

Folding polyominoes into cubes

Oswin Aichholzer^{*1}, Florian Lehner^{†2}, and Christian Lindorfer^{‡3}

¹Institute of Software Technology, Graz University of Technology, Austria

²Department of Mathematics, University of Auckland, New Zealand

³Institut für Diskrete Mathematik, Technische Universität Graz, Austria

February 29, 2024

Abstract

Which polyominoes can be folded into a cube, using only creases along edges of the square lattice underlying the polyomino, with fold angles of $\pm 90^\circ$ and $\pm 180^\circ$, and allowing faces of the cube to be covered multiple times? Prior results studied tree-shaped polyominoes and polyominoes with holes and gave partial classifications for these cases.

We show that there is an algorithm deciding whether a given polyomino can be folded into a cube. This algorithm essentially amounts to trying all possible ways of mapping faces of the polyomino to faces of the cube, but (perhaps surprisingly) checking whether such a mapping corresponds to a valid folding is equivalent to the unlink recognition problem from topology.

We also give further results on classes of polyominoes which can or cannot be folded into cubes. Our results include (1) a full characterisation of all tree-shaped polyominoes that can be folded into the cube (2) that any rectangular polyomino which contains only one simple hole (out of five different types) does not fold into a cube, (3) a complete characterisation when a rectangular polyomino with two or more unit square holes (but no other holes) can be folded into a cube, and (4) a sufficient condition when a simply-connected polyomino can be folded to a cube.

These results answer several open problems of previous work and close the cases of tree-shaped polyominoes and rectangular polyominoes with just one simple hole.

^{*}O. Aichholzer was partially supported by FWF (Austrian Science Fund) project DK W1230.

[†]F. Lehner was partially supported by FWF (Austrian Science Fund) projects P31889.

[‡]C. Lindorfer was partially supported by FWF (Austrian Science Fund) projects P31889 and DK W1230.

1 Introduction

Even in its seemingly simplest form folding a piece of paper is a fruitful source for interesting research problems. Citing from Chapter 7 (The Combinatorics of paper folding) in the book of Martin Gardner [5]: “One of the most unusual and frustrating unsolved problems in modern combinatorial theory, proposed many years ago by Stanislaw M. Ulam, is the problem of determining the number of different ways to fold a rectangular map. The map is precreased along vertical and horizontal lines to form a matrix of identical rectangles. The folds are confined to the creases, and the final result must be a packet with any rectangle on top and all the others under it.”

Surprisingly, even the version of an $1 \times n$ map is open. This is known as the stamp folding problem, as it is similar to count in how many ways a single strip of stamps can be folded along their perforated edges such that all stamps end up on top of each other. See again [5] for a history of this problem and also Example 4.2 below.

In a similar spirit, the following innocent looking problem was proposed in [2]: Which polyominoes can be folded into a cube? More precisely, given a piece of paper in the shape of a polyomino it is asked whether it has a folded state in the shape of a unit cube. Here a polyomino is a polygon in the plane build by a collection of unit squares on the integer square lattice that are connected edge-to-edge (see the next section for more formal definitions).

In the same paper different models for folding were used to tackle the problem. In addition to foldings along the creases of the lattice of the polyomino (similar to what is used for the above mentioned map problem) it can be also allowed to fold along a $\pm 45^\circ$ diagonal of a square of the polyomino, or also so-called half-grid foldings are possible. These are foldings where orthogonal and diagonal folds between half-integral points can be made. Other differences in the folding model come from the way the faces of the cube are covered. For diagonal and half-grid models it is possible that a face of the cube is only partially covered by (parts of) a square of the polyomino, and the rest is covered by other (parts of) squares. Note, however, that in all these cases it is allowed that a face is covered by multiple layers of paper (like in origami, but unlike the general polyhedron unfolding), where some of these layers might only partially cover a face. Another restriction to the folding model could be the requirement that the surface of the cube is made entirely from just one side of the paper. That is, if you have a paper which is on one side green and on the other side red, then only the green side is visible to the exterior of the cube. For further variates of folding models we refer to [2].

For the half-grid model it was shown in [2, Thm. 3] that all polyominoes of at least ten unit squares can be folded into a unit cube. Rather recently some of the remaining cases were covered in [4]. For example there it is shown that any tree-shaped polyomino of at least 9 squares can be folded to the cube in the half-grid model. Table 1 of this paper contains a nice overview of previous results in the different folding models.

In our work we focus on the grid-folding model. There foldings are allowed only along edges of the square lattice of the polyomino and the folding angles can be $\pm 90^\circ$ or $\pm 180^\circ$. Thus, every face of the cube is covered by a seamless square in the folding and we allow faces to be covered multiple times, that is, by several layers of different squares of the polyomino. Readers who are not familiar with this folding model might find it instructive to solve the three cube folding puzzles given by Nikolai Beluhov [3] to see how powerful the model is.

Previous and new results in the grid-folding model

Our first few results concern computational aspects of the problem. In [1], the authors observe that if a polyomino folds into a cube, then there is an adjacency-preserving map from the faces of the polyomino to the faces of the cube, which they call a *consistent mapping*. Clearly, if a polyomino does not admit a consistent mapping, then it cannot be folded into a cube, giving an easy brute-force algorithm to certify non-foldability. They ask whether this algorithm can also be used to certify foldability, and whether every consistent mapping can be turned into a valid folding. While the answer to the second question turns out to be negative even in the special case of trees (see Example 4.5 below), we provide a positive answer to the first question: we show that there is an algorithm which determines whether a consistent mapping can be turned into a valid folding.

The second half of this paper focuses on foldability conditions for various classes of polyominoes, namely, rectangular polyominoes with (simple) holes, tree-shaped polyominoes, and simply connected polyominoes.

Our starting point for our investigation of polyominoes with holes is a result from [1, Thm. 1] which states that any polyomino that contains a non-simple hole can be folded into a cube. There are five different simple (also called basic) holes that can occur inside a polyomino, see Figure 1. In the same paper several combinations of two simple holes were given that allow the polyomino to fold into a cube. The authors also show that rectangular polyominoes with only a single square hole ([1, Thm. 11]) or a slit of size 1 ([1, Thm. 15]) and no other hole cannot fold onto the unit cube. Extending their result and answering one of their open problems we show in Section 5 that no rectangular polyomino with only one simple hole (one of the five simple holes shown in Figure 1) folds into a cube.

Further extending this result and generalizing Theorem 12 of [1] we provide a complete characterisation when a rectangular polyomino with two or more unit square holes (but no other holes) can be folded into a cube. Roughly speaking this can be done if and only if at least two of these holes lie in the same or in adjacent rows (resp. columns) and the number of columns (resp. rows) between them is odd. See Theorem 6.2 in Section 6 for

details.

Our next group of results concerns tree-shaped polyominoes. A polyomino is called *tree-shaped* if its dual is a tree. As usual, here the dual of a polyomino consists of a vertex for every unit square of the polyomino and an edge connects two vertices whenever the corresponding unit squares are joined edge to edge (see the next section for a proper definition). In Section 7 we fully characterise all tree-shaped polyominoes which can be folded onto the cube. This is done by a combination of considering well-defined classes of small polyominoes, bounding boxes of different sizes, and some exceptional cases of constant size. See e.g. Figure 28 which shows the only tree-shaped polyomino with bounding box 4×4 that can not be folded into a cube. Besides this example we show that any tree-shaped polyomino with a bounding box of size 4×4 or larger can be folded to a cube (Theorems 7.5 and 7.6), and exhaustively characterise all tree-shaped polyominoes of smaller bounding size (Theorems 7.1 and 7.2). A similar result for bounding boxes of sizes $2 \times n$ and $3 \times n$ was presented already in [2, Thm. 4], but it seems their characterisation was incomplete, see Section 7 for details.

Finally, in Section 8 we consider simply-connected polyominoes. A polyomino is simply connected if it does not contain any holes. Clearly, all tree-shaped polyominoes are simply-connected, and simply-connected polyominoes which are not tree-shaped have to contain at least one rectangular sub-polyomino of size at least 2×2 . This makes it harder to classify foldability for simply-connected polyominoes. Using the maximal distance of a point on the boundary of the polyomino to a corner of its bounding box, we provide a sufficient condition when a simply-connected polyomino can be folded to a cube, see Theorem 8.2.

Before providing the details of these results we start by stating the underlying notation in the next section. In Section 3 we specify our folding model, and give some more rigorous definitions and general results on folding polyominoes. We then provide details on consistent mappings to the unit cube in Section 4, which will be used in the consecutive sections to derive the results mentioned above.

2 Notation

A polyomino is a plane geometric figure formed by joining one or more unit squares edge to edge. It may be regarded as a finite, connected subset of the regular square tiling. We do not require a connection between every pair of adjacent squares, i.e., we allow slits along the edges of the square tiling. As usual, the dual $D(P)$ of a polyomino P is constructed by taking a vertex for every unit square of the polyomino and connecting two vertices by an edge whenever the corresponding unit squares are joined edge to edge. We call a maximal set of missing squares and slits a hole in P , if $D(P)$ has a cycle

containing P in its interior. A hole is called simple, if it is one of the following: a unit square, a slit of size 1, an I-slit of size 2, an L-slit of size 2 or a U-slit of size 3.




Figure 1: The five simple holes: unit square hole, slit of size 1, I-slit of size 2, L-slit of size 2, U-slit of size 3. This Figure is similar to [1, Fig. 3] where these holes are called basic.

