

Unique reconstruction threshold for random jigsaw puzzles

Journal Article**Author(s):**

Nenadov, Rajko; Pfister, Pascal; Steger, Angelika

Publication date:

2017-06-30

Permanent link:

<https://doi.org/10.3929/ethz-b-000221419>

Rights / license:

[Creative Commons Attribution 3.0 Unported](#)

Originally published in:

Chicago journal of theoretical computer sciences 23(1), <https://doi.org/10.4086/cjtcs.2017.002>

Unique reconstruction threshold for random jigsaw puzzles

Rajko Nenadov Pascal Pfister Angelika Steger

Received October 21, 2016; Revised May 22, 2017, and in final form June 29, 2017; Published June 30, 2017

Abstract: A random jigsaw puzzle is constructed by arranging n^2 square pieces into an $n \times n$ grid and assigning to each edge of a piece one of q available colours uniformly at random, with the restriction that touching edges receive the same colour. We show that if $q = o(n)$ then with high probability such a puzzle does not have a unique solution, while if $q \geq n^{1+\varepsilon}$ for any constant $\varepsilon > 0$ then the solution is unique. This solves a conjecture of Mossel and Ross (*Shotgun assembly of labeled graphs*, arXiv:1504.07682).

Key words and phrases: graph reconstruction, jigsaw puzzle

1 Introduction

Recently, Mossel and Ross [9] considered the problem of reconstructing a (labelled) graph given the neighbourhood of radius r of each vertex. This problem is tightly connected to the well-known *reconstruction conjecture* of Kelly, Ulam, and Harary (see [6], [12], [5]): every graph G on at least four vertices can be reconstructed from the set of all of its *vertex-deleted subgraphs*, where a vertex-deleted subgraph is formed by deleting exactly one vertex and its incident edges from G . To make up for the lack of information of the r -neighbourhoods compared to the family of vertex deleted subgraphs, Mossel and Ross made the additional assumption that either the labels or the graph structure are random, and then asked how large the radius r needs to be so that the original graph is uniquely reconstructable. Some well-studied instances of this problem include

- *DNA shotgun assembly*: Shotgun assembly of DNA is the dominant technique used in DNA sequencing. Many randomly located short fragments of the DNA

are extracted, and the goal then is to reconstruct the original DNA sequence from these fragments. One fundamental question is what the minimum number of these randomly located fragments is to get a reliable reconstruction of the original DNA sequence (see e.g. [10], [4], [2]).

- *reconstruction of neural networks*: Shotgun assembly of neural networks was considered in recent work from applied neuroscience. The main goal is to reconstruct a bigger neural network from experimentally observed subnetworks (see e.g. [11]).

In this paper we study another instance of such a problem proposed by Mossel and Ross, namely the *random jigsaw puzzle problem*.

Suppose we are given a square wooden board 1cm thick. One way to make a puzzle out of such a board is as follows:

1. Cut the board into n^2 smaller squares of the same size. We call these squares *pieces*.
2. Arrange the pieces into an $n \times n$ grid and draw a unique, non-symmetrical piece of art on each of them. The resulting big picture is referred to as the *puzzle picture*.
3. Colour each side (other than the ones facing up and down) of a piece with one of the q available colours such that the touching sides of two adjacent pieces receive the same colour.

Suppose now that we shuffle all the pieces and give them to a friend. This friend would like to find out what the puzzle picture is, and the only way she can do it is by arranging the pieces back into an $n \times n$ grid such that the touching sides of each two adjacent pieces have the same colour (we call such an arrangement *valid*). Here we allow each piece to be arbitrarily rotated, however, once arranged in the grid, each piece should be placed such that the drawing on it is visible (i.e. flips are not allowed). This is slightly more general than the model used in [9] where no rotations were allowed.

Note that each puzzle created in the above manner contains a valid arrangement, namely the original construction. However, there might be other valid arrangements which, due to the uniqueness and non-symmetry of drawings on pieces, necessarily produce a picture different than the puzzle picture. For example, if $q = 1$, then in fact any arrangement is valid. As the puzzle picture might contain a secret message for our friend, we would like to design a puzzle such that there exists a unique valid arrangement. We make this precise in the next section. A simple way to make a puzzle with a unique solution is to have $q = 2n^2 + 2n$ colours available and use each colour exactly once. However, in this case it is computationally trivial to reconstruct the puzzle picture by simply starting with an arbitrary piece and then putting down any valid piece in each next step. Therefore, to make the puzzle interesting, we would like that (i) there exists exactly one valid arrangement and (ii) the set of colours is not ‘too large’ or, alternatively, most of the colours are used a significant number of times.

It turns out that a *random jigsaw puzzle* satisfies these two properties. A random jigsaw puzzle is created by colouring sides of pieces independently and uniformly at random using one of the q colours, respecting that once a side has been coloured the corresponding side of the adjacent piece gets the same colour. Mossel and Ross [9] studied the following question: how large does q have to be to ensure that a random

jigsaw puzzle has with high probability (w.h.p. meaning with probability $1 - o(1)$ as $n \rightarrow \infty$) only one possible arrangement? They showed that the probability of *unique vertex assembly* goes to 1 if $q = \omega(n^2)$ (as $n \rightarrow \infty$), where a puzzle has unique vertex assembly if it allows exactly one valid assembly, namely the original assembly itself. Moreover, they showed that the probability of having *unique edge assembly* goes to 0 if $q = o(n^{2/3})$, where having unique edge assembly means that in every valid assembly all *edge-positions* (e.g. the edge between the two pieces at positions (i, j) and $(i, j + 1)$), except the edge-positions along the boundary of the grid, have the same colour as in the original puzzle (note that a puzzle having unique vertex assembly immediately has unique edge assembly as well, but a puzzle with two identical pieces will not have unique vertex assembly, while it may have unique edge assembly). Furthermore, they conjectured that there exists a constant c such that for all $\varepsilon > 0$ the probability of unique reconstruction is $o(1)$ (resp. $1 - o(1)$) if $q \leq n^{c-\varepsilon}$ (resp. $q \geq n^{c+\varepsilon}$). From Mossel [8] we learned that he and Uri Feige conjecture that $c = 1$ is the correct value. We answer this conjecture in the affirmative.

