

# Tiling with Three Polygons Is Undecidable

Erik D. Demaine  

Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology,  
Cambridge, MA, USA

Stefan Langerman  

Computer Science Department, Université libre de Bruxelles, Belgium

---

## Abstract

We prove that the following problem is co-RE-complete and thus undecidable: given three simple polygons, is there a tiling of the plane where every tile is an isometry of one of the three polygons (either allowing or forbidding reflections)? This result improves on the best previous construction which requires five polygons.

**2012 ACM Subject Classification** Theory of computation → Computational geometry; Theory of computation → Problems, reductions and completeness

**Keywords and phrases** plane tilings, polygons, undecidability, co-RE

**Digital Object Identifier** 10.4230/LIPIcs.SoCG.2025.39

**Related Version** *Full Version*: <https://arxiv.org/pdf/2409.11582>

## Supplementary Material

*Software (Implementation)*: <https://github.com/edemaine/three-tiles> [4]

archived at `swh:1:dir:aef5f642bf8a40c0cecf56bf62b9e089effb265`

**Acknowledgements** This research was initiated at the Second Adriatic Workshop on Graphs and Probability, and continued at the Eleventh Annual Workshop on Geometry and Graphs. The authors thank all participants of both workshops for providing a stimulating research environment. They also thank Joshua Blömer for spotting a bug in the implementation [4] and Figure 6.

“Three Rings for the Elven-kings under the sky, . . .”  
– J. R. R. Tolkien, *The Lord of the Rings*, epigraph

## 1 Introduction

A *tiling* of the plane [9] is a covering of the plane by nonoverlapping polygons called *tiles*, isometric copies of one or more geometric shapes called *prototiles*, without gaps or overlaps. In this paper, we study the most fundamental computational problem about tilings:

▶ **Problem 1** (Tiling). *Given one or more prototiles, can they tile the plane?*

The tiling problem is *undecidable* – solved by no finite algorithm. Golomb [6] was first to prove this result, by building  $n$  polyominoes that simulate  $n$  *Wang tiles* [15] – unit squares with edge colors that must match – by adding color-specific bumps and dents to each edge. Four years earlier, Berger [1] proved that tiling with Wang tiles is undecidable (disproving Wang’s original conjecture [15]) by showing how they can simulate a Turing machine. Robinson [13] later simplified Berger’s proof. The worst-case number  $n$  of tiles (Wang or polyomino) is  $\Theta(|Q| \cdot |\Sigma|)$ , where  $|Q|$  and  $|\Sigma|$  are the number of states and symbols in the simulated Turing machine, respectively.

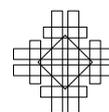
**Constant Number of Prototiles.** The first constant and previously best upper bound on the number of prototiles required to make the tiling problem undecidable is 5, as proved by Ollinger fifteen years ago [11]. Our main result, proved in Section 3, is an improvement of this upper bound to 3:



© Erik D. Demaine and Stefan Langerman;  
licensed under Creative Commons License CC-BY 4.0  
41st International Symposium on Computational Geometry (SoCG 2025).  
Editors: Oswin Aichholzer and Haitao Wang; Article No. 39; pp. 39:1–39:16  
Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



## 39:2 Tiling with Three Polygons Is Undecidable

► **Theorem 1.1.** *Given three simple-polygon prototiles, determining whether they tile the plane is undecidable.*

It remains open whether tiling with one or two given prototiles is decidable. *Periodic* tilings (tilings with two translational symmetries) can be found algorithmically by enumerating fundamental domains, as we show in the full paper. (Surprisingly, this intuitive fact does not seem to have been explicitly proved before, except in special settings like Wang tiles [15].) Thus a necessary condition for undecidability is the existence of prototile(s) with only aperiodic tilings. Recently, Smith, Myers, Kaplan, and Goodman-Strauss [14] found a single prototile with this property, so there are no obvious obstacles to undecidability.

**Tiling by Translation.** Our construction relies on rotation of the prototiles (but works independent of whether we allow reflections). If we restrict to tiling by translation only, then Ollinger’s construction can be modified to use 11 prototiles, by adding some rotations of the five polyominoes [11]. This upper bound was improved to 10 by Yang [17] and to 8 by Yang and Zhang [18]. All of these constructions use polyominoes. In higher dimensions, Yang and Zhang [19] improved the upper bound to five polycube prototiles in 3D, and four polyhypercube prototiles in 4D.

The tiling-by-translation problem also has a lower bound of 2 for undecidability: any single polygon that tiles the plane by translation can do so by periodic (even isohedral) tiling [5]. This result also holds for disconnected polyominoes [2]. If we generalize to tiling a specified periodic subset of  $d$ -dimensional space, where  $d$  is part of the input, then Greenfeld and Tao [8] recently proved tiling to be undecidable with a single disconnected polyhypercube.

**Periodic Target.** We show that Greenfeld and Tao’s generalization to tiling a specified periodic subset [8] changes the best known results also for undecidability of tiling the plane. Our 3-polygon construction and Ollinger’s 5-polyomino construction [11], and Yang and Zhang’s 8-polyomino translation-only construction [18] all have one prototile (our shurikens, and their jaws) that appear periodically in any tiling of the plane. Thus, if we remove that pattern from the target, we obtain a periodic subset of the plane which can be tiled using a reduced number of prototiles of 2, 4, and 7, respectively. In particular, we prove

► **Corollary 1.2.** *Given two simple-polygon prototiles, and given a periodic subset of the plane, determining whether the two prototiles tile the periodic subset is undecidable.*

**Logical Undecidability.** Algorithmic undecidability implies *logical undecidability* (as explained in [7] in the context of tilings). In particular, our result implies that there are three polygon prototiles that cannot be proved or disproved to tile the plane, for any fixed set of axioms (e.g., ZFC). Otherwise, we would obtain a finite algorithm to decide tileability, by enumerating all proofs.

► **Corollary 1.3.** *For any fixed set of axioms, there are three fixed simple-polygon prototiles such that both “these prototiles tile the plane” and “these prototiles do not tile the plane” have no proof.*

**Tiling Completion.** Undecidability of tiling requires the set of prototiles to depend on the Turing machine simulation. To obtain undecidability with a fixed set of prototiles, we can generalize the tiling problem as follows [13]:

► **Problem 2 (Tiling Completion).** *Given one or more prototiles, and given some already placed tiles, can this placement be extended to a tiling of the plane?*

Robinson [13] gave the first result on this problem: a set of 36 prototiles (Wang tiles or polygons) for which tiling completion is undecidable. This result applies the general Turing machine simulation to Minsky’s 4-symbol 7-state universal Turing machine, so only a finite number of tiles need to be preplaced to represent the Turing machine to simulate. Likely this result could be improved using newer smaller universal Turing machines [16]. If we allow for (countably) infinitely many tiles to be preplaced, we can use semi-universal Turing machines and simulate Rule 110, enabling undecidability with just six supertiles (Wang tile or polygons) [20]. Our main result reduces this upper bound to 3, in the stronger model of finitely many preplaced tiles:

► **Corollary 1.4.** *There are three fixed simple-polygon prototiles such that, given a finite set of already placed tiles, determining whether this placement can be extended to a tiling of the plane is undecidable.*

**Co-RE-completeness.** While past results on tiling and tiling completion have focused on undecidability, all such proofs actually show **co-RE-hardness**: the simulated Turing machine halts if and only if the prototiles fail to tile. Co-RE-hardness is a more precise statement than undecidability, so we use that phrasing here. But it raises the question: are tiling and tiling completion in co-RE? Surprisingly, this question does not seem to have been solved (or even asked) in the literature before. In Section 4, we prove the answer is “yes”:

► **Theorem 1.5.** *Given a finite set of polygon prototiles, and given a (possibly empty) connected set of already placed tiles, determining whether this placement can be extended to a tiling of the plane is in co-RE.*

This result holds in a very general model for polygons: the angles and edge lengths can be represented as **computable** numbers (meaning that a Turing machine can output the first  $n$  bits, given  $n$ ). Our three-polygon construction uses a more restricted model, where the angles are rational multiples of  $\pi$  and the edge lengths are constant-size radical expressions, showing the problem to be co-RE-complete for every model in between.