3 Folding model

In order to define our folding model rigorously, we first rigorously define what a polyomino is. An *abstract polyomino* is a square complex in which every edge and every vertex is contained in at least one face, and every edge is contained in at most two faces. In other words, an abstract polyomino is obtained from a disjoint collection of unit squares by identifying some of their edges in such a way that each edge gets identified with at most one other edge. Note that defining the faces to be unit squares allows us to view an abstract polyomino as a metric space, and therefore also as a topological space.

A *folding blueprint* of an abstract polyomino P is a map $\beta: P \rightarrow \mathbb{R}^3$ such that the restriction of β to each face A of P is an isometry. A *folded state* of an abstract polyomino P is a piecewise linear (topological) embedding $F: P \rightarrow \mathbb{R}^3$. A folded state F is said to ϵ -correspond to a folding blueprint β if $\|F(x) - \beta(x)\| < \epsilon$ for every $x \in P$; in this case we call F and ϵ -realisation of β .

We note that if F is a folded state, then the boundary of $F(P)$ in \mathbb{R}^3 is homeomorphic to a disjoint union of circles. To lighten notation, we fix some small ϵ for the rest of this section (for our purpose any $\epsilon < \frac{1}{2}$ will do), and write ‘correspond’ for ‘ ϵ -correspond’ and ‘realisation’ for ‘ ϵ -realisation’.

Let us call a folding blueprint *flat* if its image is contained in a plane and the images of different faces only overlap in edges. Note that not every abstract polyomino has a flat folding blueprint, but if it has one, then the flat folding blueprint is unique up to isometries of \mathbb{R}^3 .

Definition 3.1. A *polyomino* is an abstract polyomino P which admits a flat folding blueprint. A realisation of a flat folding blueprint is called a *trivial folding*.

Finally, we are ready to define our folding model.

Definition 3.2. Two folded states are called *equivalent* if they are ambient isotopic. A folded state F of a polyomino P is called *valid* if it is equivalent to a trivial folding; in this case we also call F a *folding* of P .

Intuitively speaking, the above definition states that two folded states F_1 and F_2 are equivalent, if we can ‘move’ the image of F_1 around in \mathbb{R}^3 to obtain the image of F_2 . In this process we allow ‘bending and stretching’ the image, but we do not allow ‘cutting it and glueing it back together’. It is not hard to see that any two trivial foldings of a polyomino P are equivalent, and therefore the same is true for any two foldings.

Definition 3.3. Let P be a polyomino, and let $Q \subseteq \mathbb{R}^3$. A *folding of P into Q* is a folding F corresponding to a folding blueprint $\beta: P \rightarrow Q$. If β is surjective, we say that F is a *folding of P onto Q* .

Usually we will consider the case where Q is the surface of the unit cube \mathcal{C} . We note that any folding blueprint $\beta: P \rightarrow \mathcal{C}$ must map faces of P to faces of \mathcal{C} . Hence, if a folding F corresponds to such a folding blueprint, then each face A of P is embedded somewhere near a face B of \mathcal{C} , and for every face B of \mathcal{C} there is a face A of P embedded near it. In particular, a folding of P into \mathcal{C} gives maps from the vertices, edges, and faces of P to the vertices, edges, and faces of \mathcal{C} . If for a folding into \mathcal{C} , all faces of \mathcal{C} are covered we call this a folding onto \mathcal{C} .

We finish this section by giving some necessary and sufficient conditions for valid foldings. Let us define a *polyhedral sphere* as a polyhedron $Q \subseteq \mathbb{R}^3$ which is homeomorphic to a sphere. An *unlink* is a collection of circles embedded in \mathbb{R}^3 which is ambient isotopic to a set of circles embedded side-by-side in a plane. The following theorem tells us that polyhedral spheres and unlinks can be used as certificates for valid foldings.

Theorem 3.4. *Let F be a folded state of a polyomino P . The following are equivalent.*

1. *The folded state F is valid, that is, F is a folding.*
2. *There is a polyhedral sphere Q such that $F(P) \subseteq Q$.*
3. *The boundary of $F(P)$ is an unlink.*

Proof. To see that condition 1 implies the other two, we note that only allowing piecewise linear movement of the polyomino leads to an equivalent definition of valid foldings, see [7, Theorem 6.2]. Since the image of any trivial folding is contained in a polyhedral sphere and has an unlink as its boundary, the same must be true for any valid folding.

The implication 2 to 1 follows from the fact that any two polyhedral spheres are ambient isotopic, for instance by the piecewise linear Schönflies theorem [8, Chapter 17, Theorem 12].

It remains to show the implication from 3 to 2. For this purpose first note that every polyomino contains a (topological) disk D such that $P \setminus \bar{D}$ is homeomorphic to a disjoint




Figure 3: Creating a polyomino from a link diagram

ϵ -neighbourhood of $x = n$ for some $n \in \mathbb{N}$, and no two crossings lie in a neighbourhood of the same x .

Let L be an arbitrary link, represented by a link diagram as above. We can create a polyomino P together with a folded state F from this diagram as follows, see Figure 3. First replace every line in the diagram by a surface strip of width 1 contained in an ϵ -neighbourhood of the x - z -plane, and resolve crossings by attaching gadgets to the strip as shown in Figure 3(a). This gives a collection of abstract polyominoes P_1, \dots, P_k (one for each connected component of L) together with folded states F_1, \dots, F_k of these abstract polyominoes. Each P_i is homeomorphic to an annulus, and each F_i corresponds to a folding blueprint $\beta_i: P_i \rightarrow Q$ where Q is a strip in the x - z -plane. Moreover, $\bigcup_i F_i(P_i)$ is ambient isotopic to a side-by-side embedding of the abstract polyominoes if and only if L is an unlink.

Next we obtain a polyomino P'_i from the abstract polyomino P_i by ‘cutting along a line in z -direction’ and ‘reconnecting the sides of the cut by a frame’, see Figure 3(b). It is not hard to see that P'_i admits a folded state F'_i corresponding to a folding blueprint $\beta'_i: P'_i \rightarrow Q$ such that $\bigcup_i F'_i(P'_i)$ is ambient isotopic to trivial side-by-side embeddings of the polyominoes P'_i if and only if L is an unlink. For instance, β'_i may map all faces of the frame we added in this step to the same unit square in Q .

We obtain a polyomino P by connecting the polyominoes P'_i as shown in Figure 3(c). Once again, it is straightforward to check that there is a folded state F of P which corresponds to a folding blueprint $\beta: P \rightarrow Q$ such that F is valid if and only if $\bigcup_i F'_i(P'_i)$ is ambient isotopic to trivial side-by-side embeddings of the polyominoes P'_i . Hence F is a folding if and only if L is an unlink.

Finally, observe that starting from F and ‘folding 180° along all edges in x -direction’ and then ‘folding 180° along all edges in z -direction’ yields a folded state F' corresponding to a folding blueprint mapping P to a unit square (and therefore to the unit cube). \square

4 Consistent mappings to the unit cube

Let us now consider polyominoes and the unit cube as (combinatorial) polygonal complexes. A *consistent mapping* from a polyomino P to the unit cube \mathcal{C} is a homomorphism φ of the corresponding polygonal complexes, i.e., it maps faces to faces, edges to edges

and vertices to vertices and preserves the incidence relation. We note that consistent mappings can be seen as a combinatorial version of folding blueprints; indeed, they carry precisely the same information. Up to isomorphism there is exactly one surjective consistent mapping from the standard cube net depicted in Figure 4 to the unit cube, which can be used to define a labelling on the unit cube \mathcal{C} . This labelling gives us a one-to-one correspondence between consistent mappings from a polyomino P to \mathcal{C} and certain labellings of the unit squares of P .




Figure 4: Standard cube net

Remark 4.1. Consistent mappings from a polyomino to \mathcal{C} are exactly the mappings yielded by the algorithm explained in [1, section 4.3]. An implementation in Python is freely available online, a link can be found in the mentioned paper.

Clearly every valid folding of a polyomino P into a cube gives a consistent mapping in the obvious way. However, not the complete information about a specific folding can be seen from its respective consistent mapping. In particular it is not possible to distinguish $+180^\circ$ folds from -180° folds. This also means that if a face of \mathcal{C} is covered by several different unit squares of P , the consistent mapping does not provide the order in which the unit squares appear on this face.

Example 4.2. Consider a strip of some integer length k , that is a rectangular polyomino of size $1 \times k$. Then one possible consistent mapping φ maps each unit square of P onto the same face **1** of \mathcal{C} , see Figure 5. It is not hard to see that there are several valid foldings yielding this specific consistent mapping φ . Clearly every edge must be folded by either $+180^\circ$ or -180° and for every given sequence of $k - 1$ angles in $\{+180^\circ, -180^\circ\}$ there is a folding of P such that the vertical edges are folded according to the sequence. Indeed, starting at one side of the strip we can consecutively fold each edge as determined by the sequence such that the next face lies on the outside of the ‘multi-layered square’ containing the squares encountered up to this points. Thus we get at least 2^{k-1} different foldings for our strip. Counting the number of possible foldings of the strip of length k into a square is known as the stamp folding problem. See Chapter 7 in Martin Gardners book [5] for the history of this long standing open problem. Its solutions are currently known up to $k = 45$ (see e.g. sequence A000136 in the online encyclopedia of integer sequences at <https://oeis.org/A000136>); for small k they are given by the sequence

1, 2, 6, 16, 50, 144, 462, 1392, 4536, 14060, 46310, 146376, 485914, 1557892, 5202690, . . .



