

Symmetric assembly puzzles are hard, beyond a few pieces

Citation for published version (APA):

Demaine, E. D., Korman, M., Ku, J. S., Mitchell, J. S. B., Otachi, Y., van Renssen, A., Roeloffzen, M., Uehara, R., & Uno, Y. (2020). Symmetric assembly puzzles are hard, beyond a few pieces. *Computational Geometry: Theory and Applications*, 90, [101648]. <https://doi.org/10.1016/j.comgeo.2020.101648>

Document license:

TAVERNE

DOI:

[10.1016/j.comgeo.2020.101648](https://doi.org/10.1016/j.comgeo.2020.101648)

Document status and date:

Published: 01/10/2020

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



Symmetric assembly puzzles are hard, beyond a few pieces

Erik D. Demaine^a, Matias Korman^b, Jason S. Ku^a, Joseph S.B. Mitchell^c,
Yota Otachi^d, André van Renssen^{e,*}, Marcel Roeloffzen^f, Ryuhei Uehara^g,
Yushi Uno^h

^a MIT, United States of America

^b Tufts University, United States of America

^c Stony Brook University, United States of America

^d Nagoya University, Japan

^e University of Sydney, Australia

^f TU Eindhoven, Netherlands

^g JAIST, Japan

^h Osaka Prefecture University, Japan



ARTICLE INFO

Article history:

Available online 9 March 2020

Keywords:

Assembly puzzle

NP-complete

Parameterized algorithms

ABSTRACT

We study the complexity of symmetric assembly puzzles: given a collection of simple polygons, can we translate, rotate, and possibly flip them so that their interior-disjoint union is line symmetric? On the negative side, we show that the problem is strongly NP-complete even if the pieces are all polyominoes. On the positive side, we show that the problem can be solved in polynomial time if the number of pieces is a fixed constant.

© 2020 Published by Elsevier B.V.

1. Introduction

The goal of a 2D *assembly puzzle* is to arrange a given set of pieces so that they do not overlap and form a target silhouette. The most famous example is the Tangram puzzle, shown in Fig. 1. Its earliest printed reference is from 1813 in China, but by whom or exactly when it was invented remains a mystery [1]. There are over 2,000 Tangram assembly puzzles [1], and many more similar 2D assembly puzzles [2]. A recent trend in the puzzle world is a relatively new type of 2D assembly puzzle that we call *symmetric assembly puzzles*. In these puzzles the target shape is not specified. Instead, the objective is to arrange the pieces so that they form a symmetric silhouette without overlap.

The first symmetric assembly puzzle, “Symmetrix”, was designed in 2003 by Japanese puzzle designer Tadao Kitazawa and was distributed by Naoyuki Iwase as his exchange puzzle at the 2004 International Puzzle Party (IPP) in Tokyo [4]. The lack of a specified target shape makes these puzzles quite difficult to solve. In this paper, we aim for arrangements that are line symmetric (reflection through a line), but other symmetries such as rotational symmetry could also be considered. We also assume that the given pieces are simple polygons, and that the line-symmetric assembly must be a simple polygon (have no holes).

* Corresponding author.

E-mail addresses: edemaine@mit.edu (E.D. Demaine), matias.korman@tufts.edu (M. Korman), jasonku@mit.edu (J.S. Ku), joseph.mitchell@stonybrook.edu (J.S.B. Mitchell), otachi@nagoya-u.jp (Y. Otachi), andre.vanrennsen@sydney.edu.au (A. van Renssen), m.j.m.roeloffzen@tue.nl (M. Roeloffzen), uehara@jaist.ac.jp (R. Uehara), uno@cs.osakafu-u.ac.jp (Y. Uno).

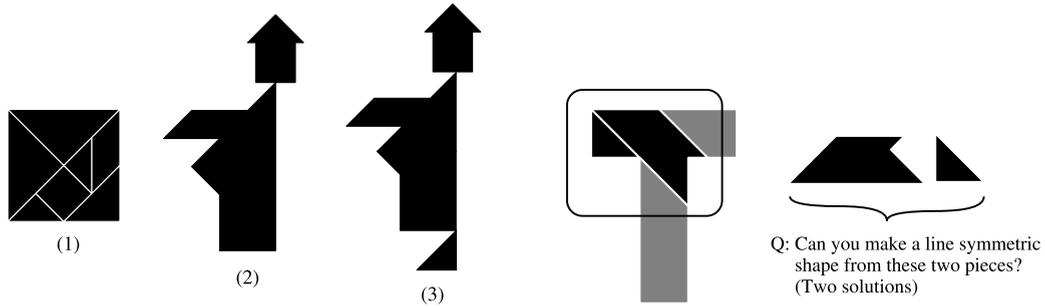


Fig. 1. [Left] The seven Tangram pieces (1) can be assembled into non-simple silhouettes (2) and (3). [Right] A symmetric assembly puzzle invented by Hiroshi Yamamoto [3]: given the two black pieces (right) from the classic T puzzle (left), make two different line symmetric shapes. (Used with permission.)

We study the computational complexity of this general form of symmetric assembly puzzle. Precisely, we define a *symmetric assembly puzzle*, or SAP, to be a set of k disjoint simple polygons $\mathcal{P} = \{P_1, P_2, \dots, P_k\}$, with $n = |P_1| + \dots + |P_k|$ the total number of vertices in all pieces. By *simple polygon* we mean a closed subset of \mathbb{R}^2 homeomorphic to a disk bounded by a closed path of a finite number of straight line segments where nonadjacent edges and vertices do not intersect. A *symmetric assembly* $f : \{p \in \mathcal{P} \mid p \in \mathcal{P}\} \rightarrow \mathbb{R}^2$, of a SAP \mathcal{P} is a planar isometric embedding of the pieces ($\{f(p) \mid p \in \mathcal{P}\}$ for each $P \in \mathcal{P}$ is a rigid transformation of P), such that their mapped interiors are disjoint and their mapped union forms a simple polygon that is line symmetric. We abuse notation slightly by using $f(P)$ to denote $\{f(p) \mid p \in P\}$ and $f(\mathcal{P})$ to denote $\{f(p) \mid p \in \mathcal{P}\}$. We refer to SAP (symmetric assembly puzzles) as the problem of deciding whether an instance \mathcal{P} has a symmetric assembly f , and we study the computational complexity of SAP. We allow pieces to flip over (reflect), but other variants of the puzzle may disallow this. Given that humans have difficulty solving SAPs with even a few low-complexity pieces, we consider two different generalizations: bounded piece complexity ($|P_i| = O(1)$) and bounded piece number ($k = O(1)$). In the former case, we prove strong NP-completeness, while in the latter case, we solve the problem in polynomial time (the exponent is linear in k).

2. Many pieces

First, we show that it is hard to solve symmetric assembly puzzles with a large number of pieces, even if each piece has bounded complexity ($|P_i| = O(1)$).

Theorem 1. *Symmetric assembly puzzles are strongly NP-complete even if each piece is a polyomino with at most six vertices with area upper-bounded by a polynomial function of the number of pieces.*

Proof. If a SAP has a solution, the location and orientation of each piece within a symmetric assembly is a solution certificate of polynomial size checkable in polynomial time, so symmetric assembly puzzles are in NP. We reduce from the RECTANGLE PACKING PUZZLE problem (in short the RPP problem), known to be strongly NP-hard [5]. Specifically, it is (strongly) NP-complete to decide whether k given rectangular pieces—sized $1 \times x_1, 1 \times x_2, \dots, 1 \times x_k$, where the x_i 's are positive integers bounded above by a polynomial in k —can be exactly packed into a specified rectangular box with given width w and height h and area $x_1 + x_2 + \dots + x_k = wh$.