► **Corollary 1.6** (Stronger form of Theorem 1.1). *Given three simple-polygon prototiles, where the edge lengths and angles in degrees are specified by computable numbers or by constant-size radical expressions, determining whether they tile the plane is co-RE-complete.*

## 2 Wang Tiling: Signed and Unsigned

We reduce from Wang tiling, which is known to be undecidable. A **Wang tile** is a square with a **glue** on each edge. Classically, Wang tiles are **unsigned**, meaning that glues match if they are equal, and **translation-only**, meaning they have a specified orientation of which edge is north, east, south, and west. In 1966, Berger proved unsigned Wang tiling undecidable:

► **Theorem 2.1** ([1]). *Given a set of Wang tiles, it is co-RE-hard to determine whether they tile the plane by translation only, matching glues of equal value.*

Our first reduction converts unsigned Wang tiling to a variant that is **signed**, meaning every glue has a sign (+ or  $-$ ) and value, and glues match if they have opposite sign and equal value, and **free**, meaning the tile can be rotated and/or reflected. This result was proved by Robinson [13, p. 179]. Figure 1 sketches the reduction; the full paper gives the details for completeness.

### 39:4 Tiling with Three Polygons Is Undecidable

► **Lemma 2.2** ([13]). *Given a set of unsigned translation-only Wang tiles  $T$ , we can construct a set of signed free Wang tiles  $T'$  that has the same tilings as  $T$  up to global isometry, allowing or forbidding reflection.*



(a) 11 unsigned translation-only Wang tiles  $T$  that tile only aperiodically, the minimum possible [10].



(b) Equivalent 11 signed free Wang tiles  $T'$ . Bumps/dents denote signs.

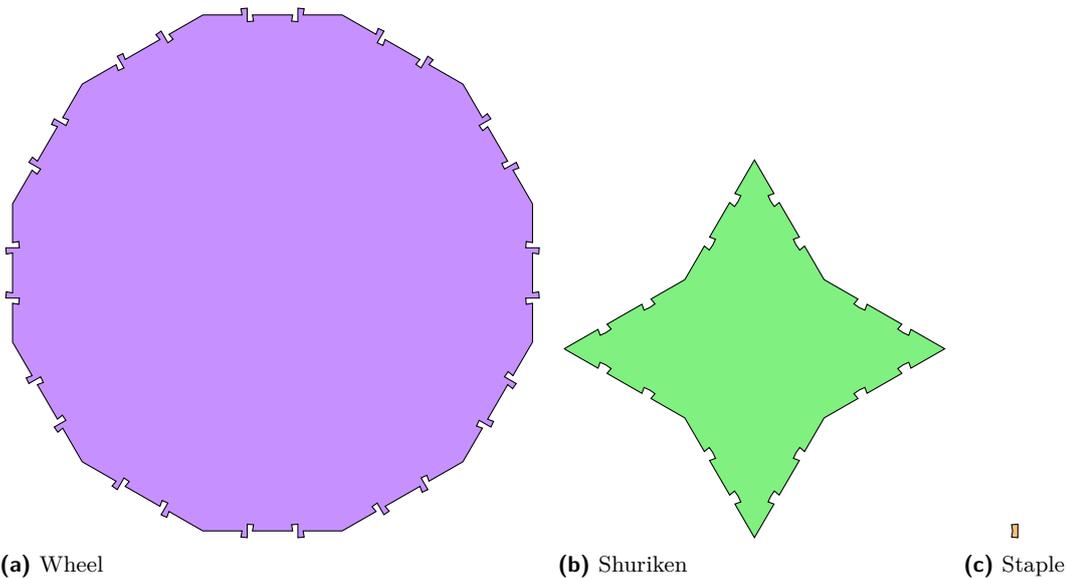
■ **Figure 1** Example of converting unsigned translation-only Wang tiles to signed free Wang tiles.

Henceforth when we say “Wang tiles” we mean signed free Wang tiles.

### 3 Three Tiles That Simulate $n$ Signed Wang Tiles

We implement any set of  $n$  Wang tiles with three tiles, illustrated in Figure 2:

1. the *wheel* which encodes all of the Wang tiles,
2. the *staple* which covers the unused Wang tiles of each wheel, and
3. the *shuriken* which fills in the remaining gaps.



(a) Wheel

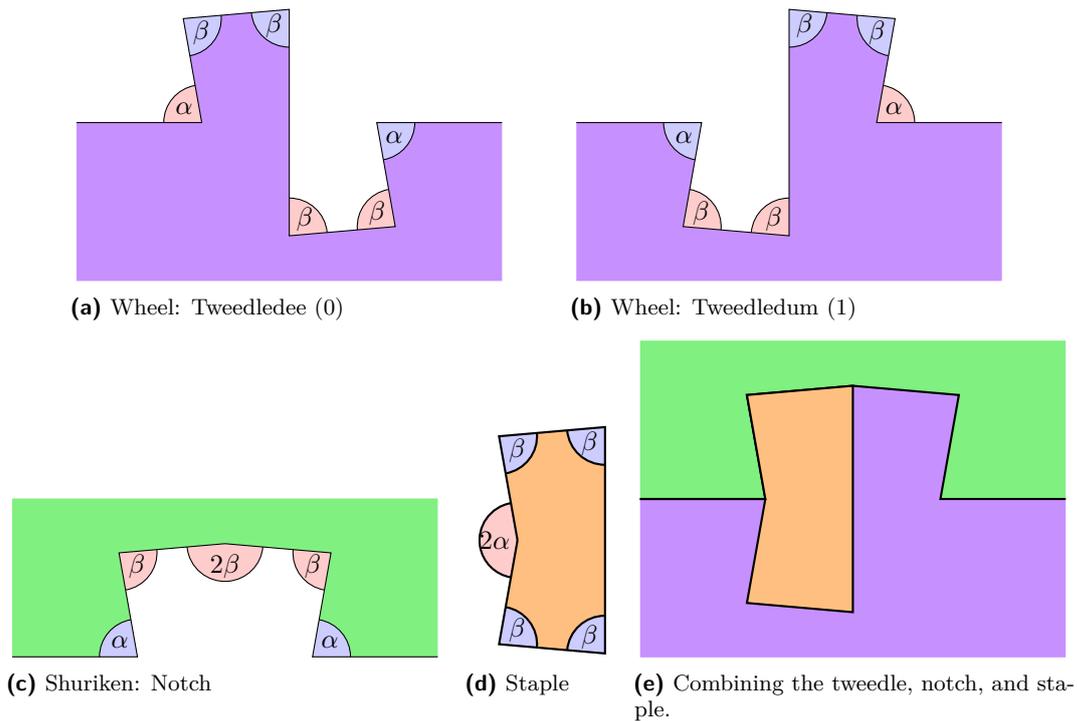
(b) Shuriken

(c) Staple

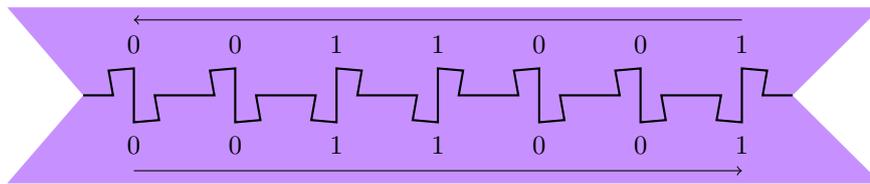
■ **Figure 2** The three tiles in our construction, to scale; Figure 3 shows zoomed details of the construction. The wheel is just an example; it depends on the  $n$  Wang tiles being simulated. The shuriken depends (only) on  $n$ .

