of these primes up to a large integer N is exactly what one expects according to the Prime Number Theorem, namely $\approx \frac{1}{\phi(k)} \frac{N}{\log N}$, where $\phi(k)$ is Euler's function, the number of positive integers up to k that are relatively prime to k. However, for us this is still not enough: we need that in *every small enough interval there is such a prime.* Such type of statements are the topic of intensive research in number theory. In particular by a result of Huxley and Iwaniec (1975) for every sufficiently large n there exists a prime $p \equiv 1 \pmod s$ such that $\sqrt{sn} - n^{1/3} < p < \sqrt{sn}$.

2. Selecting $s-1 = p-1$ would give us that *any* nonzero multiple of (a,b) is equivalent with it, that is $V(G)$ is the 1-dimensional projective space (line). The defined graph is the clique, which is then $K_{2,p}$-free indeed, since the projective line contains only $p+1$ points.

3. From the analysis of the two cases in the proof of Theorem 2.3 it could seem at first sight that, since Case 2 leads to a contradiction, every two vertices have exactly $s-1$ common neighbors. This is false however, and not only because of the presence of the loops, but also because Case 2 leads to a contradiction only if there was a common neighbor in the first place. It could happen, though rarely, that two vertices do not have a common neighbor at all. This occurs only if $(a,b) = (\lambda a, \lambda b)$ with some $\lambda \notin H$.

2.2 Forbidding $K_{3,s}$

The value of $ex(n, K_{t,s})$ for $t = 2$ is fully settled now up to lower order terms. Before going further to $t \geq 3$ we repeat our randomized approach for obtaining a dense H-free graph for arbitrary H. In particular, for $H = K_{t,s}$ we will check how the random attempt relates to the Kővári-Sós-Turán upper bound.

2.2.1 A quick detour to random graphs

The plan is again to find a probability p as large as possible such that in the random graph $G(n,p)$ we have

$$\mathbf{E}[e(G(n,p)) - \#\{\text{copies of } H \text{ in } G(n,p)\}] \geq \frac{1}{2}\binom{n}{2}p.$$

This definitely holds if

$$\frac{1}{2}\binom{n}{2}p \geq n^{n(H)}p^{e(H)} \geq \mathbf{E}[\#\{\text{copies of } H \text{ in } G(n,p)\}],$$

or if $p = O\left(n^{-\frac{n(H)-2}{e(H)-1}}\right)$.

Hence, we know that with some $p = cn^{-\frac{n(H)-2}{e(H)-1}}$ there exists a graph G such that after deleting one edge from every copy of H there are still $\frac{1}{2}\binom{n}{2}p = \Theta\left(n^{2-\frac{n(H)-2}{e(H)-1}}\right)$ edges left. We proved the following.

Proposition 2.4 *For any H,*

$$ex(n,H) = \Omega\left(n^{2-\frac{n(H)-2}{e(H)-1}}\right).$$

Applying this for $H = K_{t,s}$ we obtained a $K_{t,s}$-free graph with $\Theta\left(n^{2-\frac{t+s-2}{ts-1}}\right)$ edges. Observe, however, that $\frac{t+s-2}{ts-1}$ is *strictly* larger than $\frac{1}{t}$ for any $s \geq t \geq 2$. Hence the order of the random lower bound on $ex(n, K_{t,s})$ is *always smaller* than the order of the Kövári-Sós-Turán upper bound. So we are still out there looking for those needles in the haystacks...

2.2.2 An infinite construction

Similarly to our $K_{2,2}$-free construction we start our investigation of dense $K_{3,3}$-free graphs by an infinite geometric construction providing a good heuristic.

Motivated by our approach to the unit distance problem in the euclidean 3-space (see Exercise 2.5) we take the points of $\mathbb{R}^3$ as the vertices of our graph, and make two vertices connected with an edge if their distance is exactly one. That is, with the notation of the previous section, we consider the unit-distance graph of the 3-dimensional euclidean space. The set of neighbors of a vertex is the set of all points on the unit sphere centered in that vertex. If there were a $K_{3,3}$ in the graph, then three unit spheres would intersect in at least three points. The intersection of two unit spheres (unless it is empty) is a circle with radius strictly less than one (this radius might be 0, in case the spheres intersect in only one point). A third unit sphere can intersect this circle in at most two points, proving that G does not contain a $K_{3,3}$.