Theorem 1.1. *Let $\varepsilon > 0$ be a constant. If $q = o(n)$ (resp. $q \geq n^{1+\varepsilon}$), then the probability that a random jigsaw puzzle has a unique vertex assembly (up to rotation of the whole grid) goes to 0 (resp. 1) as $n \rightarrow \infty$.*

Note that Bordenave, Feige and Mossel independently obtained the same results [3]. Moreover, Martinsson [7] recently showed that a random jigsaw puzzle does with high probability not have unique edge assembly whenever $q \leq \frac{2}{\sqrt{e}}n - \omega(\log_2(n))$.

Our proof is elementary, essentially based only on first moment arguments. We refer the reader to [1] for an introduction to probabilistic tools. Note that a simple calculation shows that for $q = n^{1+\varepsilon}$ all colours are used asymptotically the same number of times, thus giving a simple randomized procedure for creating a puzzle satisfying (i) and (ii).

2 Notation and preliminaries

We first introduce some notation. A jigsaw puzzle is given by an $n \times n$ grid G of *puzzle pieces*, which we denote by V . We use $v_{i,j} \in V$ to denote the piece corresponding to the location (i, j) in the grid. Each piece is enclosed by four *half-edges* (top, right, bottom and left side of the piece). Two *adjacent* pieces share a pair of *adjacent* half-edges (e.g. the top half-edge of the lower piece and the bottom half-edge of the upper piece), which we call an *edge*. Observe that half-edges along the border of the grid have no other half-edge adjacent to them but, for simplicity, we assume that these ‘single’ half-edges form an edge as well. We denote by E_G the set of edges in the grid G . Given a set Q , a *colouring* of G is a function $g : E_G \rightarrow Q$ that assigns colours to edges and, in turn, to half-edges in the natural way (i.e. half-edges corresponding to an edge inherit the colour). Here we consider the case where g is a random function: the colour of each edge is chosen independently and uniformly at random from Q .

Given an $n \times n$ grid G and a colouring function g , one can do two things to reassemble the puzzle in a non-trivial way: reorder the pieces and rotate them. A reordering of the pieces corresponds to a bijection $\phi : n \times n \rightarrow n \times n$. A rotation of a piece corresponds to a cyclic shift of its edges. Therefore, we can describe rotations of all pieces with a function Π which assigns any piece $v_{i,j} \in V$ a cyclic shift $\pi_{i,j}$. Any two such functions

then define another $n \times n$ grid of pieces $H = G(\phi, \Pi)$, which we call a *reconstruction*. For a given reconstruction H , we denote by E_H the set of its edges (recall that an edge is composed of the half-edges of adjacent pieces). We say that H is a *valid* reconstruction if all edges in E_H are monochromatic, that is, every two half-edges corresponding to an edge have the same colour. Using this notation, Theorem 1.1 can be formally stated as follows.

Theorem (Theorem 1.1 restated). Let $\varepsilon > 0$ be a constant. If $q = o(n)$ (resp. $q \geq n^{1+\varepsilon}$) and $g : E_G \rightarrow Q$ is a random colouring of an $n \times n$ jigsaw puzzle G with $|Q| = q$ colours, then w.h.p there exists (resp. does not exist) a valid reconstruction $H = G(\phi, \Pi)$ such that $E_H \neq E_G$, up to rotation of the whole grid.

For any two sets $S, T \subset V$ we denote by $E_G(S, T)$ the edges for which one half-edge belongs to a piece of S and the other belongs to a piece of T , and we define $E_G(T) := E_G(T, T)$. We call T a connected set if the vertices in T induce a connected subgraph of the grid, where connectivity is defined as in the graph theoretic setting. Figure 1 illustrates this and some more basic definitions that we use throughout the paper.

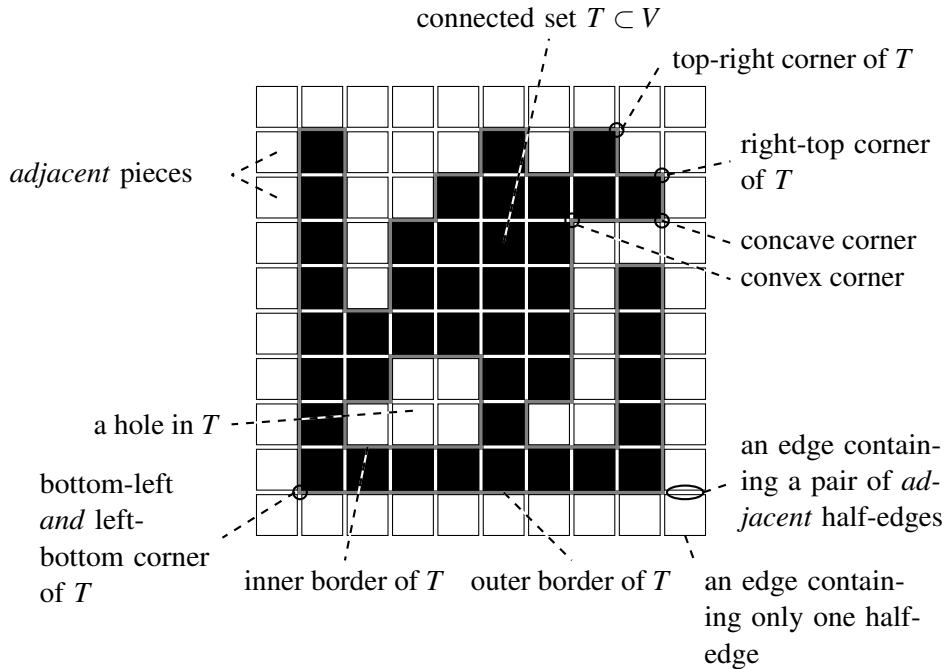


Figure 1: A 10×10 grid G with a connected set $T \subset V$

Our first lemma states some basic observations about connected sets.

