

MORE TURÁN-TYPE THEOREMS FOR TRIANGLES IN CONVEX POINT SETS

Boris Aronov[‡], Vida Dujmović[♠], Pat Morin[♣], Aurélien Ooms[‡],
and Luís Fernando Schultz Xavier da Silveira[♠]

December 10, 2018

ABSTRACT. We study the following family of problems: Given a set of n points in convex position, what is the maximum number triangles one can create having these points as vertices while avoiding certain sets of *forbidden configurations*. As forbidden configurations we consider all 8 ways in which a pair of triangles in such a point set can interact. This leads to 256 extremal Turán-type questions. We give nearly tight (within a $\log n$ factor) bounds for 248 of these questions and show that the remaining 8 questions are all asymptotically equivalent to Stein’s longstanding tripod packing problem.

[‡]Partially supported by NSF Grants CCF-11-17336, CCF-12-18791, and CCF-15-40656, and by BSF grant 2014/170. Department of Computer Science and Engineering, Tandon School of Engineering, New York University, USA.

[♠]Partially supported by NSERC and the Ontario Ministry of Research and Innovation. Department of Computer Science and Electrical Engineering, University of Ottawa, Canada.

[♣]Partially supported by NSERC. School of Computer Science, Carleton University, Canada.

[‡]Supported by the Fund for Research Training in Industry and Agriculture (FRIA). Département d’Informatique, Université libre de Bruxelles (ULB), Belgium.

Contents

1	Introduction	1
1.1	Previous Work	1
1.2	New Results	3
2	Easy Results	3
2.1	Mariposas are Irrelevant	3
2.2	Cubic-Sized Sets of Pairwise Crossing Triangles	3
2.3	Linear-Sized Sets Using Only a Single Configuration	5
3	Points of View	5
3.1	The Top/Bottom View	5
3.2	The Dot Puzzle View	7
3.3	Some Warm-Up Exercises	10
4	Results Based on Dot Puzzles	11
4.1	Linear Upper Bounds	11
4.2	Forbidding Swords	17
4.3	Monotone Matrices, Tripod Packing, and 2-Comparable Sets	20
4.4	Lower Bounds	23

1 Introduction

Let t_1 and t_2 be a pair of distinct triangles whose (4 to 6) vertices are in convex position. There are 8 combinatorially distinct ways that these triangles can interact: 2 ways in which the triangles can share an edge (\boxtimes and \diamond), 3 ways in which the triangles can share a single vertex (∇ , $\nabla\downarrow$, and $\nabla\uparrow$), and 3 ways in which the triangles can have no vertices in common ($\diamond\downarrow$, \times , and \boxtimes). Because it is difficult to keep track of nameless entities, we assign a mnemonic to each configuration (though the reader is encouraged to choose their own):

\boxtimes	\diamond	∇	$\nabla\downarrow$	$\nabla\uparrow$	$\diamond\downarrow$	\times	\boxtimes
taco	mariposa	bat	nested	crossing	ears	swords	david

We consider the following class of problems: Given a set, X , of combinatorial configurations of pairs of triangles, what is the size of a largest set, S , of triangles one can create whose vertices are n points in convex position, and such that no pair of triangles in S forms a configuration in X ? We call the size of such a set $\text{ex}(n, X)$. For example,

$$\text{ex}(n, \{\boxtimes, \nabla\downarrow, \nabla\uparrow, \times, \boxtimes\}) = n - 2. \quad (1)$$

This is because the set $X = \{\boxtimes, \nabla\downarrow, \nabla\uparrow, \times, \boxtimes\}$ forbids any form of crossings between the edges of triangles. Thus, the maximum number of triangles we can have while avoiding X is the number of triangles in a triangulation of a convex n -gon, i.e., $n - 2$.

1.1 Previous Work

Since there are eight possible forbidden configurations, there are $2^8 = 256$ sets X for which we can study $\text{ex}(n, X)$. Some of these sets have been previously studied. Braß, Rote, and Swanepoel showed that $\text{ex}(n, \{\diamond\downarrow, \times, \nabla, \nabla\downarrow\}) \leq n$ in order to solve an Erdős problem on the maximum number of maximum area/perimeter triangles determined by a point set. Braß [2] later began a systematic study in which he gave asymptotically tight bounds on $\text{ex}(n, X)$ for all singleton X and all pairs X of configurations in which two triangles share a single vertex.

Of course, upper and lower bounds are inherited through the subset relationship: $\text{ex}(n, X) \leq \text{ex}(n, Y)$ for any $X \supseteq Y$. Table 1 shows the complete set of results we obtain when we apply this exhaustively to the list of previous results. Each entry in this table presents the asymptotic behaviour of $\text{ex}(n, X)$ for the set X obtained as the union of the row and column label. Asymptotically tight bounds are coloured green, and gaps between lower and upper bounds are coloured red. Previous results imply 35 tight bounds for 256 of the possible choices of X . The configuration \diamond is omitted from the table since a simple argument (Lemma 1) shows that its inclusion in X does not change $\text{ex}(n, X)$ by more than a constant factor. Some of the results in Table 1 are marked with an F if they are easy, or folklore. Some others are marked with H if they follow from a corresponding bound for 3-regular hypergraphs. Specifically, if X includes $\{\nabla, \nabla\downarrow, \nabla\uparrow\}$ then no pair of triangles can share a vertex, so $\text{ex}(n, X) \leq n/3$ and if X includes $\{\boxtimes, \diamond\}$ then no pair of triangles can share an edge, so $\text{ex}(n, X) \leq \binom{n}{2}$.

	\bowtie \bowtie \diamond	\bowtie \diamond	\bowtie \bowtie	\bowtie	\bowtie \diamond	\diamond	\bowtie	
$\bowtie \searrow \searrow \searrow \searrow$	1 F:F	n T2:[3]	n T2:H	n T2:H	n T2:H	n T2:H	n T2:H	n T2:H
$\bowtie \searrow \searrow$	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H
$\bowtie \searrow \searrow$	n T2:[3]	n T2:[3]	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H
$\bowtie \searrow$	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H
$\searrow \searrow \searrow$	n T2:[3]	n T2:[3]	n T2:H	n T2:H	n T2:H	n T2:H	n T2:H	n H:H
$\searrow \searrow$	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	n^2 [2]:[2]
$\searrow \searrow$	n T2:[3]	n [3]:[3]	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	n^2 [2]:[2]
\searrow	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	n^2 [2]:[2]
$\bowtie \searrow \searrow \searrow$	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H
$\bowtie \searrow \searrow$	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H
$\bowtie \searrow$	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H	$n : n^2$ T2:H
$\searrow \searrow \searrow$	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	n^2 [2]:[2]
$\searrow \searrow$	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	$n : n^2$ T2:[2]	n^2 [2]:[2]
\searrow	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	$n : n^2$ [3]:[2]	$n : n^2$ T2:[2]	$n : n^3$ [3]:F	$n : n^2$ T2:[2]	n^3 [2]:[2]
	n^2 H:H	n^2 H:[2]	n^2 H:[2]	n^2 [2]:[2]	n^2 H:[2]	n^3 [2]:[2]	n^2 [2]:[2]	n^3 F:F

Table 1: Previous lower and upper bounds for $\text{ex}(n, X)$. T2 denotes a (easy) lower bound of $\Omega(n)$ that appears in Theorem 2. F denotes an obvious, or folklore result. H denotes a bound that follows from the corresponding bound on 3-regular hypergraphs.

1.2 New Results

In the current paper, we determine, up to a logarithmic factor, the asymptotics of $\text{ex}(n, X)$ for 248 sets X . These results are shown in Table 2. For the remaining 8 sets, we have determined that the asymptotics are all the same and are equivalent to a problem that appears in various contexts and under different names, including monotone matrices, tripod packing, and 2-comparable triples. We discuss this problem and its rich history in Section 4.3.