Figure 5: Strip of size $1 \times k$ for even k

This example suggests that if we want to uniquely identify a given folding, it is reasonable to also include information about the order in which the unit squares of the polyomino appear on the cube to a consistent mapping. We say that square s of polyomino P lies in layer l of a folding of P into \mathcal{C} if there are exactly $l - 1$ squares of P covering the same face as s and lying below s . A *layer map* of a polyomino P for a given consistent mapping φ is a map ℓ assigning to each unit square s of P an integer between 1 and the number of unit squares of P mapped to the face $\varphi(s)$ of \mathcal{C} such that its restriction to $\varphi^{-1}(f)$ is a bijection for every face f of \mathcal{C} . The triple (P, φ, ℓ) is called a *pseudo-folding*. It is called *surjective* if φ is surjective.

4.1 Validity of pseudo-foldings

Recall that by Definitions 3.2 and 3.3 a folding of a polyomino P into \mathcal{C} is a realisation F of a folding blueprint $\beta : P \rightarrow \mathcal{C}$ which is equivalent to a trivial folding. Clearly any realisation of a folding blueprint $\beta : P \rightarrow \mathcal{C}$ uniquely defines a pseudo-folding (P, φ, ℓ) : The consistent mapping φ is given by β , while the layer map ℓ is provided by the order of the images of the squares of P under the embedding $F : P \rightarrow \mathbb{R}^3$ for each of the faces of the cube \mathcal{C} . In this case we say that the folded state F *realises* the pseudo-folding (P, φ, ℓ) , even if it is not equivalent to a trivial folding and thus no folding.

Definition 4.3. A pseudo-folding is *valid*, if it corresponds to a folding of P into \mathcal{C} .

For the rest of this section we want to discuss obstacles preventing pseudo-foldings from being valid. Instead of giving rigorous definitions for these obstacles, we decided to present them by discussing simple examples where those occur. According to the above discussion, there are two things that can go wrong: A given pseudo-folding might not have any realisation or none of its realisations is equivalent to a trivial folding.

Let us start with the first issue and present two examples of consistent mappings which cannot have any realisations. This is due to the fact that any possible layer map forces self-intersections and thus cannot be realised by any embedding F .

Example 4.4. Let us consider the consistent mapping in Figure 6, which was implicitly already discussed in the proof of [1, Lem. 9]. In any valid folding into \mathcal{C} , whenever two unit squares with the same label have neighbours with the same label, then the order of their respective layers must be retained, otherwise there would be self-intersections. In particular, if the **1** in the top left corner is folded on top of the **1** in the bottom right corner, then the top **3** is folded on top of the bottom **3** and continuing this argument along the cycle **1** – **3** – **6** – **2** – **1** – **3** – **6** – **2** – **1** leads to a contradiction.




Figure 6: A surjective consistent mapping that cannot be realised by any folded state

Example 4.5. To show that there is no valid folding for the consistent mapping shown in Figure 7, consider the order of the five unit squares in the third (and fourth, fifth) column in which they are folded on top of each other. Note that any row ending with a T-shape cannot lie between two rows not having a ‘T’ in the same direction in this order. This is the case because the T-part would lie between those two layers and thus self-intersection would occur. In particular row 1 and 5 cannot be placed between two other rows and the same holds for 2 and 4. Thus without loss of generality rows 1 and 5 are folded on top of all other rows, and row 2 and 4 are folded on bottom. It is now easy to check that none of these permutations give a valid folding. Starting with row 1, row 2 has to be folded downwards, then row 3 upwards, row 4 again downwards and thus ends up between 2 and 3. In particular row 5 cannot be folded upwards to end up next to row 1.




Figure 7: A surjective consistent mapping of a simply connected polyomino that cannot be realised by a folded state

Self-intersections always occur at pairs of edges of our polyomino mapped to the same edge of the cube by a pseudo-folding. They are caused by layer conflicts of the two faces incident to each of the edges. Each of the edges is folded by either 90° or 180° . Thus there are three types of intersections, shown in Figure 8. Firstly, if both edges are folded by 90° , the relative order of their faces must be preserved. Secondly, a 90° folded edge cannot lie between the two faces of a 180° fold. Finally, the layers of two

180° folded edges should be well nested. While the consistent mapping of Example 4.4 forces a self-intersection at two edges folded by 90°, any pseudo-folding obtained from the consistent mapping in Example 4.5 forces an intersection of a 90° folded edge with a 180° folded edge.




Figure 8: Different types of self-intersections

But even if a given pseudo-folding has realisations, it might still not be valid. This happens if any (\Leftrightarrow each) of its realisation is not equivalent to a trivial folding. Intuitively this means that the realisation F could not have been obtained from the flat polyomino without any self-intersections occurring during the folding process. From Theorem 3.4 we know that this happens exactly if the boundary of $F(P)$ is not the unlink. Obviously this can only happen if the polyomino P contains at least one hole.

We provide two structurally different examples of two realisable pseudo-foldings which are not valid, starting with an example of a pseudo-folding containing a twist.

Example 4.6. Consider the pseudo-folding (P, φ, ℓ) given in Figure 9, where layer numbers ℓ are drawn as roman numerals in each square.




Figure 9: A pseudo-folding containing a twist

One can easily check that there are no self-intersection. However, it is still not possible to realise this pseudo-folding. In fact, we could ‘cheat’ by cutting along one edge of the polyomino, twisting one of the two resulting ends by 360° and glue them back together to obtain the requested result. In other words, the pseudo-folding is not ambient isotopic to the flat polyomino, but to a ‘doubly twisted Möbius polyomino’. A good way to see this is by looking at the linking diagram of the two boundaries, which is drawn on the left of Figure 10. It is not hard to see that this diagram is equivalent (via Reidemeister moves) to the two linked rings shown on the right.




Figure 10: Linking diagram of the two boundaries.

Twistedness only occurs by multiples of 360° , as realisations of a given consistent mapping must always be orientable surfaces. In particular, non-orientable surfaces such as Möbius strips cannot appear.

For the polyomino P and the consistent mapping φ of Example 4.6 there are in total 16 different layer maps not leading to any self-intersections. Out of these 16 maps, 4 provide valid pseudo-foldings, 8 include a 360° twist, and 4 include a 720° twist. In particular in some cases a realisable pseudo-folding can be made valid by replacing only the layer map. The following example shows that this is not always the case.

Example 4.7. Consider the polyomino P and the consistent mapping φ depicted in Figure 11. It is not hard to find a layer map ℓ such that the pseudo-folding (P, φ, ℓ) is realisable, one example can be seen in the figure. However, a brute force case-by-case analysis shows that none of the 60 layer maps which do not lead to self-intersections provides a valid pseudo-folding.




Figure 11: A consistent mapping that cannot be realised without a twist.

Finally, we provide an example of a pseudo-folding with a single hole, where the frame of the hole is non-trivially knotted.

Example 4.8. Consider the pseudo-folding (P, φ, ℓ) given in Figure 12. Again, it is not




Figure 12: A pseudo-folding containing a knot.

possible to realise this pseudo-folding. Again, we can ‘cheat’ by cutting one edge of the polyomino, thus ending up with a polyomino P' without any holes. Then one can just realise the pseudo-folding for P' and in the end glue the two sides of the cut edge back together. When trying to unfold the result, it will be clear relatively quickly, that the folded state is knotted. To make this clear, we look at the knot diagram of the outer boundary of P depicted on the left of Figure 13. As our consistent mapping only maps onto the faces **1**, **3**, **4** and **6**, it is reasonable to draw the frontal view of the cube, looking directly at the face **2** so that the layers of the four mentioned faces become visible. It is not hard to see that the knot of our boundary is the trefoil knot on the right of Figure 13, which is the simplest non-trivial knot.




Figure 13: Linking diagram of the the outer boundary.

In the proof of Theorem 3.5 we have provided a way to translate a given link diagram L consisting of $k > 0$ closed curves to a polyomino P with k holes and a folded state F corresponding to a folding blueprint mapping P to a unit square. The main idea is to first realise every curve in L as the inner boundary of a polyomino with a single hole

and then connecting the resulting k polyominoes to obtain a single polyomino P , see also Figure 3 for a reminder. In fact, the inner boundaries (the boundaries of the k holes) of the polyomino are knotted and linked in precisely the way given by L . From this we immediately obtain that every possible knot can be realised as a folding of some polyomino P with a single hole into (and also onto) \mathcal{C} .

We have seen four examples of obstacles preventing pseudo-foldings from being valid. Obviously every polyomino can be trivially folded onto a single square and thus into \mathcal{C} using only 180° folds. In fact every example polyomino in this subsection admits a valid pseudo-folding whose consistent mapping onto \mathcal{C} is surjective and thus defines a folding onto \mathcal{C} . In this context, the following important question remains open.

Question 4.9. Is there any polyomino which admits surjective pseudo-foldings, but none that are valid and define foldings?

We finish this subsection with some algorithmic considerations. By Theorem 3.5, deciding whether a given pseudo-folding is realisable is polynomially equivalent to UNLINK. Since UNLINK is known to be in NP [6], this shows that the problem CUBEFOLDING of deciding whether a given polyomino can be folded onto \mathcal{C} is in NP. In particular, there is an algorithm to solve this decision problem.

However, even if there is an efficient algorithm for UNLINK, the naive algorithm for CUBEFOLDING has exponential runtime due to the fact that there are exponentially many pseudo-foldings. We conjecture that there is an efficient algorithm.