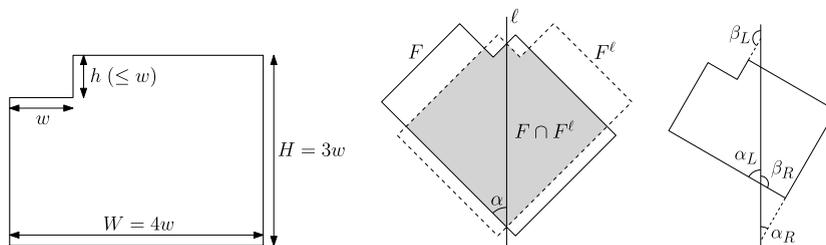


Fig. 2. [Left] The frame piece F . [Middle] If ℓ and ℓ_B form an angle of $\pi/4$, then $F \cap F^\ell$ is contained in a rectangle in an $H \times H$ and thus O^* cannot be line symmetric. [Right] The angles $\alpha_L, \beta_L, \alpha_R$, and β_R .

Let $I = (x_1, \dots, x_k, w, h)$ be a rectangle packing puzzle. Without loss of generality, we assume that $w \geq h$. Now let $I' = (P_1, \dots, P_k, F)$ be the SAP where P_i is the $1 \times x_i$ rectangle for each $i \in \{1, \dots, k\}$, and F is the polyomino in Fig. 2. We call F the *frame piece* of I' . We show that I has a rectangle packing if and only if I' has a symmetric assembly.

Clearly, if I has a rectangle packing, then the pieces P_1, \dots, P_k can be packed into the $w \times h$ hole in the frame piece creating a line symmetric $W \times H$ rectangle, solving the SAP. Now we show the reverse implication. Assume that I' has a

symmetric assembly, and let O^* be a line symmetric polygon formed by the pieces $\{P_1, \dots, P_k, F\}$. We claim that O^* must be a $W \times H$ rectangle, which will imply that I is a yes-instance of RPP. Fix a placement of the pieces of I' that forms O^* , and let ℓ be one of its lines of symmetry. Assume, without loss of generality, that ℓ is a vertical line. Let F^ℓ be the reflection of F about ℓ .

Observation 1. $\text{area}(F \cap F^\ell) \geq WH - 2wh \geq 10w^2$.

Proof. Because O^* contains F^ℓ and F , it holds that $\text{area}(F^\ell \setminus F) \leq \text{area}(O^* \setminus F) = wh$. Because $F \cup F^\ell$ is mirror-symmetric, $\text{area}(F^\ell \setminus F) = \text{area}(F \setminus F^\ell)$. Hence, it follows that $\text{area}(F \cap F^\ell) = \text{area}(F) - \text{area}(F \setminus F^\ell) \geq WH - 2wh \geq 10w^2$. \square

Observation 1 implies that ℓ passes through an interior point of F . Let ℓ_B be the line containing the segment of F with length $4w$. Let c be the center of the frame piece's bounding box.

Lemma 2. ℓ_B is either parallel to ℓ or orthogonal to ℓ .

Proof. Suppose, for contradiction, that ℓ_B is neither parallel nor orthogonal to ℓ . Let α be the smaller angle made by ℓ_B and ℓ . We partition the edges of F crossed by ℓ into two at their intersection points. Let F_L and F_R be the sets of segments on the left and right portions of F , respectively. Consider the set of counter-clockwise angles between ℓ and the lines containing segments of F_L . The assumptions that ℓ_B and ℓ are neither parallel nor orthogonal, and that F is a polyomino together imply that the set contains exactly two angles α_L and β_L , where $\alpha_L \leq \beta_L$ and $\alpha_L + \pi/2 = \beta_L$. Similarly, let α_R and β_R be the clockwise angles between ℓ and the lines containing segments of F_R , where $\alpha_R \leq \beta_R$ and $\alpha_R + \pi/2 = \beta_R$. Because $\alpha_L + \beta_R = \pi$, it holds that $\alpha_L + \alpha_R = \pi/2$. Note that $\alpha \in \{\alpha_L, \alpha_R\}$.

Two distinct pieces of I' are *connected* if the fixed placement of the two pieces to form O^* have a non-degenerate line segment on their edges in common. Let \mathcal{P} be the subset of $\{P_1, \dots, P_k, F\}$ such that each $P_i \in \mathcal{P}$ can be reached from F by repeatedly following connected pieces in O^* .

As before, consider the angles formed by ℓ and the lines containing segments in the left and right parts of \mathcal{P} . Because all pieces are polyominoes, these lines cannot make angles other than $\alpha_L, \beta_L, \alpha_R,$ and β_R with ℓ . Further note that the subset O' of O^* covered by \mathcal{P} must be mirror-symmetric with respect to ℓ , or else O^* would not be. This implies that $\alpha_L = \alpha_R$. Because $\alpha_L + \alpha_R = \pi/2$, the only solution in which ℓ is not parallel or orthogonal to ℓ_B is when $\alpha_L = \alpha_R = \pi/4$ and $\alpha = \pi/4$. However, if $\alpha = \pi/4$, then $F \cap F^\ell$ is a subset of an $H \times H$ rectangle (see Fig. 2), whose area is at most $H^2 = 9w^2$, contradicting Observation 1. \square

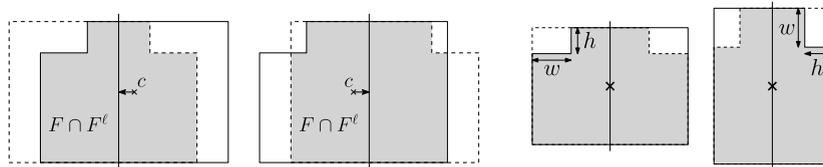


Fig. 3. [Left] When ℓ passes to the left of c , the portion of F to the left of ℓ is too small. If it passes to the right, the right portion would be too small. [Right] If ℓ passes through c , and is either orthogonal or parallel to ℓ_B , the symmetric assembly puzzle can only be completed into a rectangle.

So ℓ is either parallel or orthogonal to ℓ_B . Further, it passes through c (see Fig. 3). In either case, $F \cup F^\ell$ is a $W \times H$ rectangle, and thus $O^* = F \cup F^\ell$. This implies that $O^* \setminus F$ is a $w \times h$ rectangle that must contain the remaining pieces of I' . In particular, we have that this placement packing of P_1, \dots, P_k gives a solution to the instance I of RPP, completing the proof of Theorem 1. \square

We extend the above proof to show that the problem remains strongly NP-complete even when each piece is a convex quadrilateral.

Theorem 3. Symmetric assembly puzzles are strongly NP-complete even if each piece is a convex quadrilateral and area upper bounded by a polynomial function of the number of pieces.

Proof. We note that the only piece that is not a convex quadrilateral is the frame piece F . Hence, we split this into two convex quadrilateral pieces as shown in Fig. 4. We note that due to the dimensions of H and W , the four angles $\alpha, \beta, \gamma,$ and δ are all unique. Furthermore, only $\alpha + \delta$ and $\beta + \gamma$ do not sum up to multiples of $\pi/2$. If we show that any line symmetric solution aligns these four angles as in Fig. 4, Theorem 1 completes the proof.

