

Solving Jigsaw Puzzles by the Graph Connection Laplacian*

Vahan Huroyan[†], Gilad Lerman[‡], and Hau-Tieng Wu[§]

Abstract. We propose a novel mathematical framework to address the problem of automatically solving large jigsaw puzzles. This problem assumes a large image, which is cut into equal square pieces that are arbitrarily rotated and shuffled, and asks to recover the original image given the transformed pieces. The main contribution of this work is a method for recovering the rotations of the pieces when both shuffles and rotations are unknown. A major challenge of this procedure is estimating the graph connection Laplacian without the knowledge of shuffles. A careful combination of our proposed method for estimating rotations with any existing method for estimating shuffles results in a practical solution for the jigsaw puzzle problem. Our theory guarantees, in a clean setting, that our basic idea of recovering rotations is robust to some corruption of the connection graph. Numerical experiments demonstrate the competitive accuracy of this solution, its robustness to corruption, and its computational advantage for large puzzles.

Key words. jigsaw puzzles, graph connection Laplacian, vector diffusion maps, Z4 synchronization

AMS subject classifications. 90C20, 90C27, 90C35, 90C90

DOI. 10.1137/19M1290760

1. Introduction. Solving jigsaw puzzles is an entertaining task commonly explored by children and adults. It is also a challenging mathematical and engineering problem that occupies researchers in computer science, mathematics, and engineering. The solution of this problem is useful for several industrial applications. One example is reassembling archaeological artifacts [5, 20, 30, 37, 40], where one tries to recover the shape of an archaeological object from damaged pieces. Another example is recovering shredded documents or photographs [10, 19, 22, 24], where one tries to recover a document or a picture from small pieces of it. Additional applications appear in biology [23] and speech descrambling [45].

The automatic solution of puzzles, without having any information on the underlying image, is known to be NP hard [1, 11]. The first algorithm that attempted to automatically solve general puzzles was introduced by Freeman and Gardner [15] in 1964. It was designed to solve puzzles with 9 pieces by only considering the geometric shapes of the pieces.

In this paper we consider a setting of “jigsaw” puzzles which is common in the imaging sciences [2, 6, 16, 27, 29, 31, 38, 39]. In this setting, an image is cut into equal square pieces,

*Received by the editors October 2, 2019; accepted for publication (in revised form) July 7, 2020; published electronically October 8, 2020.

<https://doi.org/10.1137/19M1290760>

Funding: This research was supported by National Science Foundation awards DMS-14-18386 and DMS-18-30418.

[†]Department of Mathematics, The University of Arizona, Tucson, AZ 85721 USA (vahan-huroyan@math.arizona.edu).

[‡]School of Mathematics, University of Minnesota, Twin Cities, Minneapolis, MN 55455 USA (lerman@umn.edu).

[§]Department of Mathematics and Department of Statistical Science, Duke University, Durham, NC 27708 USA, and Mathematics Division, National Center for Theoretical Sciences, Taipei, Taiwan (hau.tieng.wu@duke.edu).

and the problem is to recover this image from the given pieces, which are possibly rotated and shifted along the puzzle grid. We refer to these puzzles as square jigsaw puzzles. Some examples are demonstrated in [Figure 1](#). Gallagher [16] categorized these puzzles into three types. In type 1 puzzles, the pieces are not rotated but shifted. In type 3 puzzles, the pieces are not shifted but rotated. In type 2 puzzles, the pieces are both shifted and rotated. This work aims to solve type 2 and type 3 puzzles.

Many proposals for solving the square jigsaw puzzles are based on greedy methods [2, 6, 16, 27, 29, 31, 38, 39]. However, greedy algorithms can easily get trapped in locally optimal solutions, which are not global. Some proposals also involve nongreedy constructive methods [7, 32, 34], which are often combined with greedy procedures. This work proposes a constructive framework for recovering orientations of puzzle pieces. The overall procedure for recovering both orientations and locations requires various heuristics. However, it avoids common greedy procedures in solving this problem. Our main purpose is to convey the effective use of the recent mathematical idea of the graph connection Laplacian [36] for recovering orientations of type 2 puzzles. Unlike previous methods, it is easy to understand its constructive mechanism and even guarantee some robustness to measurement errors in a clean setting. In practice, we demonstrate robustness to corruption and computational efficiency for large puzzles.

1.1. Previous work. Several algorithms have been recently proposed for the automatic solution of square jigsaw puzzles [2, 6, 7, 16, 27, 29, 31, 32, 33, 34, 35, 38, 39, 42, 43]. The problem becomes more challenging when the number of puzzle pieces increases and the sizes of puzzle pieces decrease. Some of these algorithms only consider type 1 puzzles (see, e.g., [2, 7, 31, 32, 45]), since recovering orientations increases the possible comparisons between two pieces by four and may also decrease the accuracy of solving the puzzle. The rest of these algorithms focus on type 2 puzzles; [16] also separately discusses type 3 puzzles. Other models of jigsaw puzzles and probabilistic results for their solutions are discussed in [4, 25, 28].

Cho, Avidan, and Freeman [7] proposed a probabilistic, graphical model approach to the square jigsaw puzzle problem and discussed different compatibility metrics between puzzle pieces. Yang, Adluru, and Latecki [42] proposed another probabilistic solution by using a particle filter and a state permutations framework. Pomeranz, Shemesh, and Ben-Shahar [31] proposed a greedy method, discussed a few compatibility metrics, and included some analysis on how to pick the correct compatibility metric for their method. Gallagher [16] proposed a tree-based reassembly algorithm which greedily merges components while respecting the geometric consistence constraints. It runs in three steps: building a constrained tree, trimming, and filling. Mondal, Wang, and Durocher [27] used the algorithm of Gallagher [16], but they replaced its proposed metric with a combination of two existing metrics. They claimed to achieve a more robust metric using this technique. Andaló, Taubin, and Goldenstein [2] proposed a quadratic assignment approach, which maximizes a constrained quadratic function via constrained gradient ascent. Jin et al. [18] proposed a scoring approach that, in addition to considering edge similarity, also takes into account content similarity between puzzle pieces. Paikin and Tal [29] proposed a greedy algorithm for handling puzzles of unknown size and with missing entries. Sholomon, David, and Netanyahu [32, 33, 34] proposed a genetic algorithm. In [35] they proposed a new deep neural network-based approach for the prediction of

the likelihood of correct matches.

Son, Hays, and Cooper [38] incorporated the “geometric structure” of the square jigsaw puzzle by searching for small loops (4-cycles) of puzzle pieces, which form consistent cycles, and then hierarchically combining these small loops with higher order loops in a bottom-up fashion. They argued that loop constraints could effectively eliminate pairwise matching outliers. Son et al. [39] proposed a growing consensus approach that assembles pieces by multiple modest bonds and uses a new objective function that maximizes consensus configurations. Yu, Russell, and Agapito [43] proposed a linear programming-based formulation, which combines global and greedy approaches. Their proposed solver simultaneously exploits all the pairwise matches and globally computes the location of each piece/component at each step of the algorithm. Chen, Cao, and Liu [6] proposed a greedy algorithm and combined several metrics to improve the performance of this algorithm.

The only previous procedure for solving type 3 puzzles is by Gallagher [16]. It uses a greedy and nonconstructive method. We are unaware of any previous constructive method for finding the orientations of type 2 puzzle pieces.

1.2. Our contribution. In this paper we propose a novel approach to solving type 2 and type 3 jigsaw puzzles. For type 3 puzzles, we suggest a fast, robust, constructive, and straightforward solution that uses the graph connection Laplacian (GCL) [36] (discussed in section 3.1). For type 2 puzzles we propose a novel iterative algorithm which solves the following two subproblems: The Rotation Problem (RotP) and the Location Problem (LocP):

RotP: Finding the orientations of all puzzle pieces.

LocP: Finding the locations of all puzzle pieces.

These two steps are iteratively repeated until the desired result is achieved. We solve RotP by using the GCL, where the main challenge is to construct the GCL despite the unknown locations. We solve LocP by applying any state-of-the-art solution of type 1 puzzles to the puzzle obtained from the solution of RotP. Some information inferred from the solution of LocP is further used to improve the solution of RotP.

All previous algorithms for solving type 2 puzzles simultaneously address RotP and LocP. On the other hand, this work separately solves the two subproblems, with a constructive and better understood solution of RotP. Moreover, we aim to present a principled approach and thus avoid greedy steps that are common in previous algorithms and help improve the accuracy. Empirically, the proposed method is faster than other methods for large puzzles, e.g., with thousands of patches. We can also specify more easily the overall computational complexity of our proposed components (excluding the borrowed type 1 puzzle solver). This complexity is comparable to that of the most common component of any puzzle solver. Our numerical results also demonstrate that our algorithm is more robust to corruption of the sides of puzzle patches than other algorithms.

In theory, we verify robustness under a certain mathematical setting and contribute with perturbation-type results to the general mathematical area of group synchronization. However, we would like to emphasize that the paper addresses a real applied problem, where it is unclear how to make sufficiently good measurements, which are assumed by the theory. We thus have two different components that may seem in tension. One is the clean theory that justifies recovery of orientations given a special type of measurements. The other one is

a set of methods, supported by numerical experimentation, with specific choices of measurements of initial relative orientations and graph affinities between puzzle patches. The methods include various heuristics, for example, for improving the latter measurements given better information of patch locations. This is a first attempt to incorporate a successful puzzle solver within a rigorous mathematical setting. We hope that with time the proposed mathematical foundations and algorithms can be further improved and simplified.

1.3. Structure of this paper. This paper is organized as follows: section 2 mathematically formulates the square jigsaw puzzle problem; section 3 presents a solution for RotP, given some initial measurements of relative orientations and graph affinities between puzzle patches. It also theoretically guarantees the robustness of the solution to errors in the initial measurements; section 4 explains how to measure in practice the initial relative orientations and graph affinities for type 2 and type 3 puzzles. Note that the combination of this construction with the method of section 3 provides the desired solution to RotP. While section 3 has a clean formulation with a theoretical guarantee, section 4 relies on various heuristics, which we try to motivate; section 5 describes additional heuristics for improving the initial measurements and consequently the solution of RotP given the solution of LocP. It also summarizes our full algorithm for solving square jigsaw puzzles; section 6 presents numerical experiments that test the accuracy, efficiency, and robustness to corruption of the proposed algorithm using digital images; Finally, section 7 concludes with a short discussion that includes possible extensions of this work.

2. Mathematical formulation and notation. We mathematically formulate the square jigsaw puzzle problem and introduce relevant notation in section 2.1. In section 2.2 we emphasize the main challenge in solving this problem. We remark that a more general formulation appears in the accompanying supplemental material (M129076_01.pdf [[local](#)/[web](#) 965KB]).

2.1. Setting and notation. The setting of the square jigsaw puzzle problem assumes a rectangle $M = [a_1, b_1] \times [a_2, b_2]$ in \mathbb{R}^2 and open squares $\{P_i\}_{i=1}^n$ that tile M . We refer to $\{P_i\}_{i=1}^n$ as patches and to the four nearest neighbors of a given patch (left, right, top, and bottom) as neighboring patches. The setting further assumes a function $f \in L^2(M, \mathbb{R}^k)$ and $k \geq 1$, so that the image to be recovered is the graph $\{(x, f(x)) : x \in M \subset \mathbb{R}^2\}$. One may use $k = 1$ for gray-scale images, $k = 3$ for color images, and higher k for multispectral and hyperspectral images. In this paper we use $k = 3$. Since a main challenge of the practical problem is dealing with discrete images, we further assume that f is piecewise constant with discrete values in the following way. Each patch is divided by a uniformly spaced grid into $s \times s$ subsquares, and the vector-valued f is constant on each subsquare, where each coordinate of the constant vector is discrete; for example, it lies in $0, \dots, 255$.

The setting also assumes arbitrary orientations and shuffles of puzzle pieces. It is sufficient to represent all possible orientations with rotations by 0° , 90° , 180° , and 270° , which are elements of the cyclic group \mathbb{Z}_4 , and translations of patches $\{P_i\}_{i=1}^n$ by $\{P_{\sigma(i)}\}_{i=1}^n$, where σ is a permutation of degree n . Therefore, we can write the set of image patches as $\mathcal{Q} = \{(R_{\sigma(i)}(P_{\sigma(i)}), R_{\sigma(i)} \circ f|_{P_{\sigma(i)}})\}_{i=1}^n$, where $R_{\sigma(i)}$ is an element of the cyclic group \mathbb{Z}_4 and the action \circ is defined by $R_j \circ f|_{P_j} := f(R_j^{-1})|_{P_j}$. While we use a very formal notation, an image patch is composed of a shifted and rotated patch together with the appropriately aligned value

of f . With this notation the square jigsaw puzzle problem can be expressed as follows: Given M and the above set of image patches \mathcal{Q} , recover the function $f|_M$ or equivalently the image $(M, f|_M)$.

Figure 1 demonstrates the particular instance of the square jigsaw puzzle problem we discuss in this paper. We remark that the last column of this figure illustrates the image patches $\mathcal{Q} = \{(R_{\sigma(i)}(P_{\sigma(i)}), R_{\sigma(i)} \circ f|_{P_{\sigma(i)}})\}_{i=1}^n$ discussed above. We assumed above that f is a piecewise constant function. In this figure, f has constant values on squares corresponding to image pixels. Since the resolution is relatively high, one cannot notice that f is piecewise constant. However, this is noticeable in the low-resolution demonstration of patches of another puzzle at the top right image of Figure 2.



Figure 1. Examples of puzzles with 12 patches. Left column: The original image. Central column: Division of the image into 12 square patches of the same size. Right column: The 12 patches are randomly reordered and rotated.

2.2. A challenge of square jigsaw puzzles. We recall that the formulation of the square jigsaw puzzle problem requires finding a permutation σ and rotations $\{R_i\}_{i=1}^n \subset \mathbb{Z}_4$. Equivalently, one may solve for locations $\{\mathbf{x}_i\}_{i=1}^n$ on a uniform grid, representing the centers of the patches, and rotations $\{R_i\}_{i=1}^n$. In order to estimate these from the set of image patches \mathcal{Q} with a function f , one needs to rely on the similar function values on the sides of neighboring patches. However, in our setting of digital images, f is often discontinuous in the direction from one side of a patch to a side of a neighboring patch. The top right image of Figure 2 demonstrates this phenomenon for two patches selected from the puzzles shown in top left image with lower resolution. Such discontinuity can result in loss of information for determining neighbors and may lead to ill-posed problems.

There are also special images for which the puzzle problem is ill-posed. For example, the bottom left image of Figure 2 demonstrates a case where several patches look very similar to each other and it is impossible to determine the right permutation. Nevertheless, the output of common algorithms given this particular puzzle is often visually acceptable. On the other

hand, the bottom right image of [Figure 2](#) demonstrates a case where the image consists of two parts that are disconnected by a uniform background. The background is the white sky, one part is the main scene of the image, and the other part includes two short branches of another tree at the top left corner of the image. In this case, it would be impossible to figure out the exact position of the latter part of the image.

The following definition quantifies an ideal type of metric between sides of image patches that, if it exists (i.e., if the problem is well-posed), can be used to solve the square jigsaw puzzle problem.

Definition 2.1. *Fix an image I and a set of image patches $\mathcal{Q} := \{P_i, f|_{P_i}\}_{i=1}^n$. A metric defined on all sides of image patches in \mathcal{Q} is called perfect if there exists $c > 0$ so that any two matching sides (of neighboring patches) have a distance less than c and any two nonmatching sides have a distance greater than c .*

The main challenge of solving reasonable instances of the square jigsaw puzzle problem is to find a nearly perfect metric. Empirically, we have found that the Mahalanobis Gradient Compatibility (MGC) metric, defined in [\[16\]](#) and described in section [4.1](#), is often near perfect in well-posed cases.

3. A framework for recovering rotations of puzzle pieces. This section applies the framework of [\[12, 36\]](#) for recovering the global orientations of puzzle patches. This framework requires the construction of a graph whose vertices correspond to the puzzle patches and whose edges connect neighboring patches. The rest of the section is organized as follows: section [3.1](#) forms the connection graph and its graph connection Laplacian (GCL) and explains how to estimate the rotations of puzzle patches using this Laplacian; and section [3.2](#) theoretically justifies the method described in section [3.1](#).

3.1. Estimation of orientations using the connection graph. The general connection graph [\[36\]](#) $G = (V, E, W, R)$ consists of four components: vertices V , edges E , the *affinity function* (or weight function) $W : E \rightarrow [0, 1]$, and the *connection function* $R : E \rightarrow \mathbb{G}$, where \mathbb{G} is a given group. The first three components are determined by the weighted graph, and the fourth depends on the application in which the graph is used. This formulation is most natural for the problem of group synchronization, which is carefully reviewed in the first two sections of [\[21\]](#). In this problem, one is given a graph $G = (V, E)$ with affinity function W and needs to estimate for any vertex $i \in V$ a group element $g_i \in \mathbb{G}$ from corrupted (or noisy) measurements of group ratios $g_i g_j^{-1} \in \mathbb{G}$ among all $\{i, j\} \in E$. It is natural to form the connection function by the given measurements.

