TWENTY-TWO IMPROVED LOWER BOUNDS FOR THE ZARANKIEWICZ NUMBERS $z(m, n; 3, 3)$

KOREA SUPERINTELLIGENCE LABS

Abstract. The Zarankiewicz number $z(m, n; s, t)$ is the largest number of ones in an $m \times n$ zero–one matrix with no all-ones $s \times t$ submatrix. For $s = t = 3$ the exact values are known only in a small range, and the state of the art is a table of lower and upper bounds that several groups have been improving through 2026. We give twenty-two new lower bounds in the window $11 \leq m \leq 16$, $18 \leq n \leq 23$, each certified by an explicit $K_{3,3}$-free matrix, with gains of +1 to +8 over the best previously recorded value; among them are improvements at all five cells treated in the most recent paper on the subject. We also machine-check, for all s and t, the two structural lemmas the area runs on — monotone padding $z(m, n; s, t) \leq z(m, n+1; s, t)$ and minimum-line deletion $z(m, n; s, t) \leq k \Rightarrow z(m, n+1; s, t) \leq k + \lfloor k/n \rfloor$, both of them published at this generality — and record the three upper bounds the second of these gives. Everything is verified by the Lean 4 kernel: no statement in the development delegates any step to compiled evaluation, and the reduction that makes this affordable, from a $2^m \cdot 2^n$ subset enumeration to an $n^3 m$ scan over column triples, is itself proved rather than assumed. Two cells are now bracketed within three, $120 \leq z(13, 19; 3, 3) \leq 122$ and $137 \leq z(16, 18; 3, 3) \leq 140$, and we argue that closing either needs an upper-bound argument rather than more search.

1. Introduction

For positive integers m, n, s, t let $z(m, n; s, t)$ denote the largest number of ones in an $m \times n$ matrix with entries in $\{0, 1\}$ containing no all-ones $s \times t$ submatrix, that is, no s distinct rows and t distinct columns whose st common entries are all one. Equivalently $z(m, n; s, t)$ is the largest number of edges in a bipartite graph with parts of sizes m and n containing no $K_{s,t}$ with its s-side inside the part of size m. Determining these numbers is the problem Zarankiewicz posed in 1951 [15]; the classical asymptotic upper bound is Kővári–Sós–Turán [10], and for $s = t = 2$ Reiman [12] sharpened it to

$$z(n, n; 2, 2) \;\leq\; \tfrac{n}{2}\bigl(1 + \sqrt{4n - 3}\bigr), \tag{1}$$

with equality exactly when $n = q^2 + q + 1$ and the matrix is the incidence matrix of a projective plane of order q (as stated in Theorem 3.1 of [5]).

The finite table is a different matter from the asymptotics. For $s = t = 2$ the balanced diagonal $z(n, n; 2, 2)$ has been tabulated since Guy [8]; the exact values are published for $n \leq 21$ [7] and, on Afzaly–McKay data [2] that [4] reports and tabulates, for $n \leq 31$, while [14] and [11] tabulate independently up to $n \leq 24$. For $s = t = 3$ much less is known: exact values run out quickly, and what the literature maintains is a grid of lower and upper bounds. Tan [14] prints such a grid for $3 \leq m \leq 16$, $3 \leq n \leq 23$; Collins, Riasanovsky, Wallace and Radziszowski [4] print a tighter one for $6 \leq m \leq n \leq 18$; and Davies, Gill and Horsley [6] improve the upper half over part of that range. The lower half of the grid is where the activity has been. In 2026 Bhan, Nobili and Langer [3] reported three new exact values, $z(11, 21; 3, 3) = 116$, $z(11, 22; 3, 3) = 121$ and $z(12, 22; 3, 3) = 132$, and lower bounds for 41 further cells, found by an evolutionary search driven by a large language model. Hou [9] and Afrasyab [1] then closed several cells of the frontier exactly, and Saurabh [13] improved five lower bounds — three by search and two by padding — over the values [3] compiled.

The open question this table poses is completely concrete: for each cell (m, n) in the window where lower and upper bound disagree, which is the truth? Nobody has an argument for most of these cells in either direction, and progress has come in single edges. The present paper contributes to the lower half, and does two other things.

What is settled here.

- **Twenty-two improved lower bounds**, in the window $11 \le m \le 16$, $18 \le n \le 23$, each certified by one explicit matrix. The gains over the best value we could locate in the literature range from +1 to +8; Table 1 lists them cell by cell, with the source of the value being improved. These include all five cells of [13], improved by +2, +2, +1, +5 and +6 in the order that paper states them.
- **The two structural lemmas of the area, machine-checked.** Monotone padding — appending an all-zero line — is Lemma 2 of [13]; minimum-line deletion is the Density Lemma, Lemma 3 of [4], which that paper traces back to Section 12 of [8], and which is restated as Lemma 2.1 of [1] and run by hand in [9] at $(13, 18)$. Neither is new, and both are already published at this generality; what is added here is that both are proved inside the proof assistant, and that we record what the second gives when composed with the exact values established in 2026: $z(13, 19; 3, 3) \le 122$ in place of the published 125, $z(16, 18; 3, 3) \le 140$ by a route that does not use the table of [4], and $z(16, 19; 3, 3) \le 147$ in place of 152.
- **A machine-checked development with no compiled evaluation in it.** Every statement below is proved in Lean 4, and every one of the twenty-six witness matrices is reduced by the Lean kernel itself. Checking a 16×23 witness means inspecting $\binom{23}{3} = 1771$ column triples against 16 rows inside the kernel. This is affordable only because the freeness condition is first replaced by a decidable criterion of size $n^3 m$ and the replacement is *proved* correct: the condition as defined quantifies over subsets of rows and of columns, and the decision procedure that reads straight off the definition enumerates $2^m \cdot 2^n$ of them.