Assume the angles are not matched as in Fig. 4. We first show that extending γ or δ by a multiple of $\pi/2$ is not useful. We focus on γ , but the same argument holds for δ . If we extend γ using a right angle of the other frame piece, it

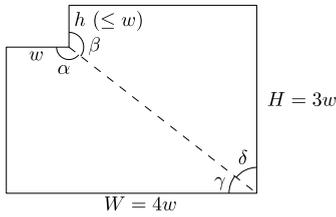


Fig. 4. Splitting the frame piece into two convex quadrilateral pieces.

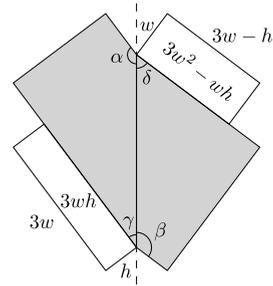


Fig. 5. Matching α to δ and β to γ .

creates an imbalance resulting from the implied line of symmetry cannot be overcome using only the remaining rectangles of combined area wh (see Fig. 6). Extending γ using the rectangles also does not lead to a line symmetric polygon, because placing the other frame piece afterwards still leads to an imbalanced shape.

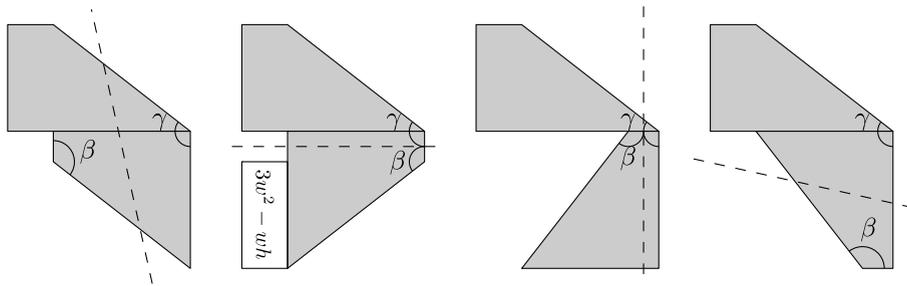


Fig. 6. The four cases when extending γ using the other frame piece.

Because the four angles are all unique and the symmetry line can pass through at most two corners of a simple polygon, at least two of these angles have to meet in a point. If α is matched to δ or if β is matched to γ , we note that the created angle is not a multiple of $\pi/2$ and thus we still have three unique angles. This implies that in this case, both α is matched to δ and β is matched to γ (see Fig. 5).

We first show that the difference between $\alpha + \delta$ and $\beta + \gamma$ cannot be a multiple of $\pi/2$, which implies that the line of symmetry still needs to pass through both of these angles. We prove this by contradiction, so assume that the difference is a multiple of $\pi/2$. We observe from Fig. 4 that $\alpha + \delta + \beta + \gamma = 2\pi$, $\beta = \gamma + \pi/2$, $\alpha = \delta + \pi/2$, and $\gamma < \delta$. Hence, we need to consider only the case where $\alpha + \delta = \beta + \gamma + \pi/2$, which implies that $\delta = \gamma + \pi/4$. Because $\delta + \gamma = \pi/2$, it follows that $\gamma = \pi/8$ and that $\tan \gamma = \sqrt{2} - 1$. However, from Fig. 4 we also observe that $\tan \gamma = (H - h)/3w$, where $2w \leq H - h < 3w$, which implies that $\tan \gamma \geq 2/3$, contradicting that γ is $\pi/8$. Thus, the difference between $\alpha + \delta$ and $\beta + \gamma$ also cannot be a multiple of $\pi/2$.

Because neither $\alpha + \delta$ nor $\beta + \gamma$ is a multiple of $\pi/2$ and the difference between $\alpha + \delta$ and $\beta + \gamma$ also cannot be a multiple of $\pi/2$, the only way to construct a line symmetric solution is for the symmetry line to pass through both created angles. However, this implies that in order to make a line symmetric shape, we need to at least add one region of area $3w^2 - wh$ and one of area $3wh$. Hence, the total area required is at least $3w^2 + 2wh$, which is more than the wh combined area of the rectangles. Therefore, the four angles have to be aligned as in Fig. 4. \square

This result raises the question of what the simplest shape is for which the problem is strongly NP-complete. We conjecture that the problem is still strongly NP-complete even when each piece is a right triangle.

Conjecture 4. Symmetric assembly puzzles are strongly NP-complete even if each piece is a right triangle with area upper-bounded by a polynomial function of the number of pieces.

While we do not have a proof of this conjecture, we do sketch an approach to a possible proof based on a reduction from the 3-PARTITION problem: It is (strongly) NP-complete to decide whether a given set of $3k$ positive integers (each integer is bounded from above by a polynomial in k) can be partitioned into k triples, such the sum of the integers in each triple is the same.

Let $\{a_1, \dots, a_{3k}\}$ be the given set of integers in increasing order. We first transform these integers into almost squares of size $1 \times 1 + \epsilon_i$, such that the $1 + \epsilon_i$ sides of each triple sum to the same length: When we want to ensure that ϵ_i is at most $1/1000$ for each square, we transform each a_i into an almost square of size $1 \times 1 + \frac{a_i}{1000a_{3k}}$. Note that this does not change triples nor the solvability of the 3-PARTITION instance.

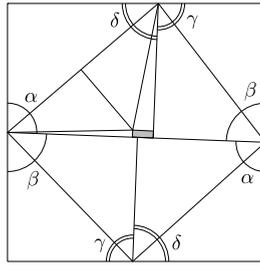


Fig. 7. The frame and its splitting lines. The hole for the 3-PARTITION instance is shown in gray.

Next, we create a big square frame that has a hole of size $k \times 3 + \frac{\sum_{i=1}^{3k} a_i/k}{1000d_{3k}}$. Note that the area of this hole is equal to the total area of the almost squares. We split the frame into right triangles as shown in Fig. 7, while ensuring that any combination of non-right angles is unique.

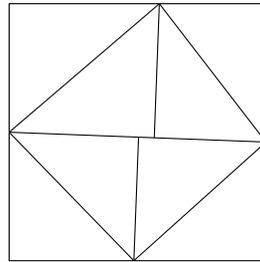


Fig. 8. Splitting an almost square.

Finally, we split the $3k$ almost squares into $24k$ right triangles. The general idea behind the splits is the same as for the frame: for each almost square, we pick four points close to the middle of its sides and split the square as shown in Fig. 8. More precisely, when s is the length of a side, we pick a point $p \in \{\frac{s}{2} + \frac{is}{2k^{20}} : i \in \{1, 2, \dots, k^{19}\}\}$ and split along the line connecting the two points on the vertical sides, along the lines from the points on the horizontal sides perpendicular to the previous splitting line, and along the lines defined by points on consecutive sides (see Fig. 8). Note that p is at most $s/2k$ away from the middle of the side. Again, we require that any combination of non-right angles is unique.

This uniqueness of angles should ensure that the triangles can only be combined to the desired frame and almost squares. Proving this formally, however, turns out to be rather intricate; thus, we leave the full proof of the above conjecture for future research.

3. Constant number of pieces

Next we analyze symmetric assembly puzzles with a constant number of pieces but many vertices, and show they can be solved in polynomial time.