In the particular case of the square jigsaw puzzle, $\mathbb{G} = \mathbb{Z}_4$ and the connection function for any edge $\{i, j\}$ assigns the rotation represented by the block $\mathbf{R}[i, j]$, which was defined in section [3.1](#). For convenience, we represent the affinity and connection functions by their corresponding matrices $\mathbf{W} \in \mathbb{R}^{n \times n}$ and $\mathbf{R} \in \mathbb{R}^{2n \times 2n}$, and thus write the connection graph as $G = (V, E, \mathbf{W}, \mathbf{R})$. In this case, the group synchronization problem is referred to as \mathbb{Z}_4 synchronization. This is the underlying mathematical problem for recovering the orientations of the patches. However, there are several practical considerations that make the problem of orientation recovery of puzzle patches more complicated than the synthetic \mathbb{Z}_4 synchronization problem. A primary issue is that one needs to find a way to measure the group ratios; that

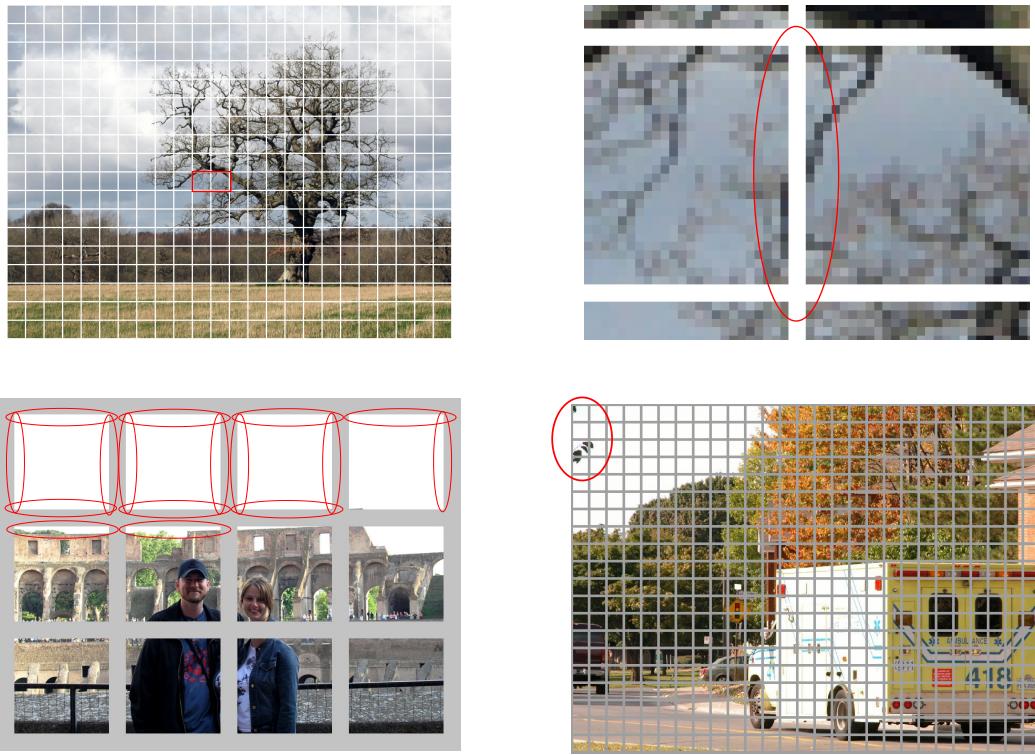


Figure 2. Examples of square jigsaw puzzles, where the comparison of two neighboring patches is challenging or impossible. The top left image shows a puzzle with 432 pieces, each of size 28×28 . The top right image demonstrates an example of 2 neighboring patches in the latter puzzle that have different pixel values around the boundaries due to the discrete nature of a digital image. These patches are circled with red in the original puzzle (top left image), and their nearby sides are circled with red in the top right image. The bottom two images demonstrate examples of puzzles that have patches with uniformly white sides (circled with red in the bottom left image) and also have some uniformly white patches. Natural solutions of the bottom left puzzle seem to yield visually correct images that may not coincide with the original assignment. However, there are natural solutions of the bottom right puzzle that result in different images than the original one. Indeed, the small component of the image circled with red can be placed in different areas within the skies.

is, a connection function needs to be estimated from the given puzzle patches. Another issue is that erroneous location assignments of patches may result in a poorly estimated connection function.

If one has a perfect metric (recall Definition 2.1) for the square jigsaw puzzle, the ideal connection graph is formed as follows. The vertices represent patches in \mathcal{Q} , the edges connect neighboring patches, and the weights are 1 for all edges and 0 otherwise. The underlying group \mathbb{Z}_4 can be represented either by the four complex numbers $\{1, i, -1, -i\}$ with complex multiplication or by the following four 2×2 matrices:

$$(1) \quad \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}$$

with matrix multiplication. Note that if $i < j$ and $\{i, j\}$ is an edge connecting patches P_i and P_j , then $R[i, j]$ is a rotation in $\mathbb{G} = \mathbb{Z}_4$ whose application to P_j , together with an appropriate plane translation, results in P_i . Throughout the paper we use the representation of \mathbb{Z}_4 in (1).

For possibly imperfect scenarios of the square jigsaw puzzles, the vertices are formed as above, but one needs to construct meaningful edges, affinity function, and connection function (with $\mathbb{G} = \mathbb{Z}_4$). A heuristic construction of these is suggested for type 2 and type 3 puzzles in section 4.3 and section 4.2, respectively. Here we propose a general heuristic that uses a given connection graph of square jigsaw puzzles to estimate the unknown orientations of the patches. This heuristic is later justified in section 3.2 under special assumptions. The main idea of this heuristic is to use the GCL for inferring global information (in the form of a certain eigendecomposition) from local information (needed to form the GCL).

Next, we review several matrices associated with a general connection graph. Recall that the functions W and R are defined on the set $\{1, \dots, n\} \times \{1, \dots, n\}$, where n is the number of puzzle pieces. Thus, from now on, we denote these functions by their corresponding matrices $\mathbf{W} \in \mathbb{R}^{n \times n}$ and $\mathbf{R} \in \mathbb{R}^{2n \times 2n}$, respectively. Note that \mathbf{R} is a block matrix whose 2×2 blocks represent two-dimensional rotations. For $1 \leq i, j \leq n$, we denote by $\mathbf{R}[i, j]$ the $[i, j]$ th 2×2 block of \mathbf{R} . We index blocks by $[i, j]$ and matrix elements by (i, j) . The connection graph is thus $G = (V, E, \mathbf{W}, \mathbf{R})$. The *connection adjacency matrix* is an $n \times n$ block matrix \mathbf{S} with 2×2 submatrices, where for $1 \leq i, j \leq n$ the (i, j) th submatrix is

$$(2) \quad \mathbf{S}[i, j] = \begin{cases} \mathbf{W}(i, j)\mathbf{R}[i, j] & \text{if } \{i, j\} \in E, \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

The *degree matrix* is an $n \times n$ block diagonal matrix \mathbf{D} , where for $1 \leq i \leq n$, its i th diagonal submatrix is

$$(3) \quad \mathbf{D}[i, i] = d(i)\mathbf{I}_2, \quad \text{where } d(i) = \sum_{j \neq i} \mathbf{W}(i, j),$$

where \mathbf{I}_2 is the 2×2 identity matrix. We define $\mathbf{C} := \mathbf{D}^{-1}\mathbf{S}$ as the graph connection weight (GCW) matrix. We define the (normalized) GCL matrix as $\mathbf{I} - \mathbf{C}$. In practice, we directly form the GCW matrix and use its eigendecomposition. Clearly, this is equivalent to using the eigendecomposition of the GCL matrix, and we thus refer to our method with the term GCL.

The GCW matrix is associated with a random walk whose transition probability matrices are $\mathbf{W}(i, j)$, $1 \leq i, j \leq n$. This can be seen by its action on a block vector $\mathbf{v} \in \mathbb{R}^{2n \times 2}$, whose n th 2×2 submatrices are

$$\mathbf{v}[j] = \begin{bmatrix} \mathbf{v}_{2j-1,1} & \mathbf{v}_{2j-1,2} \\ \mathbf{v}_{2j,1} & \mathbf{v}_{2j,2} \end{bmatrix} \in \mathbb{R}^{2 \times 2}, \quad 1 \leq j \leq n,$$

in the following way:

$$(\mathbf{C}\mathbf{v})[i] = \sum_{j: (i, j) \in E} \left[\frac{\mathbf{W}(i, j)}{\sum_{k: (i, k) \in E} \mathbf{W}(i, k)} \right] \mathbf{R}[i, j] \mathbf{v}[j].$$

That is, a block vector $\mathbf{v}[j]$ is rotated by $\mathbf{R}[i, j]$ and is assigned to the i th patch with probability $\mathbf{W}(i, j)/\sum_{k: (i, k) \in E} \mathbf{W}(i, k)$.

To recover the global orientations of the puzzle patches, we follow the procedures of [3, 36]. First, we form the block vector $\mathbf{U} \in \mathbb{R}^{2n \times 2}$ whose columns are the top 2 eigenvectors of \mathbf{C} . Then, we project each of the 2×2 blocks of \mathbf{U} onto \mathbb{Z}_4 (that is, we replace each block with its closest element, with respect to the Frobenius norm, in (1)) and use the resulting blocks as the global orientations. [Algorithm 1](#) summarizes the above straightforward procedure of recovering the unknown orientations of the image patches for a given square jigsaw puzzle.

Algorithm 1 The GCL algorithm.

Input: Connection graph: $G = (V, E, \mathbf{W}, \mathbf{R})$

- Construct the Connection Adjacency Matrix \mathbf{S} by (2)
- Construct the degree matrix \mathbf{D} by (3)
- Let $\mathbf{C} = \mathbf{D}^{-1}\mathbf{S}$
- Form $\mathbf{U} \in \mathbb{R}^{2n \times 2}$ whose columns are the 2 top eigenvectors of \mathbf{C}
- For $1 \leq i \leq n$, let $\mathbf{R}_i \in \mathbb{Z}_4$ be the projection of the i th block of \mathbf{U} onto \mathbb{Z}_4

Return: Global rotation matrices $\mathbf{R}_1, \dots, \mathbf{R}_n$

We emphasize that the GCL algorithm for recovering the orientations of patches is non-greedy. Indeed, it directly constructs the orientation of patches using the information in the connection graph via diffusion. On the other hand, other methods, such as [16, 31, 38, 39], try to greedily match pieces based on their relative orientations. We also mention that the GCL algorithm does not use any knowledge of the size of the puzzle image or, equivalently, of the number of puzzle pieces per length or width of the image.

3.2. Theoretical justification of the GCL algorithm. We show that the proposed estimation of orientations for type 2 puzzles is robust to incorrect measurements, where incorrect measurements are mistakes in estimating the connection graph. The three puzzles in [Figure 2](#) exemplify cases where incorrect measurements are expected due to indistinguishability or low resolution of patches. Incorrect measurements can also arise due to mistakes in estimating patches' locations. Indeed, such mistakes result in incorrect estimation of the connection graph.

We distinguish between the ground truth solution (or “true” solution) and the estimated one. We denote by E_{true} the set of “true edges,” that is, edges connecting neighboring patches of the true solution. We find it most natural to define (V, E_{true}) as the underlying uniform grid for the patches. Nevertheless, there is some freedom in defining (V, E_{true}) . For example, if one wants to also emphasize diagonal edges, then (V, E_{true}) can be formed by adding these edges to the uniform grid. These two different choices of (V, E_{true}) can make a slight difference in the estimates of our proposed theory described below. Using the prefixed indexing of patches, we define the true affinity function

$$(4) \quad \mathbf{W}_{\text{true}}(i, j) = \begin{cases} 1 & \text{if } \{i, j\} \in E_{\text{true}}, \\ 0 & \text{otherwise,} \end{cases}$$

and the true connection function

$$(5) \quad \mathbf{R}_{\text{true}}[i, j] = \begin{cases} \mathbf{R}_i \mathbf{R}_j^T & \text{if } \{i, j\} \in E_{\text{true}}, \\ \mathbf{0} & \text{otherwise,} \end{cases}$$

where $\mathbf{R}_1, \dots, \mathbf{R}_n$ are the rotation matrices of the rotations R_1, \dots, R_n defined in section 2.1. Note that unlike E_{true} , \mathbf{W}_{true} and \mathbf{R}_{true} are unknown to the user. Let $G_{\text{est}} = (V, E_{\text{est}}, \mathbf{W}_{\text{est}}, \mathbf{R}_{\text{est}})$ denote the estimated connection graph. We remark that the graph (V, E_{est}) may be rather different than (V, E_{true}) as the user can have some uncertainties about connecting patches. Finally, denote by \mathbf{C}_{est} the GCW matrix corresponding to G_{est} .

The following perturbation theorem states that if the estimated connection graph is a good approximation of the true connection graph and certain conditions hold, then the estimated rotations are close in some sense to the underlying rotations. To be more specific, we denote by $\mathbf{V}_{\text{est}} \in \mathbb{R}^{2n \times 2}$ the matrix whose 2 columns are the top eigenvectors of \mathbf{C}_{est} . Recall that the estimated rotations are obtained by projecting the blocks of this matrix onto \mathbb{Z}_4 . We denote the set of underlying rotations by $\mathbf{R}_1, \dots, \mathbf{R}_n \in \mathbb{Z}_4$. We further denote by $\mathbf{R}_{\text{true}}^{\text{tot}}$ the block matrix in $\mathbb{R}^{2n \times 2}$ whose n blocks are these underlying rotations. The theorem claims that under some conditions, the principal angles between the column spaces of \mathbf{V}_{est} and $\mathbf{R}_{\text{true}}^{\text{tot}}$ are sufficiently close. Note that there are only two such angles, θ_1 and θ_2 , and recall that they can be computed as follows: $\sin(\theta_1) = \cos^{-1}(\sigma_1)$ and $\sin(\theta_2) = \cos^{-1}(\sigma_2)$, where σ_1 and σ_2 are the singular values of $\mathbf{V}_{\text{est}}^T \mathbf{R}_{\text{true}}^{\text{tot}}$.

We use the following notation: $\sin(\Theta(\mathbf{R}_{\text{true}}^{\text{tot}}, \mathbf{V}_{\text{est}}))$ denotes the 2×2 diagonal matrix with diagonal elements $\sin(\theta_1)$ and $\sin(\theta_2)$ (specified above); $\|\cdot\|_F$ denotes the Frobenius norm; d_{max} denotes the maximal degree of the estimated graph (using the ∞ matrix norm, we can express it as $d_{\text{max}} = \|\mathbf{W}_{\text{est}}\|_{\infty}$); for a set $E' \subseteq E$, $\phi_{E'}$ denotes the second smallest eigenvalue of the normalized graph Laplacian of the graph $G' = (V, E')$ (a precise formula for $\phi_{E'}$ is given at the end of the proof of the theorem). Using this notation we express below the closeness of \mathbf{V}_{est} and $\mathbf{R}_{\text{true}}^{\text{tot}}$, which we further discuss and interpret after proving the theorem.

Theorem 3.1. *Let $G_{\text{true}} = (V, E_{\text{true}}, \mathbf{W}_{\text{true}}, \mathbf{R}_{\text{true}})$ be the ground-truth connection graph, where \mathbf{W}_{true} and \mathbf{R}_{true} are as defined in (4) and (5). Assume that a user estimates this connection graph from possibly corrupted data with the connection graph $G_{\text{est}} = (V, E_{\text{est}}, \mathbf{W}_{\text{est}}, \mathbf{R}_{\text{est}})$, which has maximal degree d_{max} , and that there exists a set $E' \subset E_{\text{true}} \cap E_{\text{est}}$ satisfying the following properties: (V, E') is a connected graph;*

$$(6) \quad \mathbf{R}_{\text{est}}[i, j] = \mathbf{R}_{\text{true}}[i, j] \text{ for } \{i, j\} \in E';$$

there exist $\epsilon > 0$ and $c_1, \dots, c_n > 0$ such that

$$(7) \quad \max_{1 \leq i \leq n} \frac{1}{\sqrt{d_{\text{max}}}} \sqrt{\sum_{j: \{i, j\} \in E'} \left(\mathbf{W}_{\text{true}}(i, j) - \frac{\mathbf{W}_{\text{est}}(i, j)}{c_i} \right)^2 + \sum_{j: \{i, j\} \notin E'} \left(\frac{\mathbf{W}_{\text{est}}(i, j)}{c_i} \right)^2} < \epsilon;$$

and there exists $\gamma > \epsilon$ such that

$$(8) \quad \min_{1 \leq i \leq n} \sum_{j: \{i, j\} \in E'} \mathbf{W}_{\text{true}}(i, j) / d_{\text{max}} > \gamma.$$

Then,

$$(9) \quad \|\sin \Theta(\mathbf{R}_{\text{true}}^{\text{tot}}, \mathbf{V}_{\text{est}})\|_F \leq \frac{2}{\phi_{E'}} \sqrt{\frac{2n}{d_{\text{max}}}} \frac{\epsilon}{\gamma} \left(1 + \frac{1}{\gamma - \epsilon} \right).$$

Proof. We denote the restriction of the affinity and connection functions, \mathbf{W}_{true} and \mathbf{R}_{true} , onto the edges in E' by $\mathbf{W}'_{\text{true}}$ and $\mathbf{R}'_{\text{true}}$, respectively. The corresponding connection graph, connection adjacency matrix, diagonal matrix, and GCW matrix are denoted by $G'_{\text{true}} = (V, E', \mathbf{W}'_{\text{true}}, \mathbf{R}'_{\text{true}})$, $\mathbf{S}'_{\text{true}}$, $\mathbf{D}'_{\text{true}}$, and $\mathbf{C}'_{\text{true}} \in \mathbb{R}^{2n \times 2n}$, respectively. In addition to the above notation for the estimated connection graph, we denote its connection adjacency matrix and diagonal matrix by \mathbf{S}'_{est} and \mathbf{D}'_{est} , respectively. Following (3), we denote the diagonal elements of $\mathbf{D}'_{\text{true}}$ and \mathbf{D}'_{est} by $d_{\text{true}}(i)$ and $d_{\text{est}}(i)$, $i = 1, \dots, n$. Note that the maximal degree of the estimated graph can be expressed as follows: $d_{\max} = \max_{1 \leq i \leq n} d_{\text{est}}(i)$. Let $\mathbf{V}'_{\text{true}} \in \mathbb{R}^{2n \times 2}$ denote the matrix whose 2 columns are the top eigenvectors of $\mathbf{C}'_{\text{true}}$.

