On the Turán number of ordered forests

Dániel Korándi [1,3]

Institute of Mathematics
EPFL
Lausanne, Switzerland

Gábor Tardos [2,4]

Rényi Institute
Budapest, Hungary

István Tomon [1,5]

Institute of Mathematics
EPFL
Lausanne, Switzerland

Craig Weidert [6]

Abstract

An ordered graph H is a graph with a linear ordering on its vertex set. The corresponding Turán problem, first studied by Pach and Tardos, asks for the maximum number $\mathrm{ex}_<(n, H)$ of edges in an ordered graph on n vertices that does not contain H as an ordered subgraph. It is known that $\mathrm{ex}_<(n, H) > n^{1+\varepsilon}$ for some positive $\varepsilon = \varepsilon(H)$ unless H is a forest that has a bipartition $V_1 \cup V_2$ such that V_1 totally precedes V_2 in the ordering. Making progress towards a conjecture of Pach and Tardos, we prove that $\mathrm{ex}_<(n, H) = n^{1+o(1)}$ holds for all such forests that are "degenerate" in a certain sense. This class includes every forest for which an $n^{1+o(1)}$ upper bound

was previously known, as well as new examples. For example, the class contains all forests with $|V_1| \leq 3$. Our proof is based on a density-increment argument.

1   Introduction

An *ordered graph* G is defined as a triple $(V, E, <)$ where (V, E) is a simple graph and $<$ is a linear ordering on the vertex set V. We say that an ordered graph $H = (V', E', <')$ is an *ordered subgraph* of G if there is an order preserving embedding of V' into V that maps edges into edges. If G does not contain H as an ordered subgraph then we say that G is *H-free*.

The following Turán-type question arises naturally: For a fixed ordered graph H, what is the maximum number of edges that an H-free ordered graph on n vertices can have? This maximum, called the *extremal* or *Turán number* of H, is denoted by $\mathrm{ex}_<(n, H)$. The systematic study of extremal numbers was initiated by Pach and Tardos in [5].

For usual (unordered) graphs, the Erdős-Stone-Simonovits theorem says that the extremal number of a graph is controlled by its chromatic number. As it turns out, ordered graphs exhibit a similar phenomenon. The *interval chromatic number* $\chi_<(H)$ of an ordered graph $H = (V, E, <)$ is the smallest integer r, such that V can be split into r intervals (i.e., sets of consecutive vertices in the ordering) such that no edge of H has both endpoints in the same interval. It is not hard to show (see [5]) that $\mathrm{ex}_<(n, H) = \left(1 - \frac{1}{\chi_<(H)-1}\right)\binom{n}{2} + o(n^2)$. This asymptotically determines the extremal number of H when $\chi_<(H) \geq 3$. However, much like for usual graphs, the problem becomes more difficult when $\chi_<(H) = 2$.

Our work focuses on the general problem of classifying ordered graphs that have close to linear extremal numbers, i.e., satisfy $\mathrm{ex}_<(n, H) = n^{1+o(1)}$. Note that if H contains some cycle of length k then $\mathrm{ex}_<(n, H) \geq \Omega(n^{1+1/k})$.

[3] Email: `daniel.korandi@epfl.ch`
[4] Email: `tardos@renyi.hu`
[5] Email: `istvan.tomon@epfl.ch`
[6] Email: `craig.weidert@gmail.com`

Indeed, it is well-known (see, e.g., [1]) that n-vertex (unordered) graphs not containing any k-cycle exist with $\Omega(n^{1+1/k})$ edges. [7] So if $\mathrm{ex}_<(n, H) = n^{1+o(1)}$ holds then H is acyclic with interval chromatic number 2. Pach and Tardos conjectured that the converse holds, as well. More precisely, they made the following stronger conjecture:

Conjecture 1.1 (Pach–Tardos [5]) *Let H be an acyclic ordered graph such that $\chi_<(H) = 2$. Then*

$$\mathrm{ex}_<(n, H) = n(\log n)^{O(1)}.$$

In this paper we make some progress towards proving this conjecture by showing that $\mathrm{ex}_<(n, H) = n^{1+o(1)}$ holds for a large class of ordered forests H. It will be more convenient to state our result using the language of pattern-avoiding matrices.

Let us now describe the analogous problem with $0 - 1$ matrices, i.e., matrices with all entries from $\{0, 1\}$. For a $0 - 1$ matrix A, its *weight* $w(A)$ is defined to be the number of 1-entries in it. We say that A *contains* another $0 - 1$ matrix B if either B is a submatrix of A or it can be obtained from a submatrix of A by changing some 1-entries to 0-entries. Or in other words, if we can get B by deleting some rows, columns and 1-entries from A. We denote this relation by $B \prec A$. A pattern in our context is just a fixed $0 - 1$ matrix P, and the corresponding Turán-type problem then wants to maximize the weight of a $0 - 1$ matrix A of dimensions $m \times n$ that does not contain P. Let $\mathrm{ex}(m, n, P)$ denote the maximum weight of such a matrix A. We will simply write $\mathrm{ex}(n, P)$ for $\mathrm{ex}(n, n, P)$.