The unit sphere minus a point is homeomorphic to the plane, so the degree of each vertex is of order $|\mathbb{R}|^2$. Since $n(G) = |\mathbb{R}|^3$, the number of edges in the graph is roughly $\frac{1}{2}n^{5/3}$.

2.2.3 A finite construction

In this section we "finitize" the above heuristic, but the transformation will be more problematic than it was in the case of the $K_{2,2}$-free construction. There it was clear right away what caused that the graph is $K_{2,2}$-free: two lines intersect in at most one point or a linear equation system has at most one solution. These are all properties carrying over to any field. Here it is not clear at this point what property (or rather imperfectness) of the field $\mathbb{R}$ causes that three unit spheres intersect in at most two points.

Theorem 2.5 (Brown, 1966)

$$ex(n, K_{3,3}) \gtrapprox \frac{1}{2}n^{5/3}.$$

Proof. We set $V(G) = \mathbb{F}_p^3$. For every $a \in \mathbb{F}_p^3$ we define

$$S_\alpha(a) = \{x \in \mathbb{F}_p^3 \ : \ (x_1 - a_1)^2 + (x_2 - a_2)^2 + (x_3 - a_3)^2 = \alpha\},$$

where $\alpha \in \mathbb{F}_p$ is a constant to be determined later. We define two vertices a and b to be adjacent in G if $b \in S_\alpha(a)$. Obviously the adjacency relation is symmetric.

Is G a $K_{3,3}$-free graph? Suppose not. Assume for contradiction that there exist distinct vertices $a, b, c \in V(G)$ with $|S(a) \cap S(b) \cap S(c)| \geq 3$. For $x \in S(a) \cap S(b) \cap S(c)$ we have

$$\begin{aligned}
(x_1 - a_1)^2 + (x_2 - a_2)^2 + (x_3 - a_3)^2 &= \alpha, & (2.1)\\
(x_1 - b_1)^2 + (x_2 - b_2)^2 + (x_3 - b_3)^2 &= \alpha, & (2.2)\\
(x_1 - c_1)^2 + (x_2 - c_2)^2 + (x_3 - c_3)^2 &= \alpha. & (2.3)
\end{aligned}$$

One obtains that for at least three different x,

$$\begin{aligned}
\text{(2.2) - (2.1):}\quad & 2x_1(a_1 - b_1) + 2x_2(a_2 - b_2) + 2x_3(a_3 - b_3) + \\
& +b_1^2 - a_1^2 + b_2^2 - a_2^2 + b_3^2 - a_3^2 = 0,\\
\text{(2.3) - (2.2):}\quad & 2x_1(b_1 - c_1) + 2x_2(b_2 - c_2) + 2x_3(b_3 - c_3) + \\
& +c_1^2 - b_1^2 + c_2^2 - b_2^2 + c_3^2 - b_3^2 = 0,\\
\text{(2.1) - (2.3):}\quad & 2x_1(c_1 - a_1) + 2x_2(c_2 - a_2) + 2x_3(c_3 - a_3) + \\
& +a_1^2 - c_1^2 + a_2^2 - c_2^2 + a_3^2 - c_3^2 = 0
\end{aligned}$$

holds.

The matrix of the last system of equations is

$$A = \begin{pmatrix} a_1 - b_1 & a_2 - b_2 & a_3 - b_3 \\ b_1 - c_1 & b_2 - c_2 & b_3 - c_3 \\ c_1 - a_1 & c_2 - a_2 & c_3 - a_3 \end{pmatrix}.$$

One can easily see that the rank of A is either 1 or 2. If $\mathrm{rk}(A) = 2$, then all solutions of the system lie on a line (we already assumed that there exists a solution). On the other hand, if $\mathrm{rk}(A) = 1$, then a, b, c lie on a line. Therefore, if there is a $K_{3,3}$ with partite classes $\{a, b, c\}$ and $\{d, e, f\}$ in G, then either a, b, c lie on a line, or d, e, f lie on a line.