We note that division of the i th row of the affinity matrix \mathbf{W}_{est} by the constant c_i , where $1 \leq i \leq n$, does not change the GCW matrix of $(V, E, \mathbf{W}_{\text{est}}, \mathbf{R}_{\text{est}})$. Thus, we can assume, without loss of generality, that $c_1 = \dots = c_n = 1$. Using this assumption and the definition of $\mathbf{W}'_{\text{true}}$, we rewrite (7) as

$$(10) \quad \max_{1 \leq i \leq n} \sqrt{\frac{1}{d_{\max}} \sum_{j=1}^n (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j))^2} < \epsilon.$$

We further note that the assumption in (8) can be rewritten as

$$(11) \quad d'_{\text{true}}(i)/d_{\max} \geq \gamma \quad \text{for all } 1 \leq i \leq n.$$

The proof of Theorem 3.1 consists of three steps.

Step I. We prove that

$$(12) \quad \|\mathbf{C}'_{\text{true}} - \mathbf{C}_{\text{est}}\|_F \leq \sqrt{\frac{2n}{d_{\max}}} \frac{\epsilon}{\gamma} \left(1 + \frac{1}{\gamma - \epsilon}\right).$$

This result is analogous to Lemmas 3.1 and 3.2 of El Karoui and Wu [12] but uses the Frobenius norm instead of the spectral norm and has weaker conditions in (7) and (8). El Karoui and Wu [12] also allow a very small perturbation of the connection function restricted to E' , but in the special case of \mathbb{Z}_4 , this assumption is equivalent to (6).

Using the definitions of $\mathbf{C}'_{\text{true}}$ and \mathbf{C}_{est} , we express and then bound the left-hand side (LHS) of (12) as follows:

$$(13) \quad \|\mathbf{D}'_{\text{true}}^{-1} \mathbf{S}'_{\text{true}} - \mathbf{D}_{\text{est}}^{-1} \mathbf{S}_{\text{est}}\|_F = \|\mathbf{D}'_{\text{true}}^{-1} (\mathbf{S}'_{\text{true}} - \mathbf{S}_{\text{est}}) + (\mathbf{D}'_{\text{true}}^{-1} - \mathbf{D}_{\text{est}}^{-1}) \mathbf{S}_{\text{est}}\|_F \\ \leq \|\mathbf{D}'_{\text{true}}^{-1} (\mathbf{S}'_{\text{true}} - \mathbf{S}_{\text{est}})\|_F + \left\| (\mathbf{D}'_{\text{true}}^{-1} - \mathbf{D}_{\text{est}}^{-1}) \mathbf{S}_{\text{est}} \right\|_F.$$

We follow with bounding the first term of the right-hand side (RHS) of (13):

$$(14) \quad \left\| \mathbf{D}'_{\text{true}}^{-1} (\mathbf{S}'_{\text{true}} - \mathbf{S}_{\text{est}}) \right\|_F = \left\| (\mathbf{D}'_{\text{true}}/d_{\max})^{-1} (\mathbf{S}'_{\text{true}}/d_{\max} - \mathbf{S}_{\text{est}}/d_{\max}) \right\|_F \\ \leq \left\| (\mathbf{D}'_{\text{true}}/d_{\max})^{-1} \right\|_2 \left\| (\mathbf{S}'_{\text{true}}/d_{\max} - \mathbf{S}_{\text{est}}/d_{\max}) \right\|_F.$$

We then control the first multiplicative term in the RHS of (14) using (11):

$$(15) \quad \left\| \left(\mathbf{D}'_{\text{true}} / d_{\max} \right)^{-1} \right\|_2 \leq \max_{1 \leq i \leq n} \frac{d_{\max}}{d'_{\text{true}}(i)} \leq \frac{1}{\gamma}.$$

We next control the second multiplicative term in the RHS of (14), where we follow with justification

$$(16) \quad \begin{aligned} \left\| \mathbf{S}'_{\text{true}} / d_{\max} - \mathbf{S}_{\text{est}} / d_{\max} \right\|_F^2 &= \frac{1}{d_{\max}^2} \sum_{i=1}^n \sum_{j=1}^n \left\| \mathbf{S}'_{\text{true}}[i, j] - \mathbf{S}_{\text{est}}[i, j] \right\|_F^2 \\ &= \frac{1}{d_{\max}^2} \sum_{i=1}^n \sum_{j=1}^n \left\| \mathbf{W}'_{\text{true}}(i, j) (\mathbf{R}'_{\text{true}}[i, j] - \mathbf{R}_{\text{est}}[i, j]) + (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j)) \mathbf{R}_{\text{est}}[i, j] \right\|_F^2 \\ &= \frac{1}{d_{\max}^2} \sum_{i=1}^n \sum_{j=1}^n (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j))^2 \left\| \mathbf{R}_{\text{est}}[i, j] \right\|_F^2 = \\ &\frac{2}{d_{\max}^2} \sum_{i=1}^n \sum_{j=1}^n (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j))^2 \leq \frac{2n}{d_{\max}^2} \max_{1 \leq i \leq n} \sum_{j=1}^n (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j))^2 \leq \frac{2n\epsilon^2}{d_{\max}}. \end{aligned}$$

The third equality uses (6) and the fact that $\mathbf{W}'_{\text{true}} = 0$ for $\{i, j\} \notin E'$, the fourth equality uses the fact that $\left\| \mathbf{R}_{\text{est}}[i, j] \right\|_F^2 = 2$, and the last inequality uses (10). Combining (14), (15), and (16), we bound the first term of the RHS of (13):

$$(17) \quad \left\| \mathbf{D}'_{\text{true}}^{-1} (\mathbf{S}'_{\text{true}} - \mathbf{S}_{\text{est}}) \right\|_F \leq \sqrt{\frac{2n}{d_{\max}}} \frac{\epsilon}{\gamma}.$$

We control the second term in the RHS of (13) below in (21). In order to pursue this, we need to bound several terms. We start with the following bound and then justify it:

$$(18) \quad \begin{aligned} \left\| \mathbf{D}'_{\text{true}} / d_{\max} - \mathbf{D}_{\text{est}} / d_{\max} \right\|_2 &= \max_{1 \leq i \leq n} \left| d'_{\text{true}}(i) / d_{\max} - d_{\text{est}}(i) / d_{\max} \right| \\ &= \max_{1 \leq i \leq n} \frac{1}{d_{\max}} \left| \sum_{j=1}^n (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j)) \right| \\ &\leq \max_{1 \leq i \leq n} \sqrt{\frac{1}{d_{\max}} \sum_{j=1}^n (\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j))^2} \leq \epsilon. \end{aligned}$$

The first inequality is a direct application of Cauchy–Schwarz and the fact that for each $1 \leq i \leq n$ at most d_{\max} of $\mathbf{W}'_{\text{true}}(i, j) - \mathbf{W}_{\text{est}}(i, j)$ are nonzero. In order to realize the last fact, we recall that $E' \subseteq E_{\text{est}}$ so we only need to check this fact when $\{i, j\} \in E_{\text{est}} \setminus E'$. In this case, $\mathbf{W}'_{\text{true}}(i, j) = 0$ and $\mathbf{W}_{\text{est}}(i, j)$ satisfies the required property due to (11). The last inequality follows from (10). Note that (18) implies that $d_{\text{est}}(i)/n \geq d'_{\text{true}}(i)/n - \epsilon$ for

$1 \leq i \leq n$. Applying this inequality and then (11), we obtain that

$$(19) \quad \left\| (\mathbf{D}_{\text{est}}/d_{\max})^{-1} \right\|_2 = \max_{1 \leq i \leq n} (d_{\text{est}}(i)/d_{\max})^{-1} = \frac{1}{\min_{1 \leq i \leq n} d_{\text{est}}(i)/d_{\max}} \leq \frac{1}{\min_{1 \leq i \leq n} (d'_{\text{true}}(i)/d_{\max} - \epsilon)} \leq \frac{1}{\gamma - \epsilon}.$$

Using the facts that for all $1 \leq i, j \leq n$, $0 \leq \mathbf{W}_{\text{est}}(i, j) \leq 1$, and $\|\mathbf{R}_{\text{est}}[i, j]\|_F^2 = 2$, we conclude that

$$(20) \quad \|\mathbf{S}_{\text{est}}/d_{\max}\|_F^2 \leq \max_{1 \leq i, j \leq n} \frac{n}{d_{\max}} \|\mathbf{S}_{\text{est}}[i, j]\|_F^2 = \frac{n}{d_{\max}} \max_{1 \leq i, j \leq n} |\mathbf{W}_{\text{est}}(i, j)|^2 \|\mathbf{R}_{\text{est}}[i, j]\|_F^2 \leq \frac{2n}{d_{\max}}.$$

The above first inequality is due to the fact that each column and row of \mathbf{S}_{est} contains at most d_{\max} nonzero elements, so the LHS is bounded by nd_{\max} times the maximal squared Frobenius norm of a block. The latter fact follows from the argument described below (18). Combining (15), (18), (19), and (20), we bound the second term in the RHS of (13) as follows:

$$(21) \quad \begin{aligned} \left\| \left(\mathbf{D}'_{\text{true}}^{-1} - \mathbf{D}_{\text{est}}^{-1} \right) \mathbf{S}_{\text{est}} \right\|_F &\leq \|\mathbf{S}_{\text{est}}/d_{\max}\|_F \left\| \left(\left(\mathbf{D}'_{\text{true}}/d_{\max} \right)^{-1} - \left(\mathbf{D}_{\text{est}}/d_{\max} \right)^{-1} \right) \right\|_2 \\ &\leq \|\mathbf{S}_{\text{est}}/d_{\max}\|_F \left\| \left(\mathbf{D}'_{\text{true}}/d_{\max} \right)^{-1} \left(\mathbf{D}'_{\text{true}}/d_{\max} - \mathbf{D}_{\text{est}}/d_{\max} \right) \left(\mathbf{D}_{\text{est}}/d_{\max} \right)^{-1} \right\|_2 \\ &\leq \|\mathbf{S}_{\text{est}}/d_{\max}\|_F \left\| \left(\mathbf{D}'_{\text{true}}/d_{\max} \right)^{-1} \right\|_2 \left\| \left(\mathbf{D}'_{\text{true}}/d_{\max} - \mathbf{D}_{\text{est}}/d_{\max} \right) \right\|_2 \left\| \left(\mathbf{D}_{\text{est}}/d_{\max} \right)^{-1} \right\|_2 \\ &\leq \sqrt{\frac{2n}{d_{\max}}} \frac{\epsilon}{\gamma(\gamma - \epsilon)}. \end{aligned}$$

Clearly, (13), (17), and (21) imply (12).

Step II. This step uses the Davis–Kahan sin Θ theorem [9] to prove the following inequality, where $\lambda_2(\mathbf{C}'_{\text{true}})$ and $\lambda_3(\mathbf{C}'_{\text{true}})$ denote the second and third largest eigenvalues of $\mathbf{C}'_{\text{true}}$, respectively:

$$(22) \quad \|\sin \Theta(\mathbf{V}'_{\text{true}}, \mathbf{V}_{\text{est}})\|_F \leq \frac{2\sqrt{\frac{2n}{d_{\max}}} \frac{\epsilon}{\gamma} \left(1 + \frac{1}{\gamma - \epsilon} \right)}{\lambda_2(\mathbf{C}'_{\text{true}}) - \lambda_3(\mathbf{C}'_{\text{true}})}.$$

We use a specific variant of Davis–Kahan according to [44]. This variant implies that

$$(23) \quad \begin{aligned} \|\sin \Theta(\mathbf{V}'_{\text{true}}, \mathbf{V}_{\text{est}})\|_F &\leq \frac{2 \min(\sqrt{2} \|\mathbf{C}'_{\text{true}} - \mathbf{C}_{\text{est}}\|_2, \|\mathbf{C}'_{\text{true}} - \mathbf{C}_{\text{est}}\|_F)}{\lambda_2(\mathbf{C}'_{\text{true}}) - \lambda_3(\mathbf{C}'_{\text{true}})} \leq \frac{2 \|\mathbf{C}'_{\text{true}} - \mathbf{C}_{\text{est}}\|_F}{\lambda_2(\mathbf{C}'_{\text{true}}) - \lambda_3(\mathbf{C}'_{\text{true}})}. \end{aligned}$$

The combination of (12) and (23) implies (22).

Step III. This step shows that $\mathbf{V}'_{\text{true}} = \mathbf{R}_{\text{true}}^{\text{tot}}$ and $\lambda_2(\mathbf{C}'_{\text{true}}) - \lambda_3(\mathbf{C}'_{\text{true}}) = \phi_{E'}$ and thus in view of (22) it concludes the proof of (9). We denote by $\hat{\mathbf{D}}'_{\text{true}} \in \mathbb{R}^{n \times n}$ the reduction of the matrix $\mathbf{D}'_{\text{true}} \in \mathbb{R}^{2n \times 2n}$ obtained by replacing each of its 2×2 scalar blocks, $\{d'_{\text{true}}(i)\mathbf{I}_2\}_{i=1}^n$,

with diagonal elements, $\{d'_{\text{true}}(i)\}_{i=1}^n$. The corresponding normalized weight matrix is $\widetilde{\mathbf{W}}'_{\text{true}} = \hat{\mathbf{D}}'^{-1}_{\text{true}} \mathbf{W}'_{\text{true}}$. Clearly, the largest eigenvalue of $\widetilde{\mathbf{W}}'_{\text{true}}$ is 1, it has multiplicity 1 (since E' is connected), and its eigenspace is spanned by a column vector of ones in \mathbb{R}^n , which we denote by $\mathbf{1}_{n \times 1}$.

We next relate the eigendecomposition of $\widetilde{\mathbf{W}}'_{\text{true}}$ to that of $\mathbf{C}'_{\text{true}}$. Let $\lambda_1, \dots, \lambda_n$ denote the eigenvalues of $\widetilde{\mathbf{W}}'_{\text{true}}$ in decreasing order, and let $\mathbf{l}_1, \dots, \mathbf{l}_n \in \mathbb{R}^n$ denote the corresponding orthogonal eigenvectors, each with norm \sqrt{n} . Thus, for all $1 \leq i, k \leq n$, $\sum_{j=1}^n \widetilde{\mathbf{W}}'_{\text{true}}(i, j) \mathbf{l}_k(j) = \lambda_k \mathbf{l}_k(i)$. This equation is equivalent to the following one: $\sum_{j=1}^n \widetilde{\mathbf{W}}'_{\text{true}}(i, j) \mathbf{R}_i \mathbf{R}_j^T \mathbf{R}_j \mathbf{l}_k(j) = \lambda_k \mathbf{l}_k(i) \mathbf{R}_i$. Since $\mathbf{C}'_{\text{true}}[i, j] = \widetilde{\mathbf{W}}'_{\text{true}}(i, j) \mathbf{R}_i \mathbf{R}_j^T$ for all $1 \leq i, j \leq n$, the last equation can be written as

$$\sum_{j=1}^n \mathbf{C}'_{\text{true}}[i, j] \mathbf{l}_k(j) \mathbf{R}_j = \lambda_k \mathbf{l}_k(i) \mathbf{R}_i.$$

This equation can be further written as $\mathbf{C}'_{\text{true}}[i, j] \mathbf{U}_k = \lambda \mathbf{U}_k$, where for $1 \leq k \leq n$, $\mathbf{U}_k \in \mathbb{R}^{2n \times 2}$ satisfies $\mathbf{U}_k[i] = \mathbf{l}_k(i) \mathbf{R}_i$ for $1 \leq i \leq n$. Note that the $2n$ columns of $\mathbf{U}_1, \dots, \mathbf{U}_n$ form an orthogonal system and that we obtained a one-to-one correspondence between the eigenvalues and eigenvectors of $\widetilde{\mathbf{W}}'_{\text{true}}$ and $\mathbf{C}'_{\text{true}}$.

A first implication of the above property is that 1 is also the largest eigenvalue of $\mathbf{C}'_{\text{true}}$ with multiplicity 2. Furthermore, since $\mathbf{l}_1 = \mathbf{1}_{n \times 1}$, $\mathbf{U}_1 = \mathbf{R}_{\text{true}}^{\text{tot}}$. By definition, $\mathbf{V}'_{\text{true}} = \mathbf{U}_1$ and thus, as claimed, $\mathbf{V}'_{\text{true}} = \mathbf{R}_{\text{true}}^{\text{tot}}$. Another implication of this property is that $\lambda_2(\mathbf{C}'_{\text{true}}) - \lambda_3(\mathbf{C}'_{\text{true}}) = \lambda_1(\widetilde{\mathbf{W}}'_{\text{true}}) - \lambda_2(\widetilde{\mathbf{W}}'_{\text{true}}) = 1 - \lambda_2(\widetilde{\mathbf{W}}'_{\text{true}})$. Recall that $\phi_{E'}$ is the second smallest eigenvalue of the normalized graph Laplacian, which can be written as $\mathbf{I} - \hat{\mathbf{D}}'^{-1}_{\text{true}} \mathbf{W}'_{\text{true}} = \mathbf{I} - \widetilde{\mathbf{W}}'_{\text{true}}$. Therefore, $\phi_{E'} = 1 - \lambda_2(\widetilde{\mathbf{W}}'_{\text{true}}) = \lambda_2(\mathbf{C}'_{\text{true}}) - \lambda_3(\mathbf{C}'_{\text{true}})$, as claimed. ■

For square jigsaw puzzles, [Theorem 3.1](#) implies that if one can construct a connection graph and find a connected subgraph of it that satisfies (6)–(8) with $\epsilon/\phi_{E'} = O(1/\sqrt{n})$, then [Algorithm 1](#) can nearly recover the correct orientations. More precisely, the estimated rotations of [Algorithm 1](#) are obtained by projection onto \mathbb{Z}_4 of a block matrix whose column space is sufficiently close to the column space of the block matrix of the underlying rotations.