Lemma 2.1. *Let T be a connected set of a given $n \times n$ grid G and let R be the bounding rectangle of T . Then the following holds:*

- (i) *The outer border of T contains exactly four more concave corners than convex corners.*

- (ii) Consider the left side of R and the part of T 's border corresponding to this side, i.e. the border part between the top-left and bottom-left corner (in counter-clockwise direction, see figure 2). Then this part of T 's border contains the same number of concave and convex corners.

Proof. (i): We will prove this statement by induction on the length of the outer border. As induction basis, note that a connected set consisting of one piece has an outer border with 4 concave corners and no convex corner. For the induction step, let T be a connected set with an outer border of size k . By induction assumption we know that the outer border of any connected structure with border size less than k contains four more concave than convex corners.

Let us consider the bounding rectangle R of T . Whenever the border of T ‘leaves’ the border of R and ‘comes back again’, an indentation is formed (see Figure 2). If this leaving and coming back happens on the same side of R (see Figure 2, indentation I), we know by induction that the border of the indentation contains four more concave than convex corners. Note that all convex corners of the indentation border are concave corners of the border of T and vice versa, except for the two corners incident to the border of R , which are convex corners of T and of the indentation. Hence, there are as many concave as convex corners in the border part between the two corners where the border of T leaves the border of R and comes back. We can thus fill in the indentation to get a set T' with the same difference between convex and concave corners as T but with a smaller outer border. The claim thus follows from the induction hypothesis.

If there is an indentation that leaves the border of R on one side and comes back on the next side of R (see Figure 2, indentation I'), then a similar argument shows that the border of T between the two corners where it leaves the border of R and comes back contains one more concave corner – and filling in the indentation concludes the proof similar as before.

The only case remaining is that the border of T contains no indentation at all, implying that T and R are identical and thus the claim trivially holds.




Figure 2: Indentations

- (ii): We call the border part of T between the top-left and bottom-left corner the *left boundary* and the rest of the border of T the *right boundary*. If the left boundary of T is not just a straight line, by filling all the indentations along the right boundary, we

get a set T' with four more concave than convex corners (by (i)). But, since the right boundary of T' coincides completely with the bounding rectangle R , it contains two concave and no convex corner. Thus, since the top-left and the bottom-left corner are concave corners as well, it holds that the left boundary contains the same number of concave and convex corners.

Observe that this statement is also true for the border parts of T corresponding to the right, upper and lower side of R . \square

3 Proof of Theorem 1.1, $q = o(n)$

The non-uniqueness comes essentially from the birthday paradox, cf. [8]. We add the short proof for completeness. The main idea is that w.h.p. there exist two pieces in V that have the exact same 4-tuple of colours assigned to their half-edges. These pieces can be interchanged resulting in at least two different valid reconstructions.

Note that a random colouring function g assigns the same colour to every two adjacent half-edges, thus we do not have the property that 4-tuples of colours assigned to pieces are chosen independently from Q^4 . One can circumvent this by only considering every other piece. More precisely, let

$$V' := \{v_{i,j} \in V \mid i + j \text{ is even}\}.$$

Note that $|V'| = \lceil n^2/2 \rceil$ and that V' is a set of pairwise non-adjacent pieces (it consists of all ‘black’ fields of an $n \times n$ chessboard). Since the pieces in V' are pairwise non-adjacent, it is easy to see that the corresponding 4-tuples of colours are mutually independent. Since the colouring of the half-edges of the pieces in V' is mutually independent, we can assume that the pieces receive their colour one after the other independently and uniformly at random. Conditioning on the event that there are no two identical pieces among the first k pieces, the probability that the $(k + 1)$ st piece is identical to one of the first k pieces is $(1 - \frac{k}{q^4})$. Therefore, we have that

$$\Pr[\text{no two identical pieces}] = \prod_{k=1}^{\frac{n^2}{2}-1} (1 - \frac{k}{q^4}) \leq e^{-\sum_{k=1}^{\frac{n^2}{2}-1} \frac{k}{q^4}} = e^{-\frac{n^4-2n^2}{8q^4}} = o(1),$$

where we used that $1 - x \leq e^{-x}$ and $q = o(n)$. This means that with probability $1 - o(1)$ there are two pieces in V' which can be swapped and thus the puzzle is w.h.p. not uniquely reconstructable. Note that this argument still applies if we do not allow rotations.

4 Proof of Theorem 1.1, $q \geq n^{1+\varepsilon}$

A straightforward first moment argument shows that we do not expect two identical pieces. Indeed, let $X_{(i,j),(i',j')}$ be the indicator variable for the event that $v_{i,j}$ and $v_{i',j'}$ are identical. Then the probability that this variable is one is at most $4/q^4$ if the two pieces are not adjacent in G (here the constant 4 takes care of the rotations) and at most $4/q^3$ if $v_{i,j}$ and $v_{i',j'}$ are adjacent. As we have less than n^4 such variables and only

$O(n^2)$ adjacent pairs, the expected number of identical pieces is $o(1)$. From Markov's inequality it thus follows that w.h.p. we do not have any pair of identical pieces.