Section 2 fixes notation and states the two lemmas. Section 3 sets out the $s = t = 2$ material the development rests on, all of it published. Section 4 reproduces the published $s = t = 3$ record layer, which is the control on our own verification. Section 5 is the new lower bounds. Section 6 states, as precisely as we can, how the comparison values of Table 1 were obtained, including one step that rests on reading a bitmap. Section 7 gives the negative controls, Section 8 the two cells we could not close, and Section 9 the verification.

What is not claimed. The twenty-two matrices were found by a kicked-greedy local search, which is the method [13] describes and a standard one. We claim no new method and no new theory: these are records for a public table, obtained cheaply. The total measured search cost was 2.01 CPU-hours. Nor do we claim any of the twenty-two values is optimal; each is a lower bound, and Table 1 shows how far the published upper bound still is. Finally, our comparison column is a claim about the literature, not a theorem: it says what we found, and Section 6 says how we looked.

2. Preliminaries

Definition 2.1. Fix $m, n, s, t \in \mathbb{N}$. A matrix $M : \{1, \ldots, m\} \times \{1, \ldots, n\} \to \{0, 1\}$ *contains* an all-ones $s \times t$ submatrix if there are sets R of s rows and C of t columns with $M_{ij} = 1$ for all $i \in R$, $j \in C$; M is (s, t)-*free* otherwise. Write $\|M\|$ for the number of ones in M. Then

$$z(m, n; s, t) \;=\; \max\{\, \|M\| : M \text{ is } (s, t)\text{-free}\,\}.$$

Because the maximisation is over a finite set we work throughout with its two halves separately, and never with the maximum itself. We say k is a *lower bound* at $(m, n; s, t)$ if some (s, t)-free M has $\|M\| \ge k$, and an *upper bound* if $\|M\| \le k$ for every (s, t)-free M. Every statement in this paper is one of the two, or a negation of one; an exact value is the conjunction. Keeping the halves apart is what lets an unfinished cell still refute a wrong claim, which is the point of Section 7.

Note that R and C are *sets*: the s rows are pairwise distinct, and so are the t columns. This is the hypothesis the whole subject turns on — were a single row with t ones allowed to serve as all s rows, no admissible matrix could carry t ones in any line and every value would collapse to $\min(m, n)$ — and Section 7 checks that the distinctness is doing work.

2.1. **A decidable criterion for $s = t = 3$.** Deciding $(3, 3)$-freeness from the definition means quantifying over subsets of rows and of columns. The following reformulation scans ordered triples of columns instead. For columns a, b, c let $\mathrm{rows}(a, b, c) = \{\, i : M_{ia} = M_{ib} = M_{ic} = 1 \,\}$.

Proposition 2.2. *A matrix M is $(3, 3)$-free if and only if $\big|\mathrm{rows}(j_1, j_2, j_3)\big| \le 2$ for every triple of columns $j_1 < j_2 < j_3$.*

Proof sketch. If some R, C witness an all-ones 3×3 submatrix, write $C = \{a, b, c\}$ with a, b, c distinct; then $R \subseteq \mathrm{rows}(a, b, c)$ and $|R| = 3$. Conversely, three rows covering an increasing triple of columns are exactly such a witness. The only work is that the criterion is stated for increasing triples while the witness supplies an unordered set, so the bound must be transported across the six orderings; stating the intermediate bound about the *set* $\{a, b, c\}$ rather than about the ordered triple makes each of the six a one-line rewrite. □

Proposition 2.2 is what makes the checks in Sections 4–5 affordable: the criterion is an $n^3 m$ scan where the definition is a $2^m \cdot 2^n$ enumeration. It is a restatement of the definition and carries no mathematical content, but everything decided below is decided through it, so it is proved rather than assumed.

2.2. **Padding and deletion.**

Lemma 2.3 (Monotone padding)**.** *Let $s > 0$. If k is a lower bound at $(m, n; s, t)$ then k is a lower bound at $(m, n + 1; s, t)$.*

Proof. Let M be (s, t)-free with $\|M\| \ge k$ and let M' be M with an all-zero column appended. Appending a zero column does not change the number of ones. If R, C witnessed an all-ones $s \times t$ submatrix of M' then, since $s > 0$, some row $i \in R$ has $M'_{ij} = 1$ for every $j \in C$, so C misses the new column; transporting R and C back to M gives a witness there. □

Lemma 2.3 is folklore, and is the first half of Lemma 2 of [13], which states it — for all positive m, n, s, t, with the same proof — together with its row form $z(m, n; s, t) \le z(m + 1, n; s, t)$. Only the column form is in our formal development; the row form is used once below, in Section 6, and only to grade a literature baseline. The upper-bound twin of padding is the deletion step, which we state in the form we use.