The rest of this paper is organized as follows. In Section 2 we present a few easy results that we need for completeness. In Section 3 we discuss different ways of thinking about the problem. In particular, we present a series of puzzles whose solutions determine the asymptotic growth of $\text{ex}(n, X)$. In Section 4, which represents the technical meat of the paper, we use these puzzles to derive new upper and lower bounds.

2 Easy Results

In this section we present an easy result that cuts our work in half by reducing the number of problems from 256 to 128. We then describe some easy lower bound constructions that are required for completeness.

2.1 Mariposas are Irrelevant

The following lemma shows that including the \diamond configuration in the set X of forbidden configurations has no effect on the asymptotics of $\text{ex}(n, X)$.

Lemma 1. *For any X , $\text{ex}(n, X \cup \{\diamond\}) \geq \text{ex}(n, X)/8$.*

Proof. Let S be a set of triangles that achieves $\text{ex}(n, X)$. For each pair of vertices u and w independently and uniformly choose a direction \overrightarrow{uw} or \overleftarrow{uw} . We then obtain a set $S' \subseteq S$ by removing any triangle that has a directed edge for which the triangle is to the left of the edge. Observe that the set S' does not contain a \diamond configuration.

For any particular triangle $t \in S$, the probability that $t \in S'$ is exactly $1/8$ since each of t 's three edges must be directed clockwise and edge directions are chosen independently. By linearity of expectation, $E[|S'|] = |S|/8 = \text{ex}(n, X)/8$. We conclude therefore that there exists some subset $S'' \subseteq S$ of size least $\text{ex}(n, X)/8$ that does not contain a \diamond configuration. The set S'' proves that $\text{ex}(n, X \cup \{\diamond\}) \geq \text{ex}(n, X)/8$. \square

2.2 Cubic-Sized Sets of Pairwise Crossing Triangles

Theorem 1. $\text{ex}(n, \{\diamond, \heartsuit, \clubsuit\}) \in \Omega(n^3)$.

Proof. Partition the vertices of the convex n -gon into three contiguous sets, A , B , and C , each of size $\lfloor n/3 \rfloor$ or $\lceil n/3 \rceil$, as appropriate. Consider the set, S , of all triangles having one vertex in each of A , B , and C . It is easy to check that any two triangles in S have a pair of edges that cross, thus they do not form any of \diamond , \heartsuit , or \clubsuit . Furthermore, $|S| \geq \lfloor n/3 \rfloor^3 \in \Omega(n^3)$, so $\text{ex}(n, \{\diamond, \heartsuit, \clubsuit\}) \in \Omega(n^3)$. \square

	$\begin{smallmatrix} \times \\ \times \\ \Downarrow \end{smallmatrix}$	$\begin{smallmatrix} \times \\ \Downarrow \end{smallmatrix}$	$\begin{smallmatrix} \times \\ \times \\ \Downarrow \end{smallmatrix}$	\times	$\begin{smallmatrix} \times \\ \Downarrow \end{smallmatrix}$	\Downarrow	\times	
$\begin{smallmatrix} \Downarrow \\ \Downarrow \\ \Downarrow \\ \Downarrow \end{smallmatrix}$	1 F:F	n T2:[3]	n T2:H	n T2:H	n T2:H	n T2:H	n T2:H	n T2:H
$\begin{smallmatrix} \Downarrow \\ \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T3	n^* T2:T3	n^* T2:T3	n^* T2:T3
$\begin{smallmatrix} \Downarrow \\ \Downarrow \\ \Downarrow \end{smallmatrix}$	n T2:[3]	n T2:[3]	n^* T2:T9	n^* T2:T9	n^* T2:T8	tripods	n^* T2:T8	tripods
$\begin{smallmatrix} \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T9	n^* T2:T9	n^* T2:T9	n^* T2:T9	n^* T2:T5	tripods	n^* T2:T5	tripods
$\begin{smallmatrix} \Downarrow \\ \Downarrow \\ \Downarrow \end{smallmatrix}$	n T2:[3]	n T2:[3]	n T2:H	n T2:H	n T2:H	n T2:H	n T2:H	n H:H
$\begin{smallmatrix} \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T4	n^* T2:T4	n^2 T16:[2]	n^2 [2]:[2]
$\begin{smallmatrix} \Downarrow \\ \Downarrow \end{smallmatrix}$	n T2:[3]	n [3]:[3]	n^* T2:T10	n^* [3]:T10	n^* T2:T8	n^2 T17:[2]	n^* T2:T8	n^2 [2]:[2]
$\begin{smallmatrix} \Downarrow \end{smallmatrix}$	n^* T2:T10	n^* [3]:T10	n^* T2:T10	n^* [3]:T10	n^* T2:T7	n^2 T17:[2]	n^2 T16:[2]	n^2 [2]:[2]
$\begin{smallmatrix} \Downarrow \\ \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^2 T14:H	n^2 T14:H	n^2 T14:H	n^2 T14:H
$\begin{smallmatrix} \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^2 T14:H	n^2 T14:H	n^2 T14:H	n^2 T14:H
$\begin{smallmatrix} \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T9	n^* T2:T9	n^* T2:T9	n^* T2:T9	n^2 T14:H	n^2 T14:H	n^2 T14:H	n^2 T14:H
$\begin{smallmatrix} \Downarrow \end{smallmatrix}$	n^* T2:T9	n^* T2:T9	n^* T2:T9	n^* T2:T9	n^2 T14:H	n^2 T14:H	n^2 T14:H	n^2 H:H
$\begin{smallmatrix} \Downarrow \\ \Downarrow \end{smallmatrix}$	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^2 T14:[2]	n^2 T14:[2]	n^2 T14:[2]	n^2 [2]:[2]
$\begin{smallmatrix} \Downarrow \end{smallmatrix}$	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^* T2:T6	n^2 T14:[2]	n^2 T14:[2]	n^2 T16:[2]	n^2 [2]:[2]
$\begin{smallmatrix} \Downarrow \end{smallmatrix}$	n^2 T15:[2]	n^2 T15:[2]	n^2 T15:[2]	n^2 T15:[2]	n^2 T15:[2]	n^3 T1:F	n^2 T15:[2]	n^3 [2]:[2]
	n^2 H:H	n^2 H:[2]	n^2 H:[2]	n^2 [2]:[2]	n^2 H:[2]	n^3 [2]:[2]	n^2 [2]:[2]	n^3 F:F

$$n^* = n : n \log n \quad \text{tripods} = n^{1.546} : n^2 / e^{\Omega(\log^* n)}$$

Table 2: New and previous bounds for $\text{ex}(n, X)$, up to a factor of $\log n$. New near-optimal results are in dark(er) green. TX denotes Theorem X in this paper and [X] denotes reference X in this paper. For example T16:[2] denotes a lower bound that appears in Theorem 16 and an upper bound due to Braß [2].

2.3 Linear-Sized Sets Using Only a Single Configuration

Since it is not explicitly stated in previous work, and we need it to complete our table, we now observe that for any configuration $x \in \{\boxtimes, \heartsuit, \spadesuit, \clubsuit, \diamondsuit, \heartsuit, \spadesuit, \clubsuit, \diamondsuit, \heartsuit, \spadesuit, \clubsuit, \diamondsuit\}$, one can create a linear-sized set of triangles that avoids all configurations except x .