Could it happen that $S_\alpha(a)$, for some α and some a, contains three points from a line? Or maybe even the full line? A sphere??? — By translation, we can assume that such a special line, containing three points of the sphere $S_\alpha(a)$, also passes through the origin. Let the line consist of the points τv where $v \in \mathbb{F}_p^3 \setminus \{(0,0,0)\}$ is fixed and $\tau \in \mathbb{F}_p^*$ is arbitrary. By our assumption

$$\tau^2 \sum_i v_i^2 - 2\tau \sum_i v_i a_i + \sum_i a_i^2 = \alpha$$

holds for at least three τ. This is only possible when

$$\sum_i v_i^2 = 0, \quad \sum_i v_i a_i = 0, \quad \sum_i a_i^2 = \alpha.$$

But then, assuming w.l.o.g. that $v_1 \neq 0$, we get

$$\begin{aligned} v_1^2\alpha &= v_1^2 \sum_i a_i^2 = v_1^2 a_1^2 + v_1^2(a_2^2 + a_3^2) = (-v_2a_2 - v_3a_3)^2 + v_1^2(a_2^2 + a_3^2) \\ &= (-v_2a_2 - v_3a_3)^2 + (-v_2^2 - v_3^2)(a_2^2 + a_3^2) \\ &= v_2^2a_2^2 + 2v_2a_2v_3a_3 + v_3^2a_3^2 - v_2^2a_2^2 - v_2^2a_3^2 - v_3^2a_2^2 - v_3^2a_3^2 = -(v_2a_3 - v_3a_2)^2, \end{aligned}$$

and

$$-\alpha = \left(\frac{v_2a_3 - v_3a_2}{v_1}\right)^2.$$

Hence if we select α such that $-\alpha$ is *not a square*, then we arrive at a contradiction and our graph does *not* contain a $K_{3,3}$. Recall that for an odd prime power p, the number of (non-zero) quadratic residues is $\frac{p-1}{2}$ which is the same as the number of quadratic non-residues.

Exercise 2.6 *Let k be an arbitrary field. Prove that if $-\alpha \in k$* is *a square, then the corresponding sphere-graph (defined in the 3-dimensional space over k) not only contains a $K_{3,3}$, but also a $K_{n^{1/3},n^{2/3}}$. (In case k is an infinite field we mean a $K_{|k|,|k|}$.) In particular, our heuristics would fail badly in the complex 3-space.*

To finish the proof of Theorem 2.5 we still need to see that the Brown-graph contains enough edges, i.e., the sphere $S_\alpha(a)$ (the neighborhood of vertex a) contains enough vertices — for a well-chosen α. For the $K_{2,2}$-free graphs of Subsection 2.1.3 this was a piece of cake, since to count edges we just had to solve linear equations. Here we have some difficulty, since the equations are quadratic. Of course intuitively we feel that the cardinality of $S_\alpha(a)$ is around $|\mathbb{F}_p|^2$ for any α, as $S_\alpha(a)$ is a *surface* in the three-dimensional space over $\mathbb{F}_p$.

Even more, let us take the average over all possible "radii" α. For any $a \in \mathbb{F}_p^3$ we have

$$\sum_{\alpha \in \mathbb{F}_p} |S_\alpha(a)| = p^3,$$

so by averaging there is at least one "radius" α for which $|S_\alpha(a)| \geq p^2$. For Brown's construction to work we need this for a somewhat special α: for which $-\alpha$ is not a square. These αs are abundant, half of the nonzero elements are such. Still, theoretically, it is possible that for all "bad radii" the corresponding spheres contain $2p^2$ points, while for all "good radii" the spheres are empty. In the following we will see that this is not the case and even under much more general circumstances, the number of points on a surface does not deviate much from the expected value. We include a probabilistically motivated proof due to Wolfgang Schmidt. The proof imitates the technique of the second moment method, used frequently in probabilistic combinatorics.

