

Immersive Visualization of Flat Surfaces Using Ray Marching

Fabian Lander¹ and Diaaeldin Taha²

Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

¹lander@mis.mpg.de, ²taha@mis.mpg.de

Abstract

We present an effective method for visualizing flat surfaces using ray marching. Our approach provides an intuitive way to explore translation surfaces, mirror rooms, unfolded polyhedra, and translation prisms while maintaining computational efficiency. We demonstrate the utility of the method through various examples and provide implementation insights for programmers. Finally, we discuss the use of our visualizations in outreach. We make our simulations and code available online.




Figure 1: A first-person view inside a “room” whose base is the double pentagon flat surface.

Introduction

Flat surfaces naturally arise in various mathematical contexts, from paper folding to dynamical billiards. The simplest example of a flat surface is the torus: a rectangle with opposite edges identified by translation. Walking off the right edge returns one to the left, and stepping off the top returns one to the bottom. This method of constructing surfaces by identifying parallel edges generalizes to more complex polygons, giving rise to a wide variety of flat surfaces. For example, Figure 4 shows an L-shaped polygon whose sides are identified to form one such surface. At first glance, these surfaces seem simple, as they are locally isometric to the Euclidean plane, but their global structure and dynamics can be surprisingly complex.

In this work, we visualize flat surfaces by rendering how they appear from the perspective of a viewer who inhabits a thickened three-dimensional version of the surfaces, offering an intuitive first-person experience



Figure 2: *Gluing opposite sides of a rectangle forms the simplest example of a flat surface: the torus. Though the scales are faithful, the torus appears smaller than expected because its radius equals the rectangle’s height. For a magnified view from a different angle, see Appendix D.*

of their intrinsic geometry. We build on a large body of work on rendering geometric spaces in virtual reality, including the works of Weeks [18, 19], Nelson-Segerman [13], Hart-Hawksley-Matsumoto-Segerman [9, 10], Coulon-Matsumoto-Segerman-Trettel [3, 4, 5], and many others. Additionally, we draw inspiration from Edwin Abbott’s *Flatland* and the mirror rooms of Leonardo da Vinci and Yayoi Kusama [1].

Flat Surfaces

Flat surfaces lie at the intersection of geometry, dynamics, and algebra. Their study has led to deep mathematical breakthroughs connecting complex dynamics, Teichmüller theory, ergodic theory, and number theory. Not only are mathematicians captivated by their rich structure, but these surfaces have also captured the public imagination through, e.g., their connection to billiards. Indeed, the flow of a billiard ball on a polygonal table can be unfolded to a geodesic on a flat surface, providing an elegant bridge between recreational mathematics and cutting-edge research. This connection has helped illuminate both the highly regular and chaotic behaviors that can emerge from seemingly simple geometric systems.

In what follows, a flat surface [20] will refer to a surface obtained by taking a finite collection of polygons $\{P_1, \dots, P_n\}$ in \mathbb{R}^2 and identifying their edges pairwise by isometries (translations, rotations, and reflections), resulting in a surface X with a flat metric everywhere except possibly at a finite set of singular points Σ (corresponding to some vertices of the polygons after gluing) where the cone angle may differ from 2π . Note that not all vertices necessarily become singularities after gluing, as exemplified by the standard torus construction. The metric on $X \setminus \Sigma$ is induced from the Euclidean metric on \mathbb{R}^2 .

We note that our visualization method can be readily extended to affine surfaces as well, where instead of using only isometries, we allow edges to be identified by invertible affine transformations. These surfaces carry a well-defined affine structure, although they generally do not inherit a global flat metric. Our approach can also handle certain simple infinite flat surfaces, such as infinite flat cones and cylinders.

Ray Marching

Ray marching is a class of graphics rendering methods where each pixel on the screen emits a ray into the scene, and that ray is traversed in small steps rather than solving for explicit intersection points. This approach dates back to at least the 1980s, including early work by Perlin-Hoffert [14]. Many variants of ray marching exist. For instance, *volume rendering* [6] samples density at each step, while *sphere tracing* [8] relies on signed distance functions (SDFs) to determine a safe step size.

Ray marching is particularly suited for visualizing complex geometric spaces, such as flat surfaces. It only requires local distance checks and straightforward step updates, with no need to solve complicated intersection equations. Additionally, the method maps well to modern GPUs, enabling real-time interactive visualizations of fractals, non-Euclidean rooms, and other exotic geometries.

A signed distance function (SDF) is a function that returns the signed shortest distance $d(\vec{p})$ from a

point \vec{p} to the surface of an object. By convention, if $d(\vec{p}) > 0$ then \vec{p} is outside the object, if $d(\vec{p}) = 0$ then \vec{p} lies exactly on its surface, and if $d(\vec{p}) < 0$ then \vec{p} is inside the object. A sphere provides a simple example. For a sphere centered at \vec{c} with radius r :

$$d_{\text{sphere}}(\vec{p}) = \|\vec{p} - \vec{c}\| - r.$$

Beyond this elementary example, the computer-graphics community has developed SDFs for an extensive collection of geometric primitives and complex shapes [8, 15, 16]. More complex and organic shapes can be assembled through *constructive solid geometry* (CSG) operations [7] such as *union*, *intersection*, and *difference*, and is something of an art form in its own right. For instance, the SDF for the difference of a cube and a sphere is

$$d_{\Delta}(\vec{p}) = \max(d_{\text{cube}}(\vec{p}), -d_{\text{sphere}}(\vec{p})).$$

By combining SDFs through these operations, we obtain a single SDF that describe the entire scene.