Lemma 2.4 (Minimum-line deletion)**.** *Let $n > 0$. If k is an upper bound at $(m, n; s, t)$ then $k + \lfloor k/n \rfloor$ is an upper bound at $(m, n + 1; s, t)$. Symmetrically, if $m > 0$ and k is an upper bound at $(m, n; s, t)$ then $k + \lfloor k/m \rfloor$ is an upper bound at $(m + 1, n; s, t)$.*

Proof. Let M be (s, t)-free of size $m \times (n + 1)$ and let j_0 be a column of M with the fewest ones, say c of them. Summing over columns, $(n + 1)c \le \|M\|$. Deleting column j_0 leaves an (s, t)-free $m \times n$ matrix, so $\|M\| - c \le k$; combining, $nc \le k$, hence $c \le \lfloor k/n \rfloor$ and $\|M\| \le k + \lfloor k/n \rfloor$. The row form follows by transposition, which exchanges (m, n, s, t) with (n, m, t, s) and preserves the number of ones. □

Lemma 2.4 is older still. It is the Density Lemma, Lemma 3 of [4], stated there for $s = t$ and packaged into that paper's table-filling step, Lemma 4; and [4] says of its two lemmas that they "may be found in various forms in Section 12 of" [8] and in three later papers, one of which is [7]. It reappears as Lemma 2.1 of [1], stated for all m, n, s, t with $n \ge 2$ in the equivalent form $z(m, n; s, t) \le \lfloor \frac{n}{n-1} z(m, n-1; s, t) \rfloor$, and it is the argument [9] runs by hand at $(13, 18)$. We record what it gives when composed with the exact values [9] and [1] established. The hypotheses below are published values and are carried as hypotheses, so each conclusion is exactly as good as its citation and nothing is imported silently.

Corollary 2.5. (1) *If $z(13, 18; 3, 3) \le 116$ then $z(13, 19; 3, 3) \le 122$.*
(2) *If $z(15, 18; 3, 3) \le 132$ then $z(16, 18; 3, 3) \le 140$.*
(3) *If $z(16, 18; 3, 3) \le 140$ then $z(16, 19; 3, 3) \le 147$.*

The three hypotheses are, respectively, Theorem 1 of [9] *(also Theorem 1.1 of* [1]*), Theorem 1 of* [9]*, and Table 4 of* [4]*. The best published upper bounds at the three conclusions are* 125*,* 140 *and* 152*, so (1) and (3) are below what we found in print, while (2) re-derives the entry of* [4] *from a later exact value.*

These three consequences are arithmetic on published data by a published lemma, and we claim nothing more for them. In particular (2) is not an improvement: 140 is what Table 4 of [4] already prints at $(16, 18)$, and prints undecorated, which by that table's own legend means it was obtained from its Lemmas 2–4 and not by exhaustive enumeration — so it is very likely this same deletion step applied to that table's own $z(15, 18; 3, 3) \le 132$. What is new in (2) is only the hypothesis: 132 is now known to be exact. The same arithmetic applies at other cells of the window; we have not carried those instances into the formal development, and they are not asserted here.

3. The foundation: the balanced 2×2 diagonal

The development this paper reports grew out of a formalisation of the classical $s = t = 2$ material, and the $s = t = 3$ work reuses its definitions verbatim. Nothing in this section is new; it is stated because the later sections depend on it and because it is where the definitions were tested.

Theorem 3.1 (Kővári–Sós–Turán; Reiman)**.** *For every $(2, 2)$-free M of size $m \times n$, writing r_i for the number of ones in row i,*

$$\sum_{i=1}^{m} \binom{r_i}{2} \;\le\; \binom{n}{2}, \qquad \textit{hence} \qquad \|M\|^2 \;\le\; m\,\|M\| + m\,n(n-1).$$

Consequently, if $n \le k+1$ and $n(k+1) + n \cdot n(n-1) < (k+1)^2$, then k is an upper bound at $(n, n; 2, 2)$.

The first inequality is the standard double count of pairs of ones within a row; the second is its convexity step, in the integer form that avoids square roots, and is the arithmetic version of (1). The side condition $n \le k+1$ is what makes the violation persist for every larger count, so that a single failed inequality rules out all of them. The source statement is Theorem 3.1 of Damásdi, Héger and Szőnyi [5], which attributes it to Reiman [12]; what is formalised here is its inequality half in the integer form above, not its equality clause.

Theorem 3.2 (Reiman's construction)**.** *Let $D \subseteq \mathbb{Z}/n$ be a planar difference set, that is, every nonzero residue is $d - d'$ for exactly one ordered pair $d, d' \in D$. Then the circulant matrix $M_{ij} = [\, j - i \in D\,]$ is $(2, 2)$-free, and $n\,|D|$ is a lower bound at $(n, n; 2, 2)$.*

Theorem 3.3. *$z(n, n; 2, 2)$ equals $3, 6, 9, 12, 16, 21, 52, 105, 186$ for $n = 2, 3, 4, 5, 6, 7, 13, 21, 31$ respectively.*

Proof sketch. In each case the upper half is Theorem 3.1 with two arithmetic side conditions and no search whatever; apart from the degenerate $n = 1$, these nine are exactly the $n \le 31$ at which the floor of (1) is attained. The lower half is Theorem 3.2 with the classical Singer difference sets $\{0, 1\}$, $\{0, 1, 3\}$, $\{0, 1, 3, 9\}$, $\{0, 1, 6, 8, 18\}$, $\{0, 1, 3, 8, 12, 18\}$ for $n = 3, 7, 13, 21, 31$, and an explicit matrix for $n = 2, 4, 5, 6$. Checking that a set is a planar difference set is a scan over $|D|^4$ tuples — 1296 at $n = 31$ — after which Theorem 3.2 does the rest by argument, so no n^2 search is ever run. □

These nine values are published. Each was checked against every tabulation that lists it: Table 3 of [4] has all nine; the Afzaly–McKay extremal-graph data [2] lists $\mathrm{ex}(2n; \{C_{\mathrm{odd}}, C_4\})$, which agrees with all nine; and [14] and [11] stop at $n = 24$ and so carry the first eight. There was no disagreement. The four are not independent of one another: Table 3 of [4] names [7] and the Afzaly–McKay data as its own sources. For $n = 3, 7, 13, 21, 31$ the values also follow from the equality clause of (1) together with the existence of a projective plane of order $1, 2, 3, 4, 5$, which is the argument formalised here.