There are several conditions that need to hold in order to imply the conclusion of the theorem. We review them and discuss whether they are reasonable for our proposed method. The first condition is the most restrictive one for our proposed method. It requires (V, E') to be connected. In the case of puzzles with uniform regions (such as the ones demonstrated at the bottom of [Figure 2](#)), the edges in E_{est} obtained by our proposed method, and possibly most methods, may arbitrarily connect patches in these regions while maintaining a small degree for each patch. These edges can be very different from the ones of E_{true} , and thus the intersection of (V, E_{est}) with (V, E_{true}) may result in a disconnected graph. Since $E' \subset E_{\text{true}} \cap E_{\text{est}}$, (V, E') will also be disconnected in this case of uniform regions. Nevertheless, under this setting of uniform regions, the mathematical problem does not have a unique solution and is generally ill-conditioned. None of the tested algorithms performed well in this setting when using standard reconstruction metrics; however, most solved puzzles with uniform regions that looked similar to their ground truth solutions. For nonuniform regions of natural images, this condition is

more reasonable to request from our proposed method. The other requirement is that the connection function is correctly estimated on (V, E') (see (6)) and we also find it reasonable for nonuniform regions of natural images. We note that we can obtain the requirement in (8) with $\gamma = \Theta(1)$ if we assume that d_{\max} is sufficiently small and that (V, E') is connected. Indeed, the latter assumption implies that the LHS of (8) is positive, so that $\gamma \geq 1/d_{\max}$. Our construction trims many unnecessary edges (see section 4) so that d_{\max} is either at most 8 or slightly larger than 8, where the typical 8 neighboring edges include the 4 nearest ones and 4 diagonal ones. Therefore, (8) with $\gamma = \Theta(1)$ is a reasonable requirement for our proposed method. At last, we clarify the condition in (7), where we further discuss the typical size of ϵ below. First, note that the constants c_1, \dots, c_n used in this condition are needed because $\mathbf{W}_{\text{true}}(i, j)$ obtains values in $\{0, 1\}$, whereas $\mathbf{W}_{\text{est}}(i, j)$ can obtain various nonnegative values. The underlying assumption of (7) is that the needed proportion c_i for patch i is similar for “all neighbors.” That is, when $\{i, j\} \in E'$, the scaled value of $\mathbf{W}_{\text{est}}(i, j)$ by c_i is close to the binary weight $\mathbf{W}_{\text{true}}(i, j)$. Moreover, if $\{i, j\} \notin E'$, this scaled value is close to zero. One may expect such an assumption in some practical instances. For example, if the image is continuous at a patch, then the affinities with the nearby patches are all expected to be large and comparable. Similarly, if the image is discontinuous with respect to all neighbors of a patch, then all neighboring affinities are expected to be very small. Therefore, such a constant c_i may be chosen for each patch. We remark that our proposed algorithm for solving type 2 puzzles (Algorithm 4) aims to assign very small affinities whenever there is any possible inconsistency in the construction of the graph connection Laplacian. Such an assignment aims to guarantee that the second sum in (7) is sufficiently small. Nevertheless, we cannot really guarantee that our heuristic choices work in practice. We also comment that given the order of ϵ established below, condition (7) is rather sensitive. Indeed, if for $\{i, j\} \notin E'$, $\mathbf{W}_{\text{true}}(i, j) = \Theta(1)$, then this condition is violated.

As we mentioned above, in our construction $d_{\max} = \Theta(1)$, where often $d_{\max} \leq 8$ or slightly larger than 8, and thus we can choose $\gamma = \Theta(1)$. Therefore, as long as $\epsilon/\phi_{E'} = O(1/\sqrt{n})$ the bound in (9) is sufficiently small. That is, the perturbation bound in (7) is more restrictive as the size of the puzzle increases and $\phi_{E'}$ decreases. To get a better idea of $\phi_{E'}$, we can assume a puzzle grid with lengths and widths of order \sqrt{n} . If the graph (V, E') is a lattice, then by direct application of Cheeger’s inequality, $\phi_{E'} = \Theta(1/\sqrt{n})$. Similarly, in the worst case of a path, $\phi_{E'} = \Theta(1/n)$.

A few additional remarks are in order. First, the conditions of the theorem are sufficient but not necessary. Second, one may use in the proof above another variant of the Davis–Kahan sin Θ theorem according to [13] and consequently obtain a bound similar to (9), but controlling the infinity norm, and not the Frobenius norm. Third, while we think of (V, E_{true}) as having a grid-type structure, one can assume in theory a general graph (V, E_{true}) . At last, the theorem can be easily generalized to any group synchronization and not just to \mathbb{Z}_4 synchronization. This generalization can be possibly applied to other puzzles, such as 3D puzzles or nonsquared jigsaw puzzles.

4. Connection graph construction for type 2 and type 3 puzzles. As we have discussed in section 3.1, if we are given a perfect metric, we can easily construct the connection graph. However, there is no perfect metric that would work for all images. For example, if part of

the image contains a region with a uniform color, such as sky or ocean (see the images in the second row of [Figure 2](#)), the metric between the sides of the image patches from this region will be close to zero. Thus, all these patches should be wrongly identified by a perfect metric as neighbors. Therefore, the idea of finding a perfect metric and using a threshold to identify neighbors may not lead to a correct affinity graph. Instead, we suggest iteratively updating the graph construction while identifying possibly incorrect edges and reassigning zero or small affinities to them. Nevertheless, our initial estimate of the graph affinities is based on the Mahalanobis Gradient Compatibility (MGC) of [\[16\]](#), since empirically it seems to be nearly perfect in well-posed cases. We review its construction in section [4.1](#). The constructions of the graph affinities and GCW for type 3 and type 2 puzzles are described in section [4.2](#) and section [4.3](#), respectively.

4.1. Gallagher's MGC metric. We review the construction of the MGC metric [\[16\]](#). This metric between sides of patches quantifies the proximity of the gradients computed at each side. It outputs a symmetrized version of the Mahalanobis distance between the vector of gradients of one side and the estimated distribution of gradients at the other side, where the distribution is represented by its estimated covariance.

We assume two neighboring image patches P_i and P_j of size $s \times s$. There are four different relative positions of P_i and P_j . We assume without loss of generality the left-right relative position (P_i on the left and P_j on the right) and compute the corresponding MGC, which we denote by $\text{MGC}_{\text{lr}}(P_i, P_j)$, as follows. For each color channel c (red, green, and blue) and each row r , $1 \leq r \leq s$, of the $s \times s$ patch P_i , we find the derivatives near the right side of the image patch P_i in the direction left-right as follows:

$$G_{iL}(r, c) = P_i(r, s, c) - P_i(r, s - 1, c).$$

The subscript L in the above equation indicates that patch P_i is on the left side of the patch P_j . To avoid some numerical problems, Gallagher [\[16\]](#) suggests adding the 9 “dummy gradients” $(0, 0, 0)$, $(1, 1, 1)$, $(-1, -1, -1)$, $(0, 0, 1)$, $(0, 1, 0)$, $(1, 0, 0)$, $(-1, 0, 0)$, $(0, -1, 0)$, and $(0, 0, -1)$ as additional rows of the matrix \mathbf{G}_{iL} . The extended matrix in $\mathbb{R}^{(s+9) \times 3}$ is denoted by $\tilde{\mathbf{G}}_{iL}$.

Next, for each color channel c we define

$$\boldsymbol{\mu}_{iL}(c) = \frac{1}{s} \sum_{r=1}^s \mathbf{G}_{iL}(r, c).$$

The regularized covariance matrix $\boldsymbol{\Sigma}_{iL} \in \mathbb{R}^{3 \times 3}$ between color channels is

$$\boldsymbol{\Sigma}_{iL} = \frac{1}{s+8} (\tilde{\mathbf{G}}_{iL} - \text{mean}(\tilde{\mathbf{G}}_{iL}))^T (\tilde{\mathbf{G}}_{iL} - \text{mean}(\tilde{\mathbf{G}}_{iL})),$$

where

$$\text{mean}(\tilde{\mathbf{G}}_{iL}) = \frac{1}{s+9} \sum_{r=1}^{s+9} \tilde{\mathbf{G}}_{iL}(r, c) = \frac{1}{s+9} \sum_{r=1}^s \mathbf{G}_{iL}(r, c).$$

We also define $\mathbf{G}_{ijLR}(r)$, the derivative from the left $s \times s$ image patch P_i to the right $s \times s$

image patch P_j at row r and color c , by

$$\mathbf{G}_{ij\text{LR}}(r, c) = P_j(r, 1, c) - P_i(r, s, c).$$

The left-to-right compatibility measure from P_i to P_j is defined by

$$D_{\text{LR}}(P_i, P_j) = \sum_{r=1}^s (\mathbf{G}_{ij\text{LR}}(r) - \boldsymbol{\mu}_{iL}) \boldsymbol{\Sigma}_{iL}^{-1} (\mathbf{G}_{ij\text{LR}}(r) - \boldsymbol{\mu}_{iL})^T.$$

Similarly, one can define the right-to-left compatibility measure from P_j to P_i in the same left-right setting, where P_i is to the left of P_j . The left-right MGC metric then has the symmetrized form

$$(24) \quad \text{MGC}_{\text{lr}}(P_i, P_j) = D_{\text{LR}}(P_i, P_j) + D_{\text{RL}}(P_j, P_i).$$

The right-left, top-bottom, and bottom-top MGCs, denoted by $\text{MGC}_{\text{rl}}(P_i, P_j)$, $\text{MGC}_{\text{tb}}(P_i, P_j)$, and $\text{MGC}_{\text{bt}}(P_i, P_j)$, respectively, are computed similarly.

4.2. Connection graph construction for type 3 puzzles. For type 3 puzzles, the locations of patches are given. Furthermore, edges are drawn between neighboring patches. The affinity function is set by $\mathbf{W}(i, j) = 1$ for all $\{i, j\} \in E$. One need only find the unknown orientations, that is, the unknown connection matrix \mathbf{R} .

To construct the connection function we propose using the MGC metric described in section 4.1. For all neighboring patches P_i and P_j , we calculate the possible 16 values of the MGC metric (corresponding to the 16 relative positions of P_i and P_j) and select the smallest of these numbers and its corresponding rotation $\mathbf{R}[i, j]$. If there is no unique minimum among these 16 values, we suggest assigning $\mathbf{W}(i, j) = 1/2$ (or another value smaller than 1) and letting $\mathbf{R}[i, j]$ be the mean of the candidate rotations that obtain the minimal value.

4.3. Connection graph construction for type 2 puzzles. We propose the following step-by-step procedure for constructing the affinity graph, the affinity function, and the connection function for type 2 puzzles and then summarize this procedure in [Algorithm 2](#). The rest of this section is organized as follows: section 4.3.1 discusses the initial step of constructing the connection graph; section 4.3.2 discusses the Jaccard index and explains how to use it to update the affinity function; section 4.3.3 describes how to deal with the cases when the connection graph is disconnected; section 4.3.4 describes how to find and use diagonal neighbors in order to construct a more reliable connection graph.

4.3.1. Initial step. We start with an initial construction of the directed graph $G = (V, E_{\text{est}})$. The vertex set V contains the patches in \mathcal{Q} . The edge set E_{est} is updated by the following procedure. In order to describe it, we denote by $\mathbf{R} \cdot P$ the action of the rotation $\mathbf{R} \in \mathbb{Z}_4$ on the patch P . For a patch P_i , we find the patches $P_{i_t}, P_{i_l}, P_{i_b}, P_{i_r}$ and the

corresponding rotations $\mathbf{R}[i, i_t], \mathbf{R}[i, i_l], \mathbf{R}[i, i_r], \mathbf{R}[i, i_b] \in \mathbb{Z}_4$ such that

$$(25) \quad \begin{aligned} \{P_{i_t}, \mathbf{R}[i, i_t]\} &\in \underset{P \in \mathcal{Q}, \mathbf{R} \in \mathbb{Z}_4}{\operatorname{argmin}} \text{MGC}_{\text{bt}}(P_i, \mathbf{R} \cdot P), \\ \{P_{i_l}, \mathbf{R}[i, i_l]\} &\in \underset{P \in \mathcal{Q}, \mathbf{R} \in \mathbb{Z}_4}{\operatorname{argmin}} \text{MGC}_{\text{rl}}(P_i, \mathbf{R} \cdot P), \\ \{P_{i_b}, \mathbf{R}[i, i_b]\} &\in \underset{P \in \mathcal{Q}, \mathbf{R} \in \mathbb{Z}_4}{\operatorname{argmin}} \text{MGC}_{\text{tb}}(P_i, \mathbf{R} \cdot P), \\ \{P_{i_r}, \mathbf{R}[i, i_r]\} &\in \underset{P \in \mathcal{Q}, \mathbf{R} \in \mathbb{Z}_4}{\operatorname{argmin}} \text{MGC}_{\text{lr}}(P_i, \mathbf{R} \cdot P). \end{aligned}$$

The set E_{est} of directed edges contains the edges that connect each vertex with index i , $1 \leq i \leq n$, to the vertices with indices i_t , i_l , i_b , and i_r . Note that these indices solving (25) may not be unique, and we consider all solutions of (25) when forming E_{est} .

We next modify the directed graph $G = (V, E_{\text{est}})$ into an undirected graph. We use the default parameter 0.01 and form an initial affinity matrix $\mathbf{W}_{\text{init}} \in \mathbb{R}^{n \times n}$ whose elements for $1 \leq i, j \leq n$ are

$$(26) \quad \mathbf{W}_{\text{init}}(i, j) = \mathbf{W}_{\text{init}}(j, i) = \begin{cases} 1 & \text{if both } (i, j) \text{ and } (j, i) \in E_{\text{est}}, \\ 0.01 & \text{if only one of } (i, j) \text{ or } (j, i) \text{ is in } E_{\text{est}}, \\ 0 & \text{otherwise.} \end{cases}$$

Figure 3 demonstrates the above construction for a particular patch.

Next, we enforce the constraint that, for square jigsaw puzzles, each patch can have at most one neighbor for each direction by trimming some edges that are likely not neighbors. This is done as follows. Assume without loss of generality that patch P_i has more than one neighbor in the top direction and denote these neighbors by P_{i_1}, \dots, P_{i_k} , where $k > 1$. Then we solve the minimization problem

$$(27) \quad j \in \underset{1 \leq j \leq k}{\operatorname{argmin}} \text{MGC}_{\text{bt}}(P_i, \mathbf{R}[i, i_j] \cdot P_{i_j}).$$

We remark that since (26) makes the graph undirected, P_{i_1}, \dots, P_{i_k} might contain more patches than the ones obtained from (25). Therefore, the optimization problems (25) and (27) can be different. In particular, a unique solution of (25) might become nonunique for (27). If (27) has a unique solution, we keep the edge $\{i, i_j\}$ and remove the rest of the edges. Otherwise, we remove all edges $\{i, i_j\}_{j=1}^k$ from E . The procedure is analogous if P_i has more than one neighbor from left, bottom, or right. If edges were eliminated from E_{est} , then the matrix \mathbf{W}_{init} is updated so it is zero on the corresponding indices. This process results in the following initial connection graph: $G = (V, E_{\text{est}}, \mathbf{W}_{\text{init}}, \mathbf{R})$.

The construction of this graph uses a nearest-neighbor construction. For a high-noise regime, El Karoui and Wu [12] recommend avoiding a nearest-neighbor construction. However, due to the special lattice structure of the true graph, the nearest-neighbor initial construction is natural for the square jigsaw puzzle problem.

4.3.2. Use of Jaccard index to refine the graph. Next, we refine the connection graph by trying to assess the validity of the edges and decrease the weights of edges that do not

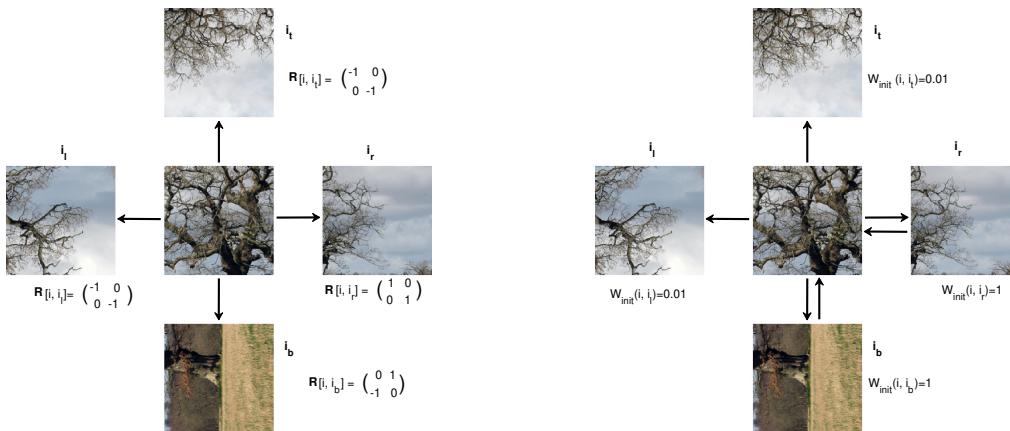


Figure 3. Demonstration of the initial step for the construction of the connection graph. The left figure demonstrates the best matches for a given patch from the four directions: top, left, bottom, and right. For each matching patch it records the rotation whose application to this patch results in correct matching with the central patch. The right figure shows the application of these rotations to the matching patches and demonstrates how to assign the weights to the undirected graph. In this example, the matching patches from top and left were originally connected by a single direction; their weights in the undirected graph are thus 0.01. On the other hand, the patches from right and bottom are connected in both directions, and thus their weights in the undirected graph are 1.

seem valid; that is, they may not appear in the true connection graph. The idea is to check after removing an edge whether its neighbors are still connected in some weak sense to each other. If so, then the edge seems to be valid, and otherwise, it may not be valid. For this purpose, we use the Jaccard index [17].