Conjecture 4.10. CUBEFOLDING can be solved in polynomial time.

5 Rectangular polyominoes with one simple hole do not fold

AICHHOLZER ET AL. showed in [1, Thm. 1] that the presence of a non-simple hole in any polyomino guarantees foldability onto \mathcal{C} . Additionally, in [1, Thm. 11] and [1, Thm. 15] the authors showed that rectangular polyominoes with only a single square hole or a slit of size 1 and no other holes cannot fold onto \mathcal{C} . The goal of this section is to extend this second statement also to the three other classes of simple holes shown in Figure 1 to obtain the following main result of this section.

Theorem 5.1. *A rectangular polyomino P with a simple hole h and no other holes does not fold onto \mathcal{C} .*

We will prove the statement separately first for I-slits of size 2 in Proposition 5.4 and then for U-slits of size 3 and L-slits of size 3 in Proposition 5.5. But first we start with some preparations.

Lemma 5.2 ([1, Lem. 7]). *In any folding of a polyomino P into \mathcal{C} , the four faces around a vertex not in the boundary of P cannot cover more than 2 faces of \mathcal{C} . In particular*

either all four faces cover the same face of \mathcal{C} or there are two adjacent faces of P covering one face of \mathcal{C} and the other two faces cover a second face of \mathcal{C} .

Corollary 5.3. *Let P be a polyomino and P' be a rectangular sub-polyomino of P whose interior does not contain any boundary of P . Then any consistent mapping φ of P to \mathcal{C} must be constant on each row of P' or constant on each column of P' .*

Let P be a rectangular polyomino. The *support* of a set H of holes of P is the smallest rectangular sub-polyomino P' of P such that each $h \in H$ is a valid hole of P' , i.e., h does not lie on the boundary of P' .

Let φ be a consistent mapping from P to \mathcal{C} . We say that φ is *trivial* for a simple hole h of P if φ is a valid consistent mapping for the polyomino obtained from P by removing h . In the case where h is a unit square hole φ is extended to the additional square as given by its boundary edges. Clearly it is sufficient to look at the support of h when checking whether φ is trivial for h .

Proposition 5.4. *A rectangular polyomino P with an I-slit of size 2 and no other holes does not fold onto \mathcal{C} .*

Proof. We will show that no surjective consistent mapping φ from P onto the unit cube \mathcal{C} can exist.

Note that φ is trivial for the I-slit h , if and only if it maps both ‘sides’ of the slit to the same edges of \mathcal{C} . By definition any trivial mapping is also a consistent mapping after removing h and it is clear that such a mapping never covers all faces of \mathcal{C} . Denote by P' the support of h , which is a 4×2 -polyomino. Using the algorithm from [1, section 4.3], we deduce that up to isomorphism and up to symmetry of P' there are 6 possible patterns a non-trivial φ can take on P' ; all of them are shown in Figure 14.




Figure 14: All non-trivial consistent mappings from the support of an I-slit of length 2 to \mathcal{C} .

We use Corollary 5.3 to show that none of these cases can extend to a surjective consistent mapping from P to \mathcal{C} . First note that it is enough to show that the restriction of φ to all unit squares of P sharing a row or a column with P' cannot map surjective onto \mathcal{C} . Then by Corollary 5.3 the image of P under φ contains the same faces as this restriction and thus not all faces of \mathcal{C} .

We make the following observations:

- If at least two unit squares in the first row of P' cover different faces of \mathcal{C} , any unit square of P lying above P' maps to the same face of \mathcal{C} as the unit square below it and the same holds for other sides of the boundary of P' . Therefore φ only maps to faces **1**, **2** and **3** of \mathcal{C} in case (a).
- In case (b) only the faces **1**, **3**, **4** and **6** of \mathcal{C} can be covered by unit squares to the left of P' , so φ does not map to face **5**.
- In cases (c), (d) and (e) only faces **1**, **2**, **5** and **6** of \mathcal{C} can be covered by unit squares above and below P' , so φ does not map to face **4**.
- In case (f) either the column of P containing the first column of P' or the row of P containing the first row of P' is mapped onto face **1** of \mathcal{C} . Thus either all unit squares of P above and below P' are mapped to face **1** and φ does not map to face **5** or all unit squares to the left of P' are mapped to face **1** and φ does not map to face **4**.

Summing up, in none of the cases φ can be surjective, so a rectangular polyomino with a single straight slit of size 2 does not fold onto \mathcal{C} . \square

Proposition 5.5. *A rectangular polyomino P with an U-slit of size 3 or an L-slit of size 2 and no other holes does not fold onto \mathcal{C} .*

Proof. Let us start with the case where P contains a single U-slit h of size 3 and no other holes. Observe that φ trivial, if both sides of the slit map to the same edge of \mathcal{C} . Additionally, if φ maps the unit square inside of the slits to the same face as the only unit square connected to it, P reduces to the case where h is a unit square hole and thus does not fold onto \mathcal{C} . Excluding these two types of consistent mappings, up to isomorphism and symmetry, only seven possible consistent mappings φ from the support P' of h to \mathcal{C} remain. They can be seen in Figure 15.

Similar as in the proof of the previous proposition, we make the following observations.

- If at least two unit squares in the first row of P' cover different faces of \mathcal{C} , any unit square of P lying above P' maps to the same face of \mathcal{C} as the unit square below it and the same holds for other sides of the boundary of P' . This handles cases (a), (b), and (c), where exactly one face of \mathcal{C} does not occur.
- In cases (d) and (e) only the faces **1**, **3**, **4** and **6** of \mathcal{C} can be covered by unit squares to the left and right of P' , so φ does not map any unit square to face **5**.
- In case (f) only faces **1**, **2**, **5** and **6** of \mathcal{C} can be covered by unit squares above and below P' , so φ does not map to face **4**.
- In case (g) either the column of P containing the first column of P' or the row of P containing the first row of P' is mapped to face **1** of \mathcal{C} . Thus either all unit squares of P above and below P' are mapped to face **1** and φ does not map to face




Figure 15: All non-trivial consistent mappings of the support of a U-slit to \mathcal{C} such that the face ‘inside’ the slit does not map to the same face as the face ‘above’ the slit.

5 or all unit squares to the left and right of P' are mapped to face **1** and φ does not map to face **4**.

As before φ cannot be surjective in any of these cases, thus P does not fold onto \mathcal{C} .

Finally note that if any P containing an L-slit of size 2 and no other holes folds onto \mathcal{C} , the polyomino with a U-slit of size 3 obtained by cutting an additional edge has to fold in the same way. We conclude that no such polyomino P can fold onto \mathcal{C} . \square

6 Folding polyominoes with two unit square holes

We have seen in Theorem 5.1 that a single hole in a rectangular polyomino P is not sufficient for P to become foldable. But what about two or more holes? In this section we fully characterise which rectangular polyominoes containing only unit square holes fold onto the cube. In particular we will show that certain pairs of unit square holes are necessary and sufficient for foldability.

Definition 6.1. Let P be a rectangular polyomino. We say that two unit square holes of P *cooperate* if the holes lie in the same or in adjacent rows (resp. columns) and the number of columns (resp. rows) between them is odd.

Observe first that by [1, Thm. 5] two cooperating unit square holes guarantee foldability. The main idea is to use 180° folds to reduce the corresponding polyomino to one of the two situations depicted in Figure 16. It is left to the reader to check that the consistent mappings shown in the figure provide a valid folding.

Let h be a unit square hole of a polyomino P and let P' be the support of h . For the upcoming Theorem it is essential to know all consistent mappings from P' to \mathcal{C} which are non-trivial for h . Following the discussion in [1, Sec. 4.1], there are two essentially




Figure 16: Consistent mappings for cooperating unit square holes

different mappings: Either the four boundary edges of h are mapped to the same edge of \mathcal{C} or they are mapped to two edges of \mathcal{C} forming an L-shape.

It is shown in [1, Lem. 9] that any mapping which maps the boundary of h to a single edge of \mathcal{C} does not yield a valid folding. The strategy in the proof is to show that in this case up to isomorphism one always ends up with the consistent mapping φ shown in Figure 17. But the discussion in Example 4.4 shows that φ does not provide a valid folding without self-intersections. In the upcoming theorem we want to exclude consistent mappings which can never yield valid foldings. Thus we call a consistent mapping φ on P *good*, if for each unit square hole h of P , the restriction of φ to the support of h is not isomorphic to the consistent mapping in Figure 17.




Figure 17: Non-trivial consistent mapping of a unit square hole onto \mathcal{C} of non-foldable type

In [1, Thm. 10] the authors show that up to isomorphisms of the cube and up to rotation and reflection of P' , there exists only a single good consistent mapping φ from P' onto \mathcal{C} which is non-trivial for h . This mapping arises from the second case and can be found in Figure 18. There is a unique face f of \mathcal{C} whose boundary contains the two edges covered by the 4 boundary edges of h . Observe that f is covered by two squares lying on different sides of the hole h and that the face opposite to f in \mathcal{C} cannot be covered by a square in P' . We say that the *type of h under φ* is f .




Figure 18: Non-trivial consistent mapping φ of the support of a unit square hole h to \mathcal{C} . The type of h under φ is **3**.

Theorem 6.2. *Let P be a rectangular polyomino with unit square holes and no other holes. Then P folds onto \mathcal{C} if and only if two of the holes cooperate.*

Proof. We have just seen that P folds if it contains two cooperating unit square holes.