Proposition 3.4. *The explicit 4×4 and 6×6 matrices used in Theorem 3.3 carry exactly 9 and 16 ones.*

The matrices of Proposition 3.4, and the 32×32 matrix of Proposition 3.6 below, were produced by a short script and are almost certainly not the published extremal graphs; a second of search reproduces them. It is the values they attain, not the matrices, that are the published data.

Proposition 3.5. $z(8, 8; 2, 2) \le 25$.

This is Theorem 3.1 instantiated, and it is where that bound stops: the true value is 24 [4, 14, 11], and the side condition at $k = 24$ reads $648 < 625$, which is false. The gap is in the bound, not in the instantiation.

Proposition 3.6. $z(32, 32; 2, 2) \geq 189$.

Proposition 3.7. $z(32, 32; 2, 2) \leq 194$.

Proposition 3.6 re-derives a published lower bound: the witness is the order-5 plane padded to 32×32 and greedily extended, gaining one entry in the new row, one in the new column and one in the corner. Proposition 3.7 is Theorem 3.1 again, five above the truth. Collins, Riasanovsky, Wallace and Radziszowski [4] record $z(32, 32; 2, 2) \in \{189, 190\}$ and describe $n = 32$ as the first open case.

Remark 3.8. That framing may now be stale. On the Afzaly–McKay extremal-graph pages [2], in the table headed $H = \{C_{\mathrm{odd}}, C_4\}$ ("the extremal bipartite graphs of girth at least 6"), the row for $nv = 64$ reads $ne = 189$, with the graph count ng given as "≥ 31", linked to the file `b4_n64e189.some.s6`. The legend of those pages says, verbatim: *If ne is not preceded by "$\geq$" but ng is preceded by "$\geq$", it means that we have the Turán number exactly but possibly there are more extremal graphs. In this case, the file name contains the word "some".* So the page asserts $\mathrm{ex}(64; \{C_{\mathrm{odd}}, C_4\}) = 189$. Every $(2, 2)$-free 32×32 matrix is the biadjacency matrix of such a graph, whence $z(32, 32; 2, 2) \leq 189$; with Proposition 3.6 that gives $z(32, 32; 2, 2) = 189$. We decoded the linked file to check that it says what the row says: it holds 31 graphs, each on 64 vertices with 189 edges, each bipartite and C_4-free, and at least one of them has a balanced 32/32 bipartition. We still record this as an inference from an unrefereed data page rather than as a citation, it is not part of the formal development, and nothing in this paper depends on it either way.

Proposition 3.9. $z(4, 4; 2, 2) \neq 10$ *and* $z(31, 31; 2, 2) \neq 185$.

Both are refutations obtained from a proved half: the first from the upper bound at $(4, 4)$, the second from the lower bound at $(31, 31)$. They are stated because a specification in which a wrong value cannot be refuted is a specification that proves nothing; see Section 7 for the same discipline at $s = t = 3$.

4. The published 3×3 record layer

Before recording anything new we reproduced the most recent published record, exactly. This is the control on our verification: if our checker disagreed with a published matrix we would learn it here rather than in Section 5.

Proposition 4.1. $z(13, 19; 3, 3) \geq 118$*,* $z(14, 19; 3, 3) \geq 126$ *and* $z(16, 18; 3, 3) \geq 136$.

These are Theorem 1 of [13], re-derived here from the three matrices that paper prints in its Appendix A. The transcription was checked against the machine-readable bitstring block [13] also supplies, and the two printings agree character for character. Each matrix was then verified twice outside the proof assistant, by an exhaustive $\binom{m}{3}\binom{n}{3}$ loop over row-triple/column-triple pairs and by an independent column-triple counter written in a different language; the three check counts 277 134, 352 716 and 456 960 reproduce those quoted in [13]. The number of ones is an equality in each case, not merely a bound: the three matrices carry exactly 118, 126 and 136 ones. Finally, each was verified by the Lean kernel through Proposition 2.2.

Proposition 4.2. $z(14, 20; 3, 3) \geq 126$ *and* $z(16, 19; 3, 3) \geq 136$.

Proof. Apply Lemma 2.3 to the second and third bounds of Proposition 4.1. □

This is Corollary 3 of [13]. That paper ships the two padded matrices as ancillary files; here the padding step is the proved lemma, so no second machine check is performed and the two padded matrices are not needed. (Both bounds are superseded in Section 5.)

Proposition 4.3. $z(10, 20; 3, 3) \geq 102$.

This one is a positive control on the search, not a record. The value 102 appears in boldface in Table z_3 of [14] and is therefore a published exact value; together with the published upper bound, Proposition 4.3 matches the literature rather than improving it. What it shows is that the search reaches a known optimum on its own — three above the value compiled in Figure 2 of [3]. We flag it because it would otherwise read as a twenty-third improvement in the same list, and it is not one.