Clearly, this argument shows that we cannot obtain a valid reconstruction by swapping only two pieces. Before we show that a similar statement holds for 'bigger components', we need some more definitions. Two half-edges which are adjacent in G are called *partners*. Let $H = G(\phi, \Pi)$ be a reconstruction. An edge $e \in E_H$ is called an *original edge* if its two half-edges are partners, i.e. if $e \in E_G$ as well. Else, we call it a *new edge*. Furthermore, let $T \subset V$ be a connected set in H . We call T *stable* if and only if all edges in $E_H(T)$ are original edges and all edges in $E_H(T, V \setminus T)$ are new edges. On the one hand this implies that T is a connected set in G as well (but not necessarily at the same place, and it might be rotated) and, on the other hand, that all pieces along the border of T in the reconstruction H are either distinct from the corresponding piece in G or are at least rotated.

We show that certain 'undesired' structures, which can be thought of as a part of the border of a stable set T , w.h.p. cannot be a part of a valid reconstruction. By defining these structures appropriately, this will in turn allow us to conclude that the puzzle has a unique reconstruction. For example, assume that T is a 1×4 subgrid of G and assume that we attach to one of its longer sides another 1×4 subgrid, different than the one which appears in G . Then the probability that all four edges are monochromatic is $1/q^4$. Thus, since there are at most $O(n^2)$ subgrids of constant size in G , the expected number of such valid configurations is bounded by $O(n^2 \cdot n^2/q^4) = o(1)$. By Markov's inequality, such valid configurations do not appear.

Note that, in general, it is not true that we get a factor of $1/q$ for each new edge. Consider, for example, two edges (a, a') and (b, b') that are rearranged to edges (a, b') and (b, a') . Clearly, the probability that they are both monochromatic is just $1/q$ and not $1/q^2$. Before we describe the undesired configurations, we first give two propositions to handle such situations. The first proposition handles the case when we can order the given new edges in such a way that we can bypass the dependencies of the colouring of the half-edges, whereas the second proposition handles the general case.

The idea of the proofs of these propositions is the so-called principle of deferred decision. That is, we do not reveal the colouring of the original edges all at once, but instead we reveal them whenever we need to know whether a new edge is monochromatic. Clearly, this implies that whenever we have so far revealed at most one of the half edges of a new edge, the probability that the edge is monochromatic is exactly $1/q$. Of course, revealing the colour of the other half-edge means that we also reveal the colour of the partner of that half-edge. Therefore, the order in which we reveal colours plays an essential part in obtaining good probability bounds.

Proposition 4.1. *Let $U = \{e_1, \dots, e_m\}$ be a set of new edges and, for all $i \in [m]$, let h_i^1 and h_i^2 be the two half-edges of e_i . Assume that the colours of all half-edges are still unknown and that at most one of the partners of h_i^1 and h_i^2 is an element of $\bigcup_{k=1}^{i-1} e_k$, for every $i \in [m]$. Then*

$$\Pr[\text{all edges in } U \text{ are monochromatic}] = \frac{1}{q^m}.$$

Proof. We consider the edges e_1, \dots, e_m one after the other, in this order. By the assumption that at most one of the partners of h_i^1 and h_i^2 is an element of $\bigcup_{k=1}^{i-1} e_k$, when

we consider e_i , the colour of at least one of its half-edges is still unknown. Thus, the probability that e_i is monochromatic is $1/q$. As all these events are independent we get that the probability that all the edges of U are monochromatic is equal to $1/q^m$. \square

In the general case, we will not be able to order the edges so that we can bypass the dependencies of the colouring. Moreover, as we will later apply Propositions 4.1 and 4.2 repeatedly (and to different sets of edges), we will not be able to always assume that the colours of all half-edges are still unknown. In Proposition 4.2 we thus allow that the colour of some half-edges are already known.

Proposition 4.2. *Let $U = \{e_1, \dots, e_m\}$ be a set of new edges and, for all $i \in [m]$, let h_i^1 and h_i^2 be the two half-edges of e_i . If for every $i \in [m]$ at most one of the colours of h_i^1 and h_i^2 is already known, then*

$$\Pr[\text{all edges in } U \text{ are monochromatic}] \leq \frac{1}{q^{\lceil m/2 \rceil}}.$$

Proof. We consider the edges in U one by one, but not necessarily in the order e_1, \dots, e_m . We call an edge *open* if the colour of both its half-edges is still unknown, *critical* if the colour of exactly one of its half-edges was already revealed, and *closed* otherwise. By assumption, all edges in U are either open or critical in the beginning. We argue that we can consider $\lceil m/2 \rceil$ edges of U in such an order that in each step we consider an open or critical edge (and thus get a factor of $1/q$) and only the considered edge and at most one additional edge change their status to closed. Clearly, this will imply the desired bound.

If there is at least one critical edge e_i in U , we consider it and reveal the colour of its so-far uncoloured half-edge. The probability that e_i is monochromatic is thus equal to $1/q$. By revealing the colour of the half-edge we also reveal the colour of its partner, which may imply that one (but only one!) more edge in U changes its status to closed.

Otherwise, if U contains only open or closed edges and at least one open edge, then revealing the colours of both half-edges of an open edge again implies that the probability that this edge is monochromatic is $1/q$ and, again, either at most one more edge can change its status to closed or at most two can change the status to critical. \square

The proof of Theorem 1.1 relies on the concept of *configurations*. We start by describing the general setup. Let C_1, \dots, C_m be disjoint sets which are connected in G such that $C_i \cup C_j$ is not a connected set (for distinct $i, j \in [m]$) and $|\bigcup_{i=1}^m C_i| \leq K(\varepsilon)^2$, where $K(\varepsilon)$ is a constant which we define later. In the cases we consider, we always reassemble the pieces into a connected set $S \cup U = \bigcup_{i=1}^m C_i$ such that

- $S \subseteq C_1$ is a connected set in G , and
- all edges between S and U are new edges.