To prove the other direction, we will show that for any polyomino P with unit square holes and no other holes such that no two of them cooperate, there cannot exist a good surjective consistent mapping. We use induction on the number of holes k .

The base cases $k = 0$ and $k = 1$ follow from Theorem 5.1. Consider now $k > 1$ and assume that the induction hypothesis holds for all polyominoes with less than k unit square holes. Let P be a polyomino with k unit square holes and no other holes and assume for a contradiction that there is a good surjective consistent mapping φ . Then φ is non-trivial for every hole h of P , otherwise the natural extension of φ to the polyomino obtained from P by removing h would be a good surjective consistent mapping from a polyomino with $k - 1$ holes to \mathcal{C} , contradicting our assumption.

Let us define a graph G as follows. The vertex set is the set of holes H of P and two holes $h_1, h_2 \in H$ are connected by an edge if they lie in the same or adjacent rows or columns of P and the support of $\{h_1, h_2\}$ in P does not contain any other holes in H .

Claim 1. The graph G is connected.

To prove this, suppose for a contradiction that G is not connected. Let h_1 and h_2 be two holes lying in different components of G such that the support P' of $\{h_1, h_2\}$ is of minimal size, that is, the sum of its rows and columns is minimal. Then P' cannot contain any other hole. Indeed, such a hole h cannot lie in the support of h_1 or h_2 and thus both the support of $\{h, h_1\}$ and $\{h, h_2\}$ have smaller size than P' . This contradicts the minimality of P' .

By construction of G the holes h_1 and h_2 cannot lie in the same or adjacent rows or columns of P , otherwise they would either cooperate or be connected by an edge in G . Assume without loss of generality that h_1 lies in row 2 and column 2 and h_2 lies in row $i \geq 4$ and column $j \geq 4$ of P' . The map φ is non-trivial for both h_1 and h_2 , thus there are two adjacent unit squares in the first three rows of column 3 mapped to different faces of \mathcal{C} . Using Corollary 5.3 we see that φ must be constant on each row of the sub-polyomino consisting of rows 1 to $i - 1$ and columns 3 to $j + 1$, see Figure 19. But this contradicts our assumption that φ is non-trivial for h_2 and thus proves Claim 1.

Claim 2. Whenever two holes $h_1, h_2 \in H$ are neighbours in G , their type coincides.

To prove this claim, assume first that the two holes lie in the same row or column of P . This case is already treated in the proof of [1, Thm. 12], we will give a brief summary. Assume without loss of generality that h_1 lies in column 2 and h_2 in column j of the support P' . The map φ is non-trivial for both h_1 and h_2 , thus there must be two adjacent unit squares in column 3 mapping to different faces of \mathcal{C} . Thus Corollary 5.3 yields that φ is constant on each row of the sub-polyomino consisting of columns 3 to




Figure 19: The holes h_1 and h_2 cannot both be non-trivial.

$j - 1$. As h_1 and h_2 do not cooperate by assumption, the number of columns between them is even. We conclude that the pattern around h_2 is the reflection of the pattern around h_1 , so in particular they are of the same type. Up to isomorphism and symmetry there are only two possible situations, see Figure 20.




Figure 20: Possible patterns for holes in the same row

Next we consider holes lying in adjacent rows or columns of P . Assume without loss of generality that h_1 lies in row 2 and column 2 and h_2 in row 3 in column j of the support P' . By assumption h_1 and h_2 do not cooperate, so j must be odd and at least 5. Additionally φ is non-trivial for h_1 and h_2 . An application of Corollary 5.3 to the sub-polyomino consisting of rows 1 and 2 and columns 3 to $j+1$ yields that φ is constant on its columns and thus in particular the squares in row 1 and 2 of column 3 are mapped to the same face of \mathcal{C} . Analogously we obtain that the squares in rows 3 and 4 of column $j-1$ are mapped to the same face of \mathcal{C} . Finally, applying Corollary 5.3 to the rectangular sub-polyomino consisting of columns 3 to $j-1$ of P' provides that the situation is up to isomorphism as shown in Figure 21. In particular the type of h_1 and h_2 under φ coincides. This concludes the proof of Claim 2.

Connectedness of G implies that the type of all unit square holes of the polyomino P coincides. Let f be this type. Then for any hole h no square of the support of h is mapped to the face f' opposite to f in \mathcal{C} . Furthermore, the restriction of φ to a rectangular sub-polyomino P'' of P containing only a single hole h cannot be surjective by the induction basis. Using that the support of h already covers all faces but f' , we

1	3	6	8		8	8	1
2		9	9		9	8	1
3	3	5	5		5		1
2	3	5	2		2	3	3

Figure 21: Possible pattern for holes in adjacent rows

conclude that no square of P'' is mapped to f' by φ . Finally, it is not hard to see that any square of the polyomino P is contained in some rectangular sub-polyomino containing precisely one hole of P , so that φ cannot be surjective on P . \square

7 Folding tree-shaped polyominoes

A polyomino P is called *tree-shaped* if its dual $D(P)$ is a tree. Note that this both forbids holes and also any rectangular sub-polyominoes of size at least 2×2 . A tree-shaped polyomino is called *path-shaped*, if additionally every square is connected to at most two other squares. The goal of this section is to fully characterise all tree-shaped polyominoes which can be folded onto the cube. The *bounding box* of a polyomino P is the smallest rectangle containing P . Note that by definition P touches each side of the rectangle. The size of the bounding box of P is called the *bounding size* of P .

Our main goal is to extend results of AICHHOLZER ET AL. in [2], who provided a characterisation of foldable tree-shaped polyominoes with a bounding box of height at most 3. However, both their statement and proof appear to be incomplete; for instance, their claimed results would imply that the polyomino sketched in Figure 25(b) below would not be foldable. Hence we decided to additionally come up with our own formulation and proof of their statement. To do this, we use similar methods.

We start by providing some useful language. Let P and Q be two tree-shaped polyominoes. We say that Q is a *minor* of P if P contains a sub-polyomino P' which can be reduced to Q by 180° folds. Note that because P is tree-shaped, containing a foldable sub-polyomino implies that P is foldable. In particular, whenever P contains a foldable minor, P must be foldable.

A *leaf square* of P is a square of P which is only connected at one side to the rest of P . In other words a leaf-square is a square whose dual is a leaf in $D(P)$. A *leaf-fold* is the 180° fold of a leaf square of P . Note that a leaf-fold transforms P into another tree-shaped polyomino P' which has one square less.

In the upcoming proofs, we often show foldability of a polyomino by proving that it contains a cube net as a minor. Up to symmetry, there are eleven different cube nets; all of them are depicted in Figure 22. Due to the number of squares in each row, the

first six cube nets are said to be of type 1-4-1, the following three are of type 2-3-1, then there is one of type 2-2-2 and the last one is of type 3-3.




Figure 22: The eleven different cube nets.

A first trivial observation is that any folding of a polyomino of bounding size $1 \times n$ can cover at most 4 faces of the cube \mathcal{C} .

7.1 Folding tree-shaped polyomino of bounding size $2 \times n$

Let us start by defining an infinite class of tree-shaped polyominoes of bounding size $2 \times n$. Class A includes all polyominoes consisting of a $1 \times n$ strip, possibly with several attachments containing only squares in columns of distance at most one to the column where it is attached to the strip. Note that each of these valid attachments is contained in the attachment shown in Figure 23. In this figure, the square at the center of the bottom row is part of the strip, which possibly continues to the left and right. In particular the squares of the attachment directly left and right of this square can only be present if the attachment is attached to the square at the respective end of the strip.




Figure 23: Valid Class A attachments are part of this attachment.

Observe that tree-shaped polyominoes of bounding size $2 \times n$ for $n \leq 3$ cannot fold onto \mathcal{C} . Indeed, any foldable polyomino has to contain at least 6 squares and none of the 11 different nets of the cube of Figure 22 fits into a 2×3 rectangle. Thus the following theorem fully classifies the set of foldable tree-shaped polyominoes of bounding size $2 \times n$.

Theorem 7.1. *Let P be a tree-shaped polyomino of bounding size $2 \times n$ for some $n \geq 4$. Then P folds onto \mathcal{C} if and only if it is not contained in Class A.*

Proof. We start by showing that a polyomino P in Class A does not fold onto \mathcal{C} . Clearly in any folding of P into \mathcal{C} , the squares contained in the strip can cover at most 4 faces of \mathcal{C} . In particular all vertices at the same (say top) side of the strip are mapped onto the corners of the same face f of \mathcal{C} . Observe that each square of the possible attachments shares one of these vertices with the strip. This implies that none of the squares of the attachments can cover the face of \mathcal{C} opposite to f .