Theorem 2. *For any $X \subsetneq \{\boxtimes, \heartsuit, \spadesuit, \clubsuit, \diamondsuit, \heartsuit, \spadesuit, \clubsuit, \diamondsuit, \heartsuit, \spadesuit, \clubsuit, \diamondsuit\}$, $\text{ex}(n, X) \in \Omega(n)$.*

Proof. Let $x \in \{\boxtimes, \heartsuit, \spadesuit, \clubsuit, \diamondsuit, \heartsuit, \spadesuit, \clubsuit, \diamondsuit, \heartsuit, \spadesuit, \clubsuit, \diamondsuit\}$ be a configuration not in X . Label the vertices of our convex n -gon $1, \dots, n$ in counterclockwise order. Depending on the value of x , we use one of the following constructions (see Figure 1):

1. For $x = \boxtimes$, we use the set of triangles $\{(1, 2, i) : i \in \{3, \dots, n\}\}$.
2. For $x = \heartsuit$, we use the set of triangles $\{(1, 2i, 2i + 1) : i \in \{1, \dots, \lfloor n/2 \rfloor - 1\}\}$.
3. For $x = \spadesuit$, we use the set of triangles $\{(1, i, n + 2 - i) : i \in \{2, \dots, \lfloor n/2 \rfloor\}\}$.
4. For $x = \clubsuit$, we use the set of triangles $\{(1, i, \lfloor n/2 \rfloor + i) : i \in \{2, \dots, \lfloor n/2 \rfloor\}\}$.
5. For $x = \diamondsuit$, we use the set of triangles $\{(3i - 2, 3i - 1, 3i) : i \in \{1, \dots, \lfloor n/3 \rfloor\}\}$.
6. For $x = \heartsuit$, we use the set of triangles $\{(i, \lfloor n/3 \rfloor + 2i - 1, \lfloor n/3 \rfloor + 2i) : i \in \{1, \dots, \lfloor n/3 \rfloor\}\}$.
7. For $x = \spadesuit$, we use the set of triangles $\{(i, \lfloor n/3 \rfloor + i, \lfloor 2n/3 \rfloor + i) : i \in \{1, \dots, \lfloor n/3 \rfloor\}\}$.

In each case, the size of the set is $\Omega(n)$ and it is straightforward to verify that each pair of triangles in the set forms the configuration x and therefore avoids all configurations in X . □

3 Points of View

In this section we describe different, but equivalent (up to a logarithmic factor), views of the problem. One of these (the dot puzzle view) will be our main line of attack for the most difficult cases.

3.1 The Top/Bottom View

It will be helpful to gain a sense of orientation by considering a top/bottom variant of $\text{ex}(n, X)$ that is defined as follows (see Figure 2). Partition the vertices of a convex n -gon using a horizontal line into a *top half* of size $\lceil n/2 \rceil$ and a *bottom half* of size $\lfloor n/2 \rfloor$. We define $\text{ex}'(n, X)$ analogously to $\text{ex}(n, X)$ except that we only count triangles having one vertex in the bottom half and two vertices in the top half. When studying ex' , each triangle we count has a naturally defined *bottom vertex* in the bottom half and a *left vertex* and *right vertex*, each in the top half.

Clearly $\text{ex}(n, X) \geq \text{ex}'(n, X)$. The following lemma shows that, without losing much precision, we can also upper bound $\text{ex}(n, X)$ by $\text{ex}'(n, X)$.

Lemma 2. *If $\text{ex}'(n, X) \in O(n^c)$, then*

$$\text{ex}(n, X) \in \begin{cases} O(n^c) & \text{if } c > 1, \\ O(n \log n) & \text{if } c = 1. \end{cases}$$

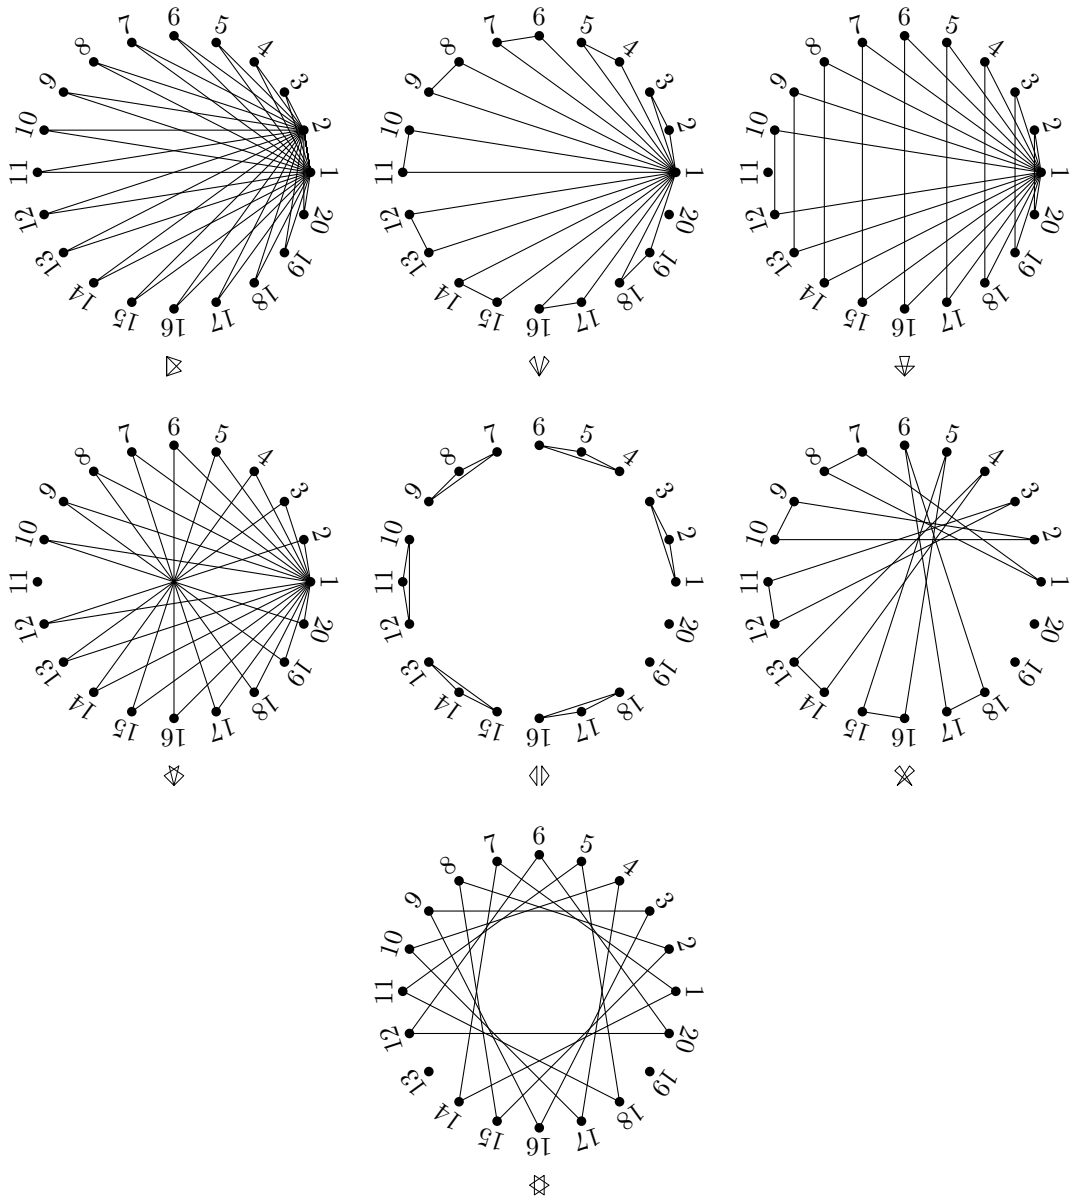


Figure 1: Constructions used in the proof of Theorem 2.

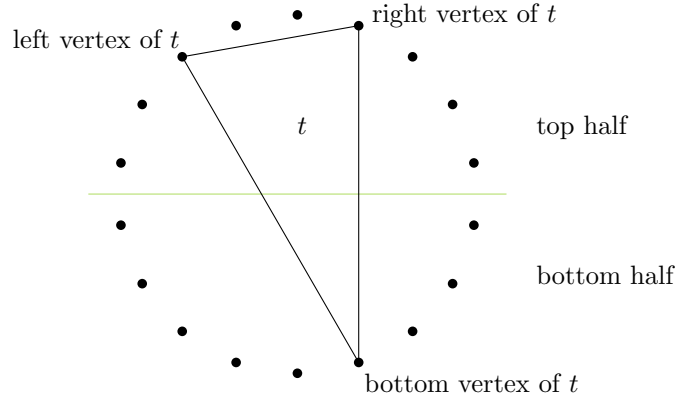


Figure 2: ex' only counts triangles with two vertices in the top half and one vertex in the bottom half.