We call an $(m+1)$ -tuple $(C_1, \dots, C_m, \text{rule})$ a *configuration*, where *rule* tells us how to rearrange and rotate the pieces of C_1, \dots, C_m . Thus *rule* consists of two functions ϕ' and Π' , where $\phi': \bigcup_{i=1}^m C_i \rightarrow K(\varepsilon) \times K(\varepsilon)$ is an injective function specifying the rearrangement of the pieces and Π' assigns any piece of $\bigcup_{i=1}^m C_i$ a cyclic shift.

In the following, we will consider five *types* of configurations, where the type of the configuration indicates the ‘shape’ of the sets S and U . Figure 3 illustrates the




Figure 3: Illustration of different types of configurations $(C_1, \dots, C_m, \text{rule})$; the black squares depict the pieces in S , while the white squares depict the pieces in U

first four types. In this figure the set S is composed of black pieces and the set U is composed of certain pieces next to S (we will make this precise for each type below). Intuitively, one can think of S as being a part of the border of a stable set and U as a set of pieces along this border that all define new edges with S . Note that for some types, namely *straightline*, *convexcorners* and *subsquare*, we additionally impose a lower bound on the size of S and/or U . Moreover, *subsquare* imposes some further technical conditions on induced stable sets, which will be discussed later. We call a configuration $(C_1, \dots, C_m, \text{rule})$ *valid* if all edges inside $S \cup U$ are monochromatic. We denote by $X_{(C_1, \dots, C_m, \text{rule})}$ the indicator random variable for this event.

Our aim is to show that the expected number of valid configurations $(C_1, \dots, C_m, \text{rule})$ of all five types in a random jigsaw puzzle is $o(1)$, since then Markov's inequality implies that w.h.p. none of these configurations can be found in a valid reconstruction. As we assumed that $|S \cup U| = |\bigcup_{i=1}^m C_i| = O(1)$, the number of tuples (C_1, \dots, C_m) is bounded by $O(n^{2m})$ and for each such tuple (and each type) we have only $O(1)$ different ways to reassemble the pieces into $S \cup U$. By linearity of expectation it thus suffices to show that $\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] = o(n^{-2m})$. This is what we will do now, details depending on the type of the configuration that we consider. Since the problem is monotone in the number of colours (if we recolour all edges of a fixed colour uniformly at random with the remaining colours, then this increases the likelihood of unique reconstructability), we can assume that $\varepsilon < 1/4$.

straightline: A configuration $(C_1, \dots, C_m, \text{rule})$ of the type *straightline* is defined as illustrated in the top left part of Figure 3. The key characteristics are:

- S is a connected *row* of length $\ell(\varepsilon) := \lceil 3(1 + 1/\varepsilon) \rceil$, and

- U is the row along the lower side of S .

We want to bound the probability that all edges within $S \cup U$ are monochromatic. To do this we apply Propositions 4.1 and 4.2 to certain sets of (new) edges within $S \cup U$. Recall that, by assumption, all edges between S and U are new.

We first consider the leftmost piece in U from every component C_i , $2 \leq i \leq m$. All but at most one of these $m - 1$ pieces have a left neighbour in U . Let $W \subseteq U$ denote the set of $m - 2$ of such pieces and let \mathcal{L} denote the set of edges between the pieces in W and their left neighbour in U together with the edges of these pieces to the set S (i.e. we look at the edges going left and up from each piece in W). Then $|\mathcal{L}| = 2(m - 2)$ and one easily checks that Proposition 4.1 implies that the probability that all these edges are valid is $q^{-|\mathcal{L}|}$.

Next consider the set \mathcal{L}' of edges between the pieces of $U \setminus W$ and S . Note that $|\mathcal{L}'| = |U| - (m - 2)$. Since W does not contain any piece from C_1 , by colouring all edges in \mathcal{L} (and thus the partners of the corresponding half-edges) we have that every edge in \mathcal{L}' still has at least one half-edge for which its colour has not yet been revealed. Indeed, the only way the lower half-edge of a piece in S adjacent to some piece in $U \setminus W$ could have already received a colour is if its partner half-edge belongs to some edge in \mathcal{L} . However, the piece containing such a half-edge then belongs to $C_1 \setminus S$ and, as W does not contain any piece from C_1 , we have not yet revealed half-edges of such pieces. Proposition 4.2 thus implies that the probability that all these edges are valid is at most $q^{-|\mathcal{L}'|/2}$. Therefore, as $q = n^{1+\varepsilon}$ and $|U| = |S| = \lceil 3(1 + 1/\varepsilon) \rceil$ we obtain

$$\begin{aligned} \Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] &\leq q^{-|\mathcal{L}|} \cdot q^{-|\mathcal{L}'|/2} \\ &= q^{-2(m-2) - (|U| - (m-2))/2} \\ &= q^{3-2m - (|U| - m)/2} \\ &\leq \begin{cases} q^{3-2m - |U|/4}, & \text{if } m \leq |U|/2 \\ q^{3-2m}, & \text{if } m > |U|/2 \end{cases} \\ &\leq \begin{cases} n^{-2m+3(1+\varepsilon) - 3(1+\varepsilon)(1+1/\varepsilon)/4}, & \text{if } m \leq |U|/2 \\ n^{-2m+3(1+\varepsilon) - 2\varepsilon m}, & \text{if } m > |U|/2 \end{cases} \\ &= o(n^{-2m}), \end{aligned}$$

where the last step follows from the assumption that $\varepsilon < 1/4$ and since $2\varepsilon m > 3(1 + \varepsilon)$ if $m > |U|/2 = \lceil 3(1 + 1/\varepsilon) \rceil/2$.

convexcorners: A configuration $(C_1, \dots, C_m, \text{rule})$ of the type **convexcorners** is defined as illustrated in the top right part of Figure 3. The key characteristics are:

- S contains $2\ell(\varepsilon)$ convex corners facing to the right,
- U is the set of pieces which are placed in these corners, and
- $|S| \leq s(\varepsilon) := 4\ell(\varepsilon)^2$.