The description of this index uses the following notation in a graph $G = (V, E)$. Given a vertex i , $1 \leq i \leq |V|$, let $N_{G,i}^1$ denote the set of vertices in V which are connected to vertex i , that is, $N_{G,i}^1 = \{j \in V \mid \{i, j\} \in E\}$. Using our terminology, $N_{G,i}^1$ contains the neighbors of i . The set $N_{G,i}^2$ contains all vertices that are at most 2 steps away from vertex i , except vertex i . That is, $N_{G,i}^2 = \bigcup_{j \in N_{G,i}^1} N_{G,j}^1 \setminus \{i\}$. Finally, let $G^{\setminus(i,j)} = (V, E \setminus \{(i,j)\})$ denote the graph with the edge (i,j) removed. The sets $N_{G,i}^1$ and $N_{G,i}^2$ are demonstrated in Figure 4.

By using this notation, we define the Jaccard index between vertices i and j as

$$(28) \quad \mu_{\text{Jaccard}}(i, j) = |N_{G^{\setminus(i,j)}, i}^2 \cap N_{G^{\setminus(i,j)}, j}^2|,$$

where $|\cdot|$ denotes the cardinality of a set. This definition is similar to the one in [17], but there are two differences. The first one is that we consider the graph $G^{\setminus(i,j)}$ instead of G to emphasize the common neighbors, while excluding the obvious pair (i,j) . The second one is that we do not divide by $|N_{G^{\setminus(i,j)}, i}^2 \cup N_{G^{\setminus(i,j)}, j}^2|$. The latter division does not matter to us as we only care about the positivity of this index. Figure 5 demonstrates calculation of the Jaccard index for a special example.

Note that the chance of two vertices i and j being neighbors in the graph (V, E_{est}) is higher if $\mu_{\text{Jaccard}}(i, j) > 0$ than if $\mu_{\text{Jaccard}}(i, j) = 0$. Thus, we propose using the Jaccard indices to

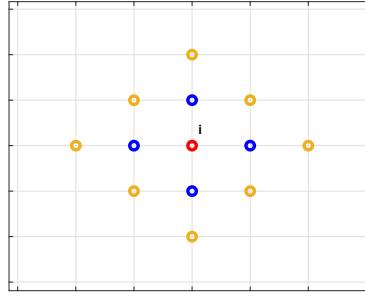


Figure 4. Demonstration of the sets $N_{G,i}^1$ and $N_{G,i}^2$. A given vertex i is colored in red, the elements of the set $N_{G,i}^1$ are colored in blue, and the elements of the set $N_{G,i}^2$ are colored in blue and orange.

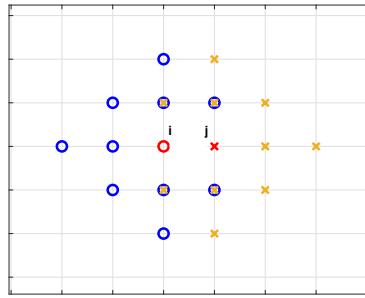


Figure 5. Demonstration of Jaccard index. Vertex i is denoted by a red circle, and vertex j is denoted by a red cross. The edge between these vertices was removed from the grid. The elements of $N_{G\setminus(i,j),i}^2$ are denoted by blue circles and the elements of $N_{G\setminus(i,j),j}^2$ by orange crosses. The Jaccard index is four since there are four elements in $N_{G\setminus(i,j),i}^2 \cap N_{G\setminus(i,j),j}^2$ (denoted by blue circles filled with orange crosses).

refine the connection graph. We use another weight matrix $\mathbf{W}_{\text{Jaccard}} \in \mathbb{R}^{n \times n}$, defined as

$$(29) \quad \mathbf{W}_{\text{Jaccard}}(i,j) = \mathbf{W}_{\text{Jaccard}}(j,i) = \begin{cases} 0, & \{i,j\} \in E_{\text{est}} \text{ and } \mu_{\text{Jaccard}}(i,j) = 0, \\ \mathbf{W}_{\text{init}}(i,j) & \text{otherwise.} \end{cases}$$

Since this procedure might also remove many correct edges by zeroing out the corresponding values of the affinity function, we propose a linear combination of \mathbf{W}_{init} and $\mathbf{W}_{\text{Jaccard}}$ with larger coefficient given to $\mathbf{W}_{\text{Jaccard}}$. In our experiments we use the default parameter 0.2 (and $0.8 = 1 - 0.2$) to set

$$(30) \quad \mathbf{W}_{\text{nb}} = 0.2 \times \mathbf{W}_{\text{init}} + 0.8 \times \mathbf{W}_{\text{Jaccard}}$$

and form the affinity graph $G = (V, E, \mathbf{W}_{\text{nb}}, R)$.

4.3.3. Making the affinity graph connected. The procedures described in section 4.3.1 and section 4.3.2 might result in a disconnected affinity graph G as demonstrated in the left



Figure 6. Demonstration of a disconnected affinity graph and the way it got connected. The left figure shows an example where the resulting affinity graph using our method is disconnected. Indeed, the two top right patches are not connected to any of the other patches. The black edges connect between true neighbors, and the only red edge is a wrongly determined edge. The right figure demonstrates the result of the simple procedure described in section 4.3.3. The connected graph has two new blue edges. While these blue edges connect between nonneighboring patches, the originally disconnected patches are uniformly white, and thus their rotations do not matter for the reconstruction of the image.

image of Figure 6. To complete G so it is connected, we first find all connected components of G . Assume that there are k connected components with corresponding vertices V_1, \dots, V_k that partition the set of vertices V . Assume further that they are labeled by descending size order, i.e., $|V_1| \geq |V_2| \dots \geq |V_k|$. Next, we find vertices $i \in V_1$ and $j \in V \setminus V_1$ that minimize the MGC metric between the patches P_i and P_j . Mathematically, we find

$$(31) \quad \{i, j, \mathbf{O}\} \in \underset{i \in V \setminus V_1, j \in V_1, \mathbf{O} \in \mathbb{Z}_4}{\operatorname{argmin}} \min\{\text{MGC}_{\text{lr}}(P_i, \mathbf{O} \cdot P_j), \text{MGC}_{\text{tb}}(P_i, \mathbf{O} \cdot P_j), \text{MGC}_{\text{rl}}(P_i, \mathbf{O} \cdot P_j), \text{MGC}_{\text{bt}}(P_i, \mathbf{O} \cdot P_j)\}.$$

If the solution of (31) is not unique, we randomly choose one solution. We then add the edge $\{i, j\}$ of the chosen solution to E_{est} and update the weight as follows:

$$(32) \quad \mathbf{W}_{\text{nb}}(i, j) = \mathbf{W}_{\text{nb}}(j, i) = 0.005, \quad \mathbf{R}[i, j] = \mathbf{O}, \quad \text{and} \quad \mathbf{R}[j, i] = \mathbf{O}^T.$$

We remark that 0.005 is a free parameter chosen to be half the size of the parameter 0.01 in (26). We iterate the procedure described above for V_2, \dots, V_k until the graph becomes connected. The number of iterations needed is $k - 1$ since there are k connected components and at each iteration we connect the largest component with a remaining component.

4.3.4. Taking advantage of 4-loops. We refine the constructed connection graph by using the following property of the square jigsaw puzzle: If two patches P_i and P_j are diagonal neighbors, then there exist exactly two other patches P_{n_1} and P_{n_2} and a cycle containing the vertices i, n_1, j , and n_2 . This idea is demonstrated in Figure 7. Such a cycle of 4 vertices is referred to as a 4-loop in [38]. In this latter work, 4-loops were used to solve the puzzle problem. We use them to define a better connection graph. As we have already discussed, for

a uniform grid, each patch can have at most 4 direct neighbors (right, top, left, and bottom). Furthermore, each patch has at most 4 diagonal neighbors. Exactly four diagonal neighbors are obtained for a patch in the interior of the puzzle, a single diagonal neighbor occurs for a corner patch, and there are 2 diagonal neighbors for a patch that lies on the boundary of the uniform grid but not on a corner.

For patches P_i and P_j we define

$$(33) \quad \delta_{\text{diag}}(i, j) = |N_{G,i}^1 \cap N_{G,j}^1|.$$

We observe that patches P_i and P_j are diagonal neighbors in the uniform grid if and only if $\delta_{\text{diag}}(i, j) = 2$. To find the diagonal neighbors for graph $G = (V, E)$ we propose a two step procedure. First, we find the set of all pairs of vertices $\{i, j\} \in V \times V$ for which $\delta_{\text{diag}}(i, j) = 2$. For each such pair $\{i, j\}$ there exists another pair $\{n_1, n_2\}$ such that

$$(34) \quad N_{G,i}^1 \cap N_{G,j}^1 = \{n_1, n_2\},$$

or, equivalently, i , n_1 , j , and n_2 are contained in a 4-loop. We set

$$(35) \quad \mathbf{W}_{\text{diag}}(i, j) = \begin{cases} 1 & \text{when } \delta_{\text{diag}}(i, j) = 2 \text{ and } \mathbf{R}[i, n_1]\mathbf{R}[n_1, j] = \mathbf{R}[i, n_2]\mathbf{R}[n_2, j], \\ 0 & \text{otherwise.} \end{cases}$$

If $\mathbf{W}_{\text{diag}}(i, j) = 1$, we add a diagonal edge between vertices i and j and assign the following value to the connection function:

$$(36) \quad \mathbf{R}[i, j] = \mathbf{R}[i, n_1]\mathbf{R}[n_1, j].$$

We remark that the condition $\mathbf{R}[i, n_1]\mathbf{R}[n_1, j] = \mathbf{R}[i, n_2]\mathbf{R}[n_2, j]$ in (35) is naturally satisfied in the ground-truth graph as demonstrated in Figure 8. Therefore, when it is satisfied and also $\delta_{\text{diag}}(i, j) = 2$, the maximal weight of 1 is assigned to the corresponding diagonal edge. We note that $\delta_{\text{diag}}(i, j) = 2$ if and only if $\delta_{\text{diag}}(n_1, n_2) = 2$, $\mathbf{R}[i, n_1]\mathbf{R}[n_1, j] = \mathbf{R}[i, n_2]\mathbf{R}[n_2, j]$ if and only if $\mathbf{R}[n_1, i]\mathbf{R}[i, n_2] = \mathbf{R}[n_1, j]\mathbf{R}[j, n_2]$ and thus $\mathbf{W}_{\text{diag}}(i, j) = 1$ if and only if $\mathbf{W}_{\text{diag}}(n_1, n_2) = 1$.

We further update blocks of the matrix \mathbf{W}_{nb} so that their weights align better with the diagonal weights. We denote by $\mathbf{W}_{\text{nb}}([i, j], [n_1, n_2])$ the 2×2 submatrix of \mathbf{W}_{nb} indexed by (i, n_1) , (i, n_2) , (j, n_1) , and (j, n_2) . We fix the default parameters $1/3$ and $2/3$ and describe the assignment of $\mathbf{W}_{\text{nb}}([i, j], [n_1, n_2])$ in the following formula (though algorithmically one does not need to implement it per each block):

$$(37) \quad \mathbf{W}_{\text{nb}}([i, j], [n_1, n_2]) = \begin{cases} \frac{\mathbf{W}_{\text{nb}}([i, j], [n_1, n_2])}{3} & \text{if } \delta_{\text{diag}}(i, j) = 2 \text{ and} \\ & \mathbf{R}[i, n_1]\mathbf{R}[n_1, j] \neq \mathbf{R}[i, n_2]\mathbf{R}[n_2, j], \\ \mathbf{1}_{2 \times 2} & \text{if } \delta_{\text{diag}}(i, j) = 2 \text{ and} \\ & \mathbf{R}[i, n_1]\mathbf{R}[n_1, j] = \mathbf{R}[i, n_2]\mathbf{R}[n_2, j], \\ \frac{2\mathbf{W}_{\text{nb}}([i, j], [n_1, n_2])}{3} & \text{otherwise.} \end{cases}$$

We note that if $\{i, j\}$ and $\{n_1, n_2\}$ are diagonal edges satisfying the rotation condition, then we enforce weight 1 for the nearest neighboring edges in the corresponding 4-loop: $\{i, n_1\}$,

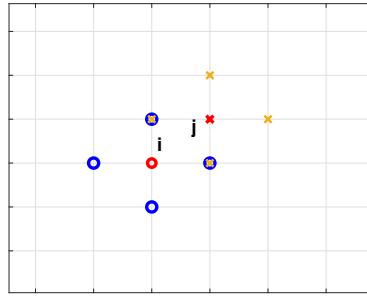


Figure 7. Demonstration of finding diagonally neighboring vertices in a uniform grid. Two vertices i and j are denoted by a red circle and a red cross, respectively. The elements of the sets $N_{G,i}^1$ and $N_{G,j}^1$ are colored by blue and orange, respectively. The intersection of these sets yields the two diagonally neighboring vertices to i and j . Together with i and j they form a 4-loop.

$\{i, n_2\}$, $\{j, n_1\}$, and $\{j, n_2\}$. If they are not diagonal and do not satisfy the rotation condition, then we decrease the existing weights of these neighboring edges by a factor of $2/3$. If the diagonal condition, $\delta_{\text{diag}}(i, j) = 2$, holds but the rotation condition does not, which is somewhat contradicting, we reduce the weights of the former neighboring edges by a factor of $1/3$.

We note that the support sets of \mathbf{W}_{nb} and \mathbf{W}_{diag} are disjoint. We set

$$(38) \quad \mathbf{W} = \mathbf{W}_{\text{nb}} + \mathbf{W}_{\text{diag}},$$

and this is the final step of constructing the connection graph $G = (V, E_{\text{est}}, \mathbf{W}, \mathbf{R})$ for square jigsaw puzzles. The full algorithm of this construction is summarized in [Algorithm 2](#).

5. Solution for type 2 puzzles via GCL and location solver. This section completes the solution of type 2 jigsaw puzzles. It assumes the formation of the connection graph according to [Algorithm 2](#), solution of the correct orientations by [Algorithm 1](#), and then application of an existing location solver. The new component is a procedure for updating the affinity function and the connection function based on the estimated rotations and locations. One can then estimate again the orientations and locations and repeat this procedure several times. This procedure and the complete solution of type 2 puzzles that uses this procedure are summarized below in section 5.1. At last, section 5.2 summarizes its time complexity.

5.1. Updating the affinity and connection functions and the resulting solution. By now there are many successful solutions to type 1 puzzles. According to our numerical tests, the stand-alone algorithms for solving type 1 puzzles of both Gallagher [16] and Yu, Russell, and Agapito [43] are highly competitive. We have often noticed a slight advantage of the latter algorithm, which applies a linear programming procedure. Therefore, we use the algorithm of Yu, Russell, and Agapito [43] (under the “fixed-rotation mode”) as a default solver for type 1 puzzles in our algorithm. One could use instead any algorithm that solves type 1 puzzles.

The basic idea for updating the values of the affinity and connection functions is that a given estimated solution for the orientations and locations can be used to infer possible mismatches. These identified mismatches could be used to reassign values for the affinity and connection functions that may lead to a more accurate solution.

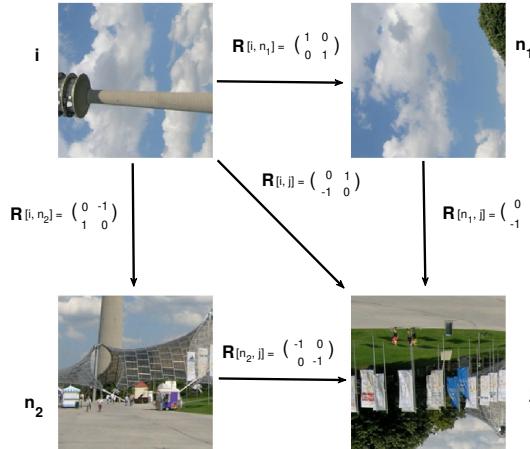


Figure 8. Intuition for the requirement in (35) and (37). The two vertices i and j are diagonal neighbors, and the vertices n_1 and n_2 satisfy (34). Thus, i , j , n_1 , and n_2 form a cycle of size 4, that is, a 4-loop. The relative rotations between vertices are indicated on the corresponding edges. We note that both $\mathbf{R}[i, n_1]\mathbf{R}[n_1, j]$ and $\mathbf{R}[i, n_2]\mathbf{R}[n_2, j]$ are equal to the relative rotation $\mathbf{R}[i, j]$ shown on edge (i, j) . In particular, $\mathbf{R}[i, n_1]\mathbf{R}[n_1, j] = \mathbf{R}[i, n_2]\mathbf{R}[n_2, j]$. The assigned weights thus try to encourage this constraint and penalize cases where it is not satisfied.