To prove the converse direction, let P be a polyomino of bounding size $2 \times n$ for some $n \geq 4$ which does not belong to Class A. In other words, it is not possible to find a strip in P such that P consists of this strip and only valid attachments of Class A as depicted in Figure 23. We conclude that there must be two adjacent squares lying in the same column i of the bounding box such that after removing the edge connecting these squares, both components are not valid attachments. Valid attachments are characterised by the fact that they do not visit any column at distance at least 2 from their root square. In particular we can assume that each of our two components visits at least one of the columns $i - 2$ and $i + 2$. In the case where one component visits column $i - 2$ and the other component visits column $i + 2$, it is not hard to see that P contains a cube net of type 3-3 as a minor and thus must be foldable. Otherwise assume that both components visit the same column at distance 2 from i , say $i + 2$. Then P contains two strips of length 3 on top of each other, which are connected only in column i . But by our assumption, the bounding box has length 4, so at least one of the strips is connected to one more square and thus must have length at least 4. But in this case we can fold this strip onto four faces of the cube and cover the other two faces with squares of the second strip in column i and $i + 2$. \square

7.2 Folding tree-shaped polyomino of bounding size $3 \times n$

Again, we start by defining an infinite class of polyominoes of bounding size $3 \times n$. Class B includes all polyominoes consisting of a $1 \times n$ strip with exactly one attachment of height two satisfying the following condition. Assume that the strip lies in the lower row of the bounding box (the other case is a horizontal reflection). Then all squares of the attachment lie in the two rows above the strip and in columns of distance at most one to the column where it is attached to the strip, see Figure 24(a). Moreover, if the square in the center of the top row is connected to its left and right neighbour, it must also be connected to its bottom neighbour. In particular any attachment containing the attachment in Figure 24(b) is forbidden because the upper central square is not connected to its bottom neighbour.

The following theorem characterises all foldable tree-shaped polyominoes of bounding size $3 \times n$ for $n \leq 5$, thus leaving only finitely many polyominoes to be treated separately.

Theorem 7.2. *Let P be a tree-shaped polyomino of bounding size $3 \times n$ for some $n \geq 5$. Then P folds onto \mathcal{C} if and only if it is not contained in Class B.*




Figure 24: Valid and non-valid attachments for Class B.

Proof. To prove the first direction, take any folding of a Class B polyomino P into \mathcal{C} . In this folding the squares of the strip of P can cover at most 4 faces of \mathcal{C} , leaving two opposite faces f and f' of \mathcal{C} unoccupied. Let s be the square of the attachment adjacent to the strip and observe that by the definition of valid attachments each square of the attachment contains one of the two vertices of s not contained in the strip. In the folding into \mathcal{C} , these two vertices are either both contained in f or in f' . We conclude that the respective second face cannot be covered by the attachment.

For the converse direction, assume that P does not fold onto \mathcal{C} . By definition of the bounding box, P contains a path-shaped sub-polyomino starting in the first column and ending in the last column of the bounding box. Assume that this path-shaped polyomino contains two adjacent squares in column i for some $3 \leq i \leq n-2$. Then clearly it contains a cube net of type 3-3 and thus folds onto \mathcal{C} . We conclude that P contains a strip of length at least $n-2$ which visits column 2 and column $n-1$. Let S be a strip in P of maximal length satisfying these conditions. Observe that each attachment attached to S in column i can visit only columns $i-1$ to $i+1$. Indeed, if this is not the case we can apply 180° folds to horizontal edges to reduce P to a polyomino of bounding size $2 \times n$ having an attachment not contained in these three columns, which is foldable by the previous theorem.

Assume first that S lies in the central row of the bounding box. If S has length at least 4, we fold all vertical edges of P not contained in the strip by 180° and the resulting structure contains a cube net of type 1-4-1. If there is an attachment visiting both the top and bottom row, we can apply 180° folds to each horizontal edge of all other attachments and end up in the previous case, where S has length 4. We are left with the case where S has length 3 and we have at least one attachment at each side of the strip. In this case column 1 must be visited by an attachment, which by the earlier observation must be attached in column 2. But then P contains a cube net of type 2-3-1 as a minor.

We conclude that S lies in an extremal row, say the bottom row and has one attachment of height 2 reaching the top row. It cannot have a second attachment, otherwise we could apply a 180° fold to the attachment of height 2 and obtain a polyomino with the same bounding size as P and S lying in the central row, which is foldable by the previous

case. Recall that if the attachment is attached in column i , it is contained in columns $i - 1$ to $i + 1$. Furthermore, it is not hard to see that no square of the attachment lies in the bottom row of the bounding box. Indeed, such a square has to be connected to its top neighbour. Figure 25(a) shows one possible polyomino satisfying these conditions. A 180° fold of the horizontal edge connecting the strip and the attachment (drawn red in the figure) yields a polyomino with the same bounding size as P having S in the central row, which is foldable by the earlier discussion.

Finally, assume for a contradiction that the square s' in the upper row of column i is connected to both its left and right neighbour, but not to its bottom neighbour. Then the shortest path-shaped sub-polyomino of P connecting s' with the strip must either visit column $i - 1$ or column $i + 1$, the first case is depicted in Figure 25(b). But then a 180° fold applied to the vertical edge at the respective side of s' yields an attachment which is not contained in columns $i - 1$ to $i + 1$ any longer, thus by the previous discussion this attachment guarantees foldability, contradicting our assumption.




Figure 25: Attachments guaranteeing foldability

□

To provide a rough idea, there are (up to symmetry) 124 tree-shaped polyominoes of bounding size 3×3 and 3942 tree-shaped polyominoes of bounding size 3×4 . Unfortunately, there seems to be no ‘nice’ characterisation for those. To keep things as simple as possible, we provide in Figures 26 and 27 all minimal foldable polyominoes (up to symmetry) of the respective bounding sizes. Here minimal foldable means that removing any leaf square either reduces the bounding size or destroys foldability. These sets were determined by computer.




Figure 26: The 7 minimal foldable tree-shaped polyominoes of bounding size 3×3 .

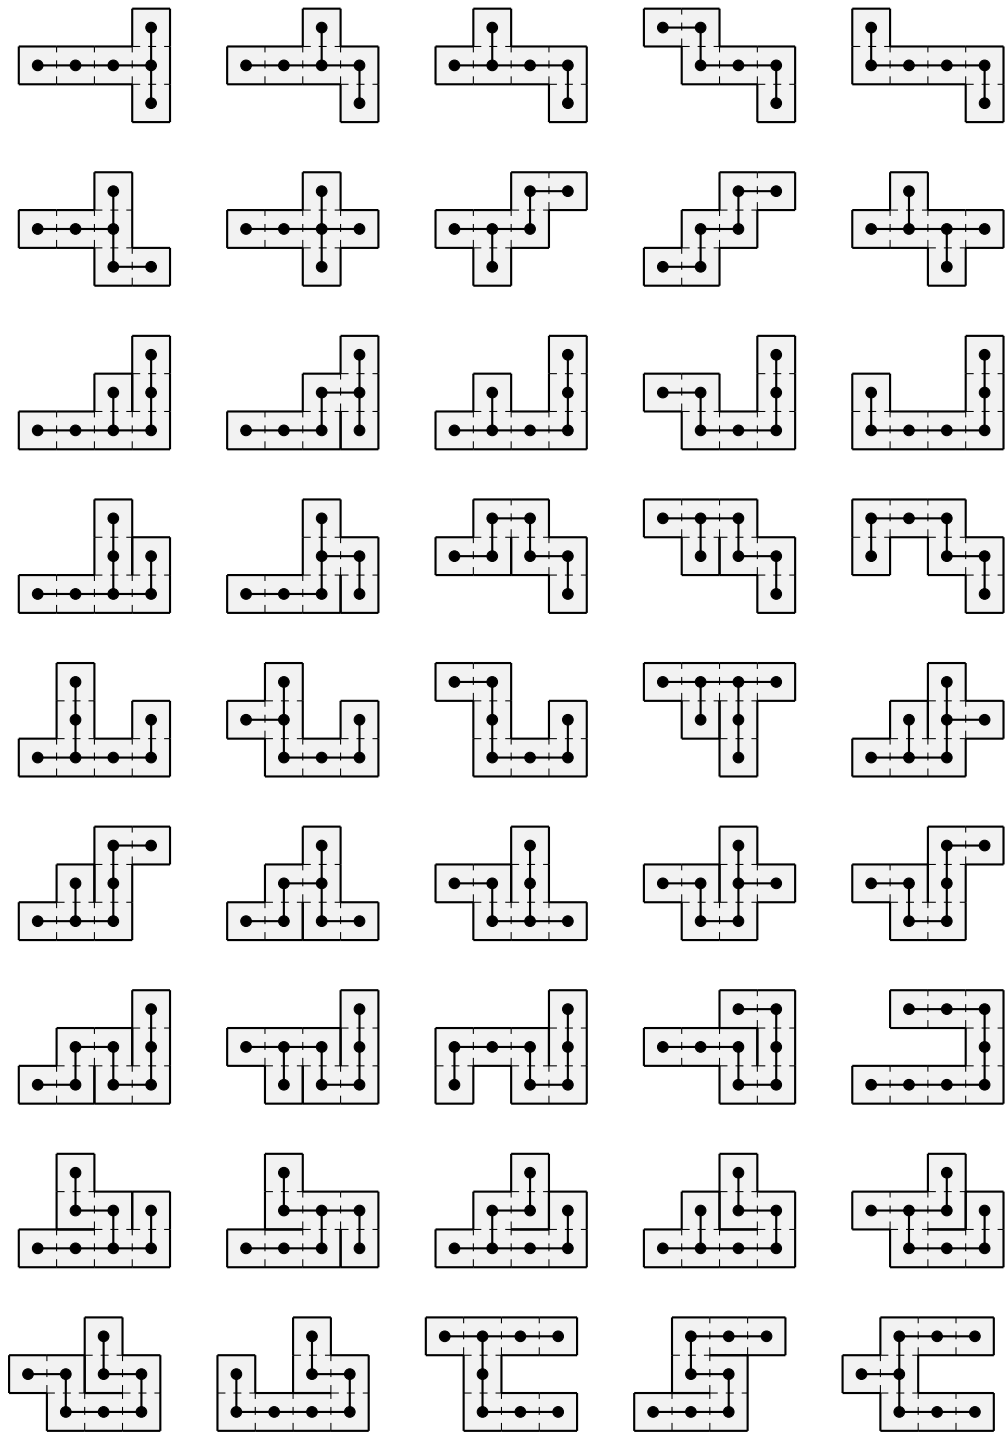


Figure 27: The 45 minimal foldable tree-shaped polyominoes of bounding size 3×4 .