Theorem 5. *Given a symmetric assembly puzzle with a constant number of pieces k containing at most n vertices in total, deciding whether it has a symmetric assembly can be decided in time that is polynomial in n .*

To prove this theorem, we present a brute force algorithm for solving a SAP that runs in polynomial time for constant k . We say two pieces in a symmetric assembly are *connected* to each other if their intersection in the symmetric assembly contains a non-degenerate line segment, and let the *connection* between two connected pieces be their intersection not including isolated points. We will call two pieces *fully* connected if their connection is exactly an edge of one of the pieces, and *partially* connected otherwise; note that two pieces may be partially connected along more than one edge. Call a piece a *leaf* if it connects to at most one piece, and a *branch* otherwise. Given a leaf, let its *parent* be the piece connected to it (if it exists), and let its *siblings* be all other pieces connected to its parent. An illustration demonstrating these terms can be found in Fig. 9.

In addition, we will need to construct simple polygons from provided simple polygons by laying them next to each other along an edge. Let E_P denote the set of directed edges (p_i, p_j) from a vertex p_i to an adjacent vertex p_j of some simple polygon P .

Given an edge $e \in E_P$, we denote its length by $\lambda(e)$. Let $e_P = (p_1, p_2)$ be a directed edge of a polygon P , let $e_Q = (q_1, q_2)$ be a directed edge of a polygon Q , and let d be a (non-negative) length strictly less than $\lambda(e_P)$. Orient P and Q such that e_P exists in a clockwise traversal of P , e_Q exists in a counter clockwise traversal of Q , e_Q is collinear and in the

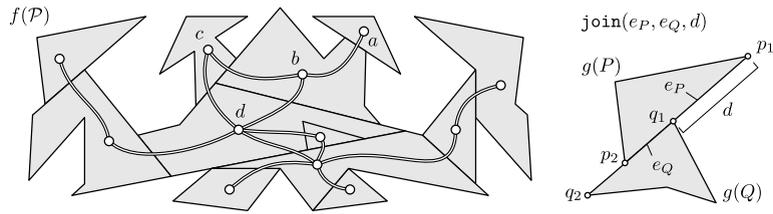


Fig. 9. [Left] Example symmetric assembly \mathcal{P} showing its connection graph. Pieces a and d are fully connected to piece b , and c is partially connected to b . Pieces b , c , and d are branches. Piece a is a leaf, with b its parent and c and d the siblings of a . [Right] Visualization of a JOIN operation.

same direction as e_P , and the distance between p_1 and q_1 is distance d . Call these transformations the mapping $g : \{p \in P \cup Q \mid P, Q \in \mathcal{P}\} \rightarrow \mathbb{R}^2$, where $\{g(p) \mid p \in P \cup Q\}$ for each $P, Q \in \mathcal{P}$ is a rigid transformation of P and Q . Then we define $\text{JOIN}(e_P, e_Q, d)$ to be $g(P) \cup g(Q)$ when $g(P) \cup g(Q)$ is a simple polygon, and otherwise the empty set. See Fig. 9.

If a SAP has a symmetric assembly, let its *connection graph* be a graph on the pieces with an edge connecting two pieces if they are connected in the symmetric assembly. Because a symmetric assembly is a simple polygon by definition, its connection graph is connected and has a spanning tree; we can then construct the assembly using a concatenation of JOIN procedures in breadth-first-search order from an arbitrary root. Because parameter d is not discrete, the total solution space of simple polygons that are constructible from the pieces of a SAP may be uncountable. However, we will exploit the structure of symmetric assemblies to search only a finite set of configurations.

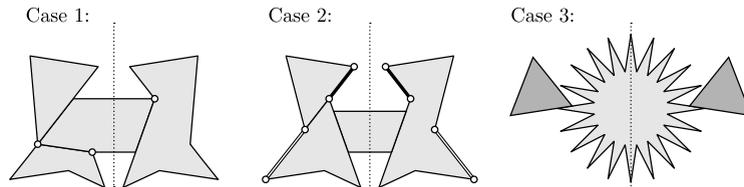


Fig. 10. Examples of symmetric assemblies belonging to each case. Case 1 highlights vertices of connected pieces that intersect. Case 2 highlights join operations using lengths of piece edges. Case 3 is constructed from one symmetric piece and a pair of congruent pieces.

In order to enumerate possible configurations, we would like to distinguish between three cases of the puzzle (see Fig. 10), specifically:

- Case 1: the puzzle has a symmetric assembly in which two connected pieces share a vertex on their connection;
- Case 2: the puzzle has no symmetric assembly satisfying Case 1, but has one in which the distance between vertices from the connecting edges between two connected pieces has the same length as another whole edge in the piece set (we say the connection between the two pieces *constructs* the length of another edge);
- Case 3: the puzzle has no symmetric assembly satisfying Case 1 or Case 2, but has one in which a nonempty set of pieces are themselves line symmetric about the line of symmetry of the symmetric assembly, and any remaining pieces are pairs of congruent pieces symmetric about the line of symmetry.

The following lemma ensures that these three cases are exhaustive.

Lemma 6. *If a SAP has a symmetric assembly, it can be described by one of the above three cases.*

To prove this lemma, we will use the following auxiliary results.

Lemma 7. *If a SAP has a symmetric assembly that is not Case 1, the connection graph of the symmetric assembly is a tree and all connections are single line segments.*

Proof. Let \mathcal{P} be a SAP with symmetric assembly $f : \{p \in P \mid P \in \mathcal{P}\} \rightarrow \mathbb{R}^2$, such that $\{f(p) \mid p \in P\}$ for $P \in \mathcal{P}$ is a rigid transformation of P , that is not Case 1. Suppose, for contradiction, that the connection graph of $f(\mathcal{P})$ is not a tree, so that there exists a cycle C in the connection graph. Let S be a simple closed curve embedded in $f(\mathcal{P})$ that traverses the piece connections from C . The region R bounded by S is completely covered by $f(\mathcal{P})$, or else it would contain a hole, contradicting that $f(\mathcal{P})$ is a simple polygon. Because $f(\mathcal{P})$ covers R and S corresponds to a cycle in the connection graph, then R contains a vertex v in R of some piece P . But then P must share vertex v along its connection with another piece, contradicting exclusion from Case 1.

So the connection graph of $f(\mathcal{P})$ is a tree. Now suppose for contradiction there exists two connected pieces P and Q whose connection is more than one line segment. Then there exists a closed curve embedded in $f(\mathcal{P})$ that crosses from P

to Q along two distinct edges. Then the region R bounded by the cycle must contain a vertex v of P . If $f(\mathcal{P})$ covers R , P must share vertex v along its connection with a vertex of Q , contradicting exclusion from Case 1. Otherwise, $f(\mathcal{P})$ does not cover R , contradicting that $f(\mathcal{P})$ is a simple polygon. \square

Lemma 8. *If a SAP has a symmetric assembly that is not Case 1 or Case 2, the reflection of any partially connected leaf is exactly another piece congruent to the leaf.*

Proof. Let \mathcal{P} be a SAP with a symmetric assembly $f : \{p \in P \mid P \in \mathcal{P}\} \rightarrow \mathbb{R}^2$, such that $\{f(p) \mid p \in P\}$ for $P \in \mathcal{P}$ is a rigid transformation of P , that is not Case 1 or Case 2. Let $s : f(\mathcal{P}) \rightarrow f(\mathcal{P})$ be an automorphism reflecting $f(\mathcal{P})$ across a line of symmetry L , and let $\mu = s \circ f$, mapping each point $p \in P_i$ of a piece $P_i \in \mathcal{P}$ to the corresponding point in the reflection of $f(P_i)$ across L .