5. Twenty-two improved lower bounds

Theorem 5.1. *$z(13,19;3,3) \geq 120$ and $z(16,18;3,3) \geq 137$.*

Theorem 5.2. *$z(14,19;3,3) \geq 128$, $z(14,20;3,3) \geq 131$ and $z(16,19;3,3) \geq 142$.*

Theorems 5.1 and 5.2 together improve all five bounds of [13]: in the order that paper states them, $z(13,19)$, $z(14,19)$, $z(16,18)$, $z(14,20)$ and $z(16,19)$ gain $+2$, $+2$, $+1$, $+5$ and $+6$. Combined with Corollary 2.5 they give $120 \leq z(13,19;3,3) \leq 122$ and, using Table 4 of [4], $137 \leq z(16,18;3,3) \leq 140$.

Theorem 5.3. *$z(11,23;3,3) \geq 123$, $z(12,23;3,3) \geq 134$, $z(13,20;3,3) \geq 125$, $z(13,21;3,3) \geq 131$ and $z(13,23;3,3) \geq 140$.*

Theorem 5.4. *$z(14,21;3,3) \geq 136$, $z(14,22;3,3) \geq 142$, $z(14,23;3,3) \geq 146$, $z(15,19;3,3) \geq 136$, $z(15,20;3,3) \geq 139$, $z(15,21;3,3) \geq 144$, $z(15,22;3,3) \geq 148$ and $z(15,23;3,3) \geq 154$.*

Theorem 5.5. *$z(16,20;3,3) \geq 147$, $z(16,21;3,3) \geq 152$, $z(16,22;3,3) \geq 156$ and $z(16,23;3,3) \geq 161$.*

Proof of Theorems 5.1–5.5. For each of the twenty-two cells we exhibit one $(3,3)$-free matrix of the stated size whose number of ones equals the stated bound. Freeness is checked by Proposition 2.2: for the cell (m,n) the check inspects all $\binom{n}{3}$ increasing column triples against all m rows and confirms that no triple is covered by three rows. This is a finite computation, and it is the only computation on which these theorems rest. The matrices are reproduced in the formal development, one bitmask per row. ☐

Table 1 gives the numbers. Column L is the best lower bound we could locate in the literature for that cell, closed under Lemma 2.3; the next column names where it comes from, "pad" marking a value obtained by padding a published value at the neighbouring cell $(m,n-1)$ or $(m-1,n)$. Column L' is the bound proved here. Column U is the smallest published upper bound we located, and $U-L'$ is the gap that remains.

The twenty-two new values are mutually consistent in the way the table forces: each is at least the maximum of its left and upper neighbours in the updated grid, and each is at most the corresponding published upper bound. Neither had to hold, and both were checked.

5.1. **The search.** The matrices were found by a kicked-greedy local search written for this work: rows are carried as bit masks together with, for each column triple, the number of rows containing it, so that $(3,3)$-freeness is the invariant "every count is at most 2", maintained incrementally. Every state the search ever records is therefore already a valid witness. The loop greedily saturates in a randomised cell order, then kicks by deleting one to four ones (more when stalled, occasionally a whole row or column), re-saturates, accepts non-worsening moves, and restarts the chain after 20 000 non-improving kicks. This is the method described in [13], Section 3, and we claim nothing new for it.

The finite search actually run was this. Twenty-four cells were attacked, each for 150 CPU-seconds from a cold start with one seed; then the same twenty-four cells were run again for 150 CPU-seconds each from a second seed, warm-started from the first sweep's best matrix with the target raised by one. The second sweep improved *nothing*: every cell returned exactly the first sweep's value after some 90 million further kicks in total. That is the evidence — and the only evidence — we have that $L'+1$ is out of reach of this method at these cells; it is not evidence that $L'+1$ is false. Total measured cost: 7 240 CPU-seconds, or 2.01 CPU-hours, eight-way parallel. Of the twenty-four cells, twenty-two improved, one reproduced a published exact value (Proposition 4.3), and one, $(10,23)$, merely tied the compiled value 112 of [3]. Two of the twenty-three matrices we kept — the $(13,19)$ one with 120 ones and the $(16,18)$ one with 137 — came from an earlier probe run of the same method under a different

cell	L	source of L	L'	gain	U	$U - L'$
$(11, 23)$	121	pad from $z(11, 22) = 121$ [3]	**123**	+2	125	2
$(12, 23)$	132	pad from $z(12, 22) = 132$ [3, 1]	**134**	+2	136	2
$(13, 19)$	118	[13], Theorem 1	**120**	+2	125	5
$(13, 20)$	120	pad from $z(12, 20) = 120$ [1]	**125**	+5	130	5
$(13, 21)$	127	[3], Figure 2	**131**	+4	135	4
$(13, 23)$	137	pad from $z(13, 22) = 137$ [1]	**140**	+3	145	5
$(14, 19)$	126	[13], Theorem 1	**128**	+2	135	7
$(14, 20)$	126	[13], Corollary 3 (pad)	**131**	+5	140	9
$(14, 21)$	131	[3], Figure 2	**136**	+5	145	9
$(14, 22)$	137	[3], Figure 2	**142**	+5	150	8
$(14, 23)$	138	[3], Figure 2	**146**	+8	155	9
$(15, 19)$	132	[3], Figure 2	**136**	+4	143	7
$(15, 20)$	138	[3], Figure 2	**139**	+1	149	10
$(15, 21)$	139	[3], Figure 2	**144**	+5	154	10
$(15, 22)$	143	[3], Figure 2	**148**	+5	160	12
$(15, 23)$	149	[3], Figure 2	**154**	+5	165	11
$(16, 18)$	136	[13], Theorem 1	**137**	+1	140	3
$(16, 19)$	136	[13], Corollary 3 (pad)	**142**	+6	152	10
$(16, 20)$	146	[3], Figure 2	**147**	+1	158	11
$(16, 21)$	147	[3], Figure 2	**152**	+5	164	12
$(16, 22)$	149	[3], Figure 2	**156**	+7	169	13
$(16, 23)$	158	[3], Figure 2	**161**	+3	175	14

TABLE 1. The twenty-two improvements. L' is proved here (Theorems 5.1–5.5); L and U are readings of the literature, described in Section 6.

implementation; before the sweep began, the program described above reproduced both values from a cold start within 20 CPU-seconds per cell.