Proof. Let S be a set of triangles that avoids X . Every triangle in S is of one of the following types:

1. It has one vertex in the top half and two in the bottom half; there are $O(n^c)$ such triangles.
2. It has two vertices in the top half and one in the bottom half; there are $O(n^c)$ such triangles.
3. It has all three vertices in the top half; there are at most $ex(\lceil n/2 \rceil, X)$ such triangles.
4. It has all three vertices in the bottom half; there are at most $ex(\lfloor n/2 \rfloor, X)$ such triangles.

Thus, we obtain the recurrence inequality

$$ex(n, X) \leq O(n^c) + ex(\lceil n/2 \rceil, X) + ex(\lfloor n/2 \rfloor, X),$$

which solves to $O(n^c)$ for $c > 1$ and $O(n \log n)$ for $c = 1$ [4, Section 4.3]. \square

3.2 The Dot Puzzle View

The top-bottom version of the problem gives us a sense of orientation, but it is still difficult to visualize the sets of triangles obtained this way. Next, we show that there is a corresponding puzzle that is easy to visualize. Refer to Figure 3.

In this puzzle, we are given $\binom{n}{2}$ points,

$$Q = \{(x, y) : y \in \{1, \dots, n-1\}, x \in \{y+1, \dots, n\}\}.$$

These points model the top/bottom view on a convex $2n$ -gon, where the point (x, y) represents a triangle whose vertices are some point on the bottom and the x th and y th points on the top, where the top vertices are labelled $1, \dots, n$ from left to right.

The dot puzzle proceeds in n rounds and during the i th round, the player selects a set $Q_i \subseteq Q$ subject to certain constraints that depend on the points selected in rounds $1, \dots, i-1$. In the top/bottom view, the i th round determines which pairs of top vertices




Figure 3: The dot puzzle view of the top/bottom view. In this example, four rounds of the Dot Puzzle have been played.

form a triangle with the i th bottom vertex, where the bottom vertices are labelled $1, \dots, n$ from right to left.

Of course, the constraints on which points can be selected during round i depend on the set of forbidden configurations and the set $\bigcup_{j=1}^{i-1} Q_j$ of points played during previous rounds. By proving bounds on $\sum_{i=1}^n |Q_i|$ we obtain bounds on the maximum number of triangles obtained in the top-bottom view on a set of $2n$ points, i.e., bounds on $ex'(2n, X)$.

Figure 4(a) shows restrictions on the locations of points placed during a single round. It is interpreted as follows: If the central point, $p = (x, y)$, is placed during round i , and we wish to avoid some particular configuration, c , then we should not place any points in the parts of the figure that have label c . For instance, if we wish to avoid the $c = \boxtimes$ configuration, then we should not place any points in the same row or column as p ; such a point creates a \boxtimes configuration in which the shared edge joins a bottom vertex to a left (same row) or right (same column) vertex.

Figure 4(b) shows the restrictions on the locations of points placed in subsequent rounds. Its interpretation is similar Figure 4(a). For example, if we wish to avoid a ∇ configuration and we place the central point, p , during round i , then, in every round $j > i$, we should not place any point directly to the left or directly below p . Any such point creates a ∇ configuration in which the shared vertex is the left vertex (to the left of p) or the right vertex (below p) of both triangles.

For any $X \subseteq \{\boxtimes, \diamond, \nabla, \nabla, \nabla, \nabla, \nabla, \nabla, \nabla\}$ and any $S \subseteq Q$, we define $\kappa(X, S)$ as the subset of Q that can no longer be played in the dot puzzle game (for configurations in X) if the points in S have been played in previous rounds (these are the points of Q killed by S). We use the complementary notation $\bar{\kappa}(X, S) = Q \setminus \kappa(X, S)$ to be the subset of points in Q that can still be played in the dot puzzle game if the points is S have been played in previous rounds.





Figure 4: The regions killed by forbidden configurations during (a) the current round and (b) subsequent rounds.

3.3 Some Warm-Up Exercises

For the remainder of the paper, we will study ex' using the dot puzzle view. Thus, all of our results are bounds on solutions to these dot puzzles.


We say that a point set is *non-decreasing* (respectively, *non-increasing*) if, when sorted lexicographically, the y -coordinates of the points form a non-decreasing (respectively, non-increasing) sequence. A point set is *increasing* (respectively, *decreasing*) if it is non-decreasing (respectively, non-increasing) and no two of its points have the same x -coordinate or the same y -coordinate.

From Figure 4, some previous upper bounds naturally fall out. Consider Braß's results [2] that $\text{ex}(n, \{\Downarrow\}) \in O(n^2)$. For the game defined by $X = \{\Downarrow\}$, we have the rules:



In particular, these rules imply that points selected during a single round of the dot puzzle must be non-decreasing, and thus at most $2n - 3$ points can be selected take part in Q_i . Thus $\sum_{j=1}^n |Q_j| \leq 2n^2 - 3n$, so $\text{ex}'(n, \{\Downarrow\}) \in O(n^2)$ and the bound $\text{ex}(n, \{\Downarrow\}) \in O(n^2)$ immediately follows from Lemma 2.

Similarly, we can almost recover the result $\text{ex}(n, \{\Leftarrow, \otimes, \Downarrow, \Downarrow\}) \leq n$ of Braß, Rote and Swanepoel [3]. Here, the rules are:



The rule for \Downarrow ensures that the set of points taken during a single round form a non-decreasing point set. The rules for points allowed in subsequent rounds ensure that, after round i any points chosen are not below or to the left of the topmost-rightmost point in Q_i . Taken together, these rules imply that

$$\text{ex}'(n, \{\Leftarrow, \otimes, \Downarrow, \Downarrow\}) \leq \sum_{i=1}^n |Q_i| \leq 3n - 4,$$


since the union of Q_i is a non-decreasing point set (whose size is therefore at most $2n - 3$), and each Q_i shares at most one point with Q_{i+1} . The bound $\text{ex}(n, \{\Leftarrow, \otimes, \Downarrow, \Downarrow\}) \in O(n \log n)$ then follows from Lemma 2.

4 Results Based on Dot Puzzles

After this warm-up, and armed with the dot puzzle view, we are ready to prove some new results. We begin by proving several linear upper bounds. From this point on, each proof of a theorem will begin with a picture, similar to Figure 4, that shows the rules of the dot puzzle considered by the theorem. For a point $q \in Q$, we use the notations $x(q)$ and $y(q)$ to denote the x- and y-coordinates of q .


4.1 Linear Upper Bounds

Theorem 3. $ex'(n, \{\boxtimes, \Downarrow, \Downarrow\}) \in O(n)$.



Proof. Taking the union of the rules for \boxtimes , \Downarrow , and \Downarrow , we obtain the rule which ensures that during subsequent rounds we can not take a point from any column or row used in a previous round. The rules for \boxtimes ensure that the set of points taken in each Q_i includes at most one point in each row (or column). Therefore each new point played can be charged to a unique row, so the total number of points played is at most n . \square

Theorem 4. $ex'(n, \{\Downarrow, \Downarrow, \Downarrow\}) \in O(n)$.