Since any piece can belong to at most two convex corners, we have $2\ell(\varepsilon) \geq |U| \geq \ell(\varepsilon)$. We bound the probability that all edges within $S \cup U$ are monochromatic.

Let us first consider one piece in U from every component C_i , $2 \leq i \leq m$ and let $W \subseteq U$ be the set of such pieces. Since all these pieces are placed in a convex corner,

each will have at least two edges going to the set S . Let \mathcal{L} denote the set of edges between the pieces in W and S . Then $|\mathcal{L}| \geq 2(m-1)$ and Proposition 4.1 implies that the probability that all these edges are valid is at most $q^{-|\mathcal{L}|}$.

Next consider the set \mathcal{L}' of edges between the pieces of $U \setminus W$ and S . Note that we have $|\mathcal{L}'| \geq 2(|U| - (m-1))$. Similarly as in the straightline case, since W does not contain any piece from C_1 , by fixing the colouring for all edges in \mathcal{L} we have that every edge in \mathcal{L}' still has one half-edge for which the colour has not yet been revealed. Proposition 4.2 thus implies that the probability that all these edges are valid is at most $q^{-|\mathcal{L}'|/2} = q^{-(|U| - (m-1))}$. Therefore, as $q = n^{1+\varepsilon}$ and $|U|/2 \geq \lceil 3(1+1/\varepsilon) \rceil/2$, by similar calculations as in the previous case we obtain

$$\begin{aligned} \Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] &\leq q^{-|\mathcal{L}|} \cdot q^{-|\mathcal{L}'|/2} \\ &= q^{-2(m-2) - (|U| - (m-1))} \\ &= q^{3-2m - (|U| - m)} \\ &\leq \begin{cases} q^{3-2m - |U|/2}, & \text{if } m \leq |U|/2 \\ q^{3-2m}, & \text{if } m > |U|/2 \end{cases} \\ &\leq \begin{cases} n^{-2m+3(1+\varepsilon) - 3(1+\varepsilon)(1+1/\varepsilon)/2}, & \text{if } m \leq |U|/2 \\ n^{-2m+3(1+\varepsilon) - 2\varepsilon m}, & \text{if } m > |U|/2 \end{cases} \\ &= o(n^{-2m}), \end{aligned}$$

where the last step follows from the assumption that $\varepsilon < 1/4 < 1/2$ and since $2\varepsilon m > 3(1+\varepsilon)$ if $m > |U|/2 \geq \lceil 3(1+1/\varepsilon) \rceil/2$.

Note that we did not use that $|S| \leq s(\varepsilon)$ in the proof above. The choice of the constant s will become clear later in the proof.

hole: A configuration $(C_1, \dots, C_m, \text{rule})$ of the type **hole** is defined as illustrated in the bottom left part of Figure 3. The key characteristic are:

- S encloses U , and
- $|U| \leq s(\varepsilon)^2$.

We bound the probability that all edges within $S \cup U$ are monochromatic.

We first consider the top-left piece in U from every component C_i , $2 \leq i \leq m$, i.e. the first piece from each C_i we encounter while traversing U row by row, from left to right. This implies that all these $m-1$ pieces have a new edge going up and left. Let $W \subseteq U$ denote the set of these pieces. Let \mathcal{L} denote the set of edges between the pieces in W and their left and upper neighbour (which are in U or in S). Then $|\mathcal{L}| = 2(m-1)$ and Proposition 4.1 implies that the probability that all these edges are valid is $q^{-|\mathcal{L}|}$.

Next, let \mathcal{L}' be the set of all the edges between U and S facing down or right. Note that $|\mathcal{L}'| = b/2$, where b is the number of edges between S and U (i.e. the border of U). By fixing the colouring for all edges in \mathcal{L} , every edge in \mathcal{L}' still has one half-edge of which the colour has not yet been revealed (again, W does not contain any piece from C_1). Proposition 4.2 thus implies that the probability that all these edges are valid is at most $q^{-|\mathcal{L}'|/2}$. Thus, if $q = n^{1+\varepsilon}$ we have

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] = q^{-|\mathcal{L}|} q^{-|\mathcal{L}'|/2} = q^{2-2m-b/4}.$$

If $b \geq 8$ then we immediately get the desired bound

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] = o(n^{-2m}).$$

Else, since the border of a connected set U is always of even size, we have that $b \leq 6$. Hence U contains at most 2 puzzle pieces and the desired probability can be obtained by a simple case analysis.

indentation: A configuration $(C_1, \dots, C_m, \text{rule})$ of the type `indentation` is defined as illustrated in the bottom right part of Figure 3. The key characteristics are:

- U is enclosed by S from either two or three sides (one can imagine that S together with one or two adjacent sides of the border of the grid enclose U),
- $|U| \leq s(\varepsilon)^2$, and
- each piece from C_1, \dots, C_m is at distance at most $s(\varepsilon)^2$ from the grid border in G .

Due to the third restriction, the number of configurations $(C_1, \dots, C_m, \text{rule})$ of the type `indentation` is $O(n^m)$, since we only have linearly many pieces available to ‘build’ the connected sets C_1, \dots, C_m . It thus suffices to show that $\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] = o(n^{-m})$.

