TWO RESULTS IN DISCRETE GEOMETRY II

SHINTARO FUSHIDA-HARDY

Abstract. Last quarter I gave a kiddie talk about classifying polyhedra, and cutting up polygons into equal area pieces. Today we'll instead be cutting up polyhedra and guarding polygons.

1. Given any two polygons of equal area, one can cut the first shape into finitely many pieces and rearrange them into the second. Does this result hold in higher dimensions? This was Hilbert's 3rd problem (in his famous list of 23 problems), and we'll recreate Dehn's remarkably simple proof that the result fails in 3 dimensions.

2. You're in charge of the security for an art gallery but have a very limited budget. What's the minimum number of guards you must employ to ensure that every artwork (every region of the gallery) is being watched at all times? We'll solve the *art gallery problem* and discuss some other variations of it.

1. Polyhedra

The theme of the first section is invariants. Let's start with the 2-dimensional version of the problem: cutting up and rearranging polygons.

Question. Given two polygons, when is it possible to cut up one of the polygons into finitely many pieces and rearrange it into the other?

We introduce some terminology: two polygons (or polyhedra in any dimension) are called *scissors congruent* if one can be cut into finitely many polygonal (polyhedral) pieces which can be glued back together to form the other.

Answer. Two polygons are scissors congruent if and only if they have the same area.

Example. A square and an equilaterial triangle are scissors congruent (with a decomposition into four pieces).

We can reformulate this in terms of invariants: Let $\mathcal{P}^2$ denote the space of all polygons (in $\mathbb{R}^2$, defined up to isometry). Then the following properties are all well defined - they can be thought of as outputs of functions $f : \mathcal{P}^2 \to X$ for some X.

(1) Area A
(2) Unordered list of numbers recording its internal angles
(3) Perimeter
(4) Count of non-convex vertices

However, we're interested in understanding the space

$$\mathcal{P}^2/\sim$$

where the equivalence relation we're modding out by is scissors congruence.

(1) Perimeter p does not descend to a well defined function $\mathcal{P}^2/\sim\to\mathbb{R}$. Cutting up and rearranging can change the perimeter.
(2) Area A defines a function $A:\mathcal{P}^2/\sim\to\mathbb{R}$.

Definition 1.1. Given some collection of objects $\mathcal{O}$, an *invariant* is a function defined on $\mathcal{O}/\sim$ where $\sim$ is some equivalence relation. A *complete invariant* is an injective invariant.

Example. Area is a complete invariant of $\mathcal{P}^2/\sim$.

Now we'll move onto the three dimensional problem.

Question. Are two polyhedra scissors congruent if and only if they have the same volume? Alternatively, writing $\mathcal{P}^3$ to denote the space of polyhedra defined up to isometry, is volume a complete invariant of $\mathcal{P}^3/\sim$?

To disprove the above result, it suffices to find an invariant which takes different values on two polyhedra of equal volume.

Proposition 1.2. (In dimension three) there exists polyhedra of equal volume that are not scissors congruent. Specifically, one can take the cube and regular tetrahedron.

Proof. Our proof is by constructing an invariant defined on $\mathcal{P}^3/\sim$ which differs on equal volume cubes and tetrahedra.

Suppose we start with a polyhedron and cut it along some plane. Typically this results in cutting some of the edges of the polyhedron, as well as cutting some of the dihedral angles of the polyhedron. However, there is a sense in which the sum of the lengths, or sum of angles, should be the same! Specifically, the local picture is as shown below. Locally there are two possible types of cuts we can make, also shown.