Proof. Observe that if Q_i contains $k > 1$ points, p_1, \dots, p_k in a single column (or row), then each of the k rows (or columns) containing one of these points is completely covered by $\kappa(\{\Downarrow, \Downarrow\}, \{p_1, \dots, p_k\})$, i.e., Q_{i+1}, \dots, Q_n can not contain any points in these rows (or columns) (see Figure 5). Therefore, when summing $\sum_{i=1}^n |Q_i|$, the contribution of points that are not alone in their row or column is at most $2n$. We therefore assume that each Q_i contains at most one point from each row and column. Since the rules for \Downarrow imply that Q_i is non-decreasing, this implies that each Q_i is an increasing set of points.

Let $S = \bigcup_{i=1}^n Q_i$ and notice that S contains at most one point from each row: each Q_i contains at most one point in each row and the first time a point p appears in some row, $\kappa(\{\Downarrow, \Downarrow\}, \{p\})$ eliminates every other point from that row. Therefore $|S| \leq n$. All that remains is to account for multiplicity; a single point in S can appear in more than one Q_i .

Now, because of the rules for \Downarrow , the condition that Q_i is increasing is quite restrictive. In particular, if we consider the last (top rightmost) point, p , of Q_i , then it must be placed so




Figure 5: A step in the proof of Theorem 4.

that $\kappa(\{\emptyset\}, \{p\})$ contains all of Q_i except p and the second-to-last point in Q_i (see Figure 6). That is, Q_i contains $|Q_i| - 2$ points that cannot appear in Q_{i+1}, \dots, Q_n . We can think of Q_i as eliminating $|Q_i| - 2$ points from S , so $\sum_{i=1}^n (|Q_i| - 2) \leq |S| \leq n$, so $\sum_{i=1}^n |Q_i| \leq 3n$. \square





Figure 6: Another step in the proof of Theorem 4.

Our next four upper bounds depend on a simple lemma about forbidden configurations of points. We say that three points $a = (x_0, y_0)$, $b = (x_0, y_1)$, and $c = (x_1, y_1)$ form a Γ -configuration if $y_0 < y_1$ and $x_0 < x_1$.

Lemma 3. *Let S be a subset of $\{1, \dots, n\}^2$ with no three points a , b , and c that form a Γ -configuration. Then $|S| \leq 2n$.*

Proof. If we remove the rightmost point from each row of S , then each column in what remains of S contains at most one point. Otherwise, we could take a to be the lowest point in a column, b to be the highest point in the same column, and c to be the removed rightmost point in b 's row. \square

Theorem 5. $ex'(n, \{\boxtimes, \boxplus, \boxminus\}) \in O(n)$.



Proof. Let $S = \bigcup_{i=1}^n Q_i$ be the set of points played in a solution to the resulting dot puzzle. Note that, by the inclusion of \mathbb{D} , each element of S appears in exactly one Q_i , so $|S| = \sum_{i=1}^n |Q_i|$ is the quantity we are interested in bounding.

Next, we claim that S does not contain any three points a , b , and c forming a Γ -configuration. Refer to Figure 7. Suppose, for the sake of contradiction, that this were not the case and that $a \in Q_i$, $b \in Q_j$, and $c \in Q_k$. We have $a \in \kappa(\{\Downarrow\}, \{b\})$ so $i \leq j$ and the rules for \mathbb{D} implies $i \neq j$, so $i < j$. We also have $b \in \kappa(\{\Downarrow\}, \{c\})$ so $j \leq k$. Finally, we have $c \in \kappa(\{\star\}, \{a\})$, so $k \leq i$. Taken together, this gives the contradiction $i < j \leq k \leq i$. Therefore, S contains no Γ -configuration and applying Lemma 3 then implies that $|S| \leq 2n$. \square





Figure 7: The proof of Theorem 5.

Theorem 6. $\text{ex}'(n, \{\Downarrow, \times\}) \in O(n)$.



Proof. Let $S = \bigcup_{i=1}^n Q_i$. We claim that S has no Γ -configuration so, by Lemma 3, $|S| \leq 2n$. To see this, assume S contains a Γ -configuration a , b , and c . Then the rules for \Downarrow imply that no set Q_i contains both a and c . However, the rules for \times and \Downarrow imply that, if $a \in Q_i$, $b \in Q_j$ and $c \in Q_k$ then

$$i \geq j \geq k \geq i.$$

(See Figure 8). But this is a contradiction, since it implies that $i = j = k$.





Figure 8: The proof of Theorem 6.

Now, for some Q_i , consider a point $p \in Q_i$ with minimum x -coordinate and, in case more than one such point exists, take the one that minimizes $y(p)$. Observe that $\kappa(\{\Downarrow, \Downarrow\}, \{p\}) \supseteq Q_i \setminus \{p\}$. Indeed, every point directly above and every point to the right of p is killed by p or cannot be included in Q_i because of the rule for \Downarrow . Therefore, Q_i eliminates at least $|Q_i| - 1$ points of S . It follows that $\sum_{i=1}^n |Q_i| \leq |S| + n \leq 3n$. \square

Theorem 7. $ex'(n, \{\Downarrow, \Downarrow, \Downarrow\}) \in O(n)$.



Proof. Let $S = \bigcup_{i=1}^n Q_i$. We will first show that, for each $i \in \{1, \dots, n\}$, there are at most two points of Q_i that appear in Q_{i+1}, \dots, Q_n . Refer to Figure 9. Define

$$Q_i^* = \{(x, y) \in Q_i : x \geq \max\{y(p) : p \in Q_i\}\},$$

and let p and r be the top rightmost and bottom leftmost points in Q_i^* , respectively. If there is more than one point in Q_i^* with x -coordinate equal to $x(r)$ (as in Figure 9(a)) then we define q to be the highest such point (note that this includes the case where $p = r$). Otherwise (as in Figure 9(b)), we define q to be the rightmost point with y -coordinate $y(r)$ (note that this includes the case where $q = r$). Now, observe that $\kappa(\{\Downarrow, \Downarrow, \Downarrow\}, \{p, q, r\}) \supseteq Q_i \setminus \{p, q\}$, so p and q are the only points of Q_i that can appear again in Q_{i+1}, \dots, Q_n . (Note that Figure 9 only illustrates the case in which $y(p) = x(r)$; if $y(p) < x(r)$, then even p is contained in $\kappa(\{\Downarrow, \Downarrow, \Downarrow\}, \{p, q, r\})$.)

We can therefore think of Q_i as eliminating $|Q_i| - 2$ points from S , so $\sum_{i=1}^n (|Q_i| - 2) \leq |S|$, which implies that $\sum_{i=1}^n |Q_i| \leq |S| + 2n$. All that remains now is to bound $|S|$.




Figure 9: A step in the proof of Theorem 7.





Figure 10: Another step in the proof of Theorem 7.

For each $i \in \{1, \dots, n\}$, let Q'_i be obtained from Q_i by removing the leftmost point in each row. Let $S' = \bigcup_{i=1}^n Q'_i$. We claim that S' contains no Γ -configuration so, by Lemma 3, $|S'| \leq 2n$. To see why this is so, suppose that S' contains a Γ -configuration a, b , and c . Then, as argued in the proof of Theorem 5, it must be that $a, b, c \in Q'_i$ for some i . However, this contradicts the fact (due to \Downarrow) that Q_i is non-decreasing since the leftmost point, b' , of Q_i in the same row as b is to the left of a (see Figure 10).

Therefore, $|S'| \leq 2n$. Now let $S'' = S \setminus S'$. We claim that S'' also satisfies the conditions of Lemma 3. Indeed, by the same reasoning as above, if there were $a, b, c \in S''$ forming a Γ -configuration, then it must be that $a, b, c \in Q_i \setminus Q'_i$ for some i . But this is a contradiction since b and c are in the same row, and $Q_i \setminus Q'_i$ contains at most one point per row.

Wrapping up, we have $|S| = |S'| + |S''| \leq 4n$ so $\sum_{i=1}^n |Q_i| \leq 6n$. \square

Theorem 8. $\text{ex}'(n, \{\Downarrow, \Downarrow, \Downarrow\}) \in O(n)$.