Consider a partially connected leaf P whose parent is Q with edge e_P connected to edge e_Q , and suppose for contradiction that $\mu(P)$ is not exactly covered by another piece congruent to P . We first show that a single piece P' contains $\mu(e_P)$ under f so that $f(P') \subset \mu(P)$, and then show that in fact $f(P') = \mu(P)$.

By Lemma 7 the partial connection is a single line segment, and $\ell_P = f(e_P) \setminus f(e_Q)$ is non-empty. $s(\ell_P)$ cannot be covered by more than one piece or else two pieces would share a vertex along their connection contradicting exclusion from Case 1. Also $s(\ell_P)$ cannot be exactly the edge of another piece or else the connection between P and Q would construct its length, contradicting exclusion from Case 2. Thus, $s(\ell_P)$ is a strict subset of an edge $e_{P'}$ from some piece P' under f . Further, $\mu(e_P) = f(e_{P'})$. Suppose for contradiction it did not, and an endpoint p of $e_{P'}$ maps to a point interior to $f(e_P) \cap f(e_Q)$. Then $f(p)$ is also an interior point of $f(\mathcal{P})$, so $\mu(p)$ is also an interior point, and $f(e_{P'})$ would share a vertex along its connection with another piece contradicting exclusion from Case 1. So $\mu(e_P)$ is exactly edge $e_{P'}$ of $f(P')$. And because P is a leaf, $f(P') \subseteq \mu(P)$.

Now we show that in fact $f(P') = \mu(P)$. Suppose for contradiction that $f(P')$ is a strict subset of $\mu(P)$, meaning that some other piece is also fully contained in $\mu(P)$. Let Q' be the first such piece connecting to P' in a clockwise traversal of P' from $e_{P'}$. Then the connection between Q' and P' must either construct the length of some edge from P under f , contradicting exclusion from Case 2, or Q' and P' must share a vertex on their connection, contradicting exclusion from Case 1. So, P' is a piece congruent to P . \square

We now use these intermediate results to prove Lemma 6.

Proof of Lemma 6. Suppose for contradiction there exists a SAP \mathcal{P} having a symmetric assembly $f : \{p \in P \mid P \in \mathcal{P}\} \rightarrow \mathbb{R}^2$, such that $\{f(p) \mid p \in P\}$ for $P \in \mathcal{P}$ is a rigid transformation of P , that does not satisfy any of the above cases, and assume \mathcal{P} has the fewest pieces among all such SAPs. We will identify a symmetric leaf or a congruent pair of leaves that can be removed from \mathcal{P} to form a SAP with fewer pieces. The new SAP must have the same classification as the original, contradicting the minimality of \mathcal{P} .

Let $s : f(\mathcal{P}) \rightarrow f(\mathcal{P})$ be an automorphism reflecting $f(\mathcal{P})$ across a line of symmetry L , and let $\mu = s \circ f$, mapping each point $p \in P_i$ of a piece $P_i \in \mathcal{P}$ to the corresponding point in the reflection of $f(P_i)$ across L . Let P be a leaf in the symmetric assembly whose siblings include at most one branch. Such a P exists, as any leaf with longest distance from an arbitrary root satisfies this property. We claim that either P is symmetric about line of symmetry L , or $\mu(P)$ is exactly covered by a second piece of the SAP congruent to P , contradicting the minimality of \mathcal{P} .

First, if P has no parent and is the only piece in the symmetric assembly, P must be a line symmetric polygon. Otherwise, let Q be the parent of P with edge e_P of P connected to edge e_Q of Q . Let e_{PQ} denote the subset of e_Q that maps to the intersection $f(e_P) \cap f(e_Q)$. We show that $f(e_{PQ})$ and $\mu(e_{PQ})$ are not the same line segment. Suppose for contradiction $f(e_{PQ}) = \mu(e_{PQ})$. Then either $f(e_{PQ})$ lies along L or is symmetric about L .

If $f(e_{PQ})$ lies along L , consider either endpoint p of e_P . $f(p)$ is either in the interior or on the boundary of $f(P)$. If $f(p)$ is interior, then the two edges of P incident to $f(p)$ must be connected to two different pieces, contradicting that P is a leaf. Alternatively, $f(p)$ is on the boundary, and a vertex of some other piece P' must contain $f(p)$, contradicting exclusion from Case 1.

Alternatively $f(e_{PQ})$ is symmetric about L . Because P is a leaf, it connects to the rest of the symmetric assembly only through $f(e_{PQ})$, so for the assembly to be symmetric, $f(P)$ must be the same as $\mu(P)$, and piece P is a line symmetric polygon.

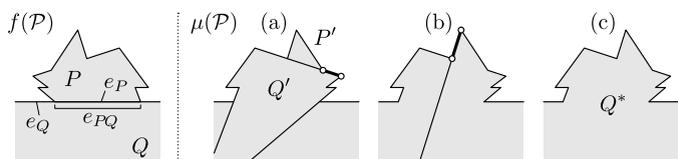


Fig. 11. Possible topological configurations of $\mu(P)$.

So $f(e_{PQ})$ and $\mu(e_{PQ})$ are not the same line segment. We claim that $\mu(P)$ is exactly covered by another piece of the SAP congruent to P . Suppose for contradiction it were not. Then by Lemma 8, P must be fully connected to Q . Then $\mu(P)$ either (a) contains a piece as a strict subset, (b) does not fully contain a piece but intersects the interiors of multiple pieces, or (c) is a strict subset of a single piece (see Fig. 11).

First suppose (a), so $\mu(P)$ contains some piece as a strict subset. We will say that piece P covers a piece P' if $f(P')$ is a strict subset of $\mu(P)$. We will identify a leaf piece P' covered by P , whose parent connection constructs the length of an edge of P , contradicting exclusion from Case 2. To find such a piece, consider any piece R that is not covered by P , and let S be a piece from among all pieces covered by P that has longest distance to R in the connection graph. This condition ensures that S is a leaf, connected to some piece Q' through edge $e_{Q'}$ from Q' . Because S is covered by P , at least one endpoint of $e_{Q'}$ maps to point contained in $\mu(P)$. Let q be such an endpoint. Point $f(q)$ is a vertex of the symmetric assembly or else the connection of Q and some other piece would share a vertex on their connection at $f(q)$. Let P' be the piece connected to $e_{Q'}$ with connection closest to q . P' is a leaf or else S would not have had a longest distance to R in the connection graph. Further, because S is covered by P , so is P' . By Lemma 8, the connection between P' and Q' must be fully connected. If $f(e_{P'}) = f(e_{Q'})$ then P' and Q' share vertices along their connection, contradicting exclusion from Case 1. If $f(e_{Q'}) \subset f(e_{P'})$, then because P' is a leaf, $f(e_{P'}) \setminus f(e_{Q'})$ constructs the lengths of two edges of P , contradicting exclusion from Case 2. So edge $e_{P'}$ fully connects P' to Q' in the assembly. And because no other piece connects to $e_{Q'}$ between vertex q and the connection between P' and Q' , the distance between them constructs the length of an edge of P , contradicting exclusion from Case 2. So $\mu(P)$ does not contain a piece as a strict subset.