#### 3.1 Construction and Intended Tiling

Suppose we are given a set of  $n$  Wang tiles, where the  $i$ th tile ( $1 \leq i \leq n$ ) has signed glues  $n_i, e_i, s_i, w_i$  on its north, east, south, and west edges respectively. Assume  $n$  is an odd integer  $\geq 5$  by possibly adding duplicate tiles.



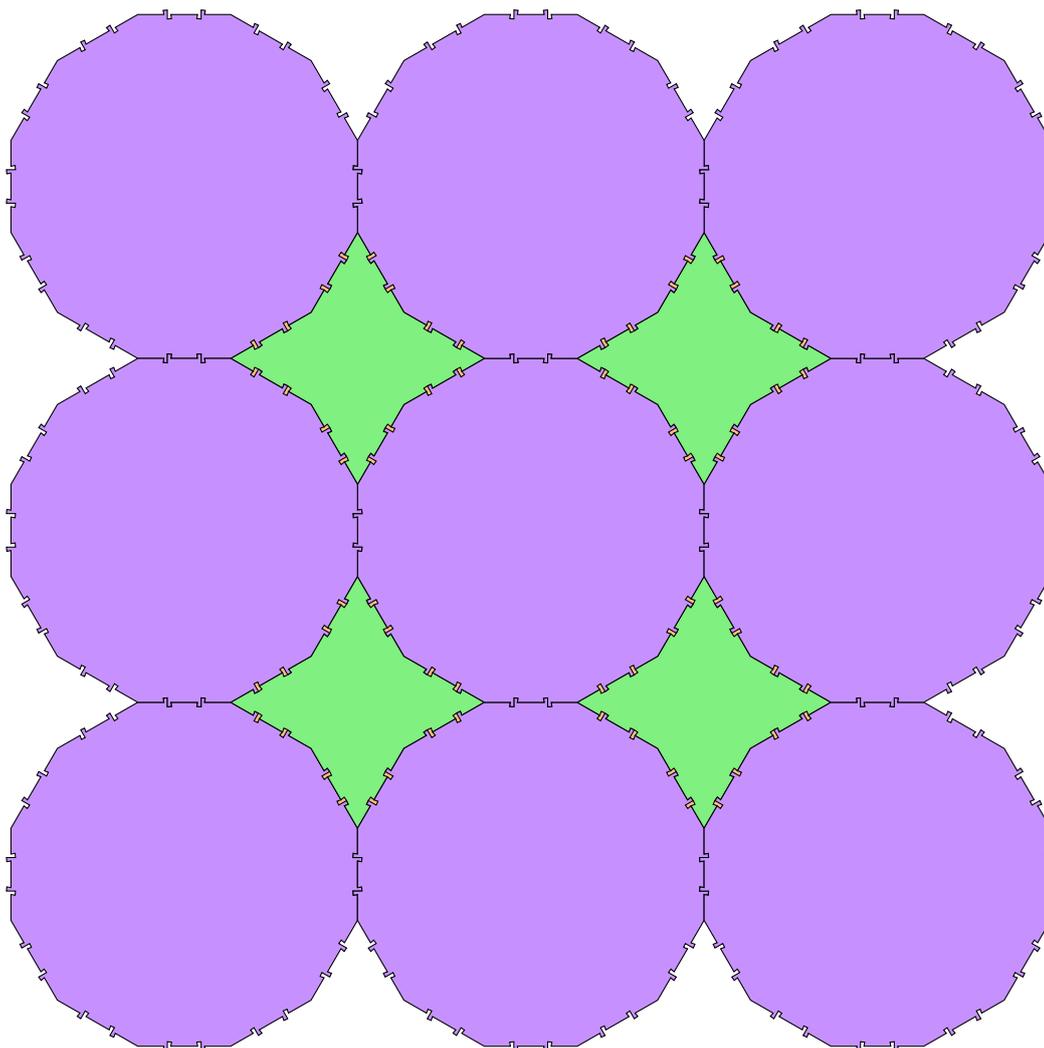
■ **Figure 3** Zoomed views of portions of the three tiles in our construction (15× scale vs. Figure 2).



■ **Figure 4** Matching a glue (top) and its negative (bottom) between two wheels.

The *wheel* is a regular  $4n$ -gon with each edge adorned by bumps and notches representing the  $4n$  glues. For tile  $i$ , the glues  $n_i, e_i, s_i, w_i$  adorn sides  $i, n + i, 2n + i, 3n + i$  of the  $4n$ -gon, respectively. To encode a glue, we encode its value in binary using  $b = O(\log n)$  bits, prepend a 00 at the beginning, and append 01 at the end. For negative glues, we reverse the order of the bits, which puts a 10 at the beginning and a 00 at the end. Then we represent each bit with a *tweedledee* (0) or *tweedledum* (1) gadget, which are rotationally symmetric zig-zags shown in Figures 3a and 3b. Both follow the sequence of angles  $\alpha, \beta, \beta, \beta, \alpha$  (where  $\alpha$  and  $\beta$  are defined below); for tweedledee, this sequence measures defect, angle, defect, defect, angle, respectively; while for tweedledum, this measures the opposite (angle, defect, defect, angle, angle, defect).<sup>1</sup> As shown in Figure 4, two adjacent glues match exactly if and only if they have the same value and opposite sign (where the opposite sign is enforced by the 00 and 01 at either end). This representation also ensures that reflecting a wheel will produce reflected glues that do not match unreflected glues: a reflection causes the bits of a

<sup>1</sup> It is also possible to use two different convex angles  $\beta_1, \beta_2$  in place of each repetition  $\beta, \beta$ , but the notation is messier, so we opt for this simpler construction.



■ **Figure 5** Example tiling with the wheel, shuriken, and staple.

glue to be reversed and negated, so the reflection of a positive glue starts with 01 and ends with 11, and the reflection of a negative glue starts with 11 and ends with 10, both of which are incompatible with unreflected glues.

By this construction, rotating the wheel so that its  $i$ th side is horizontal and at the top will have its north, east, south, and west sides represent the glues  $n_i, e_i, s_i, w_i$  of tile  $i$ . Given a tiling of the plane using this set of Wang tiles, we can place copies of the rotated wheel exactly as in the Wang tiling, and the glues will match exactly. Some space remains between the wheels, which we fill with “staples” and “shurikens”. See Figure 5.

The *shuriken* is composed of four regular concave chains of  $n - 1$  sides, matching the lengths and complementary to the angles of the regular  $4n$ -gon. Each side is adorned with  $b$  reflectionally symmetric *notches*, shown in Figure 3c, each consisting of convex angle  $\alpha$ ; reflex deficits  $\beta, 2\beta, \beta$ ; and convex angle  $\alpha$ . As shown in Figure 3e, each notch can fit a tweedle of either kind, leaving a space that is filled exactly by a *staple* (shown in Figure 3d, and consisting of convex angles  $\beta, \beta, \beta, \beta$  and reflex deficit  $2\alpha$ ). Thus each side of the shuriken can exactly match any glue, effectively hiding the unused tiles of each wheel (the glues that are not on the north, east, south, or west sides). Thus we have shown:

► **Lemma 3.1.** *Given a set of  $n$  Wang tiles and a tiling of the plane with them, we can construct a tiling of the plane with the wheel, the shuriken, and the staple constructed above.*

It remains to show that this is the only way our three tiles can tile the plane.

### 3.2 Angle Structure

We start with a few definitions and observations on the angles of the tiles.

Call an angle *clean* (and color it purple) if it is an integer multiple of  $\frac{\pi}{2n}$ . The sum of clean angles is clean, and the sum of clean angles and one nonclean angle is not clean.

