \|A\|^2 \leq |V|^2$ for all i , $\|A_{P_{i+1}}\|^2 \leq \|A_{P_i}\|^2 + \varepsilon^5|V|^2$ for some i with $1 \leq i \leq \lceil \varepsilon^{-5} \rceil$. Then, by Lemma 1, P_i is ε -regular. Moreover $|P_i| \leq f_\varepsilon^i(|P|) \leq f_\varepsilon^{\lceil \varepsilon^{-5} \rceil}(|P|)$. \blacksquare

It is important to observe that the bound on $|Q|$, though generally huge, only depends on ε and $|P|$, and not on the size of the graph. Gowers [1] showed that the bound necessarily is huge (at least a tower of powers of 2's of height proportional to $\varepsilon^{-1/16}$).

Exercise

- 1.1. Let P be an ε -balanced ε -regular partition of V , and let $C \subseteq P$ be such that all sets in C have the same size and such that $|V \setminus \bigcup C| \leq \varepsilon|V|$. Prove that at most $(\varepsilon/(1 - \varepsilon)^2)|C|^2$ pairs in C^2 are ε -irregular.

2. Δ -graphs

Call a graph $G = (V, E)$ a Δ -graph if each edge belongs to a unique triangle. For any n , let $\tau(n)$ be the maximum number of edges of any Δ -graph on n vertices.¹

Theorem 2. $\tau(n) = o(n^2)$.

Proof. Let $\varepsilon > 0$ with $\varepsilon < \frac{1}{2}$. Set $k_\varepsilon := f_\varepsilon^{\lceil \varepsilon^{-5} \rceil}(1)$. It suffices to prove:

$$(5) \quad \text{Let } G = (V, E) \text{ be a } \Delta\text{-graph with } n := |V| \geq k_\varepsilon/\varepsilon^3. \text{ Then } |E| \leq 12\varepsilon n^2.$$

By Szemerédi's regularity lemma, V has an ε -regular partition P with $|P| \leq k_\varepsilon$. Let F be the set of edges uv for which there exists $(X, Y) \in P$ with $u \in X$, $v \in Y$ such that (X, Y) is ε -irregular, or $d(X, Y) \leq 2\varepsilon$, or $|X| < \varepsilon^{-2}$. We claim

¹ $f(n) = o(g(n))$ means $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$.

(6) $E \setminus F$ contains no triangles.

If not, there exist (not necessarily distinct) $X_1, X_2, X_3 \in P$ such that (X_i, X_j) is ε -regular and $d(X_i, X_j) > 2\varepsilon$ for all distinct $i, j = 1, 2, 3$ and such that $|X_3| \geq \varepsilon^{-2}$. Let U be the set of vertices in X_1 with at most $\varepsilon|X_2|$ neighbours in X_2 . Then $d(U, X_2) \leq \varepsilon < d(X_1, X_2) - \varepsilon$, so, since (X_1, X_2) is ε -regular, $|U| \leq \varepsilon X_1$. So less than half of the vertices in X_1 have at most $\varepsilon|X_2|$ neighbours in X_2 . Similarly, less than half of the vertices in X_1 have at most $\varepsilon|X_3|$ neighbours in X_3 . Hence there exists a vertex $u \in X_1$ with more than $\varepsilon|X_2|$ neighbours in X_2 and more than $\varepsilon|X_3|$ neighbours in X_3 . Let U_i be the set of neighbours of u in X_i ($i = 2, 3$). As G is a Δ -graph, the edges spanned by $U_2 \cup U_3$ form a matching, and so $e(U_2, U_3) \leq |U_2|$. So $|U_3|^{-1} \geq d(U_2, U_3) \geq d(X_2, X_3) - \varepsilon > \varepsilon$. Hence $|U_3| < \varepsilon^{-1}$. Therefore, $|X_3| < \varepsilon^{-1}|U_3| < \varepsilon^{-2}$, a contradiction. This proves (6).