Similarly as in the previous case, we first consider the top-left piece in U from every component C_i , $2 \leq i \leq m$. Let W denote the set of such pieces and let \mathcal{L} denote the set of edges between the pieces in W and their left and upper neighbour (which are in U or in S). Then $|\mathcal{L}| = 2(m-1)$ and, as before, Proposition 4.1 implies that the probability that all these edges are valid is $q^{-|\mathcal{L}|}$. Hence, since at least one new edge is present in $S \cup U$,

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] \leq \min\{1/q, q^{-2(m-1)}\} = o(n^{-m}).$$

subsquare: The key characteristics of the configuration $(C_1, \dots, C_m, \text{rule})$ of the type `subsquare` are:

- S is an empty set,
- U is a $K(\varepsilon) \times K(\varepsilon)$ square where $K(\varepsilon) := 4s(\varepsilon)^2/\varepsilon$, and
- stable sets of U are of size at most $s(\varepsilon)^2$.

In the last property we restrict the notion of a stable set only to the rearrangement U , i.e. we say that a connected subset $U' \subseteq U$ is stable if and only if all edges in $E_U(U')$ are original edges and all edges in $E_U(U', U \setminus U')$ are new edges. Additionally, we denote by $U = \bigcup_{i=1}^z T_i$ the decomposition of U into these restricted stable sets. Note that each C_i is the union of some of the T_j ’s and that each stable set T_j belongs to exactly one C_i .

To show that $\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] = o(n^{-2m})$, we consider the following two cases:

Case 1: $m > 2K(\varepsilon)(1 + 1/\varepsilon)$:

Consider the top-left piece in U from every component C_i , $1 \leq i \leq m$. Note that all such pieces which do not touch the upper or the left border of the square have two new edges, going up and going left. Let $W \subseteq U$ be the set of these pieces and note that $|W| \geq m - 2K(\varepsilon)$. Let \mathcal{L} denote the set of edges between the pieces in W and their left

and upper neighbour. Then $|\mathcal{L}| \geq 2(m - 2K(\varepsilon))$ and Proposition 4.1 implies that the probability that all these edges are valid is $q^{-|\mathcal{L}|}$. Thus, we have that

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] \leq q^{-2(m-2K(\varepsilon))}.$$

Note that $m > 2K(\varepsilon)(1 + 1/\varepsilon)$ implies $2m < 2(1 + \varepsilon)(m - 2K(\varepsilon))$. Hence, since $q = n^{1+\varepsilon}$, it follows that

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] = o(n^{-2m}).$$

Case 2: $m \leq 2K(\varepsilon)(1 + 1/\varepsilon)$:

In this case, instead of considering the sets C_i , we will consider the sets T_i . Consider the top-left piece in U from every stable set T_i , $1 \leq i \leq z$. Note that all such pieces which do not touch the upper or the left border of U have a new edge going up and going left. So let $W \subseteq U$ be a set of these pieces and note that $|W| \geq z - 2K(\varepsilon)$. Let \mathcal{L} denote the set of edges between the pieces in W and their left and upper neighbour. Then $|\mathcal{L}| \geq 2(z - 2K(\varepsilon))$ and Proposition 4.2 implies that the probability that all these edges are valid is $q^{-|\mathcal{L}|/2}$. Thus, we have that

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] \leq q^{-(z-2K(\varepsilon))}.$$

Since $|T_i| \leq s(\varepsilon)^2$ we have $z \geq K(\varepsilon)^2/s(\varepsilon)^2$ and it is easy to check that the choice of K implies $z - 2K(\varepsilon) \geq K(\varepsilon)^2/s(\varepsilon)^2 - 2K(\varepsilon) > 4K(\varepsilon)/\varepsilon$, with room to spare. Hence, since $m \leq 2K(\varepsilon)(1 + 1/\varepsilon)$, it holds that $(1 + \varepsilon)(z - 2K(\varepsilon)) > 2m$. Therefore, if $q = n^{1+\varepsilon}$, we have that

$$\Pr[X_{(C_1, \dots, C_m, \text{rule})} = 1] \leq n^{-(1+\varepsilon)(z-2K)} = o(n^{-2m}).$$

Finally, note that $K(\varepsilon)$ is sufficiently large such that $|\bigcup_{i \in [m]} C_i| \leq K(\varepsilon)^2$ for all five types of rules. With all these preliminaries at hand we are now ready to complete the proof of Theorem 1.1.

Proof of Theorem 1.1, $q \geq n^{1+\varepsilon}$. Recall that a configuration is valid if all new edges defined by the given rule are monochromatic. We showed that the probability of a configuration being valid is $o(n^{-2m})$, respectively $o(n^{-m})$ for configurations of the type indentation. Since there are $O(n^m)$ possible configurations of the type indentation and $O(n^{2m})$ configurations of the other types, and because $m \leq |S \cup U| \leq K(\varepsilon)^2$, by Markov's inequality and a union-bound we obtain that w.h.p. none of the above mentioned configurations are valid. In other words, none of these configurations can appear in a valid reconstruction. Using this assumption, we show that a valid reconstruction is identical to G (up to rotation of the whole grid).

For the rest of the proof, let $H = G(\phi, \Pi)$ be a valid reconstruction of G . We first show that a valid reconstruction contains no stable set T of size at least $s(\varepsilon)^2$.

Let us assume, towards a contradiction, that there exists a stable set T of size at least $s(\varepsilon)^2$ and let R be the bounding rectangle of T . Let us first consider the case where at least one side of R does not coincide with the border of the grid H . Note that this implies that there has to exist a side of R which does not coincide with the border of the grid and has length at least $s(\varepsilon)$. Consider the corresponding part of the border of T , as defined in Lemma 2.1, part (ii). Then this part of the boundary of T

also has size at least $s(\varepsilon)$. Let us denote the pieces along this border part by S . Note that S cannot have a straight line segment of size $\ell(\varepsilon)$ since this would form a valid `straightline` configuration. Since the number of convex and concave corners along S is the same (Lemma 2.1, (ii)) and for every $\ell(\varepsilon)$ consecutive pieces along S we have at least one corner, we conclude that there exists a connected subset $S' \subseteq S$ of size at most $s(\varepsilon) = 4\ell(\varepsilon)^2$ which contains $2\ell(\varepsilon)$ convex corners. However, this forms a valid `convexcorners` configuration, which is again a contradiction.