Proof. Consider the set Q_i played during some round i . The rules for \Downarrow imply that Q_i is non-decreasing. Consider the following subsets of Q_i :

1. the set L_i of points in the leftmost column of Q_i ;
2. the set T_i of points in the topmost row of Q_i ;
3. the set B_i of points in the bottommost row of Q_i ; and
4. the set $N_i = Q_i \setminus (L_i \cup T_i \cup B_i)$.

Let p_i denote the lowest leftmost point of Q_i (the unique point in $B_i \cap L_i$). The rule for \Downarrow implies that every point of N_i is contained in $\kappa(\{\Downarrow\}, \{p_i\})$.

Refer to Figure 11. Observe that, for any point $p \in N_i$, the entire row containing p is killed in the sense that it is contained in $\kappa(\{\Downarrow\}, \{p_i\}) \cup \kappa(\{\Downarrow\}, \{p\})$. Next, let t_i be the topmost point in L_i and observe that, for any point $p \in L_i \setminus \{t_i, p_i\}$, the entire row containing p is killed by $\kappa(\{\Downarrow\}, \{p_i\}) \cup \kappa(\{\Downarrow\}, \{t_i, p\})$.

Consider the operation of removing the leftmost point from each row of Q_i , for each $i \in \{1, \dots, n\}$. We claim that this removes a total of at most $4n$ points from Q_1, \dots, Q_n . Indeed, by the preceding discussion, if we remove a point $p \in N_i$ or $p \in L_i \setminus \{t_i, p_i\}$ then this point can be charged to a row that is never used again in Q_{i+1}, \dots, Q_n . For each round $i \in \{1, \dots, n\}$, there are only three other choices for p : $p = p_i$, $p = t_i$, or p is the leftmost point in T_i . Thus, we can charge each row for removing at most one point and each round for removing at most 3 points.

For each $i \in \{1, \dots, n\}$, let Q'_i be the subset of Q_i obtained by removing the leftmost point in each row and observe that Q'_i is an increasing set of points. By the preceding discussions $\sum_{i=1}^n |Q_i| \leq 4n + \sum_{i=1}^n |Q'_i|$. Let $S = \bigcup_{i=1}^n Q'_i$. We claim that S contains no Γ -configuration. To see why this is so, observe that, since each Q'_i is an increasing set, if two points $a \in Q'_i$ and $b \in Q'_j$ of S are in the same column then $i \neq j$. Assume b is above a , then $a \in \kappa(\{\Downarrow\}, \{b\})$ so $i < j$. Now, if $c \in Q'_k$ is to the right of b , then $b \in \kappa(\{\Downarrow\}, \{c\})$, so $j \leq k$. Therefore, $i < j \leq k$, but this is not possible since $c \in \kappa(\{\Downarrow\}, \{a\})$, so $k \leq i$. Therefore, by Lemma 3, $|S| \leq 2n$.

All that remains is to account for points in S that are played multiple times. In each Q_i there are at most two points that can be played in subsequent rounds: the rightmost point in B_i and the rightmost point in T_i . We charge each occurrence of such repeated points to the rounds in which they are these extreme points. In this way, each round is charged for at most two such points and the total contribution of these points to $\sum_{i=1}^n |Q_i|$ is at most $2n$.




Figure 11: The proof of Theorem 8.




Figure 12: The proof of Lemma 4.

In summary,

$$\sum_{i=1}^n |Q_i| \leq 4n + \sum_{i=1}^n |Q'_i| \leq 4n + 2n + |S| \leq 8n. \quad \square$$


4.2 Forbidding Swords

Next we focus on the configuration \bowtie and give linear upper bounds on $\text{ex}'(n, \{\bowtie, \boxtimes\})$ and $\text{ex}'(n, \{\bowtie, \boxplus\})$. We begin with another lemma about forbidden configurations of points that is similar to Lemma 3. We say that a point (x_i, y_i) *SE-dominates* a point (x_j, y_j) if $x_i > x_j$ and $y_i < y_j$. We say that three points $a = (x_0, y_0)$, $b = (x_0, y_1)$, and c form an *obtuse-L-configuration* if $y_1 < y_0$ and c SE-dominates b .

Lemma 4. *Let S be a subset of $\{1, \dots, n\}^2$ with no three points $a, b,$ and c forming an obtuse-L-configuration. Then $|S| \leq 3n$.*

Proof. Refer to Figure 12. Consider the *Pareto boundary* $P \subseteq S$ containing each point of S that is not SE-dominated by any other point in S . The set P is non-decreasing, so it has size at most $2n$. We claim that the set $S \setminus P$ has at most one point in each column, so $|S| \leq 3n$. To see why this claim is true, observe that if some column of $P \setminus S$ contains two points a and b with a above b , then at least one point c in P SE-dominates b , so that a, b , and c would form the obtuse- L -configuration. \square


Theorem 9. $\text{ex}'(n, \{\boxtimes, \boxtimes\}) \in O(n)$.



Proof. Let $S = \bigcup_{i=1}^n Q_i$. Then the rules for \boxtimes imply that $|S| = \sum_{i=1}^n |Q_i|$, so it suffices to bound $|S|$. We claim that S contains no obtuse- L -configuration so, by Lemma 4, $|S| \leq 3n$.

Suppose there were $a \in Q_i$, $b \in Q_j$, and $c \in Q_k$ forming an obtuse- L -configuration. Now, $a \in \kappa(\{\boxtimes\}, \{c\})$ and $c \in \kappa(\{\boxtimes\}, \{a\})$, so it must be that $i = k$. The same argument, applied to b and c , implies that $j = k$, so $i = j = k$. But this is a contradiction since the rules for \boxtimes imply that $i \neq j$. \square

Theorem 10. $\text{ex}'(n, \{\boxtimes, \boxtimes\}) \in O(n)$.



Proof. Let $S = \bigcup_{i=1}^n Q_i$. We claim that S is non-decreasing. Indeed, the rule for \boxtimes prevents two decreasing points from being played in the same round, while the rule for \boxtimes prevents two decreasing points from being played in different rounds. This implies that $|S| \leq 2n$. What remains is to account for points of S that are played in multiple rounds.

Consider the graph G with vertex set S that contains an edge uw if and only if the x -coordinate of u is equal to the y -coordinate of w . We claim that G is 4-colourable. To prove this, we partition S into two sets A and B and show that each of the graphs $G[A]$ and $G[B]$ induced by A and B is 2-colourable (in fact, $G[A]$ and $G[B]$ are each forests). Thus, if we colour A with colours $\{1, 2\}$ and B with colours $\{3, 4\}$, then we obtain a 4-colouring of G .

Refer to Figure 13. Remove the bottom-most point of S from each column and what remains is the set A . Observe that, since S is non-decreasing, A contains at most one point per row. Imagine directing the edges of $G[A]$ from right to left (top to bottom). This directed graph is obviously acyclic and, since each row contains at most one point of A , has

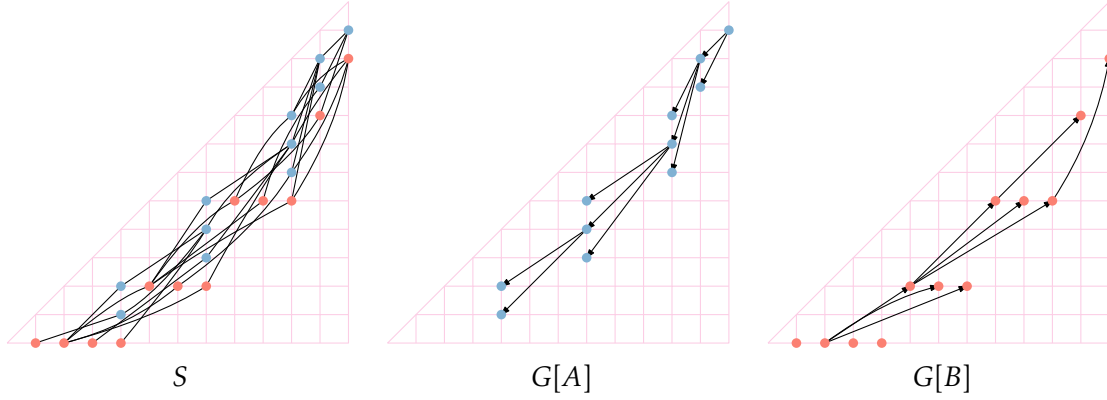


Figure 13: Four-colouring the graph G in the proof of Theorem 10.

maximum in-degree 1. Therefore $G[A]$ is a forest and can be 2-coloured using the colours $\{1, 2\}$. By a similar argument, using the fact that each column contains at most one point of B , the graph $G[B] = G[S \setminus A]$ can be 2-coloured using the colours $\{3, 4\}$.