The vertices of the convex  $4n$ -gon, which we call *corners*, have clean convex angle  $\pi(1 - \frac{1}{2n})$ . The matching shuriken reflex *anticorners* have a matching defect  $\pi(1 - \frac{1}{2n})$ , and each of the four concave chains are connected at the convex *tip* vertices, which have a clean convex angle of  $\frac{\pi}{n}$ .

Define angles  $\alpha = \frac{\pi}{2} - 2\varepsilon$  and  $\beta = \frac{\pi}{2} - \varepsilon$ , and pick  $\varepsilon = \frac{\pi}{16}$ .<sup>2</sup> These angles and their combinations are not clean:

► **Property 1.** *For any  $\theta_1, \theta_2 \in \{\alpha, \beta\}$ , neither  $\theta_1$  nor  $\theta_1 + \theta_2$  is clean.*

**Proof.** The relevant angles are  $\alpha = \frac{\pi}{2} - 2\varepsilon = \frac{3\pi}{8}$ ,  $\beta = \frac{\pi}{2} - \varepsilon = \frac{7\pi}{16}$ ,  $2\alpha = \frac{3\pi}{4}$ ,  $2\beta = \frac{7\pi}{8}$ , and  $\alpha + \beta = \frac{13\pi}{16}$ , which all have a doubly even denominator (divisible by 4). Because  $n$  is odd, these numbers cannot be equal to  $\frac{i\pi}{2n}$  for any integer  $i$ , so none of these angles are clean. ◀

■ **Table 1** Angles used in the wheel, shuriken, and staple. For convex vertices, we give the interior angle, while for reflex vertices, we give the defect ( $2\pi$  minus the interior angle).

| Shape    | angle of convex vertices               | defect of reflex vertices              |
|----------|--|--|
| staple   | $\beta$                                | $2\alpha$                              |
| shuriken | $\alpha, \frac{\pi}{n}$                | $\beta, 2\beta, \pi(1 - \frac{1}{2n})$ |
| wheel    | $\alpha, \beta, \pi(1 - \frac{1}{2n})$ | $\alpha, \beta$                        |

Table 1 lists the angles used by each polygon, colored to indicate which are clean. Angles  $\alpha, \beta$  are all a bit less than  $90^\circ$ , so a sum of two of them is a bit less than  $180^\circ$ , and a sum of three of them is a bit less than  $270^\circ$ . More precisely:

► **Lemma 3.2.** *For any  $\theta_1, \theta_2, \theta_3 \in \{\alpha, \beta\}$ ,  $\theta_1 \in (\frac{3}{8}\pi, \frac{7}{16}\pi)$ ,  $\theta_1 + \theta_2 \in (\frac{3}{4}\pi, \frac{7}{8}\pi)$  which is  $< \pi$ , and  $\theta_1 + \theta_2 + \theta_3 \in (\frac{9}{8}\pi, \frac{21}{16}\pi)$  which is  $> \pi$ . Note that these intervals are disjoint, so the value of a sum  $\sum_i \theta_i$  with at most three terms determines the number of terms in the sum.*

### 3.3 Edge Lengths

We design the edge lengths of the tweedledee and tweedledum in Figures 3a and 3b so that the near-vertical and near-horizontal edges all have the same length, which we call 1, and the total horizontal traversal is exactly 4. Equivalently, we choose the sequence of edge lengths for either tweedle to be

$$\begin{aligned} &\langle 2 - \sin \beta + \cos \alpha, 1, 1, 2(\sin \alpha + \cos \beta), 1, 1, 2 - \sin \beta + \cos \alpha \rangle \\ &= \langle 2 - \cos \varepsilon + \sin 2\varepsilon, 1, 1, 2(\cos 2\varepsilon + \sin \varepsilon), 1, 1, 2 - \cos \varepsilon + \sin 2\varepsilon \rangle. \end{aligned}$$

<sup>2</sup> Other choices of  $\varepsilon$  also work; the choice here is so that all edge lengths can be expressed by radical expressions, and all angles are rational multiples of  $\pi$  (or equivalently, rational numbers of degrees).

### 39:8 Tiling with Three Polygons Is Undecidable

Given our choice of  $\varepsilon = \frac{\pi}{16}$ ,

$$\begin{aligned} \cos \varepsilon &= \frac{1}{2} \sqrt{2 + \sqrt{2 + \sqrt{2}}}, & \sin \varepsilon &= \frac{1}{2} \sqrt{2 - \sqrt{2 + \sqrt{2}}}, \\ \cos 2\varepsilon &= \frac{1}{2} \sqrt{2 + \sqrt{2}}, & \sin 2\varepsilon &= \frac{1}{2} \sqrt{2 - \sqrt{2}} \end{aligned}$$

are all radical expressions, so all our edge lengths can be expressed by radical expressions.

To prevent the notches from intersecting each other at the tip of the shuriken (with the sharp angle of  $\pi/n$ ), we place the notches of the shuriken and the tweedles of the wheel at a distance of  $n$  from the anticorners of the shuriken and the corners of the wheel respectively. That is, each side of the shuriken is composed of an edge of length  $n$ , followed by the  $b$  notches, followed by an edge of length  $n$ ; and each side of the wheel is composed of an edge of length  $n$ , followed by the  $b$  tweedles composing the glue, followed by an edge of length  $n$ .

#### 3.4 Forced Tiling Structure

► **Lemma 3.3.** *The staple does not tile the plane (even allowing reflections).*

**Proof.** The staple has convex vertices just of angle  $\beta$  and one reflex vertex of defect  $2\alpha$ . In any tiling of staples, every reflex vertex must have its defect  $2\alpha$  filled by some convex angles. But  $2\alpha = \frac{3}{4}\pi$ , so by Lemma 3.2, it must be filled by exactly two convex angles. But  $2\beta = \frac{7}{8}\pi > 2\alpha$ , so two of the convex angles do not fit. Thus it is impossible to exactly fill the deficit. ◀

► **Lemma 3.4.** *The staple and shuriken do not tile the plane (even allowing reflections).*

**Proof.** By Lemma 3.3, any tiling with staples and shurikens has a shuriken. Any shuriken has a reflex anticorner of defect  $\pi(1 - \frac{1}{2n})$ , which is clean and less than  $\pi$ . This defect must be filled by some convex angles, of which we have three:  $\alpha, \beta, \frac{\pi}{n}$ . By Lemma 3.2, this defect can be filled by at most two angles  $\in \{\alpha, \beta\}$ . But by Property 1, summing one or two of these angles is not clean, while the remaining convex angle  $\frac{\pi}{n}$  and the target sum  $\pi(1 - \frac{1}{2n})$  are. Thus we cannot use any angles  $\in \{\alpha, \beta\}$  to fill the deficit, leaving only the convex angle  $\frac{\pi}{n}$ . But  $\frac{\pi}{n}$  is an even multiple of  $\frac{\pi}{2n}$ , while the target sum is an odd multiple of  $\frac{\pi}{2n}$ . Thus it is impossible to exactly fill the deficit. ◀

► **Lemma 3.5.** *Any tiling of the plane with staples, shurikens, and wheels (even allowing reflections) must consist of an infinite square grid of wheels corresponding to a Wang tiling.*

**Proof.** By Lemma 3.4, any tiling with staples, shurikens, and wheels has a wheel. Any wheel has a convex corner of angle  $\pi(1 - \frac{1}{2n})$ , which has a deficit of  $\pi(1 + \frac{1}{2n})$ , which is clean. We claim that this deficit can be filled in exactly two ways: one reflex anticorner of a shuriken of deficit  $\pi(1 - \frac{1}{2n})$ , or one convex tip vertex of a shuriken of angle  $\frac{\pi}{n}$  and one convex corner of a wheel of angle  $\pi(1 - \frac{1}{2n})$ .