First, we figure out which patches are wrongly placed in the assembled puzzle and remove them from the grid. For this purpose we use the following kinds of metrics, which we refer to as NAM (Neighbor-Averaged Metric). For each patch i , $1 \leq i \leq n$, with neighbors i_t , i_l , i_b , and i_r , from top, left, bottom, and right, respectively, we utilize the MGC metric (see (24) and text below it) to define

$$(39) \quad \text{NAM}_{\text{all}}(i) = (\text{MGC}_{\text{bt}}(i, i_t) + \text{MGC}_{\text{rl}}(i, i_l) + \text{MGC}_{\text{tb}}(i, i_b) + \text{MGC}_{\text{lr}}(i, i_r))/4.$$

If a patch i is at the edge or corner of the puzzle grid, then it has 3 or 2 neighbors, respectively. In this case, we only sum up the respective MGC values and divide the sum by the number of neighbors. Similarly we define the following four metrics:

$$(40) \quad \begin{aligned} \text{NAM}_{\text{ltr}}(i) &= (\text{MGC}_{\text{rl}}(i, i_l) + \text{MGC}_{\text{bt}}(i, i_t) + \text{MGC}_{\text{lr}}(i, i_r))/3, \\ \text{NAM}_{\text{trb}}(i) &= (\text{MGC}_{\text{bt}}(i, i_t) + \text{MGC}_{\text{lr}}(i, i_r) + \text{MGC}_{\text{tb}}(i, i_b))/3, \\ \text{NAM}_{\text{blt}}(i) &= (\text{MGC}_{\text{tb}}(i, i_b) + \text{MGC}_{\text{rl}}(i, i_l) + \text{MGC}_{\text{bt}}(i, i_t))/3, \\ \text{NAM}_{\text{lbr}}(i) &= (\text{MGC}_{\text{rl}}(i, i_l) + \text{MGC}_{\text{tb}}(i, i_b) + \text{MGC}_{\text{lr}}(i, i_r))/3. \end{aligned}$$

Again, if the patch is at the edge or corner of the puzzle grid, then we sum the appropriate MGC values and divide by the corresponding number of neighbors.

If the puzzle is correctly assembled, the NAM values of all patches are relatively small as demonstrated for NAM_{all} values in last figure of the first row of Figure 9. Otherwise, if there are some wrongly placed patches, their corresponding NAM values should be relatively higher. This is demonstrated in the first row of Figure 9, where the spikes of NAM_{all} values

Algorithm 2 Connection graph construction for type 2 puzzles.**Input:** Puzzle Patches: $\{P_i\}_{i=1}^n \subset \mathbb{R}^{s \times s \times 3}$

- For all $1 \leq i < j \leq n$ calculate the 16 MGC metric values between patches P_i and P_j as explained in section 4.1
- Construct $G = (V, E_{\text{est}}, \mathbf{W}_{\text{init}}, \mathbf{R})$ according to the procedure described in section 4.3.1 with the following three stages: nearest-neighbors construction based on (25), symmetrization of \mathbf{W}_{init} , and pruning extra neighbors with the use of (27)
- For all $\{i, j\} \in E_{\text{est}}$, calculate $\mu_{\text{Jaccard}}(i, j)$ according to (28)
- For all $\{i, j\} \in E_{\text{est}}$, if $\mu_{\text{Jaccard}}(i, j) = 0$, set $\mathbf{W}_{\text{Jaccard}}(i, j) = 0$; otherwise, $\mathbf{W}_{\text{Jaccard}}(i, j) = 1$
- Set $\mathbf{W}_{\text{nb}} = 0.8 \times \mathbf{W}_{\text{Jaccard}} + 0.2 \times \mathbf{W}_{\text{init}}$
- If the graph G is disconnected, iteratively connect the largest connected component to smaller connected components as explained in section 4.3.3
- For all $i, j \in V$, calculate $\delta_{\text{diag}}(i, j)$ according to (33), and if $\delta_{\text{diag}}(i, j) = 2$, calculate n_1 and n_2 according to (34)
- Form \mathbf{W}_{diag} according to (35) and update \mathbf{R} and \mathbf{W}_{nb} according to (36) and (37)
- Set $\mathbf{W} = \mathbf{W}_{\text{nb}} + \mathbf{W}_{\text{diag}}$

Return: $G = (V, E_{\text{est}}, \mathbf{W}, \mathbf{R})$, MGC values for all pairs of patches

correspond to wrongly orientated or placed patches. Based on this observation, we suggest finding all patches for which the corresponding NAM_{all} value and at least one of the NAM_{ltr} , NAM_{trb} , NAM_{blt} , and NAM_{lbr} values exceeds 1.5 times the median of all corresponding NAM values. We remove the corresponding edges from the grid. For example, for NAM_{ltr} we remove the edges connecting vertex i with its left, top, and right neighbors. We refer to a location as empty if all edges connecting the patch in this location to its top, bottom, left, and right neighbors were removed. Patches at empty locations at each iteration of this procedure are demonstrated in the second row of Figure 9. Their removal, which literally creates empty locations, is demonstrated in the last row of this figure.

Next, we select all patches at empty locations for which at least 2 of 4 neighboring locations in the puzzle grid (including edges removed in the current process) are not empty. For each selected patch, we denote the set of neighboring patches by S_{nb} . Note that by our selection criterion, S_{nb} contains 2, 3, or 4 indices. For each selected patch, we find an oriented patch in S_{nb} that minimizes the averaged MGC metric. We denote this minimal value by $\text{NAM}_{S_{\text{nb}}}$ and also denote by $\text{med}(\text{NAM}_{\text{all}})$ the median of all NAM_{all} values. For a selected patch i and a patch $j \in S_{\text{nb}}$, we fix the default parameters 0.3 and 0.6 and update the affinity function as follows:

$$(41) \quad \mathbf{W}_{\text{est}}(i, j) = \mathbf{W}_{\text{est}}(j, i) = \begin{cases} 0.6 & \text{if } \text{NAM}_{S_{\text{nb}}} < \text{med}(\text{NAM}_{\text{all}}), \\ 0.3 & \text{if } \text{med}(\text{NAM}_{\text{all}}) < \text{NAM}_{S_{\text{nb}}} < 2\text{med}(\text{NAM}_{\text{all}}). \end{cases}$$

We also update the connection function as follows, where we denote by $\mathbf{R}_{i,j}$ the current solution of the pairwise orientation of patch i with respect to patch j :

$$(42) \quad \mathbf{R}_{\text{est}}[i, j] = \mathbf{R}_{i,j} \text{ and } \mathbf{R}_{\text{est}}[j, i] = \mathbf{R}_{i,j}^T \text{ if } \text{NAM}_{S_{\text{nb}}} < 2\text{med}(\text{NAM}_{\text{all}}).$$

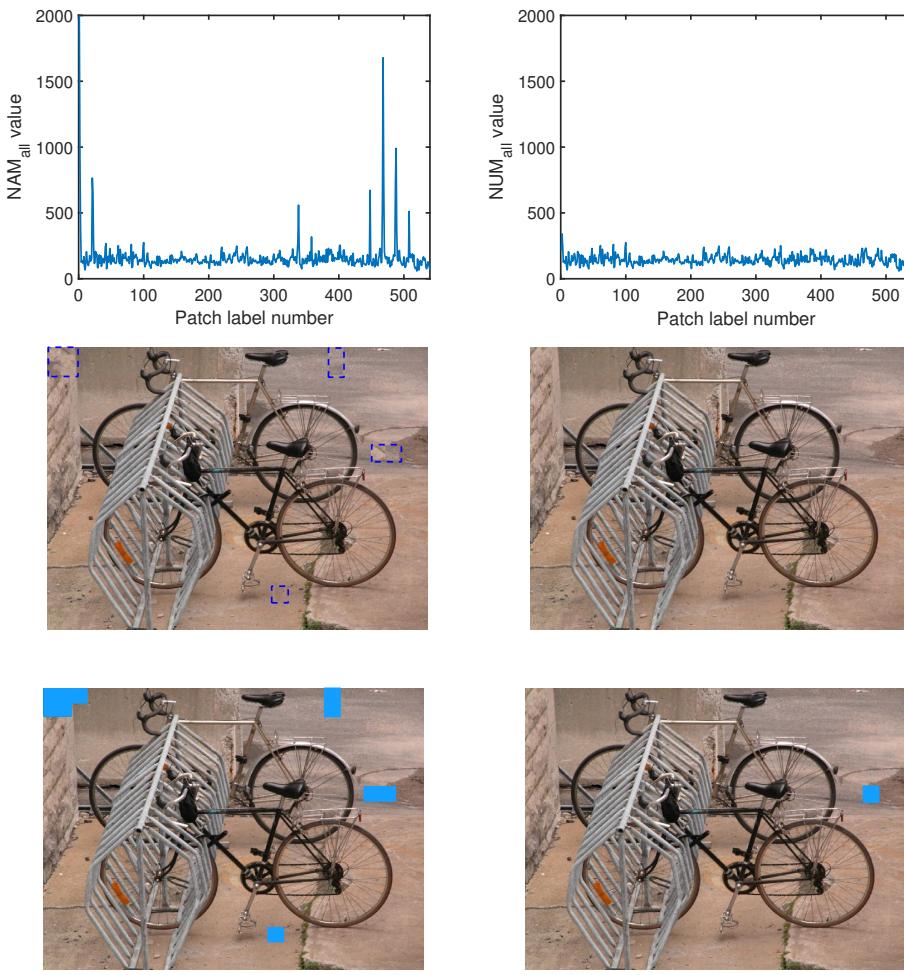


Figure 9. Demonstration of the update step for 2 iterations, described in section 5.1, of a type 2 puzzle with 540 pieces each with sizes of 28×28 . The first row shows the histograms of the NAM_{all} metric values for all patches, defined in (39). The second row shows the solution of the puzzle after each iteration of assembling the puzzle, and the third row shows the remaining patches of an assembled puzzle after removing the patches that are wrongly placed or oriented.

Algorithm 3 summarizes this update procedure.

Finally, our proposed algorithm for reassembling square jigsaw puzzles is summarized in Algorithm 4. It iteratively solves the puzzle by repeating the following 3 steps: finding the orientations of all patches, finding the locations of all patches, and updating the connection function and the affinity function. To measure how good the solution is at each iteration, we

Algorithm 3 Updating the affinity and connection functions at a given iteration.

Input: MGC metric values between all patches, current solution to the puzzle problem

- Calculate the NAM values for all patches according to (39) and (40)
- Remove all edges from the solution grid for which the corresponding NAM_{all} value and at least one of the NAM_{ltr} , NAM_{trb} , NAM_{blt} , and NAM_{lbr} values exceeds 1.5 times the median of all corresponding NAM values
- For any patch at an empty location (that is, location whose edges were all removed) which has at least two nonempty neighboring locations (according to the natural grid of locations), find the patch with the correct rotation which best fits in that position and update the affinity function and the connection function according to (41) and (42)

Return: $G = \{V, E_{\text{est}}, \mathbf{W}_{\text{est}}, \mathbf{R}_{\text{est}}\}$

recommend using the following metric:

$$(43) \quad \text{Err}(\{R_i\}_{i=1}^n, \sigma) = \sum_{i=1}^n (\text{MGC}_{\text{lr}}(R_i \cdot P_i, R_{i_{\sigma,r}} \cdot P_{i_{\sigma,r}}) + \text{MGC}_{\text{tb}}(R_i \cdot P_i, R_{i_{\sigma,b}} \cdot P_{i_{\sigma,b}}) + \text{MGC}_{\text{rl}}(R_i \cdot P_i, R_{i_{\sigma,l}} \cdot P_{i_{\sigma,l}}) + \text{MGC}_{\text{bt}}(R_i \cdot P_i, R_{i_{\sigma,b}} \cdot P_{i_{\sigma,b}})),$$

where $i_{\sigma,t}$, $i_{\sigma,l}$, $i_{\sigma,b}$, and $i_{\sigma,r}$ are the indices of the neighbors of patch i from top, left, bottom, and right, respectively, according to the solution σ . If patch i is at the edge or corner of the puzzle grid, we only sum the respective MGC values.

Algorithm 4 Solution of type 2 puzzles.

Input: Puzzle patches: $\{P_i\}_{i=1}^n \subset \mathbb{R}^{s \times s \times 3}$

- Apply [Algorithm 2](#) with $\{P_i\}_{i=1}^n$ to construct the affinity graph $G = (V, E, \mathbf{W}, \mathbf{R})$ and obtain MGC values between all patches
- Run [Algorithm 1](#) with $G = (V, E, \mathbf{W}, \mathbf{R})$ to find the orientations $\{R_i\}_{i=1}^n$
- Apply the type 1 jigsaw puzzle solver of [43] to solve the type 1 puzzle with patches $\{R_i \cdot P_i\}_{i=1}^n$ and obtain their estimated permutation vector σ
- Compute and record $\text{Err}(\{R_i\}_{i=1}^n, \sigma)$ by (43)
- **for** iterations 1:5 **do**
 - Apply [Algorithm 3](#) with σ , $\{R_i\}_{i=1}^n$ and the MGC values to obtain the updated connection graph $G = (V, E, \mathbf{W}, \mathbf{R})$
 - Apply [Algorithm 1](#) with $G = (V, E, \mathbf{W}, \mathbf{R})$ to recover the orientations $\{R_i\}_{i=1}^n$
 - Apply the type 1 jigsaw puzzle solver of [43] to solve the type 1 puzzle with patches $\{R_i \cdot P_i\}_{i=1}^n$ and obtain their estimated permutation vector σ
 - Compute and record $\text{Err}(\{R_i\}_{i=1}^n, \sigma)$ by (43)
- **end for**

Return: $\{R_i\}_{i=1}^n$ and σ , which minimize $\text{Err}(\{R_i\}_{i=1}^n, \sigma)$ among all the above choices

We remark that most state-of-the-art methods use a greedy step to make final corrections to the solved puzzle. On the other hand, the step discussed here only updates the connection

graph and is thus nongreedy. It is possible to incorporate greedy procedures that may improve the performance of our algorithm; however, we would like to show that a more principled method can be competitive.

The above description of our algorithm mentions various parameters, and we further review them in section [SM5](#) of the supplemental material. All parameters, except the number of iterations (which can be rather small), are used in forming the graph affinities. That is, throughout the algorithm, various tests are performed, and based on these, the affinities of the estimated graph are changed. While the relative order of these parameters was often clear to us, we performed some experiments to get the actual effective range of these parameters and noticed some stability to changes of the chosen parameters. We remark that the use of unknown parameters is common in other puzzle solvers and other graph-based algorithms.

5.2. Time complexity of Algorithm 4. Empirically, the most time-consuming step is finding the MGC metric between all puzzle pieces. This step is vital for all jigsaw puzzle solvers. Its order of operations is $O(n^2d)$, where n is the number of image patches and d is the number of pixels in the side of each square image patch. One may parallelize this procedure and achieve faster computation.

After finding the MGC metric between all puzzle pieces, our proposed procedure constructs the connection graph by following [Algorithm 2](#). The main computation of this step is in finding the nearest neighbors of each point in the lattice, which is of order $O(n^2)$. Next, our procedure finds the orientations of all patches. Here, the main computation is in finding the top two eigenvectors of a sparse symmetric matrix with 4 nonzero elements in each column and row. Using the Lanczos algorithm, the time complexity of this step is of order $O(n)$. Once the orientations are computed, the type 2 puzzle problem becomes type 1. The complexity of solving the type 1 puzzle depends on which algorithm is used.

The last step of the proposed procedure is the update step, which is summarized in [Algorithm 3](#). It first requires finding a median of an array of size n , whose computation is of order $O(n \log(n))$. It then requires filling up the empty locations, whose computation is of order $O(\#(\text{empty locations}) \times n)$, where the $\#$ of empty locations is at most $n/2$. This step recurs at most 5 times.

In summary, the complexity of the whole procedure, excluding the chosen type 1 solver, is of order $O(n^2)$. A main computation is that of the MGC metric, where the constant multiplying n^2 seems to be large in practice. As we mention below, in theory, the worst-case complexity of the type 1 solver of Yu, Russell, and Agapito [\[43\]](#) can be higher than $O(n^2)$, but we did not see any evidence for this in our numerical experiments in section [6](#).

At last, we discuss the complexity of the type 2 solvers of Gallagher [\[16\]](#) and Yu, Russell, and Agapito [\[43\]](#). We first mention that they both implement greedy procedures, so it is hard to bound their complexity. Moreover, Yu, Russell, and Agapito [\[43\]](#) use multiple iterations (until convergence) with no guarantees of convergence. Nevertheless, we discuss two of their main algorithms. Gallagher [\[16\]](#) uses Kruskal's algorithm to find the minimum spanning tree, whose complexity is of order $O(n^2 \log(n))$ [\[8\]](#), though another implementation for the minimum spanning tree has complexity of order $O(n^2 \log^*(n))$ [\[14\]](#). Yu, Russell, and Agapito [\[43\]](#) solve a linear program at each iteration, where the worst-case complexity of linear programs is $O(n^{2.5})$ [\[41\]](#). We remark that the same bound on the order of complexity holds for the

type 1 solver, which we use in our algorithm. However, our numerical experiments in section 6 may indicate different orders of complexity for the type 1 and type 2 puzzle solvers of Yu, Russell, and Agapito [43]. One component which makes their type 2 puzzle solver slower is the artificial enlargement of the number of puzzle patches by 4 in the former one, but this only increases the complexity in a constant factor.