Let us fix positive integers $d_1, \ldots, d_n \in \mathbb{N}$ and consider the following equation

$$a_1x_1^{d_1} + a_2x_2^{d_2} + \cdots + a_nx_n^{d_n} = a_0, \qquad a_0, \ldots, a_n \in \mathbb{F}_p.$$

Let $N(a_0, \ldots, a_n)$ denote the number of solutions $(x_1, \ldots, x_n) \in \mathbb{F}_p^n$ of this equation. We have

$$\sum_{a_1,\ldots,a_n} \sum_{a_0 \in \mathbb{F}_p} N(a_0, \ldots, a_n) = \sum_{a_1,\ldots,a_n} p^n = p^{2n}.$$

Then, the average value of $N(a_0, \ldots, a_n)$ over $(a_0, \ldots, a_n)$ is p^{n-1}. The next theorem shows that the function $N(\cdot)$ *never* deviates from its average by much.

Theorem 2.6 *Let $c_0, c_1, \ldots, c_n \in \mathbb{F}_p \setminus \{0\}$. Then we have*

$$|N(c_0, \ldots, c_n) - p^{n-1}| \le \left[\left(\frac{p}{p-1} \right)^{n/2} \prod_{i=1}^{n} (d_i, p-1) \right] p^{(n-1)/2}.$$

Remark: 1. Estimating the number of solutions of higher degree equations over finite fields is a classic and well-studied area of mathematics full of beautiful ideas and hard theorems. There is an even more general theorem about the number of solutions to high degree equations over finite fields (due to Weil), but its proof well exceeds the possibilities of our course.
2. A more precise and elementary proof of what is needed for the Brown-graph is sought for in Exercise 2.7.
3. Assuming $c_1, \ldots, c_n \neq 0$ is necessary otherwise the theorem is (easily) not valid. The assumption $c_0 \neq 0$ is not crucial, see Exercise 2.8.
Proof. (W. Schmidt) Before starting, let us note that we can assume w.l.o.g. that $d_i | p-1$ for every $i = 1, \ldots, n$, since for every $\gamma \in \mathbb{F}_p$,

$$\#\{x \in \mathbb{F}_p \ : \ x^{d_i} = \gamma\} = \#\{x \in \mathbb{F}_p \ : \ x^{(d_i, p-1)} = \gamma\}.$$

We will look at the following sum, which is appropriate to measure the *average* deviation from the average.

$$\sum_{a_0,\ldots,a_n} (N(a_0, \ldots, a_n) - p^{n-1})^2$$

For a moment, consider N as a random variable determined by $a_0, \ldots, a_n$, which are chosen independently, uniformly at random from $\mathbb{F}_p$. Then the above sum (divided by p^{n+1}) is its variance $var(N)$, measuring how much the values of the random variable can deviate from its average; exactly what we are interested in.

Our plan is first to bound the variance from above, i.e., to show that the deviation of N from p^{n-1} is not too large *on the average*. This still does not say *anything* about the deviation of each individual term of the sum. But then we prove that N takes on only a few (in fact constantly many) values, moreover any such value occurs on a large (i.e., constant) fraction of the $(n+1)$-tuples $(c_0, \ldots, c_n)$. Hence we will conclude that $N(c_0, \ldots, c_n)$ cannot deviate from p^{n-1} too much, because otherwise N would deviate from p^{n-1} by much on a large fraction of its domain, which would imply that the variance would be too large.

$$\begin{aligned}
p^{n+1} \cdot var(N) &= \sum_{a_0,\dots,a_n} \left(N^2(a_0,\dots,a_n) - 2N(a_0,\dots,a_n)p^{n-1} + p^{2(n-1)}\right) \\
&= \sum_{a_0,\dots,a_n} N^2(a_0,\dots,a_n) - 2p^{2n} \cdot p^{n-1} + p^{n+1} \cdot p^{2(n-1)} \\
&= \sum_{a_0,\dots,a_n} N^2(a_0,\dots,a_n) - p^{3n-1}
\end{aligned}$$

We estimate the above sum of squares similarly to calculating $\sum N(a_0,\dots,a_n)$: by double-counting.

$$\begin{aligned}
\sum_{a_0,\dots,a_n} N^2(a_0,\dots,a_n) &= \sum_{a_0,\dots,a_n} \Bigg(\sum_{\substack{x_1,\dots,x_n \\ a_1x_1^{d_1}+\dots+a_nx_n^{d_n}=a_0}} 1 \Bigg)\Bigg(\sum_{\substack{y_1,\dots,y_n \\ a_1y_1^{d_1}+\dots+a_ny_n^{d_n}=a_0}} 1 \Bigg) \\
&= \sum_{x,y\in\mathbb{F}_p^n} \sum_{\substack{(a_0,\dots,a_n) \\ a_1x_1^{d_1}+\dots+a_nx_n^{d_n}-a_0=0 \\ a_1y_1^{d_1}+\dots+a_ny_n^{d_n}-a_0=0}} 1
\end{aligned}$$