Therefore, we conclude that if such a set T exists, then the bounding rectangle R has to touch all four borders of the grid H , implying that R is identical to the grid H . Note that in this case T touches all four borders of the grid H as well (see Figure 4 below).




Figure 4: Example of a ‘giant’ stable set T

Consider an indentation U in T and let $S \subseteq T$ be the set of pieces along the border of U (see Figure 4). Note that U has four more concave than convex corners (Lemma 2.1, part (i)). Moreover, there are at most three corners of U which touch the boundary of the grid and thus there are more concave corners lying ‘inside’ of the grid than the total number of convex corners of U . Since every such concave corner ‘seen’ from the set U is a convex corner ‘seen’ from the set S we conclude that S contains more convex than concave corners along its border with U . A similar argument as in the previous case then shows that the set S is of size at most $s(\varepsilon)$ as otherwise there exists either a valid `straightline` or a valid `convexcorners` configuration. Therefore, every indentation U is of size at most $s(\varepsilon)^2$. Similarly, the set S of pieces around a hole is of size at most $s(\varepsilon)$, and thus every hole of T has to be of size at most $s(\varepsilon)^2$. However, since there are no valid `hole` configurations, we conclude that T has no hole at all. Therefore, the pieces of an indentation (like U in Figure 4) cannot originate from the ‘inside’ of the grid, but only from a zone around the border of the grid of depth at most $s(\varepsilon)^2$. Finally, the existence of such indentations is excluded since there are no valid `indentation` configurations. Therefore, since a stable set T touches all four borders of the grid and has no holes and indentations, we conclude that T is the whole grid G .

To summarize, we showed that a valid reconstruction which is not identical to G (up to rotation of the whole grid) contains only stable sets of size smaller than $s(\varepsilon)^2$. However, since there are no valid `subsquare` configurations we in turn conclude that no $K(\varepsilon) \times K(\varepsilon)$ subsquare in a valid reconstruction contains a new edge (as otherwise

it would give a valid subsquare configuration). By shifting a subsquare of size $K(\varepsilon) \times K(\varepsilon)$ through the whole puzzle we thus deduce that E_H contains no new edge at all. That is, H is identical to the original puzzle G up to rotation of the whole grid. \square

To finish, let us make two observations about possible extensions of the proof. First, we see no reason why our arguments should not generalize to higher dimensions. Second, note that the above proof for the upper bound can be slightly adapted so that everything holds for suitable (though very slowly decreasing) $\varepsilon = o(1)$ as well. It would be interesting to determine this more precisely. In particular, it is tempting to conjecture that actually $q = n \cdot \omega(n)$ suffices for unique reconstruction, where $\omega(n)$ is some slowly growing function in n .

References

- [1] NOGA ALON AND JOEL H SPENCER: *The probabilistic method*. John Wiley & Sons, 2015. 3
- [2] RICHARD ARRATIA, DANIELA MARTIN, GESINE REINERT, AND MICHAEL S WATERMAN: Poisson process approximation for sequence repeats, and sequencing by hybridization. *Journal of Computational Biology*, 3(3):425–463, 1996. 2
- [3] CHARLES BORDENAVE, URIEL FEIGE, AND ELCHANAN MOSSEL: Shotgun assembly of random jigsaw puzzles. *arXiv preprint arXiv:1605.03086*, 2016. 3
- [4] MARTIN DYER, ALAN FRIEZE, AND STEPHEN SUEN: The probability of unique solutions of sequencing by hybridization. *Journal of Computational Biology*, 1(2):105–110, 1994. 2
- [5] FRANK HARARY: On the reconstruction of a graph from a collection of subgraphs. In *Theory of Graphs and its Applications (Proc. Sympos. Smolenice, 1963)*, pp. 47–52, 1964. 1
- [6] PAUL J KELLY ET AL.: A congruence theorem for trees. *Pacific J. Math*, 7(1):961–968, 1957. 1
- [7] ANDERS MARTINSSON: Shotgun edge assembly of random jigsaw puzzles. *arXiv preprint arXiv:1605.07151*, 2016. 3
- [8] ELCHANAN MOSSEL: Personal communication, 2015. 3, 6
- [9] ELCHANAN MOSSEL AND NATHAN ROSS: Shotgun assembly of labeled graphs. *arXiv preprint arXiv:1504.07682*, 2015. 1, 2
- [10] ABOLFAZL S MOTAHARI, GUY BRESLER, AND NC DAVID: Information theory of dna shotgun sequencing. *IEEE Transactions on Information Theory*, 59(10):6273–6289, 2013. 2
- [11] DANIEL SOUDRY, SURAJ KESHRI, PATRICK STINSON, MIN-HWAN OH, GARUD IYENGAR, AND LIAM PANINSKI: A shotgun sampling solution for the common input problem in neural connectivity inference. *arXiv preprint arXiv:1309.3724*, 2013. 2

- [12] STANISLAW M ULAM: *A collection of mathematical problems*. Volume 8. Interscience Publishers, 1960. 1

AUTHORS

Rajko Nenadov
research fellow
School of Mathematical Sciences, Monash University, Australia
rajko.nenadov@monash.edu
<https://www.inf.ethz.ch/personal/rnenadov/>

Pascal Pfister
Ph.D. student
Institute of Theoretical Computer Science, ETH Zürich, Switzerland
ppfister@inf.ethz.ch
<http://www.as.inf.ethz.ch/people/members/ppfister/index.html>

Angelika Steger
professor
Institute of Theoretical Computer Science, ETH Zürich, Switzerland
steger@inf.ethz.ch
<http://www.as.inf.ethz.ch/people/prof/index.html>