Because the target deficit is  $> \pi$ , we need to consider both convex and reflex angles as well as flat edges (of angle  $\pi$ ) for possible fillings. First consider the unclean angles starting with reflex angles of deficit  $\alpha, \beta, 2\alpha, 2\beta$ . But  $\alpha < \beta < 2\alpha < 2\beta = \frac{7}{8}\pi = \pi(1 - \frac{1}{8}) < \pi(1 - \frac{1}{2n})$  for any  $n \geq 5$ . Thus none of the unclean reflex angles can be used to fill the deficit. The only remaining reflex angle that can fill the deficit is a shuriken anticorner of deficit  $\pi(1 - \frac{1}{2n})$ , and if we use that angle, we are done.

Next consider the unclean convex angles  $\alpha$  and  $\beta$ . Because  $\pi(1 + \frac{1}{2n}) < \frac{9}{8}\pi$  for any  $n \geq 5$ , by Lemma 3.2 this deficit can be filled by at most two angles  $\in \{\alpha, \beta\}$ . But by Lemma 1, summing one or two of these angles is not clean, while the remaining convex angles  $\frac{\pi}{n}, \pi(1 - \frac{1}{2n})$ , flat edge of angle  $\pi$ , and the target sum  $\pi(1 + \frac{1}{2n})$  are all clean. Thus we cannot use any angles  $\in \{\alpha, \beta\}$  to fill the deficit.

This leaves only the flat edge of angle  $\pi$  and two convex angles: the shuriken **tip** of angle  $\frac{\pi}{n}$  and the wheel **corner** of angle  $\pi(1 - \frac{1}{2n})$ . If we used only copies of the **tip**  $\frac{\pi}{n}$ , we would get an even multiple of  $\frac{\pi}{2n}$ , but the target sum  $\pi(1 + \frac{1}{2n})$  is an odd multiple of  $\frac{\pi}{2n}$ . Using a flat edge  $\pi$  would leave a gap of angle  $\frac{\pi}{2n}$ , which is too small to fill with any of the available angles. Finally, using the wheel corner  $\pi(1 - \frac{1}{2n})$  (gluing a second wheel to the first) will leave a gap of angle  $\frac{\pi}{n}$ , which can only be filled by the **tip** angle  $\frac{\pi}{n}$ .

Thus we have shown that, in all cases, the wheel's convex angle  $\pi(1 - \frac{1}{2n})$  has to be matched with an **anticorner** or a **tip** of the shuriken. Furthermore, an edge of the wheel adjacent to a **corner** must be glued to an edge of the concave chain adjacent to a **tip** or **anticorner** of the shuriken. We can follow the path of both the wheel and the concave chain of the shuriken and observe that the  $n - 2$  **anticorners** and the two **tips** of that concave chain will be glued to consecutive **corners** of the wheel.

At the end of the concave chain of the shuriken, we find a **tip** glued to the **corner** of a wheel, leaving a deficit of  $\pi(1 - \frac{1}{2n})$  to fill. As shown in Lemma 3.4, this can only be filled by another wheel **corner**. Therefore, the surround of a wheel must be filled by an alternating sequence of wheels and shurikens (omitting the small gaps left between the tweedles and the notches, which are filled by staples).

Pick one wheel  $T$  in the tiling, translate the tiling so that its center<sup>3</sup> is at the origin  $(0, 0)$ , and rotate the plane so that edge  $i$  of its  $4n$ -gon is glued to another wheel at a horizontal edge of the  $4n$ -gon, with  $T$  below that edge. Also rescale so that the width of a wheel's  $4n$ -gon (the distance between parallel edges, without adornments) is 1. Then the wheel adjacent to edge  $i$  of  $T$  will have its center at coordinate  $(0, 1)$ . Following the boundary of  $T$  clockwise, we find a shuriken glued to the edges  $i + 1, \dots, i + n - 1$  of  $T$ , and a wheel glued to the edge  $i + n$  of  $T$ . The wheel glued to edge  $i + n$  has its center at coordinates  $(1, 0)$ . Continuing this reasoning, we find that the tiling is a grid of wheels with centers on all integer lattice points. The shurikens and staples ensure that all spaces are filled, and the tweedles ensure that the tiles are compatible as in a Wang tiling. Therefore, for any tiling with the three tiles, we can produce a tiling of the plane with the original Wang tiles. ◀

### 3.5 Undecidability

The angles of the polygons, as listed in Table 1, are all rational multiples of  $\pi$ . The edge lengths of the polygons, as listed in Section 3.3, can all be expressed as radical expressions. We call polygons with such angles and edge lengths *nice*.

► **Theorem 3.6.** *Given  $n$  Wang tiles, we can construct three nice polygons that can tile the plane (allowing or forbidding reflections) if and only if the Wang tiles can tile the plane.*

**Proof.** Combine Lemma 3.1 (if) and Lemma 3.5 (only if). ◀

Combining this with membership of the tiling problem in co-RE (to be proved in the next section), we obtain:

<sup>3</sup> Define the *center* of a wheel to be the center of gravity of its  $4n$ -gon, without adornments.

## 39:10 Tiling with Three Polygons Is Undecidable

► **Corollary 3.7** (Nice form of Corollary 1.6). *Given three nice polygons in the plane, deciding whether they tile the plane is co-RE-complete and thus undecidable.*

In a recent paper, Greenfeld and Tao [8] consider a generalized version of the tiling problem, where only a periodic subset of space needs to be covered by the tiles. In our reduction, the union of the shirikens form a periodic subset of  $\mathbb{R}^2$ , and so does its complement. Thus, tiling the complements of the shirikens with the two remaining tiles is undecidable:

► **Corollary 3.8** (Stronger form of Corollary 1.2). *Given two nice polygons, deciding whether they tile a periodic subset of the plane is co-RE-complete and thus undecidable.*

By plugging in Wang tiles that simulate a universal Turing machine, such as Robinson’s 36 Wang tiles [13], we also obtain undecidability of tiling completion with three specific tiles:

► **Corollary 3.9** (Stronger form of Corollary 1.4). *There exist three fixed tiles for which completing a given finite partial tiling is co-RE-complete and thus undecidable.*

We implemented our construction in an open-source web application [4], where the user can input any set of Wang tiles and see the resulting polygons. Figure 6 shows the output for the 11 Wang tiles from Figure 1.

## 4 Membership in Co-RE

The previous sections show the co-RE-hardness of the tiling problem for three given tiles, and the tiling completion problem for three fixed tiles. We counterbalance these intractability results by showing that the tiling problem and tiling completion problem are in co-RE, for any finite set of prototiles. Refer to the full paper for omitted proofs in this section.