Next we show:

(7) $|F| \leq 4\varepsilon n^2$.

The number of edges connecting any ε -irregular pair (X, Y) is at most εn^2 , by (2). The number of edges spanned by those $(X, Y) \in P^2$ with $d(X, Y) \leq 2\varepsilon$ is at most $\sum_{X, Y \in P} 2\varepsilon|X||Y| \leq 2\varepsilon n^2$. The number of edges intersecting those X_i with $|X_i| < \varepsilon^{-2}$ is at most $k_\varepsilon \varepsilon^{-2} n \leq \varepsilon n^2$. So we have (7).

By (6), each triangle of G contains an edge in F . Hence, by (7), G has at most $4\varepsilon n^2$ triangles, and hence, as G is a Δ -graph, at most $12\varepsilon n^2$ edges. This proves (5). \blacksquare

Note that ε -balancedness of partition P of V is not used in this proof.

3. Arithmetic progressions

An *arithmetic progression of length k* is a sequence of numbers a_1, \dots, a_k with $a_i - a_{i-1} = a_2 - a_1 \neq 0$ for $i = 2, \dots, k$. For any k and n , let $\alpha_k(n)$ be the maximum size of a subset of $[n]$ containing no arithmetic progression of length k . (Here $[n] := \{1, \dots, n\}$.)

We can now derive the theorem of Roth [3], which implies that any set X of natural numbers with $\limsup_{n \rightarrow \infty} |X \cap [n]|/n > 0$ contains an arithmetic progression of length 3.

Corollary 2a. $\alpha_3(n) = o(n)$.

Proof. We show that $\alpha_3(n) \leq \tau(9n)/3n$. Then Theorem 2 gives the corollary.

Choose $S \subseteq [n]$ with $|S| = \alpha_3(n)$ such that S contains no arithmetic progression of length 3. Let $V := [3n] \times [3]$ and for $i \in [n]$ and $s \in S$, let $T_{i,s}$ be the triangle spanned by $(i, 1), (i + s, 2), (i + 2s, 3)$. Let E be the set of edges spanned by these $T_{i,s}$. We show that $G = (V, E)$ is a Δ -graph.

Let T be any triangle in E . Let T be spanned by $(i, 1), (j, 2)$ and $(k, 3)$. Then $j = i + s$, $k = j + t$, and $k = i + 2u$ for some $s, t, u \in S$. So $u = \frac{1}{2}(s + t)$. If $T \neq T_{i,s}$, then $t \neq s$, and hence s, u, t is an arithmetic progression of length 3, contradicting our assumption. So G is a Δ -graph.

Therefore, $3n\alpha_3(n) = 3n|S| = |E| \leq \tau(|V|) = \tau(9n)$. \blacksquare

This was extended to $\alpha_k(n) = o(n)$ for any k by Szemerédi [4]. Recently, Green and Tao [2] proved that there exist arbitrarily long arithmetic progressions of primes.

References

- [1] W.T. Gowers, Lower bounds of tower type for Szemerédi's uniformity lemma, *Geometric and Functional Analysis* 7 (1997) 322–337.
- [2] B. Green, T. Tao, The primes contain arbitrarily long arithmetic progressions, *Annals of Mathematics* (2) 167 (2008) 481–547.
- [3] K. Roth, Sur quelques ensembles d'entiers, *Comptes Rendus des Séances de l'Académie des Sciences Paris* 234 (1952) 388–390.
- [4] E. Szemerédi, On sets of integers containing no k elements in arithmetic progression, *Acta Arithmetica* 27 (1975) 199–245.
- [5] E. Szemerédi, Regular partitions of graphs, in: *Problèmes combinatoires et théorie des graphes* (Proceedings Colloque International C.N.R.S., Paris-Orsay, 1976) [Colloques Internationaux du C.N.R.S. N° 260], Éditions du C.N.R.S., Paris, 1978, pp. 399–401.