The point of exchanging the summation is obvious: now, instead of having equations of high degree, we have systems of *linear* equations (with the x_i and y_i being fixed and the a_i being the variables). We denote the matrix of this homogeneous system by

$$A := \begin{pmatrix} -1 & x_1^{d_1} & \dots & x_n^{d_n} \\ -1 & y_1^{d_1} & \dots & y_n^{d_n} \end{pmatrix}.$$

If $\mathrm{rk}(A) = 2$, then there are p^{n-1} solutions $a = (a_0,\dots,a_n)$ to $Aa^T = (0,0)^T$.

If $\mathrm{rk}(A) = 1$, then there are p^n solutions $(a_0,\dots,a_n)$.

Now we only have to count how many times we encounter the second case. How often do we have $\mathrm{rk}(A) = 1$? In other words, for how many vectors $x, y \in \mathbb{F}_p^n$ do we have $x_i^{d_i} = y_i^{d_i}$ for all $i = 1,\dots,n$? For a fixed i, with $x_i \neq 0$ there are exactly d_i solutions y_i satisfying $x_i^{d_i} = y_i^{d_i}$. If $x_i = 0$, then of course $y_i = 0$ as well. In any case, for any fixed $(x_1,\dots,x_n)$ we have *at most* $d_1 \cdots d_n$ appropriate $(y_1,\dots,y_n)$ giving us a matrix A of rank 1.

The sum can then be estimated by

$$\begin{aligned}
\sum_{x,y\in\mathbb{F}_p^n} \sum_{\substack{(a_0,\dots,a_n) \\ a_1x_1^{d_1}+\dots+a_nx_n^{d_n}-a_0=0 \\ a_1y_1^{d_1}+\dots+a_ny_n^{d_n}-a_0=0}} 1 &\leq p^{2n}p^{n-1} + p^n d_1 \cdots d_n (p^n - p^{n-1}) \\
&= p^{3n-1} + p^{2n-1}(p-1)d_1 \cdots d_n.
\end{aligned}$$

Finally, we can upper bound the variance,

$$p^{n+1} \cdot var(N) = \sum_{a_0,\dots,a_n} (N(a_0,\dots,a_n) - p^{n-1})^2 \leq p^{2n-1}(p-1)d_1 \cdots d_n. \tag{2.4}$$

At this point we can conclude that the *average deviation* of $N(a_0,\ldots,a_n)$ from p^{n-1} is at most $\sqrt{d_1\cdots d_n}p^{(n-1)/2}$. This is roughly just the square-root of p^{n-1}, a promising sign.

Now consider the $(n+1)$-tuple $(c_0,\ldots,c_n)$ from the statement of the theorem. We claim that the variable N equals $N(c_0,\ldots,c_n)$ on a constant (roughly $\frac{1}{d_1\cdots d_n}$) fraction of its domain $\mathbb{F}_p^{n+1}$. Indeed,

$$N(c_0,c_1,\ldots,c_n) = N(tc_0,tc_1b_1^{d_1},\ldots,tc_nb_n^{d_n})$$

for any $t\neq 0$ and $b_1,\ldots,b_n\neq 0$, since the two equations

$$\begin{aligned} c_1x_1^{d_1}+\cdots+c_nx_n^{d_n}-c_0 &= 0\\ tc_1(b_1x_1)^{d_1}+\cdots+tc_n(b_nx_n)^{d_n}-tc_0 &= 0\end{aligned}$$