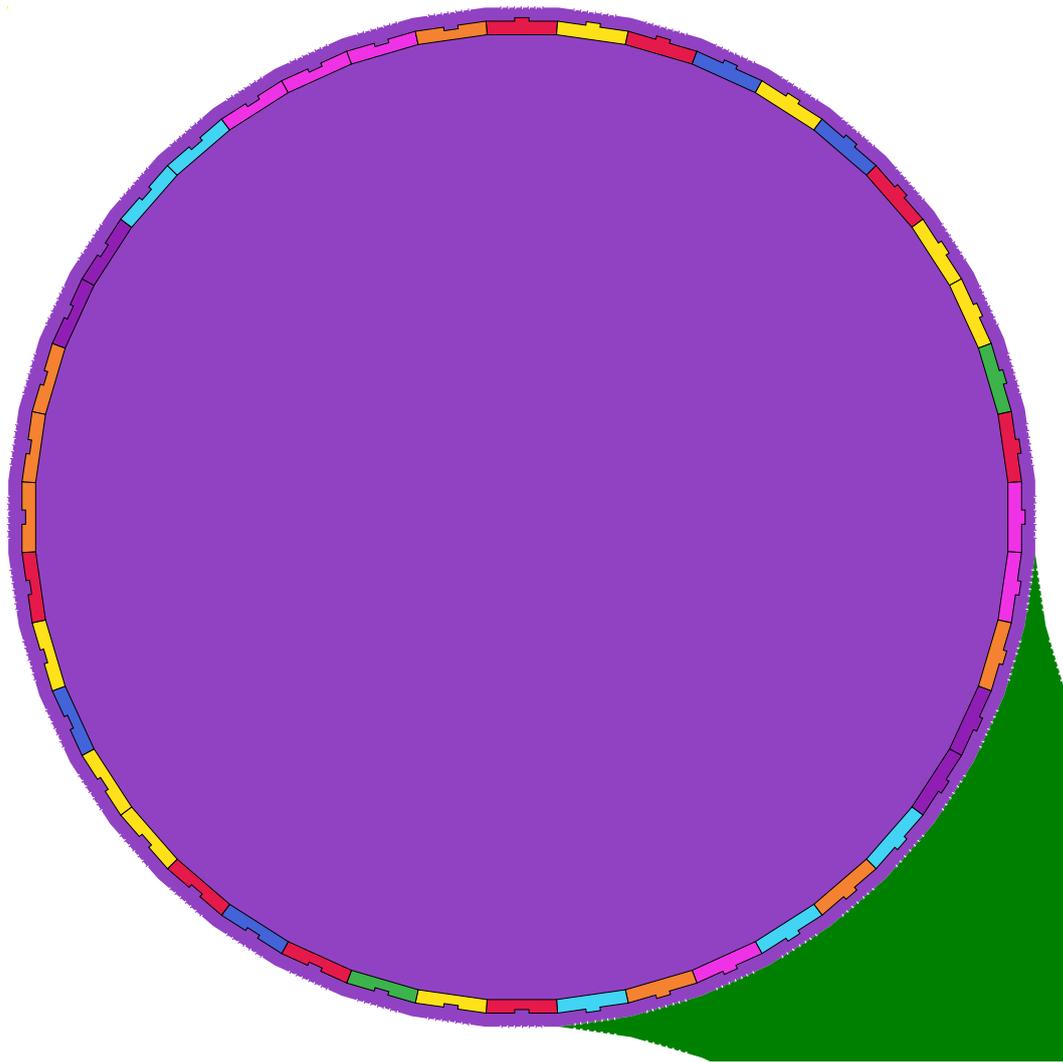
### 4.1 Model: Computable Polygons

To make these positive results as strong as possible, we use the weakest possible model of computation. We use a standard Turing machine, and represent polygons by their sequences of angles and edge lengths (equivalently, instructions in the Logo/Turtle graphics language), which need not be given explicitly but can be computed to any desired precision. More precisely, we assume the input polygons are “computable” in the following sense:

► **Definition 4.1.** *A real number  $a$  is **computable** [3] if there is a Turing machine  $T_a$  that, given a natural number  $n$ , outputs an integer  $T_a(n)$  such that  $\frac{T_a(n)-1}{n} \leq a \leq \frac{T_a(n)+1}{n}$ . A polygon is **computable** if it is promised to be simple and closed, and its  $n$  angles and edge lengths are computable.*

Computability is likely the most general representation of real numbers that is still usable for our problem. Computable numbers include all rational numbers, algebraic numbers, and transcendental numbers that can be computed to any desired precision. In particular they are closed under addition, subtraction, multiplication, division, integer roots and even trigonometric functions:

► **Theorem 4.2** ([3, Theorem 4.14]). *If  $x$ ,  $y$ , and  $z$  are computable real numbers with  $z > 0$ , then  $x + y$ ,  $x - y$ ,  $xy$ ,  $x/z$ ,  $|x|$ ,  $\min(x, y)$ ,  $\max(x, y)$ ,  $\exp(x)$ ,  $\sin(x)$ ,  $\cos(x)$ ,  $\log(z)$ , and  $\sqrt{z}$  are computable as well.*



■ **Figure 6** Output of our implemented reduction [4] for the 11 Wang tiles from Figure 1, focusing on the wheel and how the shuriken fits around the bottom-right corner. The wheel is annotated with edge colors from Figure 1b.

Note that some basic operations can be intractable for computable numbers. For instance, determining whether a computable number is zero, or the equality between two computable numbers, is undecidable (in fact, co-RE-complete). To see this, given a Turing machine  $T$ , define a number  $a_T$  to have its  $n$ th bit after the binary point be 1 if  $T$  halts after  $n$  steps, and 0 otherwise. The number  $a_T$  is computable, and is zero if and only if  $T$  does not halt. For our problem, we can show a similar undecidability:

► **Theorem 4.3.** *Given a single computable pentagon, determining whether it tiles the plane is co-RE-hard and thus undecidable.*

**Proof.** Pick a generic quadrilateral, and glue a very flat isosceles triangle to one of its edges, where the apex angle is  $180^\circ - r/100$  for a given constructible number  $r \in [0, 1]$ . The other angles and edge lengths of the triangle are constructible via trigonometric functions.

## 39:12 Tiling with Three Polygons Is Undecidable

The resulting pentagon tiles the plane only in the degenerate case where the triangle is degenerate, i.e., a line segment, which happens exactly when the obtuse angle of the triangle is exactly  $180^\circ$ . Thus the tiling problem is equivalent to testing whether  $r = 0$  for a given constructible number  $r \in [0, 1]$ , which is co-RE-hard as shown above. ◀

Of course, this result is quite unsatisfying, as the reason for undecidability stems from the extreme weakness of the model and generality of the representation of the polygons, and the inability to even check locally that a tiling is valid. Yet surprisingly, in this same model, we are able to show membership in co-RE.

### 4.2 Co-RE Algorithm

The high-level idea of our algorithm is to try to build partial tilings that cover a larger and larger disk. If we ever fail to cover a disk, then we know that the plane cannot be tiled; and if we never fail, the well-known Extension Theorem (Theorem 4.11 below) guarantees that the plane can be tiled. To determine whether we can tile enough to cover a disk, we bound the number of tiles that could possibly intersect the disk, then enumerate all possible combinatorial ways for these tiles to fit together, and for each, check whether the tiles fit together properly. Checking fit is limited to tiles that share vertices, however, so we need to take care to handle the case that there are seam lines in the tiling where no tiles on opposite sides of the line share a vertex (as in, e.g., the classic brick tiling). We also avoid checking for global intersection between tiles (because doing so is tricky in our model), opting instead to check just locally that angles add up correctly at vertices and that edge lengths add up correctly along edges. Our notion of “neat carpet” handles both of these issues by forbidding only local self-overlap, and guaranteeing that every boundary vertex is either outside the specified disk or has total angle  $180^\circ$  so potentially forms a seam boundary. We are then able to show that arbitrarily large neat carpets imply the existence of a plane tiling.

Our co-RE algorithm will in particular need to repeatedly test for equality among constructible numbers, a co-RE-complete problem. Thus we need a way to compose co-RE decisions. We use the following standard result (mentioned, e.g., in [12]):

► **Lemma 4.4.** *Finite disjunctions and recursively enumerable conjunctions of co-RE decision problems are in co-RE.*

Let  $\mathcal{T}$  be a set of prototiles, where each tile is a (simple closed) computable polygon. Define a *carpet* to be a topological disk produced by gluing together a finite collection of tiles from  $\mathcal{T}$ , where every interior vertex has  $360^\circ$  total angle from incident tiles. We assume that the carpet is laid out in the Euclidean plane, that is, every point in a tile has real coordinates, but we allow the surface to be self-overlapping, that is, a point of the plane might be covered by more than one tile. A *patch* is a carpet whose embedding in the plane is not self-overlapping.