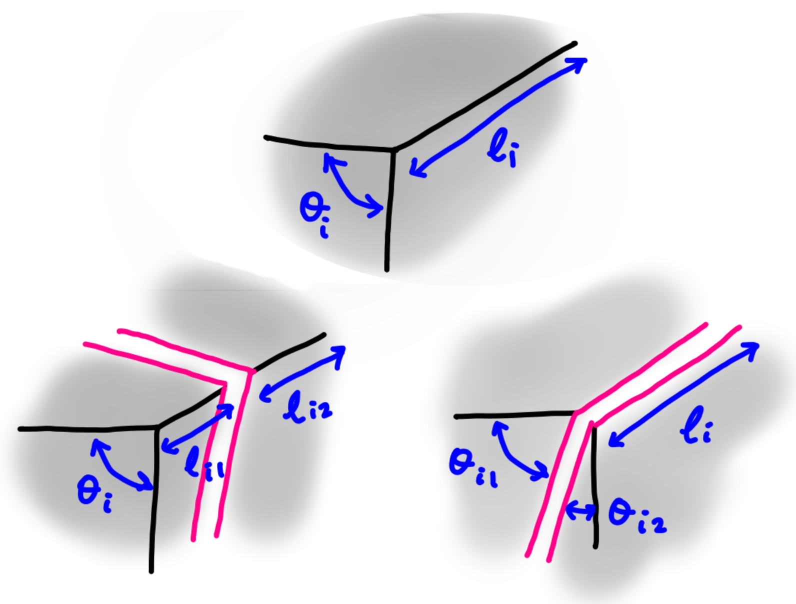

In one case, the cut takes us from (θ_i, ℓ_i) to $(\theta_i, \ell_{i1}), (\theta_i, \ell_{i2})$ where $\ell_{i1} + \ell_{i2} = \ell_i$. The other cut is similar, but instead the angles are additive. Maybe some of you have noticed that this is very tensor producty!

Definition 1.3. Given a polyhedron P, we define its *Dehn invariant* to be

$$D(P) = \sum_{i \in \text{edges}} \ell_i \otimes \theta_i \in \mathbb{R} \otimes_{\mathbb{Q}} \mathbb{R}/2\pi\mathbb{Z}.$$

Recall that the tensor product has these properties:

(1) $\ell_1 \otimes \theta + \ell_2 \otimes \theta = (\ell_1 + \ell_2) \otimes \theta$
(2) $\ell \otimes \theta_1 + \ell \otimes \theta_2 = \ell \otimes (\theta_1 + \theta_2)$
(3) $q(\ell \otimes \theta) = (q\ell) \otimes \theta = \ell \otimes (q\theta)$ for $q \in \mathbb{Q}$.

These three properties ensure that the Dehn invariant is truly an invariant of scissors congruence! Whenever we cut a polyhedron, existing edges are changed in the two ways described earlier. In those instances, properties 1 and 2 ensure that the invariant is unchanged. The only other thing that happens is that new edges can form - but these necessarily form on faces or the interior of the polyhedron. In each case, the sum of the angles about the new edge is π or 2π, but these are zero in the tensor product.

Now that we've defined a new invariant and shown that it's an invariant, let's compute it for the cube and regular tetrahedron.

(1) Given any cube, all dihedral angles are rational multiples of pi, so terms involving them are automatically 0 in the tensor product. Therefore the Dehn invariant vanishes.
(2) Next consider any regular tetrahedron. These all have edge length ℓ and dihedral angle $\arccos(1/3)$. This means their Dehn invariant is

$$6\ell \otimes \arccos(1/3).$$

Since $\arccos(1/3)$ is *not* a rational multiple of 2π, this proves that every regular tetrahedron has non-zero Dehn invariant.

In summary, cubes and tetrahedra are never scissors congruent. □

Question. Is there a complete set of invariants for scissors congruence of polyhedra?

Answer. Two polyhedra are scissors congruent if and only if they have the same volume and Dehn invariant.

Remark. There's a version of the Dehn invariant for 4 dimensional polytopes, in which the above result also holds. I'm not aware of anything in higher dimensions.

2. Polygons

Now it's time for my second topic. I'll start by posing a question:

Question. Given an art gallery of some shape, how many guards do we need to defend every part of the gallery?

This is pretty vague, but we can draw some pictures to come up with a more formal question.

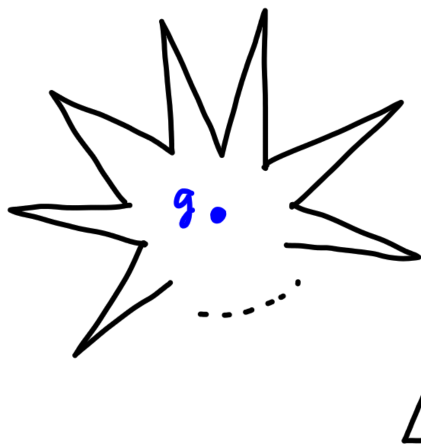

On the left, we see that we can build an art gallery with arbitrarily many walls for which one guard is sufficient. This mathematically corresponds to a "star shaped domain" which is a close relative of convex domains.