have the same number of solutions. More precisely, $(y_1,\ldots,y_n)$ is a solution of the first equation if and only if $(y_1/b_1,\ldots,y_n/b_n)$ is a solution of the second equation. It remains to determine the number of different $(n+1)$-tuples $(tc_0,tc_1b_1^{d_1},\ldots,tc_nb_n^{d_n})$. Each of $t,b_1,\ldots,b_n$ can be chosen in $p-1$ different ways, which gives $(p-1)^{n+1}$ choices altogether. However, some of the $(n+1)$-tuples they give rise to are identical. Such a thing happens if and only if $b_i^{d_i}=b_i'^{d_i}$ for every $i=1,\ldots,n$ (Here we used that $c_i\neq 0$ for $i=0,\ldots,n$). For fixed b_i there are exactly d_i such b_i', so the number of different $(n+1)$-tuples $(tc_0,tc_1b_1^{d_1},\ldots,tc_nb_n^{d_n})$ is $(p-1)^{n+1}/(d_1\cdots d_n)$. The fact that all these $(n+1)$-tuples have the same N-value $N(c_0,c_1,\ldots,c_n)$ combined with inequality (2.4) gives

$$\begin{aligned}\frac{(p-1)^{n+1}}{d_1\cdots d_n}\left(N(c_0,\ldots,c_n)-p^{n-1}\right)^2 &\leq \sum_{a_0,\ldots,a_n}\left(N(a_0,\ldots,a_n)-p^{n-1}\right)^2\\ &\leq p^{2n-1}(p-1)d_1\cdots d_n,\end{aligned}$$

which proves the theorem. □

Corollary 2.7 $|S_\alpha(a)|\geq p^2-16p$

Proof. Substitute $n=3$, $d_1=d_2=d_3=2$, $a_0=\alpha$, $a_1=a_2=a_3=1$ in Theorem 2.6, and use $(\frac{p}{p-1})^{3/2}\leq 2$. □

To complete the proof of Theorem 2.5 note that indeed the Brown graph has the claimed number of edges, since by Corollary 2.7 the degree of each vertex is (roughly) p^2. □

In fact, for the special case of $S_\alpha(a)$ the deviation from the average p^2 can be calculated exactly. Let $QR(p)$ be the set of quadratic residues of $\mathbb{F}_p$.

Exercise 2.7 *Give an elementary proof that for any $a\in\mathbb{F}_p^3$ the sphere $S_\alpha(a)$ contains either p^2-p or p^2+p points depending on whether α and -1 are quadratic residues*

or not. Note that four cases need to be considered.
For this purpose recall that the equation $x^2+y^2=\beta$, where $\beta \neq 0$ is fixed, has $p-1$ solutions $x, y \in \mathbb{F}_p$ if -1 is a quadratic residue in $\mathbb{F}_p$, and $p+1$ solutions if -1 is not a quadratic residue; furthermore, $x^2+y^2=0$ has $2p-1$ solutions if $-1 \in QR(p)$, or 1 single solution if $-1 \in QNR(p)$.
Give a general exact formula for the number of solutions to $x_1^2+\cdots+x_k^2=\beta$, for any fixed $k \in \mathbb{N}, \beta \in \mathbb{F}_p$.

Exercise 2.8 *Prove that $N(c_0, c_1, \ldots, c_n)$ cannot deviate by "much" from the average even if $c_0=0$. The point is of course how much is "much"? (The assumption that all other $c_i \neq 0$ is still needed, otherwise the problem becomes smaller dimensional and the error term gets be worse.)*

The next exercise exhibits some of the difficulties one faces when stepping from dense $K_{3,3}$-free graphs to $K_{4,4}$-free graphs.

Exercise 2.9 *A natural thought to extend the idea of the Brown graph to $K_{4,4}-$ or $K_{4,1000}-$avoiding dense graphs is the following. Instead of three dimensions let us take four, i.e., our vertex set is $\mathbb{F}_p^4$. Let the neighborhood of a vertex x be determined by a four-dimensional sphere around it, in particular a vertex y is adjacent to x if and only if $\sum_i (y_i-x_i)^2=1$. According to Theorem 2.6, our graph has roughly $cn^{7/4}$ edges — the conjectured truth. Prove, however, that this graph contains a $K_{p,p}$. Prove that even taking a higher degree surface of the form $\sum_i (y_i-x_i)^{1000}=1$ as the neighborhood of x instead of the sphere would not help us. (Note that Theorem 2.6 ensures that this graph also has roughly the correct number $cn^{7/4}$ of edges.)*