Now suppose (b), so that two connected pieces intersect $\mu(P)$. The edges connecting these two pieces must overlap in $\mu(P)$ to construct a length equal to an edge of P , contradicting exclusion from Case 2. So $\mu(P)$ does not intersect the interior of multiple branch pieces.

Finally suppose (c), and let $\mu(P)$ be the strict subset of some piece Q^* . Let ℓ be the line collinear with segment $f(e_{PQ})$, and let e_ℓ be the subset of Q that maps to the largest connected subset of $\ell \cap f(Q)$ containing $f(e_{PQ})$. Consider the two disconnected sections of the boundary of Q between an endpoint of e_{PQ} and an endpoint of e_ℓ , which must each be more than an isolated point or exclusion from Case 1 would be violated. Piece P has at most one branch sibling, so at most one of these sections can be connected to a branch. Let q be an endpoint of e_ℓ in a section not connected to a branch.

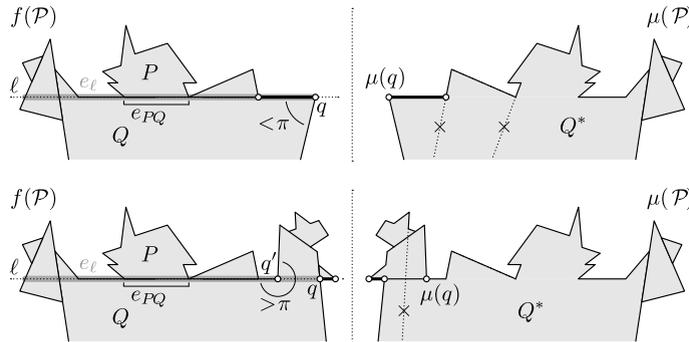


Fig. 12. Considering if $\mu(P)$ is a strict subset of Q^* and the boundary between e_{PQ} and q is a [Left] straight line or [Right] not a straight line.

Consider the boundary of Q between e_{PQ} and q . Suppose this boundary were a line segment subset of e_Q , implying the internal angle of Q at q is less than π ; see Fig. 12. Point q is not included in the connection between Q and another piece through e_Q . If it were, it would be a partial connection with a leaf piece, which by Lemma 8 would be part of a partially connected pair contradicting the minimality of $f(\mathcal{P})$. Further, $\mu(q)$ is a vertex of $f(Q^*)$ or else Q^* would connect to another piece somewhere on the segment between $\mu(e_{PQ})$ and $\mu(q)$, and their connection would construct an edge of the same length as an edge from a leaf connected to $f(e_Q)$, contradicting exclusion from Case 2. The edge of Q^* adjacent to $\mu(q)$ contained in $\mu(e_Q)$ will have the same length as the subset of e_Q between q and a vertex of a leaf, contradicting exclusion from Case 2.

Thus, the boundary of Q between e_{PQ} and q is not a line segment, so $f(Q)$ must cross ℓ , and the endpoint q' of e_Q in this section is a vertex of Q with internal angle greater than π ; see Fig. 12. By the same argument as we applied to q in the preceding paragraph, $\mu(q')$ must be in $f(Q^*)$, and if it were a vertex, we would have the same contradiction as with q before. However this time $\mu(q')$ need not be a vertex of $f(Q^*)$ because $f(Q^*)$ may extend past $\mu(q')$, with Q^* connecting to another piece on the other side of e_ℓ . However, the connection between these pieces will construct an edge that is the same length as an edge in either Q or a leaf connected to Q , and we have arrived at our final contradiction. So if P is not line symmetric, $\mu(P)$ is itself a piece of the SAP congruent to P , contradicting the minimality of \mathcal{P} . \square

Because every symmetric assembly can be classified as one of these cases, we can check for each case to decide whether the SAP has a symmetric assembly. Given a SAP that does not satisfy Case 1 or Case 2, by Lemma 6 it must satisfy Case 3 if it has a symmetric assembly. However, satisfying Case 3 is not sufficient to ensure a symmetric assembly. For example,

two congruent regular polygons with many sides and a single regular star with many spikes cannot by themselves form a symmetric assembly, though they satisfy Case 3, because no pair of edges can be joined without making the pieces overlap (for example, the Case 3 example from Fig. 10, exchanging to the triangular pieces for large regular hexagons). Thus given a SAP in Case 3, we must search the configuration space of possible connected arrangements of the pieces for an arrangement that forms a simple polygon.

Recall that the connection graph for a symmetric assembly not in Case 1 must be a tree. For a SAP with k pieces, consisting of at most n vertices in total, Cayley's formula says the number of distinct connection trees is k^{k-2} [6]. However, even if two pieces are connected, they could be connected through $O(n^2)$ different pairs of directed edges, so the number of different *edge distinguishing connection trees*, connection trees distinguishing between which pairs of edges are connected, can be no more than $n^{2k}k^k = O(n^{2k})$ (k is constant). As an instance of Case 3, \mathcal{P} consists of one or more symmetric pieces, with the rest being congruent pairs. Let $\mathcal{D}_{\mathcal{P}}$ and $\mathcal{D}'_{\mathcal{P}}$ be maximal disjoint subsets of \mathcal{P} such that there exists a matching $\eta : \mathcal{D}_{\mathcal{P}} \rightarrow \mathcal{D}'_{\mathcal{P}}$ between pieces in $\mathcal{D}_{\mathcal{P}}$ and $\mathcal{D}'_{\mathcal{P}}$ such that matched pairs are congruent. Let $\mathcal{S}_{\mathcal{P}}$ be the set of symmetric pieces in \mathcal{P} not in $\mathcal{D}_{\mathcal{P}}$ or $\mathcal{D}'_{\mathcal{P}}$. Let \mathcal{D}_s denote some subset of the symmetric pieces contained in $\mathcal{D}_{\mathcal{P}}$, and define a *trunk* to be a subset of symmetric pieces $\mathcal{R} = \mathcal{S}_{\mathcal{P}} \cup \mathcal{D}_s \cup \eta(\mathcal{D}_s)$ that can be connected into a simple polygon without overlap while aligning each of their lines of symmetry to a common line L (see Fig. 13). Define a *half tree* T to be an edge distinguishing connection tree on $\mathcal{R} \cup \mathcal{D}_{\mathcal{P}}$ such that every piece in $\mathcal{D}_{\mathcal{P}}$ connected to a piece R in \mathcal{R} connects through an edge of R intersecting the same half-plane bounded by L . We call this half-plane the *connecting half-plane*, with the other half-plane the *free half-plane*. The reason we define half trees is if we can find a point in their configuration space for which pieces do not intersect and for which pieces in $\mathcal{D}_{\mathcal{P}}$ not in the trunk do not intersect the free half-plane, we can place the remaining congruent pieces in $\mathcal{D}_{\mathcal{P}} \setminus \mathcal{D}_s$ at the mirror image of their respective matched pairs to complete a symmetric assembly. If a symmetric assembly exists satisfying Case 3, the assembly will correspond to a point in the constructed configuration space by definition.