The resulting 4-colouring of G partitions S into 4 colour classes S_1, \dots, S_4 . We now argue that each of these colour classes contributes $O(n)$ to $\sum_{i=1}^n |Q_i|$. Consider a (new) directed graph $H_j = (S_j, E_j)$ that contains the edge \overrightarrow{uw} if and only if u kills w , i.e., $w \in \kappa(\{\nabla, \otimes\}, \{u\})$. We claim that this graph is *complete*, i.e., for any $u, w \in S_j$ at least one of \overrightarrow{uw} or \overrightarrow{wu} is in E_j . To see why this is so, consider any two distinct points $u, w \in S_j$ with $u = (x_0, y_0)$ and $w = (x_1, y_1)$. Since S (and hence S_j) is non-decreasing, we may assume without loss of generality that $x_0 \leq x_1$ and $y_0 \leq y_1$. There are five cases to consider:

1. $y_0 = y_1$, so $x_0 < x_1$. In this case, $u \in \kappa(\{\nabla\}, w)$, so $\overrightarrow{wu} \in E$.
2. $x_0 = x_1$, so $y_0 < y_1$. In this case, $u \in \kappa(\{\nabla\}, w)$, so $\overrightarrow{wu} \in E$.
3. $x_0 < x_1$, $y_0 < y_1$, and $y_1 > x_0$. In this case, $w \in \kappa(\{\otimes\}, \{u\})$, so $\overrightarrow{uw} \in E$.
4. $x_0 < x_1$, $y_0 < y_1$, and $y_1 < x_0$. In this case, $u \in \kappa(\{\otimes\}, \{w\})$, so $\overrightarrow{wu} \in E$.
5. $x_0 < x_1$, $y_0 < y_1$, and $y_1 = x_0$. This case cannot occur since, in this case, the graph G contains the edge uw , so u and w are assigned different colours and at most one of them is j .

Suppose now that H_j contains a directed cycle $C = u_0, \dots, u_{\ell-1}$. Since the points in this cycle all kill each other, i.e., $u_{k+1 \bmod \ell} \in \kappa(\{\nabla, \otimes\}, \{u_k\})$, it must be the case that all the vertices of C are played in the same round $i' \in \{1, \dots, n\}$ and never played again, i.e., $V(C) \subseteq Q_{i'}$ and $V(C) \cap Q_{j'} = \emptyset$ for all $j' \in \{1, \dots, n\} \setminus \{i'\}$.

Now, if we repeatedly find a cycle in H_j and remove its vertices, we will eventually be left with an acyclic subgraph H'_j with vertex set $S'_j \subseteq S_j$. From the preceding discussion, we know that each cycle vertex we remove contributes only 1 to $\sum_{i=1}^n |Q_i|$:

$$\sum_{i=1}^n |(S_j \setminus S'_j) \cap Q_i| = |S_j \setminus S'_j|.$$

Finally, we are left with the complete acyclic subgraph H'_j with vertex set S'_j and

whose topological sort order we denote by \prec . Now, if Q_i contains vertices $v_1 \prec \dots \prec v_k$ of H'_j , then v_1 kills all of v_2, \dots, v_k so that these vertices can not appear in any $Q_{i'}$ with $i' > i$. This implies that


$$\sum_{i=1}^n (|S'_j \cap Q_i| - 1) \leq |S'_j|,$$

so $\sum_{i=1}^n |S'_j \cap Q_i| \leq |S'_j| + n$. Putting everything together, we have

$$\begin{aligned} \sum_{i=1}^n |Q_i| &= \sum_{j=1}^4 \left(\sum_{i=1}^n |Q_i \cap S_j| \right) \\ &= \sum_{j=1}^4 \left(\sum_{i=1}^n (|Q_i \cap (S_j \setminus S'_j)| + |Q_i \cap S'_j|) \right) \\ &\leq \sum_{j=1}^4 (|S_j \setminus S'_j| + |S'_j| + n) \\ &= |S| + 4n \leq 6n. \end{aligned} \quad \square$$

4.3 Monotone Matrices, Tripod Packing, and 2-Comparable Sets

In this section, we discuss $\text{ex}(n, \{\boxtimes, \boxplus\})$:



Determining the asymptotics of $\text{ex}(n, \{\boxtimes, \boxplus\})$ was given explicitly as an open problem in the conclusion of Braß's paper. We spent more than a year working on this problem and this work included computer searches for a variant of the problem played on the square grid $\{1, \dots, n\}^2$.¹ Using the results of these computer searches in the Online Encyclopedia of Integer Sequences [9], we discovered that this problem, when played on the square grid, is equivalent to several other known problems. See Figure 14.

1. *Monotone matrix problem*: How many values from $\{1, \dots, n\}$ can one write in an $n \times n$ matrix, so that each row is increasing from left-to-right, each column is increasing from bottom-to-top, and for each $i \in \{1, \dots, n\}$, the positions of i in the matrix form an increasing sequence?
2. *Tripod packing problem*: A tripod with top $p = (x, y, z) \in \mathbb{R}^3$ is the union of three closed

¹Playing on the square grid does not change the asymptotics of the problem. Any solution for $\{1, \dots, \lfloor n/2 \rfloor - 1\}^2$ can be used as a solution for the triangular grid Q and any upper bound for $\{1, \dots, n\}^2$ is also an upper bound for the triangular grid Q .




Figure 14: The dot puzzle induced by excluding the taco and nested configurations has already been studied under several equivalent formulations.

rays originating at p and directed in the positive x -, y -, and z -directions, i.e.,

$$\text{tripod}(x, y, z) = \bigcup_{0 \leq t < \infty} \{(x+t, y, z), (x, y+t, z), (x, y, z+t)\}.$$

How many disjoint tripods can be packed with tops in $\{1, \dots, n\}^3$?

3. *2-comparable sets of triples problem.* Two triples of integers (a_1, a_2, a_3) and (b_1, b_2, b_3) are *2-comparable* if $a_i < b_i$ for at least two values of $i \in \{1, 2, 3\}$ or $a_i > b_i$ for at least two values of $i \in \{1, 2, 3\}$. What is the largest set, S , of pairwise 2-comparable triples one can make whose entries come from $\{1, \dots, n\}$?

Several simple and natural constructions give lower bounds of $\Omega(n^{3/2})$ for these problems. However, this bound is not tight. A sequence of recursive constructions has steadily raised this lower bound [6, 10, 11, 12, 13]. The current record is held by Gowers and Long [6], who describe a construction of size $\Omega(n^{1.546})$.

Theorem 11 (Gowers and Long [6]). $ex'(n, \{\boxtimes, \boxplus\}) \in \Omega(n^{1.546})$.

The only known upper bound for this problem comes from the fact that a solution to this problem gives a solution to the Ruzsa-Szemerédi induced matching problem [8].

4. *Induced-matching problem:* What is the maximum number of edges in a bipartite graph $G = (A, B, E)$ with $|A| = |B| = n$ such that E can be partitioned into n induced matchings M_1, \dots, M_n ? That is, each M_i is a matching, and for any two edges $e, f \in M_i$ there is no edge in E that joins an endpoint of e to an endpoint of f .