A carpet can be described by its *combinatorial gluing*, which specifies (1) the set of tiles, each of which is an instance of a prototile; (2) a partition of the tile vertices into coincident (glued together) points; and (3) for each tile edge, the sequence of other tile edges and/or boundary that the edge has positive-length overlap with, in order along the edge. Call a carpet or a patch *seamless* if the position of all tiles in the carpet is fully determined by the combinatorial gluing and the position of its first tile. This notion forbids carpets whose tiles can be separated by a line along which the two sides could slide (causing an uncountable infinity of solutions). To verify seamlessness, build the *incidence graph* on the tiles of the carpet, where two tiles are connected by an edge if they share a vertex, and

check that the graph is connected. For the case of tiling completion, we can ensure that the given patch is seamless by adding a vertex along each seam, common to both adjacent tiles; as this modification to the preplaced tiles does not change their shape, it does not change the outcome of the decision.

First we show how to verify that a carpet is valid:

► **Lemma 4.5.** *Given a combinatorial gluing of a possible seamless carpet, deciding whether it corresponds to a seamless carpet is co-RE-complete.*

Define the *distance* between two points in a carpet to be the Euclidean distance between those two points when the carpet is laid out in the Euclidean plane (note that this embedding may be self-overlapping, and this distance is no larger than the intrinsic distance within the carpet). Call a vertex of a carpet *neat* if it is either interior to the carpet and surrounded by tiles summing up to an angle of  $2\pi$ , or is on the boundary of the carpet and is surrounded by contiguous tiles summing up to an angle of  $\pi$ . A carpet is *neat within radius  $< r$*  if every vertex at distance  $< r$  from the origin is neat. In an *anchored* carpet, we assume one *anchored* tile in the combinatorial gluing has been chosen to be placed with its center of gravity at the origin, and with a canonical rotation (say, matching its prototile).

► **Lemma 4.6.** *Given an anchored carpet and a computable positive number  $r$ , deciding whether it is neat within radius  $< r$  is in co-RE.*

Let  $\rho$  be the maximum “radius” of tiles in  $\mathcal{T}$ , meaning that a disk of radius  $\rho$  centered at the center of gravity of each tile in  $\mathcal{T}$  covers that tile. Let  $A_{\min}$  be the minimum area of a tile in  $\mathcal{T}$ . Let  $D_r$  denote the disk of radius  $r$  centered at the origin.

► **Lemma 4.7.** *If  $\mathcal{T}$  can tile the plane, then for any  $r > 0$ , (a) there is an anchored patch with a finite number  $N(r)$  of tiles that covers the disk  $D_r$ , and (b) there is an anchored seamless patch with  $\leq N(r)$  tiles that is neat within radius  $< r$ .*

*If a seamless patch  $P$  can be extended to tile the plane, then for any  $r > 0$ , (a) there is an anchored patch containing  $P$  with a finite number  $|P| + N(r)$  of tiles that covers the disk  $D_r$ , and (b) there is a seamless patch containing  $P$  with  $\leq |P| + N(r)$  tiles that is neat within radius  $< r$ .*

**Proof.** For the (a) statements, translate (and rotate) the tiling so that one of its tiles is anchored; and if we are given a seamless patch  $P$ , choose to anchor one of its tiles. Consider the disk  $D_{r+2\rho}$  of radius  $r + 2\rho$  centered at the origin (the anchored tile’s center of gravity). Take all the tiles in the plane tiling that are fully inside  $D_{r+2\rho}$ , and take all the tiles of  $P$  (if given). Because the tiles do not overlap and are each of area  $\geq A_{\min}$ , there are at most  $\pi(r + 2\rho)^2/A_{\min}$  tiles inside  $D_{r+2\rho}$ . Take the connected component that contains the origin (and thus the tiles of  $P$ , if given), which only decreases the number of tiles. This is an anchored patch that covers the smaller disk  $D_r$ .

For the (b) statements, build the incidence graph on the tiles of the carpet, where two tiles are connected by an edge if they share a vertex. If this graph is connected, then the patch is seamless. If this graph is disconnected, it is because of seams. Seam lines cannot intersect: otherwise, their intersection point is a vertex common to both sides of the seams. Thus, cutting the patch along all seams, or equivalently taking one connected component of the incidence graph, will create neat vertices on the new boundary, where the seams were. Take the component that contains a tile covering the origin. This is a seamless patch that is neat within radius  $< r$ . ◀

### 39:14 Tiling with Three Polygons Is Undecidable

► **Lemma 4.8.** *Given prototiles  $\mathcal{T}$  and a computable radius  $r$ , deciding whether there is an anchored carpet that is neat within radius  $< r$  is in co-RE. If there is no such carpet, then  $\mathcal{T}$  cannot tile the plane.*

*Given prototiles  $\mathcal{T}$ , a seamless patch  $P$ , and a computable radius  $r$ , deciding whether  $P$  can be extended to a carpet that is neat within radius  $< r$  is in co-RE. If there is no such carpet, then  $P$  cannot be extended to tile the plane.*

► **Lemma 4.9.** *If an anchored carpet is neat within radius  $< r + 2\rho$ , then the carpet contains an anchored patch that is neat within radius  $< r$ .*

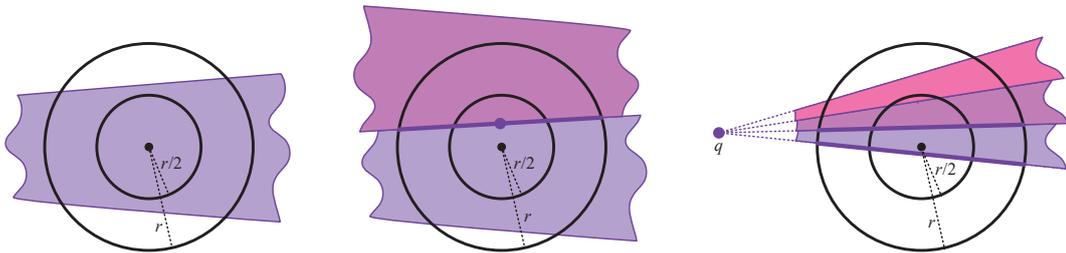
**Proof.** Consider the intersection between the carpet  $C$  and the disk  $D_{r+2\rho}$ , which might have multiple connected components and/or be self-overlapping. Pick the component  $S$  that contains the anchored tile. The boundary of  $S$  is composed of straight lines (the edges of the tiles connected by neat vertices) and circular arcs (portions of the boundary of  $D_{r+2\rho}$ ), connected together at convex angles. Thus  $S$  is convex, so non-self-overlapping.

Back in the carpet, remove all tiles that are not entirely contained in  $D_{r+2\rho}$ . Again, this possibly results in several connected components. Retain the component  $C'$  that contains the anchored tile, where connectivity is defined by interior paths, so that  $C'$  does not have pinch points. By construction, this component is contained in  $S$  and thus is not self-overlapping. Also  $C'$  is a topological disk: a hole in  $C'$  could only come from a removed tile, which must touch the outside of  $D_{r+2\rho}$ , contradicting that it is surrounded by  $C'$  (by planarity). Therefore,  $C'$  is a patch.

Finally, because all tiles removed have diameter  $\leq 2\rho$ , none of the deleted tiles intersect the disk  $D_r$ , and therefore all vertices within radius  $< r$  remain untouched. Therefore the patch  $C'$  is neat within radius  $< r$ . ◀

► **Lemma 4.10.** *If there exists an anchored carpet  $C$  that is neat within radius  $< r + 2\rho$ , then there exists a patch that covers the disk  $D_{r/2}$ . Furthermore, that patch contains all tiles of  $C$  that intersect  $D_{r/2}$ .*

**Proof.** Suppose we have a carpet that is neat within radius  $< r + 2\rho$ . By Lemma 4.9, we have a patch that is neat within radius  $< r$ . The intersection of the patch and the disk  $D_r$  is a disk cut by noncrossing chords; refer to Figure 7. Note that any chord of  $D_r$  that intersects  $D_{r/2}$  cuts off an arc of angle  $\frac{2}{3}\pi$  from  $D_r$ . Because the chords defined by the patch do not intersect in  $D_r$ , at most two of them intersect  $D_{r/2}$ .