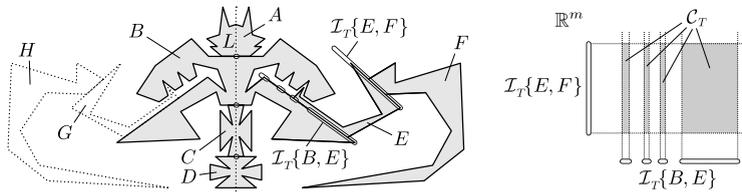


Fig. 13. An example showing a SAP \mathcal{P} satisfying Case 3, with $\mathcal{S}_{\mathcal{P}} = \{A, B\}$, $\mathcal{D}_{\mathcal{P}} = \{C, E, F\}$, $\mathcal{D}'_{\mathcal{P}} = \{D, G, H\}$, $\mathcal{D}_s = \{C\}$, $\eta(\mathcal{D}_s) = \{D\}$, and trunk $\mathcal{R} = \{A, B, C, D\}$. \mathcal{I}_T for two connected pieces in the trunk is just a single point as shown by the midpoint of their connection. Pieces not in the trunk have a degree of freedom sliding along their connection. $\mathcal{I}_T\{E, F\}$ is a single interval where F can attach to E , while $\mathcal{I}_T\{B, E\}$ is four intervals. The right diagram shows \mathcal{C}_T the Cartesian product of each \mathcal{I}_T .

Let $\mathcal{T}_{\mathcal{P}}$ be the set of possible half trees of \mathcal{P} . Let \mathcal{L}_T be the set of undirected edges $\{P, Q\}$ where piece P is connected to piece Q in tree $T \in \mathcal{T}_{\mathcal{P}}$, and let $m = |\mathcal{L}_T| \leq k$. For a fixed edge distinguishing connection tree, the orientation of each piece is fixed as pieces may only translate along their specified connection. We want to define a set of intervals $\mathcal{I}_T\{P, Q\}$ where we could join pieces P and Q along respective edges e_P to e_Q that are connected in tree T , while together forming a simple polygon without overlap. For each $\{P, Q\} \in \mathcal{L}_T$ with e_P and e_Q the respective connecting edges of P and Q with $\lambda(e_P) \geq \lambda(e_Q)$, let $\mathcal{I}_T\{P, Q\}$ be defined as follows. If P and Q are both in \mathcal{R} , let $\mathcal{I}_T\{P, Q\}$ be the empty set if $\text{JOIN}(e_P, e_Q, d_{PQ})$ is the empty set and $\{d_{PQ}\}$ otherwise, where we use d_{PQ} to denote $|\lambda(e_P) - \lambda(e_Q)|/2$, the distance d would need to be in order to align the midpoints of e_P and e_Q . Alternatively if P or Q are not both in \mathcal{R} , let $\mathcal{I}_T\{P, Q\}$ be the closure of the set of distances d for which $\text{JOIN}(e_P, e_Q, d)$ is a simple polygon for which P and Q do not share a vertex along their connection. Note that if P or Q are not both in \mathcal{R} , $\mathcal{I}_T\{P, Q\}$ will be a sequence of positive length closed intervals. Each interval endpoint represents a point at which P and Q would just intersect, no longer forming a simple polygon. The number of such points is upper bounded by the $O(n^2)$ possible intersections between some edge of P and some edge of Q when sliding the two pieces along their connection; so the number of distinct intervals in $\mathcal{I}_T\{P, Q\}$ is at most quadratic in the number of vertices, $O(n^2)$. Any fixed arrangement of the pieces consistent with an edge distinguishing connection tree T joins each pair of pieces by fixing one point in every $\mathcal{I}_T\{P, Q\}$, so the set of configurations is a subset of \mathbb{R}^m . Ignoring overlap between pieces that are not connected, the configuration space \mathcal{C}_T of possible arrangements is defined as the Cartesian product of $\mathcal{I}_T\{P, Q\}$ for every $\{P, Q\} \in \mathcal{L}_T$. Thus \mathcal{C}_T is a set of $O(n^m)$ disjoint m -dimensional hyperrectangles in \mathbb{R}^m .

We now describe the subset of \mathbb{R}^m where intersection occurs between two pieces that are not connected in T . If two pieces in a configuration overlap, by continuity there exist two edges e_P and e_Q from two distinct pieces P and Q that also intersect. The positions of e_P and e_Q are translations parameterized by a point in \mathcal{C}_T and the region in which the two edges intersect is a convex region $\mathcal{X}_T\{e_P, e_Q\} \subset \mathbb{R}^m$ bounded by four hyperplanes forming the m -dimensional parallelogram representing the intersection of the two edges. For each $O(n^2)$ pair of edges from distinct pieces that are not connected, we can subtract each $\mathcal{X}_T\{e_P, e_Q\}$ from \mathcal{C}_T to form \mathcal{C}'_T .

If \mathcal{C}'_T is empty, there will certainly be no symmetric assembly satisfying Case 3. If \mathcal{C}'_T is a single point, tree T places all pieces in the trunk \mathcal{R} to form a symmetric assembly. Lastly, if \mathcal{C}'_T is non-empty and contains a point in its interior,

then there exists a symmetric assembly because it will be a point in the configuration space avoiding overlap between pieces. Points on the boundary of C'_T correspond to configurations that are non-simple (the symmetric assembly is not homeomorphic to a disc), as the boundaries of each \mathcal{I}_T not between two pieces in \mathcal{R} and the boundaries of each \mathcal{X}_T correspond to configurations which produces a hole in the assembly or a cycle in the connection graph. Thus, if \mathcal{P} has a symmetric assembly satisfying Case 3, C'_T will have a point on its interior or be a single point.

Consider the function HASASSEMBLYCASE3 described in Algorithm 1.

```

1 Function HASASSEMBLYCASE3( $\mathcal{P}$ )
2   input : Symmetric assembly puzzle  $\mathcal{P}$ .
3   output: TRUE if  $\mathcal{P}$  has a Case 3 symmetric assembly, // FALSE otherwise.
4   for  $T \in \mathcal{T}_{\mathcal{P}}$  do
5      $C'_T \leftarrow C_T$ 
6     for  $\{P, Q\} \in \mathcal{L}_T$  do
7        $C'_T \leftarrow C'_T \setminus \mathcal{X}_T\{e_P, e_Q\}$ 
8       if  $\text{interior}(C'_T) \neq \emptyset$  then
9         return TRUE
10      else if  $C'_T \neq \emptyset$  and  $\mathcal{D}_{\mathcal{P}} = \emptyset$  then
11        return TRUE
12   return FALSE

```

Algorithm 1: Pseudocode for function HASASSEMBLYCASE3(\mathcal{P}).

Lemma 9. Given symmetric assembly puzzle \mathcal{P} that satisfies Case 3, function HASASSEMBLYCASE3(\mathcal{P}) returns TRUE if and only if \mathcal{P} has a symmetric assembly, and terminates in $O(n^{6k})$ time.