It is simple to verify that if one takes a 2-comparable set of triples $S = \{(a_i, b_i, c_i) : i \in \{1, \dots, m\}\}$ then the bipartite graph $G = (A, B, E)$ with $A = B = \{1, \dots, n\}$ and

$$E = \{(b_j, c_j) : j \in \{1, \dots, m\}\}$$

satisfies the conditions of the induced matching problem, with the partition into matchings given by

$$M_i = \{(b, c) : (i, b, c) \in S\}.$$

Thus, any upper bound for the induced-matching problem is also an upper bound on the size of a 2-comparable set of triples.


Known upper bounds for the induced matching problem are barely subquadratic, with the current record being held by Fox's improved version of the triangle removal lemma [5], which gives an upper bound of $n^2/e^{\Omega(\log^* n)}$. See the discussion, for example, in Gowers and Long [6]. Lower bounds for the induced matching problem are surprisingly high; a result of Behrend [1] can be used to construct n vertex graphs with $n^2/e^{O(\sqrt{\log n})}$ edges that can be decomposed into induced matchings.

Theorem 12 (Fox [5]). $ex'(n, \{\boxtimes, \boxplus\}) \in n^2/e^{\Omega(\log^* n)}$.

While discovering these results, we noticed that the relationships between some of these problems have gone unnoticed. Here we make a few bibliographic notes.

- Braß [2] seems to have been unaware that the question he posed was equivalent to tripod packing and monotone matrices (or, like us, had never heard of these problems).
- Tiskin [13], apparently unaware of the relation between tripod packing and induced matchings, proved an upper bound of $o(n^2)$ for tripod packing. His proof does not depend on any properties of tripod packing that are not also true for induced matchings, and uses the same tools (namely Szemerédi’s Regularity Lemma) as the original upper bounds for the induced matching problem.
- Gowers and Long [6] seem to be unaware that the problem on 2-comparable sets of triples was studied under other names.
- Gowers and Long [6] arrived at 2-comparable sets as a relaxation of a problem (the size of the largest 2-increasing sequence of triples) proposed by Loh [7]. In his discussion of this problem, Loh formulates a restricted version of the induced matching problem, in which the matching must satisfy a certain Σ -free property and expresses hope [7, remark on page 9] that this restricted version has an $O(n^{3/2})$ upper bound. However, solutions for tripod packing correspond to Σ -free induced matchings, so Σ -free induced matchings of size $\omega(n^{3/2})$ are already known.

In our context, the only new observation we have pertains to $\text{ex}'(n, \{\boxtimes, \Downarrow, \Psi, \Downarrow\})$:




We observe that $\text{ex}'(n, \{\boxtimes, \Downarrow\}) \in O(\text{ex}'(n, \{\boxtimes, \Downarrow, \Psi, \Downarrow\}))$ so these two functions therefore have the same asymptotic growth. This comes from the fact that a solution for the dot puzzle of size n resulting from $X = \{\boxtimes, \Downarrow\}$ can be used as a solution for the dot puzzle of size $2(n+1)$ resulting from $X = \{\boxtimes, \Downarrow, \Psi, \Downarrow\}$ by only playing in the lower-right quadrant. When played this way, the extra restrictions caused by Ψ and \Downarrow do not affect the lower-right quadrant.

Theorem 13. $\text{ex}'(n, \{\boxtimes, \Downarrow\}) \in \Theta(\text{ex}'(n, \{\boxtimes, \Downarrow, \Psi, \Downarrow\}))$.

4.4 Lower Bounds


Finally, we finish up with some $\Omega(n^2)$ lower bound constructions. In each case, a matching upper bound follows from one of the results in Braß [2]. The following are essentially “proofs by figure” in which a brief description of the solution is illustrated alongside the rules of each dot puzzle. In order to avoid floors and ceilings, we assume n is even.

Theorem 14. $\text{ex}(n, \{\boxtimes, \boxtimes, \Downarrow, \Downarrow, \Downarrow\}) \in \Theta(n^2)$.




Proof. For each $i \in \{1, \dots, n/2\}$, we take Q_i to be all the points of $Q' = \{n/2, \dots, n\}^2$ on the line $\{(x, y) : y = 3n/2 - x - i + 1\}$. □

Theorem 15. $\text{ex}(n, \{\bowtie, \Downarrow, \Leftarrow, \star\}) \in \Theta(n^2)$.




Proof. We take $Q_1 = \{(x, y) \in Q : x > n/2, y < n/2\}$ and set $Q_2 = Q_3 = \dots = Q_n = \emptyset$. □

Theorem 16. $\text{ex}(n, \{\star, \Downarrow, \Leftarrow\}) \in \Theta(n^2)$.



Proof. We repeatedly take every second point on the diagonal $y = x - 1$, i.e., for each $i \in \{1, \dots, n\}$, $Q_i = \{(2j, 2j - 1) : j \in \{1, \dots, n/2\}\}$. □

Theorem 17. $\text{ex}(n, \{\Downarrow, \Leftarrow, \Leftarrow\}) \in \Theta(n^2)$.



Proof. We repeatedly take points on the diagonal of the lower-right quadrant, i.e., for each $i \in \{1, \dots, n\}$, we take $Q_i = \{(n/2 + i, i) : i \in \{1, \dots, n/2\}\}$. □

Acknowledgement

Some of this work was carried out at the Fourth Annual Workshop on Geometry and Graphs, held at the Bellairs Research Institute in Barbados, March 6–11, 2016. The authors are grateful to the organizers and to the other participants of this workshop for providing a stimulating working environment.

References

- [1] F. Behrend. On sets of integers which contain no three terms in arithmetic progression. *Proceedings of the National Academy of Sciences USA*, 32:331–332, 1946.
- [2] Peter Braß. Turán-type extremal problems for convex geometric hypergraphs. *Contemporary Mathematics*, 342:25–34, 2004.
- [3] Peter Braß, Günter Rote, and Konrad J. Swanepoel. Triangles of extremal area or perimeter in a finite planar point set. *Discrete & Computational Geometry*, 26(1):51–58, 2001. URL: <http://dx.doi.org/10.1007/s00454-001-0010-6>, doi: 10.1007/s00454-001-0010-6.
- [4] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press and McGraw–Hill, second edition, 2001.
- [5] J. Fox. A new proof of the graph removal lemma. *Annals of Mathematics*, 174:561–579, 2011.
- [6] W. T. Gowers and J. Long. The length of an s -increasing sequence of r -tuples. arxiv:1609.08688, 2016.
- [7] Po-Shen Loh. Directed paths: from Ramsey to Ruzsa and Szemerédi. arxiv:1505.07312v2, 2016.
- [8] I. Z. Ruzsa and E. Szemerédi. Triple systems with no six points carrying three triangles. *Combinatorics (Keszthely, 1976)*, 18:939–945, 1978.
- [9] N. J. A. Sloane. The online encyclopedia of integer sequences. Sequence number A070214. URL: <http://oeis.org/A070214>.
- [10] Sherman Stein. Combinatorial packing of R^n by certain error spheres. *IEEE Transactions on Information Theory*, 30(2):364–368, 1984.
- [11] Sherman K. Stein. Packing tripods. *Mathematical Intelligencer*, 17(2):37–39, 1995.
- [12] Sherman K. Stein and Sándor Szabó. *Algebra and Tiling: Homomorphisms in the Service of Geometry*, volume 25 of *The Carus Mathematical Monographs*. The Mathematical Association of America, 1994.
- [13] Alexandre Tiskin. Packing tripods: Narrowing the density gap. *Discrete Mathematics*, 307(16):1973–1981, 2007. URL: <http://dx.doi.org/10.1016/j.disc.2004.12.028>, doi: 10.1016/j.disc.2004.12.028.