Two further checks were run on all twenty-three landed matrices outside the proof assistant. Each was verified by three implementations sharing no code — a row-triple by column-triple double loop, a column-triple counter, and a third check inside the search binary that scans row triples rather than column triples — all reporting no violation. And each was verified *maximal*: of the 98 to 207 zero entries in each matrix, not one can be flipped to a one without creating an all-ones 3×3 submatrix. Maximality is the cheapest sharp evidence that a matrix is a local optimum rather than a stopping point, and it does not follow from the search having terminated.

6. Where the comparison values come from

Grading a new lower bound requires the maximum over all published sources *closed under Lemma 2.3*, since a cell whose neighbour is already better gains nothing from a search. Assembling that closed maximum was the larger part of the work behind Table 1, and we describe it in full because the column L is the one claim in this paper that no machine checks.

Seven sources touch the window $11 \leq m \leq 16$, $18 \leq n \leq 23$. Tan [14], Table z_3, gives the full grid $3 \leq m \leq 16$, $3 \leq n \leq 23$ with boldface marking exact values. Collins, Riasanovsky, Wallace and Radziszowski [4], Table 4, gives $6 \leq m \leq n \leq 18$, so it meets the window only in the column $n = 18$; at the one cell of Table 1 in that column it supplies the smallest upper bound we found. Davies, Gill and Horsley [6], Table 2 of the preprint, improves upper bounds in the range $10 \leq m \leq 16$, $17 \leq n \leq 23$. Bhan, Nobili and Langer [3], Figure 2, is the grid of lower bounds everything since mid-2026 has been improving. Hou [9] gives $z(12, 18) = 108$, $z(13, 17) = 110$, $z(13, 18) = 116$, $z(14, 17) = 118$, $z(14, 18) = 124$, $z(15, 17) = 126$, $z(15, 18) = 132$; Afrasyab [1] gives $z(12, n) = 6n$ for $18 \leq n \leq 22$, $z(13, 22) = 137$, the frontier equalities at $(13, 18)$, $(14, 17)$, $(14, 18)$, $(15, 17)$ and $(15, 18)$, and $132 \leq z(16, 17) \leq 133$. Saurabh [13] gives the five values of Propositions 4.1 and 4.2.

6.1. **Figure 2 of** [3] **is a bitmap.** The grid of lower bounds in [3] is published as an image. There is no prose transcription of it in the paper, so the entries had to be read off the picture: we extracted the figure's PNG from the paper's own source archive and read the grid from it by eye. This is the load-bearing step in column L of Table 1 — thirteen of the twenty-two rows are graded directly against it — and we state it plainly rather than leave it implicit.

The read was cross-checked against prose. Two later papers quote entries of that figure in running text: [9] quotes seven of its lower bounds (108, 106, 115, 118, 124, 125, 132 at $(12, 18)$, $(13, 17)$, $(13, 18)$, $(14, 17)$, $(14, 18)$, $(15, 17)$, $(15, 18)$) and [13] quotes five more (114, 121, 130, 125, 132 at its five cells) as well as four of the figure's upper bounds; the ancillary files of [13] repeat its five in machine-readable form. All twelve agree with what we read off the image. Two independent papers thus confirm the reading at twelve cells. None of those twelve, however, is one of the thirteen cells that Table 1 grades directly against the figure: all twelve have $n \leq 20$, and eight of them $n \leq 18$.

Because the read is load-bearing we also pinned the image itself. The archive first read was the May 2026 source of arXiv:2605.01120, while the current version is v3 (25 Aug 2026); we downloaded the v3 e-print on 2026-08-29 and compared. The file `zarankiewicz_results.png` is byte-identical in the two archives (same MD5), and the only change to the paper's TeX source between them is an added author and one deleted space. Figure 2 is therefore literally the same picture in the version current at the time of writing, and the grid was re-read from it cell by cell.

6.2. **The padded floors.** Monotone padding raises the floor under several cells above their printed entry, and three of the twenty-two are graded against such a padded floor coming from an exact value established after [3] appeared:

- $(12, 23)$: from $z(12, n) = 6n$ we get $z(12, 22) = 132$, which Lemma 2.3 raises to $z(12, 23) \geq 132$, so 132 and not the printed entry is what a search here has to beat.
- $(13, 20)$: the same family gives $z(12, 20) = 120$, and padding in the row direction — the second half of Lemma 2 of [13], which is Lemma 2.3 transposed and is the one direction our formal development does not carry — gives $z(13, 20) \geq 120$, against a printed entry of 119.
- $(13, 23)$: Afrasyab's $z(13, 22) = 137$ with Lemma 2.3 gives $z(13, 23) \geq 137$.