We can think of a pattern P as the bipartite adjacency matrix of some ordered graph H_P of interval chromatic number 2, where the order of the vertices coincides with the order of the corresponding rows and columns of P, and row vertices precede the column vertices. Then $\mathrm{ex}(m, n, P)$ translates to the maximum number of edges in an H_P-free ordered graph G_A (corresponding to A) on $m + n$ vertices, such that all edges of G_A connect the first m vertices to the last n. If not for this last condition, this would be the exact same quantity as $\mathrm{ex}(m + n, H_P)$. And indeed, as it was observed in [5], the two functions are closely related:

$$\mathrm{ex}(\lfloor n/2 \rfloor, P) - O(n) \leq \mathrm{ex}_<(n, H_P) = O(\mathrm{ex}(n, P) \log n). \tag{1}$$

---

[7] of course, for odd k we even have $\mathrm{ex}_<(n, H) \geq \Omega(n^2)$

Let us call a pattern P *acyclic* if H_P does not contain any cycle. It is easy to see that P is acyclic if and only if it contains no submatrix P' such that every row and column of P' contains at least two 1-entries. Then (1) shows that Conjecture 1.1 can be stated in the following, equivalent form:

Conjecture 1.2 (Pach–Tardos [5]) *Let P be an acyclic pattern. Then*

$$\mathrm{ex}(n, P) = n(\log n)^{O(1)}.$$

It is worth mentioning that a stronger variant of this conjecture had earlier been proposed by Füredi and Hajnal [2], who thought $\mathrm{ex}(n, P) = O(n \log n)$ holds for every acyclic pattern P. However, this conjecture was refuted by Pettie [6], who constructed a pattern P_0 such that $\mathrm{ex}(n, P_0) = \Theta(n \log n \log \log n)$.

There are several patterns that are known to satisfy Conjecture 1.2. For example, it is known [5] that if P is obtained from a pattern P' by appending a new last column with a single 1-entry, then

$$\mathrm{ex}(n, P) = O(\mathrm{ex}(n, P') \log n) \tag{2}$$

(provided $w(P') \geq 2$). Using this and similar operations, the conjecture has been verified for a large family of matrices [2,5,7], including all patterns of weight up to 6, with essentially two exceptions (omitting 0-entries for clarity):

$$Q_1 = \begin{pmatrix} 1 & & 1 & \\ 1 & & & 1 \\ & 1 & & 1 \end{pmatrix} \quad \text{and} \quad Q_2 = \begin{pmatrix} & 1 & & 1 \\ 1 & & 1 & \\ 1 & & & 1 \end{pmatrix}.$$

In certain special cases, e.g., when P is a permutation matrix [4] or a double permutation matrix [3], the even stronger bound $\mathrm{ex}(n, P) = O(n)$ is known to hold.

Let us now define a broad class of patterns that includes, up to transposing, all matrices that are known to satisfy Conjecture 1.2. We say that a $0-1$ matrix P of dimensions $l \times k$ is *horizontally separable* if it can be cut along a horizontal line without separating the 1-entries in more than one column. In other words, there is a row $a \in [l-1]$ such that for all but at most one column $y \in [k]$, we have $P(x, y) = 0$ either for every $x \leq a$ or for every $x > a$. We call a matrix P *horizontally degenerate* if every submatrix P' of P is either a single row or a horizontally separable matrix. Note that horizontally degenerate matrices are always acyclic.

The reader might find it helpful to visualize patterns as graphs whose vertices are the 1-entries, and two 1-entries are connected if they are in the same row or column with only 0-entries between them. A pattern is acyclic if and only if this graph is acyclic. Horizontally separable means that the matrix can be split into two parts along a horizontal cut that only cuts through at most one (vertical) edge. A pattern is horizontally degenerate if it can be split into its rows by applying a series of such cuts.

Our main result is the following theorem that says that every horizontally degenerate pattern has close to linear extremal number. Some preliminary results from this work establishing $\text{ex}(n, Q_1) \leq n^{1+o(1)}$ have previously appeared in the master's thesis of the fourth author [8].

Theorem 1.3 *Let P be a horizontally degenerate matrix. Then $\text{ex}(n, P) \leq n^{1+o(1)}$.*

This theorem can be thought of as a common generalization of all previously known results about acyclic patterns, albeit with a somewhat worse upper bound. However, it also applies to many new matrices, including, for example, all $3 \times k$ acyclic patterns, as they are easily verified to be horizontally degenerate. In particular, we get $\text{ex}(n, Q_1), \text{ex}(n, Q_2) \leq n^{1+o(1)}$, improving the previous best bound of $O(n^{5/3})$. As discussed above, this also implies that every pattern P of weight up to 6 satisfies $\text{ex}(n, P) \leq n^{1+o(1)}$.

It is clear that transposing a matrix does not change its extremal number: $\text{ex}(n, P) = \text{ex}(n, P^T)$. In particular, Theorem 1.3 holds for the analogously defined *vertically degenerate* patterns, as well. The smallest acyclic pattern that is neither horizontally, nor vertically degenerate (and hence not covered by our theorem) is the following "pretzel"-like matrix:

$$
R = \begin{pmatrix} 1 & & & 1 \\ & 1 & & \\ 1 & & 1 & \\ & 1 & & 1 \end{pmatrix}.
$$

It would be interesting to obtain nontrivial estimates on $\text{ex}(n, R)$.