■ **Figure 7** A neat patch within  $< r$ , and how it can interact with the smaller disk  $D_{r/2}$ . From left to right: no chords, one chord, and two chords.

If no chord intersects  $D_{r/2}$ , then the patch covers  $D_{r/2}$ , and we are done.

If exactly one chord intersects  $D_{r/2}$ , then the patch is contained in a half-plane that contains the origin and is bounded by the extension of the cord. Rotate a copy of the patch by  $180^\circ$  about the center of the cord, and glue the two pieces together. This produces a patch that covers  $D_{r/2}$ .

If two chords intersect  $D_{r/2}$ , then the patch is contained in a wedge defined by the extensions of the two chords. Let  $q$  be the apex of the wedge, and assume  $q$  is on the  $x$  axis, with one of the chords above the  $x$  axis and the other chord below. Assume without loss of generality that the chord below is the longest one. Repeatedly stitch copies of the patch by rotating it about  $q$ , upward, gluing the longer chord of the previous copy to the shorter chord of the next copy, until the upper half of  $D_{r/2}$  is covered. Now repeat the procedure for a single chord. The resulting patch covers the disk of radius  $r/2$ . ◀

► **Theorem 4.11** (Extension Theorem [9, p. 151]). *Given a finite collection  $\mathcal{T}$  of prototiles, if they tile arbitrarily large disks, then they admit a tiling of the plane.*

*Given a finite collection  $\mathcal{T}$  of prototiles and patch  $P$  using tiles of  $\mathcal{T}$ , if  $P$  can be extended to cover arbitrarily large disks centered in  $P$ , then  $P$  can be completed to tile the plane.<sup>4</sup>*

We translate the above theorem to seamless anchored patches.

► **Lemma 4.12.** *Given a collection  $\mathcal{T}$  of prototiles, if there exist anchored carpets that are neat within radius  $< r$  for arbitrarily large  $r$ , then  $\mathcal{T}$  admits a tiling of the plane.*

*Given a collection  $\mathcal{T}$  of prototiles, and given an anchored patch  $P$  using tiles of  $\mathcal{T}$ , if  $P$  can be extended to a carpet that is neat within radius  $< r$  for arbitrarily large  $r$ , then  $P$  can be completed to tile the plane.*

► **Theorem 4.13** (Precise form of Theorem 1.5). *Given a set  $\mathcal{T}$  of  $k$  polygons in our model, deciding whether they tile the plane is in co-RE. Also given a patch  $P$  of tiles from  $\mathcal{T}$ , deciding whether  $P$  can be completed to tile the plane is in co-RE.*

**Proof.** For every positive integer  $k$ , set  $r = k\rho$  and use Lemma 4.8 to determine whether there exists an anchored carpet that is neat within radius  $< r$ , or whether  $P$  can be completed to produce such a carpet. This is a recursively enumerable disjunction of co-RE problems, so by Lemma 4.4 is in co-RE. By Lemma 4.12, if all of these problems output true, then the polygons tile the plane or complete  $P$  to tile the plane. By Lemma 4.8, if any of these problems outputs false, then the polygons do not tile or complete  $P$  to tile the plane. ◀

---

## References

- 1 Robert Berger. The undecidability of the domino problem. *Memoirs of the American Mathematical Society*, 66:1–72, 1966. doi:10.1090/memo/0066.
- 2 Siddhartha Bhattacharya. Periodicity and decidability of tilings of  $\mathbb{Z}^2$ . *American Journal of Mathematics*, 142(1):255–266, 2020. doi:10.1353/ajm.2020.0006.
- 3 Vasco Brattka, Peter Hertling, and Klaus Weihrauch. A tutorial on computable analysis. In S. Barry Cooper, Benedikt Löwe, and Andrea Sorbi, editors, *New Computational Paradigms*, pages 425–491. Springer, New York, NY, 2008. doi:10.1007/978-0-387-68546-5\_18.
- 4 Erik D. Demaine and Stefan Langerman. edemaine/three-tiles, 2024. Software, swId: swh:1:dir:aef5f642bf8a40c0cecf56bfb62b9e089effb265 (visited on 2025-06-05). URL: <https://github.com/edemaine/three-tiles>, doi:10.4230/artifacts.23339.
- 5 D. Girault-Beauquier and M. Nivat. Tiling the plane with one tile. In *Topology and Category Theory in Computer Science*. Oxford University Press, August 1991. doi:10.1093/oso/9780198537601.003.0012.
- 6 Solomon W. Golomb. Tiling with sets of polyominoes. *Journal of Combinatorial Theory*, 9(1):60–71, 1970. doi:10.1016/S0021-9800(70)80055-2.

---

<sup>4</sup> Although [9, p. 151] states only the first version of the theorem, the same proof establishes both.

## 39:16 Tiling with Three Polygons Is Undecidable

- 7 Rachel Greenfeld and Terence Tao. Undecidable translational tilings with only two tiles, or one nonabelian tile. *Discrete & Computational Geometry*, 70(4):1652–1706, 2023. doi:10.1007/s00454-022-00426-4.
- 8 Rachel Greenfeld and Terence Tao. Undecidability of translational monotilings. *Journal of the European Mathematical Society*, 2024. To appear. arXiv:2309.09504.
- 9 Branko Grünbaum and G. C. Shephard. *Tilings and Patterns*. W. H. Freeman, 1987.
- 10 Emmanuel Jeandel and Michaël Rao. An aperiodic set of 11 Wang tiles. *Advances in Combinatorics*, January 2021. doi:10.19086/aic.18614.
- 11 Nicolas Ollinger. Tiling the plane with a fixed number of polyominoes. In *Proceedings of the 3rd International Conference on Language and Automata Theory and Applications (LATA 2009)*, volume 5457 of *Lecture Notes in Computer Science*, pages 638–647, Tarragona, Spain, April 2009. Springer. doi:10.1007/978-3-642-00982-2\_54.
- 12 H. G. Rice. Classes of recursively enumerable sets and their decision problems. *Transactions of the American Mathematical Society*, 74(2):358–366, 1953. doi:10.2307/1990888.
- 13 Raphael M. Robinson. Undecidability and nonperiodicity for tilings of the plane. *Inventiones Mathematicae*, 12(3):177–209, September 1971. doi:10.1007/BF01418780.
- 14 David Smith, Joseph Samuel Myers, Craig S. Kaplan, and Chaim Goodman-Strauss. An aperiodic monotile, 2023. doi:10.48550/arXiv.2303.10798.
- 15 Hao Wang. Proving theorems by pattern recognition – II. *Bell Systems Technical Journal*, 40(1):1–41, 1961. doi:10.1002/j.1538-7305.1961.tb03975.x.
- 16 Damien Woods and Turlough Neary. The complexity of small universal turing machines: A survey. *Theoretical Computer Science*, 410(4):443–450, 2009. doi:10.1016/j.tcs.2008.09.051.
- 17 Chao Yang. Tiling the plane with a set of ten polyominoes. *International Journal of Computational Geometry & Applications*, 33(03n04):55–64, 2023. doi:10.1142/S0218195923500012.
- 18 Chao Yang and Zhujun Zhang. A proof of Ollinger’s conjecture: undecidability of tiling the plane with a set of 8 polyominoes, 2024. doi:10.48550/arXiv.2403.13472.
- 19 Chao Yang and Zhujun Zhang. Undecidability of translational tiling of the 4-dimensional space with a set of 4 polyhypercubes, 2024. doi:10.48550/arXiv.2409.00846.
- 20 Jed Chang-Chun Yang. *Computational Complexity and Decidability of Tileability*. PhD thesis, University of California, Los Angeles, 2013.