A fourth cell, $(11, 23)$, is graded against $z(11, 22) = 121$, one of the three values [3] determined exactly, padded in n; and the floors at $(14, 20)$ and $(16, 19)$ are themselves the padded corollaries of [13]. In every case we grade against the larger, padded value, so the gains reported in Table 1 are the conservative ones. Note that the padded floors enter the paper only through the column L, which is a reading of the literature; no theorem proved here depends on any of them, and in particular nothing proved here depends on the row form of padding.

6.3. **What was searched for competing values.** A full pass over a local mirror of the arXiv text corpus, matching the notations $z(m, n)$, `\zar`$\{m\}\{n\}$ and $Z(m, n)$ for $10 \leq m \leq 16$ and $18 \leq n \leq 23$, produced, when re-run on 2026-08-29, 67 matching lines in three files: 39 inside [1] and [9], 18 inside [13], and 10 inside one unrelated paper in which $Z(m, n)$ denotes a Zolotarev graph. No fourth paper states a value in that window in any of those notations. (The four table-shaped sources print unlabelled grids and were read directly instead, which is why they do not appear as hits.) A live query of the arXiv API for all papers matching "Zarankiewicz", sorted by submission date, confirmed on 2026-08-29 that [13], posted 27 August 2026, was then the most recent paper on the subject. We also read and excluded a 2026 paper of Sadhu on a congruence obstruction to Roman's bound (arXiv:2608.07607, not in the bibliography since nothing here uses it): its conclusions are stated for $n = T - c$ with $1 \leq c \leq \frac{s}{s+2}T$, where $T = (t-1)\binom{m}{s}/(s+1)$ is the design threshold, so for $s = t = 3$ they need $n \geq \frac{2}{5}T = \frac{1}{5}\binom{m}{3}$, which is $n \geq 58$ at $m = 13$. Our cells have $n \leq 23$, so it cannot affect any of them.

Under the convention we use, "not found" is not "new". What Table 1 asserts is that we located no published lower bound as large as any of the twenty-two, having looked in the places described above.

7. Negative controls

Every check in Sections 4 and 5 asserts that something is *absent*, and a freeness predicate that were accidentally true of everything would pass all of them while proving nothing. The following are stated and proved for that reason.

Proposition 7.1. (1) *The all-ones 3×3 matrix is not $(3,3)$-free, and neither is the all-ones 16×18 matrix — at the exact size where Theorem 5.1 claims a record — which has 288 ones, so the check is not being passed by an empty matrix.*

(2) *The 16×18 matrix of Theorem 5.1 has a zero in position $(1,3)$; the matrix obtained by turning that entry into a one has 138 ones and is* not *$(3,3)$-free.*

(3) $z(3,3;3,3) = 8$*; consequently the claim $z(3,3;3,3) \geq 9$ is false. Both halves are decided over the whole space of 2^9 matrices, so this is the one cell where a too-large claim is refuted rather than merely left unproved.*

(4) *The distinctness hypothesis inside the proof of Proposition 2.2 is load-bearing: the 3×3 matrix whose first two columns are all ones and whose third is all zero is $(3,3)$-free, yet all three of its rows carry ones in the column list $(1,1,2)$, whose underlying set has only two elements. So the count bound stated for a triple of columns fails as soon as two of them coincide.*

(5) *The column Lemma 2.3 appends is identically zero, and padding the 16×18 matrix of Theorem 5.1 leaves the count at 137 — so padding cannot be manufacturing ones.*

Item (2) is the sharpest of these. A witness one short of the truth would admit some single flip; this one admits none, and the failure is explicit — an all-ones 3×3 submatrix created by the flip sits on rows $\{1,3,4\}$ and columns $\{3,11,14\}$. Item (3) is the only place in the development where an upper bound at $s=t=3$ is decided rather than cited. The value 8 agrees with the first entry of Table z_3 of [14].

8. What remains open

Two cells are now bracketed within three and are, we think, the natural next targets:

$$120 \ \leq\ z(13,19;3,3) \ \leq\ 122, \qquad 137 \ \leq\ z(16,18;3,3) \ \leq\ 140,$$

the lower halves by Theorem 5.1, the upper halves by Corollary 2.5(1) and Table 4 of [4].

We do not expect more search to close either. At $(13,19)$, a further 150 CPU-seconds of warm-started search produced 4.7 million kicks and never left 120; at $(16,18)$ the same budget produced 4.5 million kicks and never left 137, and twelve independent seeds in an earlier run all plateaued there. What these cells appear to want is an upper-bound argument of the kind [9] and [1] use — a certificate excluding $e+1$ ones — rather than a larger search budget. The same is true of $(16,17)$, where [1] already has $132 \leq z(16,17;3,3) \leq 133$ and only the upper half is in question.

Beyond the window, every cell with $m \leq 8$ is exact in Table z_3 of [14], as is every cell with $m=9$ except $(9,23)$; and cells with $n \geq 24$ lie outside every compiled table we found, so a lower bound there would have nothing to be graded against.