7.3 Folding tree-shaped polyomino of bounding size at least 4×4

The goal of this section is to show that the polyomino P_W depicted in Figure 28 is the only tree-shaped polyomino of bounding size at least 4×4 which does not fold onto \mathcal{C} .




Figure 28: Polyomino of bounding size 4×4 which does not fold onto \mathcal{C} .

Lemma 7.3. P_W does not fold onto \mathcal{C} .

Proof. Consider the square in the center of the polyomino P_W . If one of the two edges connecting it to the rest of P_W is folded by 180° , then the resulting shape has the form of the letter ‘L’. It consists of 6 squares and is not a cube net, so it cannot be folded onto \mathcal{C} . Thus we may assume that both boundary edges of the central square are folded by 90° . But then no folding of P_W onto \mathcal{C} can cover the face directly opposite of the face covered by the central square. \square

Lemma 7.4. Let P be the polyomino P_W with one additional leaf-square added. Then P folds onto \mathcal{C} .

Proof. By symmetry of P_W , we may assume without loss of generality that the additional leaf shares a vertical edge with P_W . All possible locations for this additional leaf are shown as red squares in Figure 29, where some sub-figures contain multiple locations as they are treated in the same way. It is not hard to check that each of the polyominoes is foldable. After applying 180° folds along the coloured edges of each polyomino, we see that each of the resulting shapes contains a 4×1 sub-polyomino, either with leaves attached on both sides or with a 3×1 sub-polyomino attached on one of its sides. In both cases folding everything by 90° yields the cube \mathcal{C} .




Figure 29: Possible positions of the additional square and possible 180° folds to reduce the polyomino to a foldable shape.

\square

Theorem 7.5. *Let P be a tree-shaped polyomino of bounding size 4×4 which is not the polyomino P_W depicted in Figure 28. Then P folds onto \mathcal{C} .*

Proof. Given a polyomino P of bounding size 4×4 , we start by applying leaf-folds in arbitrary order to reduce the polyomino as long as possible without reducing the bounding size. Clearly P folds onto \mathcal{C} whenever the reduced shape folds onto \mathcal{C} . Additionally, by Lemma 7.4 the polyomino P folds onto \mathcal{C} if the reduced shape equals P_W and we folded away at least one leaf during the process.

Thus it is enough to prove the statement of the theorem for *reduced* polyominoes, that are polyominoes, where any leaf-fold reduces the bounding size.

Let P be a reduced polyomino of bounding size 4×4 which does not fold onto \mathcal{C} . We show that P must be equal to P_W . A leaf square of P is said to *point towards* the top/left/bottom/right boundary, if it lies in the topmost row/leftmost column/bottommost row/rightmost column of the 4×4 bounding box and its neighbour does not lie in this row or column. The following two observations follow easily from the reducedness of P .

Observation 1. Every leaf square of P points towards a boundary.

Observation 2. If a leaf square points towards a boundary, then the respective row or column belonging to this boundary does not contain any other squares of P .

We organize our proof as a sequence of claims.

Claim 1. The polyomino P does not contain a 4×1 strip as a sub-polyomino.

Indeed, if it contains a 4×1 strip not lying in an extremal row or column of the bounding box, this 4×1 strip has two neighbours on different sides and thus contains a cube net of type 1-4-1. If it contains a 4×1 strip lying in an extremal row, say the bottom row, then there must be a sequence of squares connected to it and reaching the top row. We can now use 180° folds along vertical edges to reduce this sequence to a 3×1 strip attached at the 4×1 strip, which can be folded onto \mathcal{C} with only 90° folds.

Claim 2. The polyomino P must be path-shaped.

Assume for a contradiction that P is not path-shaped. Then P contains a square having at least 3 neighbouring squares and thus has a sub-polyomino with a T-shape. This square will henceforth be called the *central square*. Assume without loss of generality that the central square has a neighbour to the left, at the bottom and to the right and that the left square lies in the leftmost column of the bounding box of P , see Figure 30(a), otherwise rotate or reflect P . Removing the central square from P splits P into 3 connected components, which will be called left, bottom and right arm of P . Clearly each arm has to contain at least one leaf square of P , which must point towards one of the four boundaries.

Assume first that the right arm contains a square of the rightmost column of the bounding box. Let s denote the square of the right arm lying in the rightmost column which is closest to the central square with respect to the graph metric in the dual of P . If s lies above the row containing the central square, then it contains a 2-3-1 cube net as a minor. Otherwise s lies below the row containing the central square. Note that in this case, the bottom arm has to contain a leaf square pointing toward the bottom boundary, thus we can apply 180° folds along horizontal edges of the right arm to reduce P to a polyomino containing a 4×1 strip without reducing the bounding size, which is foldable. We conclude that the right arm cannot visit the rightmost column and thus contains only a single leaf square pointing towards the top boundary. In particular the square to the right of the central square is connected to its upper neighbour.

We are left with the case where the bottom arm visits the rightmost boundary. Then the left arm cannot visit the bottom row, otherwise we could apply 180° folds along horizontal edges of the bottom arm to reduce P to a polyomino containing a cube net of type 1-4-1, which makes P foldable. Thus the bottom arm also has to visit the bottom row. If the central square lies in the second row from top, we apply 180° folds along vertical edges of the bottom arm to see that the cube net of type 2-3-1 shown in Figure 30(b) is a minor of P . Otherwise the central square lies in the third row from the top and we apply 180° folds along vertical edges of the right arm to see that the cube net of type 1-4-1 shown in Figure 30(c) is a minor of P . This proves the claim that P is path-shaped.

Claim 3. The polyomino P contains a 3×1 strip as a sub-polyomino.

Indeed, if P does not contain a 3×1 strip, it is not hard to check that P either contains a cube net of type 2-2-2 or the shape in Figure 30(d), both of which are foldable.

We finally end up in the case where P is a path and contains a 3×1 strip. Assume again without loss of generality that the 3×1 strip lies horizontally in the bounding box of P and that the leftmost square lies in the leftmost column. Furthermore, as the rightmost square cannot be a leaf square by Observation 1 we may assume that it is connected to its upper neighbour. Again the square in the center of the 3×1 strip will be called central square, and the two components of P after removing this square are called left and right arm of P . Clearly the left arm cannot visit the rightmost column, otherwise P contains the foldable polyomino in Figure 30(e) or its reflection as a minor. Thus the right arm has to visit the right column.

Consider the square to the left of the central square. If it is not a leaf square it is either connected to its lower or upper neighbour. In the first case we can apply 180° folds along horizontal edges of the right arm to see that P contains a cube net of type 2-3-1 as a minor. In the other case the unique leaf square of P contained in the left arm must point towards the top boundary. By Observation 2 the central square cannot lie in the second row from the top, thus P can be reduced to the polyomino in Figure 30(f)

by only 180° folds, which with the additional fold shown in the figure reduces to a cube net of type 2-3-1.

We are left with the case where the square to the left of the central square is a leaf square. If the right arm touches both the top and the bottom boundary, we can use 180° folds along vertical edges to reduce it to a polyomino containing a vertical 4×1 strip, which is foldable. Thus the central square lies in the bottom row of the boundary box of P . By Observation 1 the square at the top right from the central square cannot be a leaf square. If it is not connected to its right neighbour, P contains one of the polyominoes in Figure 30(g) and Figure 30(h) as a minor, which both contain a cube net and thus are foldable. Otherwise P either contains the polyomino in Figure 30(i) as a minor, which is foldable because one additional 180° fold yields the polyomino in Figure 30(h), or it is equal to the polyomino P_W .




Figure 30: Polyominoes encountered in the proof of Theorem 7.5.

□

Theorem 7.6. *Let P be a tree-shaped polyomino of bounding size at least 4×5 . Then P folds onto \mathcal{C} .*

Proof. For a given tree-shaped polyomino P of bounding size at least 4×5 we consider the following 3-step approach.

Similar to the proof of the previous theorem, we start by applying leaf-folds in arbitrary order such that the resulting flat tree-shaped polyomino has bounding size at least 4×4 . If no such leaf-fold is possible any more, we call the polyomino *reduced*. Clearly any reduced polyomino P' is still tree-shaped and has bounding size $4 \times k$ for some $k \geq 4$. If P' is equal to the polyomino P_W in Figure 28, we undo the last leaf-fold and see that the resulting polyomino and thus also the original polyomino P is foldable by Lemma 7.4.

Thus it is enough to show that any reduced polyomino which is not equal to P_W can be folded onto the cube. We use induction on $k \geq 4$.

For the induction basis assume first that the resulting tree-shaped polyomino P' has bounding size 4×4 . If this shape is not the polyomino P_W depicted in Figure 28 we are done by Theorem 7.5 and can obtain a proper cube folding.

Assume now that the reduced tree-shaped polyomino P' has bounding size $4 \times k$ for some $k > 4$ (the case $k \times 4$ is symmetric, and bounding boxes where both dimensions are larger than 4 for sure have a leaf-fold.) We make a sequence of useful observations.