Finally, it may be worth translating our result back to graphs. In terms of ordered graphs of interval chromatic number 2 (with interval partition $V = A \cup B$), horizontally separable means that the vertices of the first interval A can be split into two subintervals A_1 and A_2 such that there is at most one

vertex in the second interval B that has neighbors in both A_1 and A_2. An ordered graph H is horizontally degenerate if every ordered subgraph of H is horizontally separable. Our result then says that $\mathrm{ex}_<(n, H) \leq n^{1+o(1)}$ for every horizontally degenerate H.

The proof of Theorem 1.3 is based on a density-increment argument. We provide a sketch in the next Section.

2 Proof outline

Let P be an $l \times k$ horizontally degenerate matrix. We want to show that $\mathrm{ex}(n, P) \leq n^{1+o(1)}$. The general proof strategy is very simple: For fixed $\varepsilon > 0$ and large enough n, we take an arbitrary $n \times n$ matrix A with weight at least $n^{1+\varepsilon}$, and then show that either we can embed P in A, or A contains a large and significantly denser submatrix. More precisely, if our embedding algorithm fails then A has an $n' \times n'$ submatrix A' with weight $w(A') \geq (n')^{1+\varepsilon+\varepsilon'}$, where ε' is about ε^l and n' is close to n. Then we can repeat the same argument on A'. After about ε^{-l} iterations, we must find an embedding of P, otherwise we would get some $n_0 \times n_0$ matrix with weight over n_0^2, which is impossible.

Let us now describe our embedding scheme. Take a partition of A into k equal-sized $n \times \frac{n}{k}$ *blocks*, formed by consecutive columns. Via some technical considerations, we may assume that each row in each block contains the same number of 1-entries. We want to find a *block-respecting* embedding of P that maps the i'th column of P to the i'th block of A.

To embed P, we will apply induction on the number of rows l. The horizontal degeneracy of P means that we can cut it into two matrices $\overline{P}$ and $\underline{P}$ that are *almost* independent from each other: there is only one column c that might contain 1-entries in both $\overline{P}$ and $\underline{P}$. We will try to find block-respecting embeddings of the two halves relatively far from each other ($\overline{P}$ above $\underline{P}$) such that their c'th columns are mapped to the same column of A (in the c'th block).

More precisely, we split the rows of A into β equal-sized $\frac{n}{\beta} \times n$ *parts* $A_1, \ldots, A_\beta$, where β is an appropriately chosen integer. We will embed $\overline{P}$ and $\underline{P}$ in different parts, one after the other. Once we found an embedding of $\underline{P}$ in some A_j using our induction hypothesis, the image of column c is fixed. By our technical assumptions on A, if this column has a 1-entry in a part A_i above A_j, then a critical row of $\overline{P}$ (one that has a 1-entry in its c'th column) can be embedded in A_i. Using an appropriate induction argument

we can extend this to an embedding of $\overline{P}$ in A_i. Then the images of $\overline{P}$ and $\underline{P}$ together will form an embedding of P.

What we need in the above argument is that A_j is not the topmost part that contains 1-entries in the image of the c'th column. To ensure this, we delete the 1-entries from the topmost nonempty block of each column before applying the argument on the matrix. One can check that either A has a dense submatrix or the weight of the new matrix $\tilde{A}$ is not significantly smaller than the weight of A (and hence we can apply the previous paragraph on $\tilde{A}$). But then when we embed $\underline{P}$ into part A_j of $\tilde{A}$, the image of the c'th column will certainly have the deleted 1-entries in a part A_i of A above A_j. This concludes our sketch.

References

[1] Bollobás, B., "Extremal Graph Theory," Academic Press, London, 1978.

[2] Füredi, Z., and P. Hajnal, *Davenport-Schinzel theory of matrices,* Discrete Mathematics **103** (1992), 233–251.

[3] Geneson, J. T., *Extremal functions of forbidden double permutation matrices,* Journal of Combinatorial Theory, Series A **116** (2009), 1235–1244.

[4] Marcus, A., and G. Tardos, *Excluded permutation matrices and the Stanley-Wilf conjecture,* Journal of Combinatorial Theory, Series A **107** (2004), 153–160.

[5] Pach, J., and G. Tardos, *Forbidden paths and cycles in ordered graphs and matrices,* Israel Journal of Mathematics **155** (2006), 359–380.

[6] Pettie, S., *Degrees of nonlinearity in forbidden $0-1$ matrix problems,* Discrete Mathematics **311** (2011), 2396–2410.

[7] Tardos, G., *On 0-1 matrices and small excluded submatrices,* Journal of Combinatorial Theory, Series A **111** (2005), 266–288.

[8] Weidert, C., "Extremal problems in ordered graphs," Master's thesis, Simon Fraser University, Burnaby, BC, Canada, 2009.