6. Numerical experiments. We apply our proposed algorithm to solve square jigsaw puzzles of the following standard image datasets: the MIT dataset from Cho, Avidan, and Freeman [7], which contains 20 images, each with 432 patches, and three datasets from Pomeranz, Shemesh, and Ben-Shahar [31], where the first two, which are referred to as McGill and Pomeranz, include 20 images with 540 and 805 patches, respectively, and the third one has 3 images with 3300 patches, which is also referred to as large Pomeranz. For all datasets, the patches are of size 28×28 . Figure 10 demonstrates the application of our proposed algorithm to four images that represent the four datasets. We remark that our algorithm perfectly solves the last 3 puzzles, and we later discuss the solution of the first puzzle.

To test the accuracy of our proposed algorithm we use the following four metrics, defined in Gallagher [16] and Cho, Avidan, and Freeman [7]: the direct comparison, the neighbors comparison, the largest component, and the perfect reconstruction. The direct comparison measures the percentage of image patches whose location and orientation are correct. The neighbors comparison calculates the percentage of pairs of image patches that are matched correctly. The largest component calculates the percentage of patches in the largest correctly assembled component of the solved puzzle. Finally, the perfect reconstruction of a puzzle is 1 if it is solved correctly and 0 otherwise.

We compared our algorithm with those of Gallagher [16] and Yu, Russell, and Agapito [43] for type 2 puzzles since they were the only algorithms with available codes (we requested codes from all authors of relevant published algorithms). Note that we use the type 1 and type 2 solvers of Yu, Russell, and Agapito [43] in two different ways: We apply the former one as a component of Algorithm 4 and the latter one for comparing performance with Algorithm 4. One of the many procedures in our algorithm is random and is described in section 4.3.3. We have also noticed some randomness in the results of the other two algorithms. Therefore, for each puzzle we run each algorithm 20 times and report the averaged result. To get an idea of the randomness of the three algorithms, we summarize here the averaged standard deviations of the neighbors comparison metric when applying each algorithm 20 times to each of the 20 puzzles in the MIT dataset and averaging over the 20 puzzles. These averaged standard deviations for Gallagher [16], Yu, Russell, and Agapito [43], and our algorithm are 6.5, 1.5, and 0.17, respectively. In this and other experiments, we notice that the randomness of our algorithm is not significant.

Table 1 compares the four metrics of the three algorithms. For the first three metrics of percentages, we report the means and standard deviations among the images in each dataset and among the 20 instances per image. We remark that in the above paragraph, the standard deviations were different as they were computed among 20 instances per image and then averaged among all images in a dataset. We only report the value of the fourth metric, that is, the number of perfectly solved images. One can use this reported metric and the total number of images to compute the mean and standard deviation of perfect reconstruction

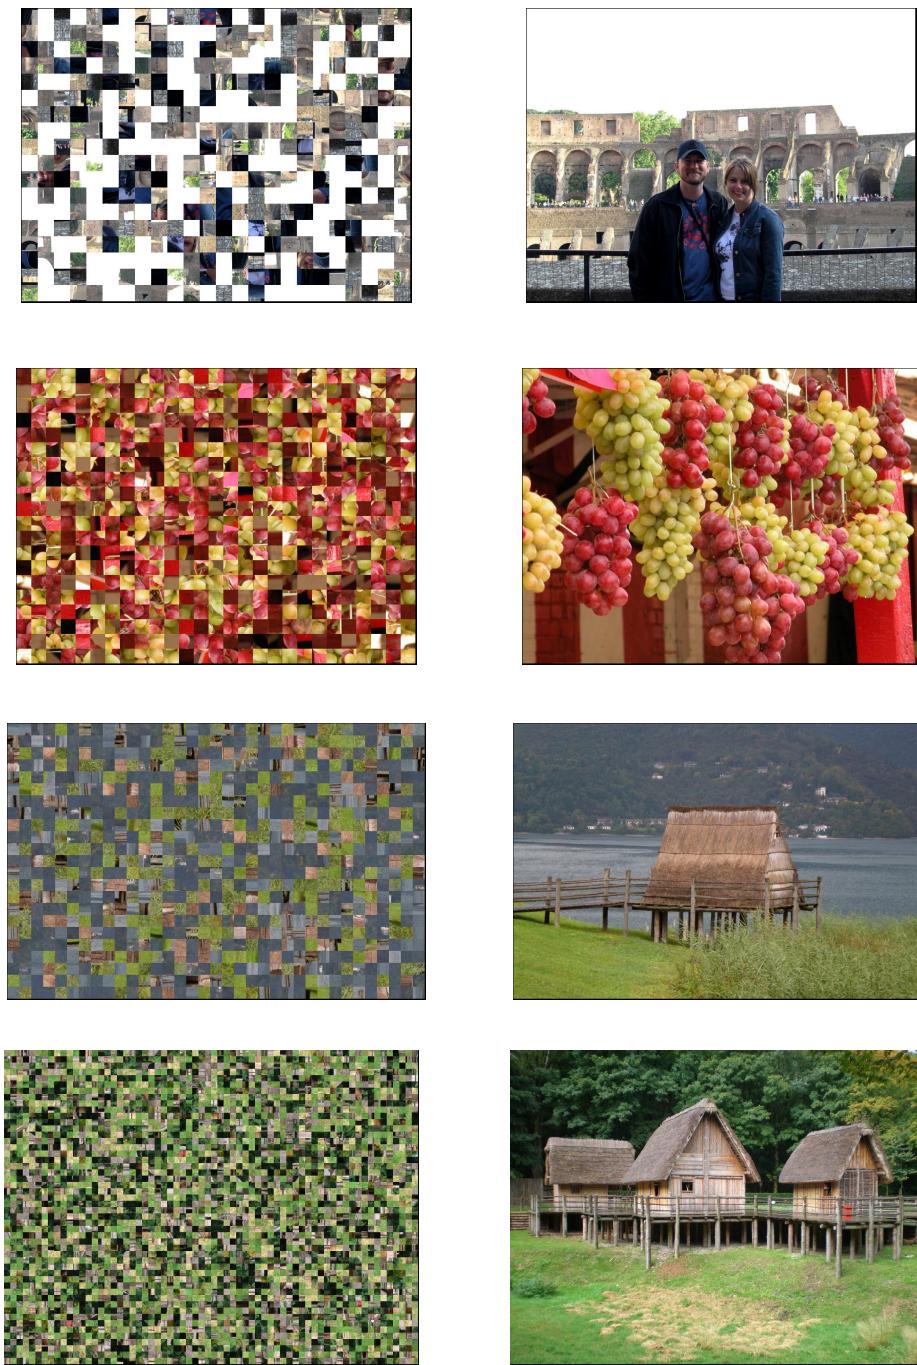


Figure 10. Solutions by *Algorithm 4* of type 2 puzzles representing the four datasets. The images in the left column are the inputs for the algorithm, and the ones in the right column are the outputs generated by our proposed algorithm. The puzzle in first row is from the MIT dataset with 432 patches, the puzzle in the second row is from the McGill dataset with 540 patches, the puzzle in the third row is from the Pomeranz dataset with 805 patches, and the puzzle in the fourth row is from the large Pomeranz dataset with 3300 patches.

Table 1

Comparison of results for type 2 puzzles. For the first three metrics (direct, neighbor, and largest), we report the mean values and standard deviations over all the images in a dataset and over 20 instances of solving each puzzle. For the fourth metric, we report the total number of perfectly solved images (which was identical to all 20 instances per puzzle).

Dataset	Method	Direct		Neighbor		Largest		Perfect
		mean	std	mean	std	mean	std	
MIT (20 images, 432 patches)	Gallagher [16]	84.2	19.7	89.1	12.4	87.2	14.3	9
	Yu et al. [43]	95.5	13.0	95.4	8.7	95.4	13.2	13
	Our method	94.8	11.3	95.2	9.2	95.4	9.1	13
McGill (20 images, 540 patches)	Gallagher [16]	77.2	35.3	85.8	19.8	84.6	21.3	7
	Yu et al. [43]	92.9	24.6	93.5	14.8	93.1	15.4	13
	Our method	88.3	25.6	92.2	15.2	91.4	17.2	13
Pomeranz (20 images 805 patches)	Gallagher [16]	77.5	27.8	85.3	15.5	79.3	22.6	5
	Yu et al. [43]	91.8	14.2	92.7	13.0	91.7	14.2	9
	Our method	86.8	21.4	90.0	14.2	89.3	15.4	9
Large Pomeranz (3 images 3300 patches)	Gallagher [16]	82.9	15.6	84.2	14.2	82.8	15.7	1
	Yu et al. [43]	89.7	12.3	90.2	11.0	89.7	12.3	1
	Our method	86.4	14.0	88.1	11.7	86.4	14.0	1

among the images in a dataset, and we thus do not find it necessary to include it. We also remark that the standard deviations of this metric among the 20 instances were always zero. [Figure 11](#) presents histograms of the neighbors comparison metric for the first three datasets and the three different algorithms. The fourth dataset is excluded from this figure since it only has three images. Histograms for the other metrics, which are not provided here, indicate similar comparisons of the three algorithms.

The results presented in [Table 1](#) and [Figure 11](#) indicate that the accuracy of our algorithm is comparable to the state-of-the-art methods when tested on previously suggested datasets. In general, we noted that when most of the patches have nonzero gradients around their boundaries, our algorithm obtained perfect recovery. On the other hand, we noted that puzzles with low percentages of recovery by all algorithms have large portions of patches with the same uniform color. For example, we demonstrate such a puzzle from the MIT dataset, together with its solution by our algorithm, in the first row of [Figure 10](#). For this image and all three methods, the neighbors comparison metric was 65%, where the errors occurred in the top part of the image of uniform white background. Furthermore, this was the minimal value of this metric among all puzzles in the MIT dataset for all methods. The solutions obtained by the three algorithms are visually identical to the original one. On the other hand, the solutions of any of the three methods to the two images with minimal neighbors comparison metric of either the McGill or the Pomeranz datasets (the corresponding two images are the same for all three methods) do not look visually identical to the original images.

We further test the robustness of the three algorithms to corruption. We fix a corruption rate of value 0.02, 0.04, 0.06, 0.08, 0.10, 0.12, or 0.14. We arbitrarily fix a 28×28 patch and then a side of this patch and change the values of this side (we uniformly sample 28×3 values from all pixel values of the given image and assign them to this side) with probability that equals the corruption rate. Note that if, for example, the corruption rate is 0.14, then a given edge (which may arise from two different patches) is corrupted with probability $1 - 0.86^2 \approx 0.26$.

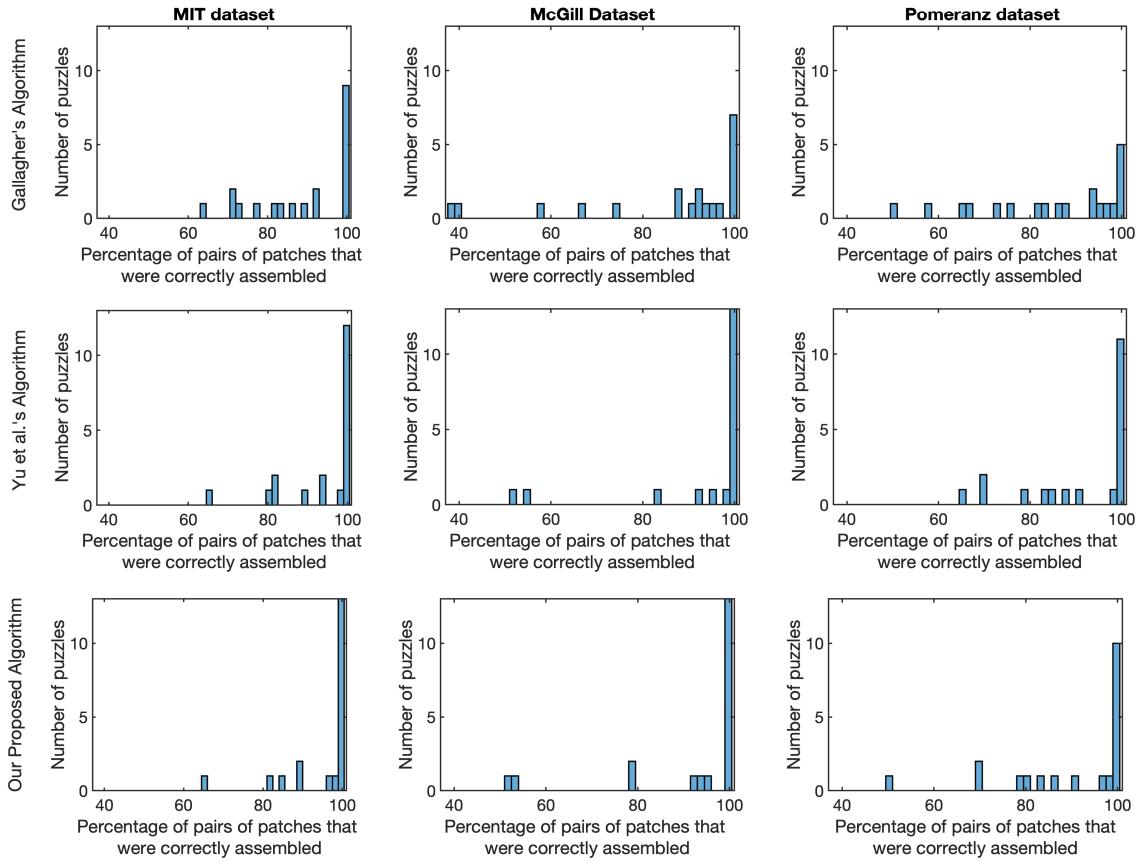


Figure 11. Histograms of the neighbors comparison metric (percentage of the recovered pairs of patches). The three rows correspond to results by the algorithm of [16], the algorithm of [43], and our proposed algorithm, respectively. The three columns correspond to the MIT, McGill, and Pomeranz datasets, respectively.

There are several reasons for this model. First, since gradients are sensitive to noise, the MGC metric is sensitive to noise, and thus application of standard noise models to the whole image will result in similar degradation of performance by all tested methods. Second, this model of corrupting sides of patches is somewhat similar to adversarially corrupting edges in group synchronization [21, 26], where despite really bad corruption, one has some clean information that can help in solving the synchronization problem. At last, we hope that for images without uniform regions, the conditions of [Theorem 3.1](#) may hold under small corruption in this model, but of course we cannot verify this as there is a gap between the clean theory and the applied problem.

[Figure 12](#) demonstrates some examples of puzzles with corrupted patches. [Figure 13](#) shows graphs of the averaged neighbors comparison metric as a function of the corruption rates computed for puzzles of the MIT dataset and the McGill dataset. For each dataset, there are three graphs for the three algorithms: [Algorithm 4](#), Gallagher [16], and Yu, Russell, and Agapito [43]. We note that our algorithm obtains the highest averaged accuracy for all nonzero corruption rates. Nevertheless, the algorithm of Yu, Russell, and Agapito [43] is still

somewhat comparable to our algorithm; however, Gallagher [16] is clearly less accurate than both methods.

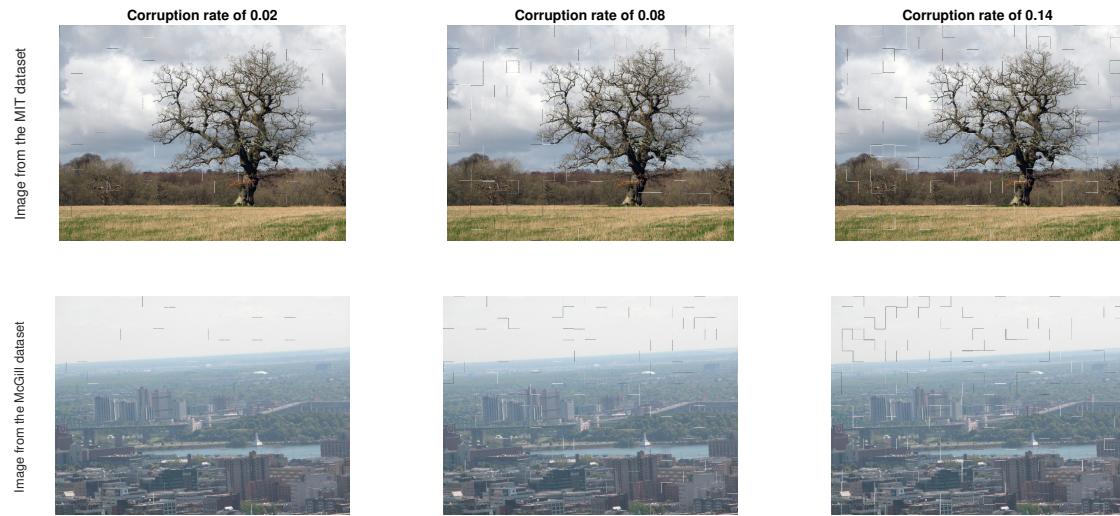


Figure 12. Examples of puzzles with corrupted patch sides. The puzzle in the first row is from the MIT dataset, and the puzzle in the second row is from the McGill dataset. The three columns correspond to the following corruption rates: 0.02, 0.08, and 0.14.

Next, we compare the running times of Gallagher [16], Yu, Russell, and Agapito [43], and our algorithm, where we use a MacBook Pro with a 2.3 GHz Intel Core i5 processor and an 8 GB 2133 MHz LPDDR3 memory. Table 2 reports results for the four different datasets. It includes results for both the faster implementation of our algorithm with a single iteration (and no update) and our default implementation with 5 iterations described in Algorithm 4. It also includes the time of calculating the MGC metric, which is shared by all algorithms and is rather slow. We can see from Table 2 that Gallagher [16] is the fastest algorithm for the smaller puzzles, and our algorithm is the fastest for the largest puzzles. Yu, Russell, and Agapito [43] is the slowest algorithm.