On the right, we have an art gallery which is much worse! The top of each spike is only visible from the triangle formed by the spike, and none of these triangles overlap. This means the number of guards required is the number of spikes: For each $w \in 3\mathbb{Z}$, we can find a gallery with w walls for which

$$g = w/3$$

guards are required to defend the gallery. With these examples in mind, we can formulate a problem statement.

Question. For each w in $\mathbb{N}$, is there an upper bound on the number of guards needed to defend an art gallery with w walls?

Answer. Yes! We can take

$$g \leq \lfloor w/3 \rfloor.$$

This bound is sharp because of the example earlier.

Proof. An *art gallery* is any polygon P in $\mathbb{R}^2$. Suppose P has w edges. Our proof strategy is as follows:

(1) Triangulate the polygon (without adding new edges).
(2) Three-colour the triangulation.
(3) Place guards at the vertices with the least frequent colour.

Step 1. I won't prove this, because it's intuitively clear enough: given any polygon, you can draw edges between vertices to triangulate it. One can describe an algorithm that will take in a polygon and output a triangulation.

Step 2. Here things get interesting! Why would our triangulation be three-colourable? We consider the dual graph of the triangulation.

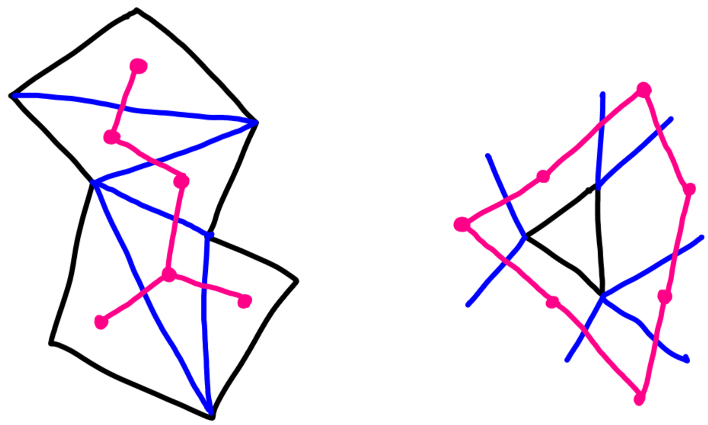

The dual graph is guaranteed to be a tree, because the only way a cycle can form is if the polygon has a hole (and then it's not a polygon). We now proceed by induction: A tree with one vertex corresponds to a single triangle, and this can be trivially three-coloured. Now suppose trees with n vertices correspond to three-colourable triangulations. Then a tree with $n+1$ vertices corresponds to adding an extra vertex with degree 1 to a tree with n vertices. On the triangulation side, this corresponds to gluing a triangle onto a triangulation along a single edge. This adds one extra vertex, which can be chosen to be the third colour not already used by the gluing edge.

Step 3. We're basically done! The number of vertices of the polygon is the number of edges (walls). The three-colouring defines a partition on the vertices, so the least frequently occuring colour occurs at most $\lfloor w/3 \rfloor$ times. Place the guards on these vertices. Since every triangle in the triangulation meets eeach of the three colours, every triangle has a guard on exactly on vertex. This guard can see the entire triangle (since triangles are convex), so the collection of guards can see the entire polygon. □

Now I'll give a bit of background, and describe a few different versions of the problem. The *art gallery problem* was first formally stated and proved in the mid 1970s after a conversation at a combinatorics conference. This lead to rapid interest in other versions of the problem in subsequent years. I'll now state several of them.

Example. Art gallery problem with holes: in practice, most art galleries aren't polygons - they often have obstructions (holes). An art gallery with w walls and h holes can always be guarded by at most

$$g = \lfloor (w+h)/3 \rfloor$$

guards.