2.2.4 An upper bound

The above construction of Brown could also be called the "unit-distance graph" of $\mathbb{F}_p^3$ if we choose $\alpha=1$, which we can do if $p \equiv 3 \pmod 4$ (since then -1 is a quadratic non-residue). After this, here is where we are standing in terms of the Turán number of $K_{3,3}$:

$$\frac{1}{2}n^{5/3} \lessapprox ex(n, K_{3,3}) \lessapprox \frac{1}{2}2^{1/3}n^{5/3}.$$

In 1996, Füredi proved that the upper bound can be improved to match the lower bound, so the construction of Brown (which is from 1966) is asymptotically optimal. This is the first (and so far lone) improvement on the classical $(s-1)^{1/t}n^{2-1/t}$ upper bound (which is from 1954). Remember, this simple bound was asymptotically tight for $t=2$ and s arbitrary.

Before we start proving Füredi's upper bound we state a technical lemma which will be convenient later. Behind its artificial appearance it is quite easy: it only formalizes the same convexity calculation we made when proving the Kővári-Sós-Turán upper bound.

Lemma 2.7.1 *Let $v, k \geq 1$ be integers, and let $c, x_0, x_1, \ldots, x_k \geq 0$ be integers. Then*

$$\sum_{i=1}^{v} \binom{x_i}{k} \leq c\binom{x_0}{k}$$

implies

$$\sum_{i=1}^{v} x_i \leq x_0 c^{1/k} v^{1-1/k} + (k-1)v.$$

Proof. The small technical hurdle in proving this lemma is that the function $\binom{x}{k}$ is not convex on the full interval $[0, \infty)$. Hence we define $\widetilde{\binom{x}{k}}$ to be zero if $x \leq k-1$ and to be equal to $\binom{x}{k}$ otherwise. With this new definition $\widetilde{\binom{x}{k}}$ agrees with $\binom{x}{k}$ on integers, but it is a convex function not only on $\{x : x \geq k-1\}$, but also on $\{x : x \geq 0\}$.

So our assumption can be written as

$$\sum_{i=1}^{v} \widetilde{\binom{x_i}{k}} \leq c\binom{x_0}{k}$$

Case 1. If $\sum x_i \leq (k-1)v$, then the statement is clear (the "error term" takes care).
Case 2. Otherwise, let $S = \frac{\sum x_i}{v} > k-1$. Then by convexity and Jensen's inequality

$$v\widetilde{\binom{S}{k}} \leq \sum_{i=1}^{v} \widetilde{\binom{x_i}{k}} \leq c\binom{x_0}{k}.$$

Our assumption for Case 2 ensures that $\widetilde{\binom{S}{k}} = \binom{S}{k}$, so we have

$$v\frac{S(S-1)\cdots(S-k+1)}{k!} \leq c\frac{x_0(x_0-1)\cdots(x_0-k+1)}{k!}.$$

Estimating both sides trivially, we have

$$v(S-k+1)^k \leq cx_0^k.$$

Solving this inequality for S we obtain

$$S \leq c^{1/k} v^{-1/k} x_0 + k - 1,$$

which is exactly what we wanted. □

Let us now turn to the proof of Füredi. Once we decide that Brown's construction should be asymptotically optimal and we try to prove that, we better first find the inaccuracy of the Kővári-Sós-Turán argument. That argument has two estimates: The one applying Jensen's inequality is tight for all graphs that are more or less regular, and the Brown graph is. So the only problem can occur with the crucial combinatorial idea, when we say that in a $K_{3,3}$-free graph each triple of vertices can be the endpoints of at most *two* $K_{1,3}$. This in fact fails badly in Brown's graph, where roughly half of the triples is not the endpoint of *any* $K_{1,3}$ at all! Such a phenomenon is not so surprising once we think back to our motivating infinite $K_{3,3}$-free unit-distance graph in $\mathbb{R}^3$. There the number of neighbors of a triple depends on the radius of its circumcircle: if the radius is less than 1, then there are two common neigbors, if the radius is larger than 1, then there are no common neighbors, and only those "degenerate" triples have one common neighbor whose circumcircle has radius exactly one.