Observation 1. In the top (bottom) row of the bounding box there is only one leaf square and no other polyomino squares. Otherwise we could make a leaf-fold without decreasing the dimension of the bounding box below 4.

Observation 2. Except the two leaf squares on top and bottom, there is no other leaf square. Otherwise we could leaf-fold it without decreasing the dimension of the bounding box below 4.

Observation 3. Since there are only two leaf squares the polyomino P' is path-shaped.

Observation 4. If there is no leaf square in the leftmost column, then the two squares in the left most column are in rows 2 and 3 and are connected in a C-shaped way together and to the right.

Observation 5. If there is a leaf square in the leftmost column, then this is connected in an L-shaped way to the right. Note that there are several possibilities. The L could be 2 or 3 rows high, having the leaf square on top or bottom, or there could even be both leaf squares in the leftmost columns with two L-shaped connections of height 2 each.

In all cases we can again fold the leftmost column with a 180° fold into another path-shaped polyomino P'' , reducing the bounding size to $4 \times (k - 1)$. See Figure 31 for the different possible foldings. Note that for the third polyomino in the figure, we first apply the 180° fold drawn in the figure to reduce it to the fourth polyomino in the figure.




Figure 31: Possible shapes of the two leftmost columns.

It is not hard to check that in all cases P'' is reduced, thus we can apply the induction hypothesis and see that P' and thus also P can be folded onto the cube. \square

8 Folding simply connected polyominoes

A polyomino is called simply connected if it does not contain any holes. Clearly this includes all tree-shaped polyominoes. However, it appears to be much harder to find a full characterisation for the set of all foldable simply connected polyominoes. Lemma 5.2 leads to a first naive approach. Clearly whenever a polyomino is simply connected but not tree-like, it has to contain at least one rectangular sub-polyomino of size at least 2×2 . By the mentioned lemma each of these large rectangular sub-polyominoes has to be folded by 180° to either horizontal or vertical strips of width 1. The result then is in some sense tree-shaped. However, some squares might have multiple layers, which might allow more possible foldings than the respective normal tree-shaped polyomino with the same shape. The following example should make this clear.

Example 8.1. Consider the leftmost polyomino in Figure 32. It contains a rectangular sub-polyomino of size 4×2 . By the previous discussion either all horizontal or all vertical edges must be folded by 180° . It clearly does not make sense to fold the vertical edges as this would collapse the polyomino to a 1×6 strip, which obviously cannot be folded onto \mathcal{C} . Horizontally folding along the red line yields the shape in the center of the figure. When considering this a normal single layered polyomino, it is not foldable. However we can use the two layers and just fold one of the attached arms along the red edge to obtain the shape at the right, which contains the standard cube net and thus is foldable.




Figure 32: Multiple layers can make non-foldable shapes foldable.

We have seen in Section 7 that the bounding size of non-foldable tree-shaped polyominoes is bounded from above, in the sense that they either fit into a bounding box of size $3 \times n$ or of size 4×4 . Clearly a similar upper bound for the bounding box cannot be given for simply connected polyominoes because a simply connected rectangular polyomino can never fold onto \mathcal{C} . However, the upcoming result shows that in fact a non-foldable polyomino cannot differ a lot from a rectangular polyomino.

Let P be a simply connected polyomino and let B be its bounding box. For the upcoming arguments we orient both the boundary ∂P of P and the bounding box B in clockwise direction. By definition ∂P touches each of the four sides B . Let v be a vertex of ∂P which does not lie on B . We denote by $\pi_1(v)$ and $\pi_2(v)$ the unique paths obtained by starting at v and following ∂P in negative and positive direction until reaching B , respectively. Let c_1 and c_2 be the not necessarily different closest corners of B in positive direction from the endpoint of $\pi_1(v)$ and in negative direction from the endpoint of $\pi_2(v)$, respectively. Then we say that c_1 and c_2 are *valid corners* for v .

This definition should become clear when looking at the example polyomino in Figure 33. Note that $p_1(v)$ and $p_2(v)$ both end at the left boundary of the boundary box, thus the top left and the bottom left corner are valid for v . On the other hand, $p_1(w)$ ends at the right boundary, while $p_2(w)$ ends at the upper boundary. Thus by definition only the corner between these two boundaries, that is the top right corner, is valid for w .




Figure 33: Valid corners for two vertices on the boundary of P

Theorem 8.2. *Let P be a simply connected polyomino containing a boundary point whose horizontal and vertical distance to one of its valid corners is at least 3. Then P folds onto \mathcal{C} .*

Proof. Assume without loss of generality that the top left corner of the bounding box of P is valid for the point v on the boundary of P and that v has horizontal and vertical distance at least 3 to this corner. Let s denote the unit square having v at its top left corner, which does not have to be present in P . Finally let $\pi_1(v)$ and $\pi_2(v)$ be the unique sub-paths of ∂P connecting v with the next points where P touches the bounding box as defined above.

We fold P in several steps. To assist the reader, in Figure 34 these steps are applied to the example polyomino of Figure 33. Start by rolling up all rows below the row

containing s . Next roll to the left all columns to the right of the column containing s . Then s is the bottom right square of the resulting shape P' . Additionally, while $\pi_1(v)$ and $\pi_2(v)$ might have been deformed by this transformation in the case they visit rows below or columns to the right of s , they still connect the top left corner of s with the respective boundaries. In particular, after removing s , the shape P' has two connected components: the *left component* containing the square to the left of s and the *upper component* containing the square above s . Note that because the top left corner is valid for v , the left component touches the left boundary and the upper component touches the top boundary of the bounding box of P .

We roll down all squares of the left component lying in a row above s and roll to the right all squares of the upper component lying in columns left of s . Clearly all squares of the resulting shape either lie in the same row or the same column as s , thus the shape has the form of the letter 'L'. Furthermore by construction the left component still touches the left boundary and the upper component still touches the upper boundary, thus the two components must have length at least 3. We conclude that it folds onto \mathcal{C} , and thus the same is true for P .




Figure 34: Folding strategy in the proof of Theorem 8.2

□

Observe that by definition every point of the boundary of a polyomino P not lying inside the boundary box must have some valid corner. Thus we immediately see that whenever the boundary of P contains a point at horizontal and vertical distance at least 3 from each of the four corners of the bounding box, P folds onto \mathcal{C} . We obtain the following corollary.

Corollary 8.3. *Let P be a simply connected polyomino of bounding size $m \times n$ which does not fold onto \mathcal{C} . Then it contains a rectangular sub-polyomino of size $(m - 4) \times (n - 4)$ at distance 2 from each boundary.*

As mentioned before, this result has a close relation to the bounded size of non-foldable tree-shaped polyominoes. In fact, it directly implies that any non-foldable simply connected polyomino with a bounding box of size at least 6×6 contains a 2×2 rectangular sub-polyomino, or in other words, that every tree-shaped polyomino with a bounding box of size at least 6×6 folds onto \mathcal{C} .

Conclusion and open problems

For polyominoes with holes we showed in Section 5 that any rectangular polyomino with only one simple hole does not fold into a cube. Moreover, we provided a complete characterisation when a rectangular polyomino with two or more unit square holes (but no other holes) can be folded into a cube. In [1] several combinations of two simple holes were given that allow the polyomino to fold into a cube, and also examples for simple hole(s) were identified that cannot be folded into a unit cube. Restricting considerations to rectangular polyominoes we ask the following question.

Open Problem 8.1. *For which combinations of (two or more) simple holes can a rectangular polyomino be folded into a cube?*

Concerning tree-shaped polyominoes we gave a full characterisation which of them can be folded onto the cube (Section 7). This settles the case of tree-shaped polyominoes.

Finally, in Section 8 we provided a sufficient condition when a simply-connected polyomino can be folded to a cube. Here a complete classification is still open.

Open Problem 8.2. *Provide a complete characterisation when a simply-connected polyomino can be folded into a cube.*

References

- [1] O. Aichholzer, H. A. Akitaya, K. C. Cheung, E. D. Demaine, M. L. Demaine, S. P. Fekete, L. Kleist, I. Kostitsyna, M. Löffler, Z. Masárová, K. Mundilova, and C. Schmidt. Folding polyominoes with holes into a cube. *Computational Geometry*, 93:101700, 2021.
- [2] O. Aichholzer, M. Biro, E. D. Demaine, M. L. Demaine, D. Eppstein, S. P. Fekete, A. Hesterberg, I. Kostitsyna, and C. Schmidt. Folding polyominoes into (poly)cubes. *International Journal of Computational Geometry & Applications*, 28(03):197–226, 2018.
- [3] N. Beluhov. Cube folding. <https://nbpuzzles.wordpress.com/2014/06/08/cube-folding>, 2014.

- [4] K. Y. Czajkowski, E. D. Demaine, M. L. Demaine, K. Eppling, R. Kraft, K. Mundilova, and L. Smith. Folding small polyominoes into a unit cube. In *CCCG*, pages 95–100, 2020.
- [5] M. Gardner. *Wheels, Life and Other Mathematical Amusements*. W.H. Freeman and Company, 1983. University of Chicago Lecture Notes prepared with the assistance of J. L. Shaneson and J. Lees.
- [6] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999.
- [7] J. F. P. Hudson. *Piecewise linear topology*. W. A. Benjamin, Inc., New York-Amsterdam, 1969. University of Chicago Lecture Notes prepared with the assistance of J. L. Shaneson and J. Lees.
- [8] E. E. Moise. *Geometric topology in dimensions 2 and 3*, volume Vol. 47. Springer-Verlag, New York-Heidelberg, 1977.