Comparing the ratio of the time of each algorithm over the time of computing the MGC metric, one can notice the following: For our single iteration implementation, these ratios are very similar across different puzzle sizes. Thus, the complexity of this implementation seems to match order $O(n^2)$ on these datasets (as the complexity of computing the MGC metric is of order $O(n^2)$). For our full implementation according to Algorithm 4, these ratios slightly decrease. That is, relatively, fewer updates are needed for larger puzzles. These ratios slightly increase for Gallagher [16], and they significantly increase for Yu, Russell, and Agapito [43]. In view of this observation and the discussion in subsection 5.2, it is possible that for typical puzzles the order of complexity of Yu, Russell, and Agapito [43] is higher than $O(n^2)$.

We remark that for most images, our algorithm obtains competitive accuracy with either one or two iterations. However, there are a couple of images for which more iterations are needed to achieve competitive accuracy.

Finally, we would like to mention that most state-of-the-art algorithms, in particular

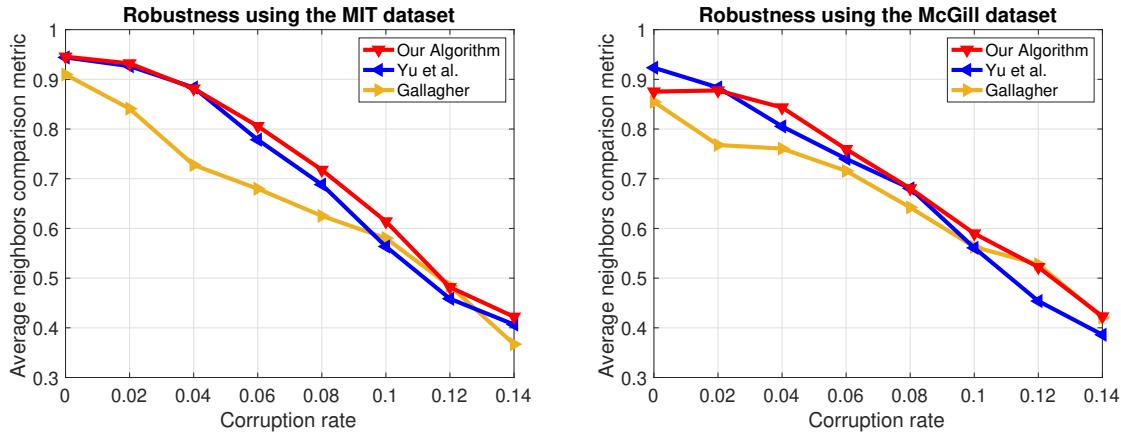


Figure 13. Demonstration of the robustness of *Algorithm 4* in case of corrupted puzzle sides. We compare our proposed algorithm with the algorithms of Gallagher [16] and Yu, Russell, and Agapito [43] when the patches are corrupted. The x axis is the corruption rate, and the y axis is the average neighbors comparison metric for the puzzles under different noise levels. We report the mean values of the neighbors comparison metric for 5 instances and 20 puzzles of the MIT dataset in the first row and the McGill dataset in the second row.

Table 2
Comparison of running times for type 2 puzzles for the four datasets.

Dataset	Method	Time (seconds)
MIT (20 images, 432 patches, 28×28)	MGC metric calculation	32.5
	Gallagher [16]	36.4
	Yu et al. [43]	60.2
	Our method with 1 iteration	38.1
	Our method with 5 iterations	57.6
McGill (20 images, 540 patches, 28×28)	MGC metric calculation	51.3
	Gallagher [16]	58.5
	Yu et al. [43]	100.1
	Our method with 1 iteration	58.5
	Our method with 5 iterations	87.5
Pomeranz (20 images, 805 patches, 28×28)	MGC metric calculation	112.3
	Gallagher [16]	135.5
	Yu et al. [43]	234
	Our method with 1 iteration	128
	Our method with 5 iterations	178
Large Pomeranz (3 images, 3300 patches, 28×28)	MGC metric calculation	1908.5
	Gallagher [16]	3288
	Yu et al. [43]	6120
	Our method with 1 iteration	2214
	Our method with 5 iterations	2857

[16, 38, 43], use a greedy step to make final corrections. We believe that by using that final step of corrections we could further improve our results; however, we would like to avoid greedy procedures.

We refer the reader to <https://github.com/vahanhuroyan/PuzzleDemoGCL> for all the codes necessary to duplicate these results.

7. Conclusion. This paper introduces a novel and constructive mathematical approach for solving square jigsaw puzzles. It first suggests a procedure for recovering the unknown orientations of type 2 puzzle patches and guarantees its robustness to measurement errors assuming a special clean setting. Furthermore, it also suggests a principled strategy for updating the full puzzle solution based on the latter strategy for solving orientations. Some components of the proposed algorithm, in particular the strategy for recovering orientations, are relatively fast. Nevertheless, the main bottleneck in the computational complexity of our algorithm, that is, calculating the MGC metric, is shared by all existing algorithms. Numerical experiments on datasets of square jigsaw puzzles indicate that the accuracy of our algorithm is comparable to that of state-of-the-art methods. Furthermore, on average, our algorithm seems to outperform the existing methods in the presence of corruption. It is also more computationally efficient for large puzzles.

We expect some possible extensions of the proposed algorithm. First, we believe that the ideas pursued in this work could be extended to puzzles that come from more complicated manifolds, such as the two-dimensional sphere or a three-dimensional cube jigsaw puzzle, or puzzles with more complicated shapes of patches, such as tangrams. The GCL algorithm should be the same; however, instead of considering the group \mathbb{Z}_4 , one needs to consider the corresponding rotation group. Two challenges though are defining a good metric between puzzle pieces and constructing the connection graph. By doing this, one will extend the applicability of this work to various real-world applications, such as three-dimensional image reconstruction from two-dimensional images.

In terms of theory, it is interesting to analyze our proposed GCL algorithm with more complicated perturbations. Additional theoretical questions arise from different ideas discussed in the accompanying supplemental material (M129076_01.pdf [[local](#)/[web](#) 965KB]) that we cannot make practical. For example, we are interested to find out if one can effectively utilize the vector diffusion distances or a modification of them. Moreover, we would like to know if one can better estimate the locations of the patches by using a quadratic assignment problem formulation.

Acknowledgments. We would like to thank Rui Yu for kindly sending us his code of the algorithm presented in [43] and Andrew Gallagher for posting the code used in [16]. We are very grateful to the anonymous reviewers for their very careful reading of the manuscript and their valuable comments. We are also thankful to Dr. Brendt Wohlberg for his professional handling of the manuscript.

REFERENCES

- [1] T. ALTMAN, *Solving the JIGSAW puzzle problem in linear time*, Appl. Artificial Intell., 3 (1989), pp. 453–462, <https://doi.org/10.1080/08839518908949937>.
- [2] F. A. ANDALÓ, G. TAUBIN, AND S. GOLDENSTEIN, *Solving image puzzles with a simple quadratic programming formulation*, in Proceedings of the 25th SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI), IEEE, Washington, DC, 2012, pp. 63–70.
- [3] A. S. BANDEIRA, A. SINGER, AND D. A. SPIELMAN, *A Cheeger inequality for the graph connection Laplacian*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1611–1630, <https://doi.org/10.1137/120875338>.
- [4] C. BORDENAVE, U. FEIGE, AND E. MOSEL, *Shotgun Assembly of Random Jigsaw Puzzles*, preprint, <https://arxiv.org/abs/1605.03086>, 2016.

[5] B. J. BROWN, C. TOLER-FRANKLIN, D. NEHAB, M. BURNS, D. P. DOBKIN, A. VLACHOPOULOS, C. DOUMAS, S. RUSINKIEWICZ, AND T. WEYRICH, *A system for high-volume acquisition and matching of fresco fragments: Reassembling Theran wall paintings*, ACM Trans. Graph., 27 (2008), pp. 84:1–84:9.

[6] L. CHEN, D. CAO, AND Y. LIU, *A new intelligent jigsaw puzzle algorithm base on mixed similarity and symbol matrix*, Int. J. Pattern Recognit. Artif. Intell., 32 (2018), 1859001.

[7] T. S. CHO, S. AVIDAN, AND W. T. FREEMAN, *A probabilistic image jigsaw puzzle solver*, in Proceedings of the 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Washington, DC, 2010, pp. 183–190.

[8] T. H. CORMEN, C. E. LEISERSON, R. L. RIVEST, AND C. STEIN, *Introduction to Algorithms*, 3rd ed., MIT Press, Cambridge, MA, 2009, <https://mitpress.mit.edu/books/introduction-algorithms>.

[9] C. DAVIS AND W. M. KAHAN, *The rotation of eigenvectors by a perturbation. III*, SIAM J. Numer. Anal., 7 (1970), pp. 1–46, <https://doi.org/10.1137/0707001>.

[10] A. DEEVER AND A. GALLAGHER, *Semi-automatic assembly of real cross-cut shredded documents*, in Proceedings of the 19th IEEE International Conference on Image Processing (ICIP), IEEE, Washington, DC, 2012, pp. 233–236.

[11] E. D. DEMAINE AND M. L. DEMAINE, *Jigsaw puzzles, edge matching, and polyomino packing: Connections and complexity*, Graphs Combin., 23 (2007), pp. 195–208.

[12] N. EL KAROUI AND H.-T. WU, *Graph connection Laplacian methods can be made robust to noise*, Ann. Statist., 44 (2016), pp. 346–372.

[13] J. FAN, W. WANG, AND Y. ZHONG, *An l_∞ eigenvector perturbation bound and its application*, J. Mach. Learn. Res., 18 (2017), 207, <https://jmlr.org/papers/v18/16-140.html>.

[14] M. L. FREDMAN AND R. E. TARJAN, *Fibonacci heaps and their uses in improved network optimization algorithms*, J. ACM, 34 (1987), pp. 596–615, <https://doi.org/10.1145/28869.28874>.

[15] H. FREEMAN AND L. GARDER, *Apictorial jigsaw puzzles: The computer solution of a problem in pattern recognition*, IEEE Trans. Electron. Comput., 13 (1964), pp. 118–127.

[16] A. C. GALLAGHER, *Jigsaw puzzles with pieces of unknown orientation*, in Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (Providence, RI), IEEE, Washington, DC, 2012, pp. 382–389, <https://doi.org/10.1109/CVPR.2012.6247699>.

[17] P. JACCARD, *Étude comparative de la distribution florale dans une portion des alpes et des jura*, Bulletin de la Société Vaudoise des Sciences Naturelles, 37 (1901), pp. 547–579.

[18] S.-Y. JIN, S. LEE, N. A. AZIS, AND H.-J. CHOI, *Jigsaw puzzle image retrieval via pairwise compatibility measurement*, in Proceedings of the 2014 International Conference on Big Data and Smart Computing (BIGCOMP), IEEE, Washington, DC, 2014, pp. 123–127.

[19] E. JUSTINO, L. S. OLIVEIRA, AND C. FREITAS, *Reconstructing shredded documents through feature matching*, Forensic Sci. Internat., 160 (2006), pp. 140–147.

[20] D. KOLLER AND M. LEVOY, *Computer-aided reconstruction and new matches in the Forma Urbis Romae*, Bullettino Della Commissione Archeologica Comunale di Roma, 15 (2006), pp. 103–125.

[21] G. LERMAN AND Y. SHI, *Robust Group Synchronization via Cycle-Edge Message Passing*, preprint, <https://arxiv.org/abs/1912.11347>, 2019.

[22] H. LIU, S. CAO, AND S. YAN, *Automated assembly of shredded pieces from multiple photos*, IEEE Trans. Multimedia, 13 (2011), pp. 1154–1162.

[23] W. MARANDE AND G. BURGER, *Mitochondrial DNA as a genomic jigsaw puzzle*, Science, 318 (2007), pp. 415–415.

[24] M. A. MARQUES AND C. O. FREITAS, *Reconstructing strip-shredded documents using color as feature matching*, in Proceedings of the 2009 ACM symposium on Applied Computing, ACM, New York, 2009, pp. 893–894.

[25] A. MARTINSSON, *Shotgun Edge Assembly of Random Jigsaw Puzzles*, preprint, <https://arxiv.org/abs/1605.07151>, 2016.

[26] T. MAUNU AND G. LERMAN, *A Provably Robust Multiple Rotation Averaging Scheme for $SO(2)$* , preprint, <https://arxiv.org/abs/2002.05299>, 2020.

[27] D. MONDAL, Y. WANG, AND S. DUROCHER, *Robust solvers for square jigsaw puzzles*, in Proceedings of the 2013 International Conference on Computer and Robot Vision, IEEE, Washington, DC, 2013, pp. 249–256.

[28] E. MOSEL AND N. ROSS, *Shotgun assembly of labeled graphs*, IEEE Trans. Network Sci. Engrg., 6 (2019), pp. 145–157, <https://doi.org/10.1109/TNSE.2017.2776913>.

[29] G. PAIKIN AND A. TAL, *Solving multiple square jigsaw puzzles with missing pieces*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE, Washington, DC, 2015, pp. 4832–4839.

[30] R. PINTUS, K. PAL, Y. YANG, T. WEYRICH, E. GOBBETTI, AND H. RUSHMEIER, *A survey of geometric analysis in cultural heritage*, Comput. Graphics Forum, 35 (2015), pp. 4–31.

[31] D. POMERANZ, M. SHEMESH, AND O. BEN-SHAHAR, *A fully automated greedy square jigsaw puzzle solver*, in Proceedings of the 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011 (Colorado Springs, CO), IEEE, Washington, DC, 2011, pp. 9–16, <https://doi.org/10.1109/CVPR.2011.5995331>.

[32] D. SHOLOMON, O. E. DAVID, AND N. S. NETANYAHU, *A generalized genetic algorithm-based solver for very large jigsaw puzzles of complex types*, in Proceedings of the 28th AAAI Conference on Artificial Intelligence, Québec City, Québec, Canada, 2014, pp. 2839–2845, <http://www.aaai.org/ocs/index.php/AAAI/AAAI14/paper/view/8650>.

[33] D. SHOLOMON, O. E. DAVID, AND N. S. NETANYAHU, *Genetic algorithm-based solver for very large multiple jigsaw puzzles of unknown dimensions and piece orientation*, in Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation, GECCO '14 (New York, NY), 2014, ACM, New York, pp. 1191–1198.

[34] D. SHOLOMON, O. E. DAVID, AND N. S. NETANYAHU, *An automatic solver for very large jigsaw puzzles using genetic algorithms*, Genetic Program. Evolv. Mach., 17 (2016), pp. 291–313, <https://doi.org/10.1007/s10710-015-9258-0>.

[35] D. SHOLOMON, O. E. DAVID, AND N. S. NETANYAHU, *DNN-buddies: A deep neural network-based estimation metric for the jigsaw puzzle problem*, in Artificial Neural Networks and Machine Learning—ICANN 2016, A. E. Villa, P. Masulli, and A. J. Pons Rivero, eds., Springer, New York, 2016, pp. 170–178.

[36] A. SINGER AND H.-T. WU, *Vector diffusion maps and the connection Laplacian*, Comm. Pure Appl. Math., 65 (2012), pp. 1067–1144.

[37] E. SIZIKOVA AND T. FUNKHOUSER, *Wall painting reconstruction using a genetic algorithm*, J. Comput. Cultural Heritage (JOCCH), 11 (2017), p. 3.

[38] K. SON, J. HAYS, AND D. B. COOPER, *Solving square jigsaw puzzles with loop constraints*, in European Conference on Computer Vision, Springer, New York, 2014, pp. 32–46.

[39] K. SON, D. MORENO, J. HAYS, AND D. B. COOPER, *Solving small-piece jigsaw puzzles by growing consensus*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE, Washington, DC, 2016, pp. 1193–1201.

[40] C. TOLER-FRANKLIN, B. BROWN, T. WEYRICH, T. FUNKHOUSER, AND S. RUSINKIEWICZ, *Multi-feature matching of fresco fragments*, ACM Trans. Graphics (Proc. SIGGRAPH Asia), 29 (2010), 185.

[41] P. M. VAIDYA, *Speeding-up linear programming using fast matrix multiplication (extended abstract)*, in Proceedings of the 30th Annual Symposium on Foundations of Computer Science (Research Triangle Park, NC), IEEE, Washington, DC, 1989, pp. 332–337, <https://doi.org/10.1109/SFCS.1989.63499>.

[42] X. YANG, N. ADLURU, AND L. J. LATECKI, *Particle filter with state permutations for solving image jigsaw puzzles*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Washington, DC, 2011, pp. 2873–2880.

[43] R. YU, C. RUSSELL, AND L. AGAPITO, *Solving jigsaw puzzles with linear programming*, in Proceedings of the 2016 British Machine Vision Conference, BMVC 2016, York, UK, 2016, 139.

[44] Y. YU, T. WANG, AND R. J. SAMWORTH, *A useful variant of the Davis–Kahan theorem for statisticians*, Biometrika, 102 (2015), pp. 315–323.

[45] Z. YU-XIANG, S. MU-CHUN, C. ZHONG-LIE, AND L. JONATHAN, *A puzzle solver and its application in speech descrambling*, in Proceedings of the 2007 WSEAS International Conference on Computer Engineering and Applications, Gold Coast, Australia, 2007, pp. 171–176.