Proof. We can test all pieces for line symmetry or congruence in $O(nk)$ time [7]. If \mathcal{P} has a symmetric assembly satisfying Case 3 with nonempty $\mathcal{D}_{\mathcal{P}}$, C'_T will have a point on its interior for some tree T as argued above; or if $\mathcal{D}_{\mathcal{P}}$ is empty, C'_T must be nonempty, i.e., a single point corresponding to constructing a trunk from all the pieces. There are $O(n^{2k})$ elements of $\mathcal{T}_{\mathcal{P}}$. There are $m = O(k)$ interval sets $\mathcal{I}_T\{P, Q\}$ each having computational complexity $O(n^2)$, so we can construct C_T naively in $O(n^{2k})$ time. The union X_T of the $O(n^2)$ regions $\mathcal{X}_T\{e_P, e_Q\}$, which are m -dimensional convex regions, has computational complexity at most $O(n^{2m})$, so the final computational complexity of $C'_T = C_T \setminus X_T$ is at most $O(n^{4m})$ and can be computed in as much time. Checking each of the $O(n^{2k})$ elements of $\mathcal{T}_{\mathcal{P}}$ in this way yields the running time for HASASSEMBLYCASE3 bounded by $O(n^{6k})$. \square

Our brute force algorithm HASASSEMBLY(\mathcal{P}) is described in Algorithm 2.

```

1 Function HASASSEMBLY( $\mathcal{P}$ )
2   input : Symmetric assembly puzzle  $\mathcal{P}$ .
3   output: TRUE if  $\mathcal{P}$  satisfies Case 1 or Case 2 or Case 3, FALSE otherwise.
4   for  $e_P \in E_P, e_Q \in E_Q, \{P, Q\} \subset \mathcal{P}$  do
5      $S \leftarrow \text{JOIN}(e_P, e_Q, 0)$ 
6      $\mathcal{P}' \leftarrow (\mathcal{P} \setminus \{P, Q\}) \cup \{S\}$ 
7     if  $S \neq \emptyset$  and HASASSEMBLY( $\mathcal{P}'$ ) then
8       return TRUE // Case 1
9     for  $e_R \in E_R, R \in \mathcal{P}$  do
10      if  $\lambda(e_R) < \lambda(e_P)$  then
11         $S \leftarrow \text{JOIN}(e_P, e_Q, \lambda(e_R))$ 
12         $\mathcal{P}' \leftarrow (\mathcal{P} \setminus \{P, Q\}) \cup \{S\}$ 
13        if  $S \neq \emptyset$  and HASASSEMBLY( $\mathcal{P}'$ ) then
14          return TRUE // Case 2
15   return HASASSEMBLYCASE3( $\mathcal{P}$ ) // Case 3

```

Algorithm 2: Pseudocode for function HASASSEMBLY(\mathcal{P}).

Lemma 10. Function HASASSEMBLY(\mathcal{P}) returns TRUE if and only if \mathcal{P} has a symmetric assembly that satisfies either Case 1, Case 2, or Case 3, and terminates in $O(n^{6k})$ time.

Proof. We prove by induction. For the base case, \mathcal{P} consists of only a single piece satisfying Case 3, which will drop directly to the last line of the algorithm checking Case 3 which, by Lemma 9 will evaluate correctly. Now suppose HASASSEMBLY returns a correct evaluation for SAPs containing $k - 1$ pieces. Then we show HASASSEMBLY returns a correct evaluation for SAPs containing k pieces.

The outer **for** loop of HASASSEMBLY cycles through every pair of directed edges $e_P = (p_1, p_2)$ and $e_Q = (q_1, q_2)$ taken from different pieces P and Q . For each pair, HASASSEMBLY first checks to see if there exists a symmetric assembly for which e_P is connected to e_Q with p_1 coincident to q_1 , which would satisfy Case 1. If one exists, then joining P and Q into

one piece as described would produce a SAP \mathcal{P}' with one fewer piece that also has a symmetric assembly. Then evaluating HASASSEMBLY on the smaller instance will return correctly by induction. Because the outer **for** loop checks every possible pair of edges that could be joined in a symmetric assembly satisfying Case 1, HASASSEMBLY will return TRUE if \mathcal{P} satisfies Case 1.

Next HASASSEMBLY checks to see if there exists a symmetric assembly for which e_P is connected to e_Q with p_1 and q_1 separated by a distance equal to the length of some other edge e_R in \mathcal{P} , which would satisfy Case 2. In the same way as with Case 1, both **for** loops check every possible pair of edges and that could be joined at every possible length that could produce a symmetric assembly satisfying Case 2, so HASASSEMBLY will return TRUE if \mathcal{P} satisfies Case 2.

Otherwise, no symmetric assembly exists satisfying Case 1 or Case 2. By Lemma 9, HASASSEMBLYCASE3 correctly evaluates if \mathcal{P} is in Case 3, so HASASSEMBLY returns a correct evaluation for SAPs containing k pieces.

Let $T(k)$ be the running time of HASASSEMBLY on an instance with k pieces. Then the recurrence relation for HASASSEMBLY is $T(k) = O(n^3)T(k-1) + O(n^{6k})$, where $O(n^{6k})$ is the running time given by Lemma 9. Running time for Case 3 at the leaves dominates the recurrence relation so HASASSEMBLY terminates in $O(n^{6k})$. \square

Now we can determine whether a symmetric assembly puzzle with a constant number of pieces has a symmetric assembly in polynomial time.

Proof of Theorem 5. By Lemma 6, if the SAP has a symmetric assembly, it satisfies either Case 1, Case 2, or Case 3, and by Lemma 10 HASASSEMBLY(\mathcal{P}) can correctly determine if it has a symmetric assembly satisfying one of the cases in polynomial time, proving the claim. \square

4. Conclusion

Several open questions remain. It may be interesting to consider SAPs for special classes of shapes $P_i \in \mathcal{P}$. We conjecture that SAPs remain hard for instances in which the shapes P_i are right triangles (Conjecture 4). Are SAPs hard for a constant number $k = O(1)$ of pieces if the target shape is allowed to be nonsimple (a polygon with holes)? Are SAPs fixed-parameter tractable with respect to the number k of pieces? (We conjecture W[1]-hardness.)

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Many of the authors were introduced to symmetric assembly puzzles during the 30th Winter Workshop on Computational Geometry at the Bellairs Research Institute of McGill University, March 2015. Korman was supported by MEXT KAKENHI No. 17K12635 and the NSF award CCF-1422311. Mitchell is supported in part by the National Science Foundation (CCF-1526406). Van Renssen and Roeloffzen were supported by JST ERATO Grant Number JPMJER1201, Japan. Uno is supported in part by JSPS KAKENHI No. 17K00017 and by JST CREST Grant Number JPMJCR1402, Japan.

References

- [1] J. Slocum, *The Tangram Book: The Story of the Chinese Puzzle with Over 2000 Puzzle to Solve*, Sterling Publishing, 2004.
- [2] E. Fox-Epstein, K. Katsumata, R. Uehara, The convex configurations of “Sei Shonagon Chie no Ita,” tangram, and other silhouette puzzles with seven pieces, *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* 99 (6) (2016) 1084–1089.
- [3] H. Yamamoto, Personal communication, 2014.
- [4] N. Iwase, Symmetrix, in: 24th International Puzzle Party (IPP 24), IPP24 Committee, unpublished, 2005, p. 54.
- [5] E.D. Demaine, M.L. Demaine, Jigsaw puzzles, edge matching, and polyomino packing: connections and complexity, *Graphs Comb.* 23 (Supplement) (2007) 195–208, <https://doi.org/10.1007/s00373-007-0713-4>.
- [6] A. Cayley, A theorem on trees, *Quart. J. Math.* 23 (1889) 376–378.
- [7] J.D. Wolter, T.C. Woo, R.A. Volz, Optimal algorithms for symmetry detection in two and three dimensions, *Vis. Comput.* 1 (1) (1985) 37–48.