9. Verification

Every statement in this paper has been formally verified in Lean 4 against *Mathlib*, and the development contains no `sorry`. What is unusual about it, and worth saying explicitly, is that *nothing is delegated to compiled evaluation*: every matrix check — the three published matrices of Proposition 4.1, the twenty-two of Theorems 5.1–5.5, the positive control of Proposition 4.3, and the matrices appearing in Proposition 7.1 — is reduced by the Lean kernel itself, and every headline declaration depends on no axioms beyond `propext`, `Classical.choice` and `Quot.sound`. There is no use of `native_decide`, no external certificate file, and no trusted evaluator anywhere in the $s=t=3$ development. Two design choices make that affordable. The first is Proposition 2.2, which replaces a $2^m \cdot 2^n$ enumeration by an n^3m scan and is proved rather than postulated. The second is a measured detail: stating the scanned predicate with an explicit three-element set of columns cost 93.8 seconds and 8.45 GB of memory on the 16×19 witness, while stating it as a plain Boolean conjunction $M_{ia} \wedge M_{ib} \wedge M_{ic}$, proved equal to the set form, cost 13.3 seconds and 4.02 GB — building a three-element finite set inside a kernel loop is what dominates. Padding (Lemma 2.3) and deletion (Lemma 2.4) are proved by argument for all s and t and are not decided at all, so Proposition 4.2 and Corollary 2.5 cost no machine check. All numerical values were first produced by independent implementations in C and Python, as described in Sections 4 and 5.1, and the formal statements were written against them afterwards. The repository is private at the time of writing: repository reference to be added at submission.

References

[1] K. Afrasyab, *Exact Zarankiewicz Values On Two Finite Frontier Slices*, arXiv:2608.08154v1 (8 Aug 2026). Preprint; arXiv listed no journal reference and no further version as of 2026-08-29.

[2] N. Afzaly and B. D. McKay, *Extremal graphs* (data pages), `https://users.cecs.anu.edu.au/~bdm/data/extremal.html`, accessed 2026-08-29. The pages carry no date and no accompanying paper; they refer to work in preparation.

[3] J. Bhan, N. Nobili and P. Langer, *New Bounds for Zarankiewicz Numbers via Reinforced LLM Evolutionary Search*, arXiv:2605.01120 (v1 1 May, v2 7 May, v3 25 Aug 2026). The title page of every version is marked "Working Paper"; no journal reference is listed. The author list here is the one arXiv records; the v3 PDF itself carries a fourth author, S. Raghuraman, whom the arXiv metadata does not list. Figure 2, which is all we use, is the same image file in v2 and v3.

[4] A. F. Collins, A. W. N. Riasanovsky, J. C. Wallace and S. P. Radziszowski, *Zarankiewicz numbers and bipartite Ramsey numbers*, Journal of Algorithms and Computation **47** (2016), no. 1, 63–78; `doi:10.22059/jac.2016.7943`; arXiv:1604.01257.

[5] G. Damásdi, T. Héger and T. Szőnyi, *The Zarankiewicz problem, cages, and geometries*, Annales Universitatis Scientiarum Budapestinensis, Sectio Mathematica **56** (2013), 3–37.

[6] S. Davies, P. Gill and D. Horsley, *Improved upper bounds on Zarankiewicz numbers*, Discrete Mathematics **349** (2026), no. 5, Paper No. 114924; `doi:10.1016/j.disc.2025.114924`; arXiv:2411.18842. Table and theorem numbers are cited from the preprint, v2, where the improved $z(m,n;3,3)$ bounds are Table 2.

[7] J. Dybizbański, T. Dzido and S. Radziszowski, *On some Zarankiewicz numbers and bipartite Ramsey numbers for quadrilateral*, Ars Combinatoria **119** (2015) 275–287; arXiv:1303.5475.

[8] R. K. Guy, *A many-facetted problem of Zarankiewicz*, in: G. Chartrand and S. F. Kapoor (eds.), *The Many Facets of Graph Theory*, Lecture Notes in Mathematics **110**, Springer, Berlin, 1969, 129–148. We did not read this paper. The bibliographic record above is from zbMATH (Zbl 0186.01801) and Crossref; everything attributed to it here is attributed on the authority of [4] and [14]. Note that [14] recomputed Guy's tables independently and reports that they contain "only eight such errors, all in the $z_2(m,n)$ table, and all are too low by just one".

[9] S. Hou, *Seven Exact Finite Zarankiewicz Numbers from a Single* 13×18 *Core*, arXiv:2608.08549v1 (9 Aug 2026). Preprint; arXiv listed no journal reference and no further version as of 2026-08-29.

[10] T. Kövári, V. T. Sós and P. Turán, *On a problem of K. Zarankiewicz*, Colloquium Mathematicum **3** (1954), no. 1, 50–57; `doi:10.4064/cm-3-1-50-57`.

[11] OEIS Foundation Inc., *Sequence A072567*, The On-Line Encyclopedia of Integer Sequences, `https://oeis.org/A072567`, accessed 2026-08-29. The listed terms are $a(1), \ldots, a(24)$.

[12] I. Reiman, *Über ein Problem von K. Zarankiewicz*, Acta Mathematica Academiae Scientiarum Hungaricae **9** (1958), no. 3–4, 269–273; `doi:10.1007/BF02020254`. We did not read this paper; the record above is from Crossref, and the inequality and its equality clause are quoted in Section 1 as they are stated in Theorem 3.1 of [5].

[13] A. Saurabh, *Five improved lower bounds for Zarankiewicz numbers* $z(m,n;3,3)$, arXiv:2608.26603v1 (27 Aug 2026). Preprint; v1 was still the only version, with no journal reference, as of 2026-08-29. Its theorem numbering, used above, is Theorem 1, Lemma 2, Corollary 3.

[14] J. Tan, *An attack on Zarankiewicz's problem through SAT solving*, arXiv:2203.02283v2 (19 Apr 2022). Preprint; we could not locate a journal version. Neither arXiv, zbMATH nor dblp records one.

[15] K. Zarankiewicz, *Problem P 101*, Colloquium Mathematicum **2** (1951) 301.