Example. Half guard theorem: In practice, a guard doesn't have 360 degree vision at all times. A more realistic model is to restrict the guard's vision. For example, how many "half guards" do we need to secure an art gallery? That is, how many guards with 180 degree vision are necessary? Theorem:

$$g = \lfloor w/3 \rfloor$$

half guards are sufficient to guard a polygonal art gallery with w walls. This is surprising! We'd expect more than this number to be required, considering $\lfloor w/3 \rfloor$ was a tight bound for 360 vision guards.

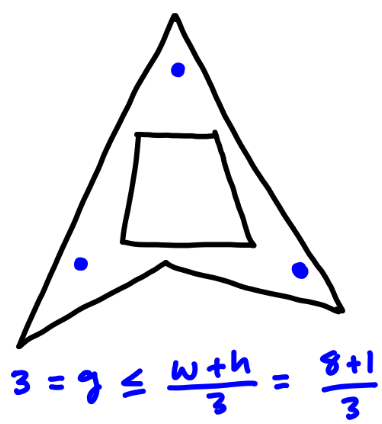

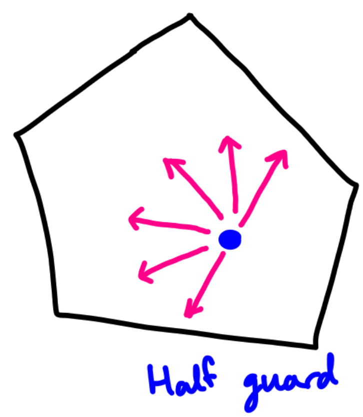

Example. Another version of the art gallery problem considers guarded guards: given a gallery with w walls, how many guards do you need to see the whole gallery, and ensure that every guard is seen by at least one other guard?

Clearly we need at least the bound for the usual art gallery problem, and it's sufficient to have twice the number, by placing guards in pairs:

$$\lfloor w/3 \rfloor \leq g \leq 2\lfloor w/3 \rfloor.$$

We can improve the lower bound by considering the following depicted family of galleries:

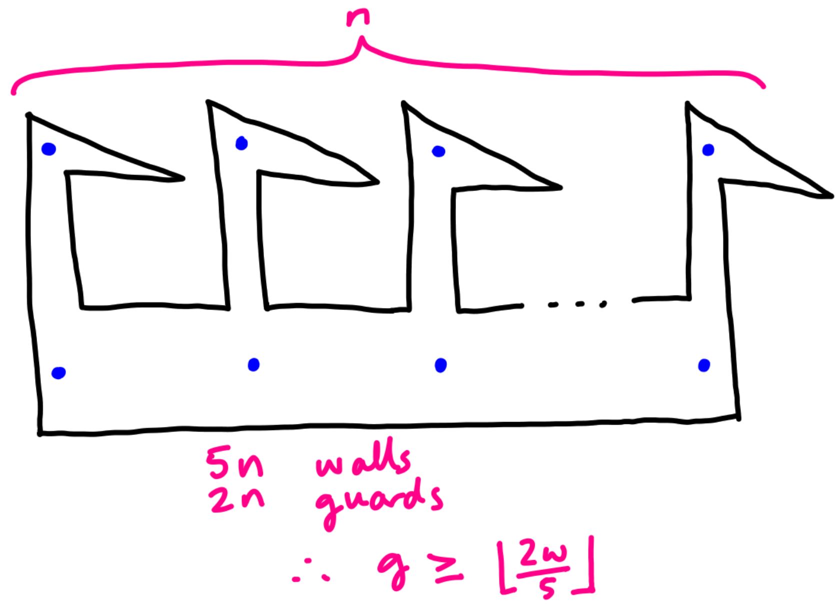

In fact, the true answer for "the number of guards that is sufficient for any gallery with w walls" is

$$g = \lfloor (3w - 1)/7 \rfloor.$$

Example. Fortress problem: the final version of the the problem I'll describe is the fortress problem, in which we want to find the minimum number of guards to defend a fortress (from the outside). That is, place guards exterior to a polygon so that evey edge of the polygon can be seen.

It turns out that

$$g = \lceil w/3 \rceil$$

guards are sufficient and sometimes necessary to guard a fortress with w walls from its exterior, and

$$g = \lceil w/2 \rceil$$

guards are sufficient and sometimes necessary if we place the guards on the walls of the fortress.