Exercise 2.10 *Prove that in Brown's graph roughly half of the triples has two common neighbors and the other half has none. Even more: describe explicitly those triples of the Brown graph which do not have a common neighbor!*

Keeping this in mind, we will repeat the Kővári-Sós-Turán-argument *only for the triples, that do have a common neighbor*, i.e. have a chance to have two of them. This turns out to be the key observation we need to improve on the upper bound.

Theorem 2.8 (Füredi, 1996)

$$ex(n, K_{3,3}) \lessapprox \frac{1}{2}n^{5/3}.$$

Proof. Suppose G is a $K_{3,3}$-free graph. We fix a vertex $x \in V(G)$. There are $\binom{d(x)}{3}$ triples of vertices in the neighborhood of x. Since G is $K_{3,3}$-free, each of these triples can be fully contained in the neighborhood of at most *one other* vertex. This implies the following inequalities,

$$\sum_{y \neq x} \binom{|N(x) \cap N(y)|}{3} \leq \binom{d(x)}{3}, \quad \forall\, x \in V(G).$$

Applying Lemma 2.7.1 we have

$$\sum_{y \neq x} |N(x) \cap N(y)| \leq d(x)(n-1)^{2/3} + 2(n-1).$$

Summing up over all x we get

$$\sum_{x \in V(G)} \sum_{y \neq x} |N(x) \cap N(y)| \leq \sum_{x \in V(G)} \left[d(x)(n-1)^{2/3} + 2(n-1)\right].$$

The left side counts twice the cherries $K_{1,2}$ in G by their endpoints. Counting at their midpoints and introducing $\bar{d}(G) = \sum_{x \in V(G)} d(x)/n$ for the average degree, we have

$$2 \cdot \sum_{z \in V(G)} \binom{d(z)}{2} \leq n\bar{d}(G)(n-1)^{2/3} + 2(n-1)n.$$

By Jensen's inequality we infer

$$2n\binom{\bar{d}(G)}{2} \leq n\bar{d}(G)(n-1)^{2/3} + 2(n-1)n,$$

implying

$$\bar{d}(G) - 1 \leq (n-1)^{2/3} + \frac{2(n-1)}{\bar{d}(G)}.$$

If $\bar{d}(G) \leq (n-1)^{2/3}$, then we are done. Otherwise by the above we have

$$\bar{d}(G) \leq (n-1)^{2/3} + \frac{2(n-1)}{\bar{d}(G)} + 1 \leq (n-1)^{2/3} + 2(n-1)^{1/3} + 1.$$

Hence, with $e(G) = n\bar{d}(G)/2$ we concluded the proof of

$$ex(n, K_{3,3}) \leq \frac{1}{2}n^{5/3} + n^{4/3} + \frac{1}{2}n.$$

□

Exercise 2.11 *Improve the KST-upper-bound a bit. Show that for arbirary $s \geq 3$, we have $ex(n, K_{3,s}) \lessapprox \frac{\sqrt[3]{s-2}}{2}n^{5/3}$.*

Exercise 2.12 *Generalize the proof above and show that $ex(n, K_{4,4}) \lessapprox \frac{1}{2}n^{7/4}$. (Hint: Instead of lower bounding $\sum \binom{x_i}{3}$ in terms of $(\sum x_i)^3$ (which follows from the convexity of $\binom{x}{3}$) you might want to bound it from below in terms of the product of $\sum \binom{x_i}{2}$ and $\sum x_i$.)*

Open Problem. The asymptotics of $K_{3,s}$ is not known for any $s > 3$. There are infinitely many values of s for which the upper and lower bounds are within a constant factor of $\sqrt[3]{2}$ of each other (we will discuss these results later), but there are also infinitely many values s where this constant factor separation is $\sqrt[3]{s-2}$.
Any improvement would be very interesting. The value of $ex(n, K_{3,4})$ is the first unknown.

2.3 Forbidding $K_{t,f(t)}$

So far in all we've seen on the front of dense $K_{t,s}$-free constructions, the smaller of the parameters t and s was at most 3. So what does become so hard when t and s are both at least 4? Phrasing it mysteriously, besides us being not creative enough, the problem is that $4 = 2 + 2$.

Typical $K_{t,s}$-free constructions live in the t-dimensional space (over a finite field). The vertex set is usually chosen to be the space itself, and the neighborhood of each vertex