Each pixel is rendered by casting a ray from the camera into the scene and advancing it step-by-step according to the signed distance function of the scene. The ray direction \vec{v} is computed by normalizing the vector from the camera to the location of the pixel in the view plane. The ray is then advanced iteratively according to the signed distance function of the scene. We start at the camera position \vec{p}_0 and repeatedly update

$$\vec{p}_{n+1} = \vec{p}_n + d(\vec{p}_n) \vec{v},$$

where $d(\vec{p}_n)$ is the distance to the nearest surface at position \vec{p}_n . If $d(\vec{p}_n)$ falls below a small scene-scale-dependent threshold (e.g., $\varepsilon = 10^{-4}$ units of length), we consider the ray to have hit an object and assign the corresponding pixel a color based on the object's material or shading model. If the ray travels beyond a predefined maximum distance without hitting any object, it is assumed to have missed all surfaces, and the pixel is assigned a background color.




Figure 3: Illustration of the extended ray marching algorithm. (a) The core ray marching algorithm, where a ray iteratively approaches the surface of an object represented as a signed distance function (SDF). (b) Rays are translated and rotated when they encounter walls, with an offset to prevent immediate reintersection.

Once a ray hits an object, the final color at the corresponding pixel can be computed using a desired lighting or shading model. In our simulations, we use a simple coloring based on normal vectors. We approximate $\nabla d(\vec{p})$ at the hit point \vec{p} , where $\nabla d(\vec{p})$ is the gradient of the signed distance function of the object and thus points in the direction normal (perpendicular) to the object's surface $d(\vec{p}) = 0$. These normal vectors can then be mapped to RGB color values directly or combined with classic shading models like Blinn-Phong. More advanced methods such as physically based rendering can also be used.

Ray Marching on Flat Surfaces

We present an adaptation of ray marching to flat surfaces. For our primary example, we consider a specific translation surface L , shown in Figure 4a, whose shape resembles the letter “L”. The surface is defined by identifying the opposite edges by translations. While we focus on this concrete case, our approach readily generalizes to other flat surfaces such as unfolded polyhedra and mirror rooms.




Figure 4: Visualization of the translation surface L and its three-dimensional thickening. (a) The translation surface L with identified edges marked in matching colors. Parallel opposite edges are identified by translation. (b) Three-dimensional visualization of L with identified opposite walls shown in matching colors.

To create an immersive visualization, we lift L into three dimensions by adding a Euclidean height component to thicken the surface. This transforms our surface into a three-dimensional “room,” where parallel walls are identified by translations, leaving the vertical component unchanged. For simplicity, we continue to refer to this three-dimensional space as a “translation surface” and use the same notation L .

In our ray marching setup, the scene includes both a test object—a cube with a sphere removed—and the walls of L , as shown in Figure 4b. Rays are traced until they either intersect the test object or encounter a wall. For wall hits, the ray’s position is updated based on the gluing rules of the surface.

This process introduces a technical challenge: standard signed distance functions (SDFs) only provide the distance to the nearest surface but do not indicate which object the ray is interacting with. To address this, we augment the SDF with a flag that distinguishes between the walls and the test object.

When a wall is hit, the appropriate translation is applied, and the ray’s position is adjusted to avoid immediate re-intersection. Specifically, the translated point is placed at a safe distance greater than ε from the new wall, as illustrated in Figure 3b.

Our implementation targets geometries obtained by thickening a finite collection of Euclidean polygons and assigning boundary rules (translations, rotations, reflections) to the vertical walls and optionally the floors and the ceilings of the thickened polygons:

- *Translation surfaces:* matching parallel walls are identified by horizontal translations.
- *Unfolded polyhedra:* matching walls are identified by horizontal translations and rotations.
- *Mirror rooms:* walls act as ideal mirrors.
- *Translation prisms:* translation surface rooms in which the floors and ceilings are identified by a vertical translation, so that the height coordinate becomes periodic.

Visualizations and more intuitive descriptions of translation surfaces and unfolded polyhedra are presented

in the *Gallery* section, and mirror rooms and translation prisms are presented in the *Extended Gallery* section (Appendix A).

Our interactive visualizations were implemented using WebGL through the three.js framework, which provides an efficient WebGL environment. The visualization interface includes custom JavaScript code for handling keyboard and mouse inputs, allowing users to control the camera and explore the flat surfaces from different perspectives. The core visualization logic, including the ray marching algorithm and the geometric representations of the flat surfaces, was implemented in fragment shaders following the GLSL specification. Our shader-based implementation can be easily ported to any modern graphics or game engine that supports programmable fragment shaders and user input, such as Unity, Unreal, or Godot. We make our simulations and source code available online [11, 12]. To run the simulation locally, the code repository can be cloned and hosted as a local server (e.g., using Python’s `http.server` module).

Simulation Controls

Our simulation implements the following controls:

<i>W / A / S / D</i>	move forward / left / backward / right
<i>E / Q</i>	move up / down
<i>Mouse</i>	rotate the camera

We note that the key and mouse mappings mirror standard first-person video game controls, and in our experience, feel immediately intuitive to many users.

Use in Outreach and Public Engagement

We use our simulations at science fairs, math festivals, and other public events to illustrate flat surfaces and their surprisingly rich mathematics. We begin with the simplest example: forming a torus by gluing opposite edges of a rectangle, and making a comparison to Pac-Man, where crossing one side of the screen makes one reappear on the other. Next, we highlight more complex surfaces and mirror rooms, drawing parallels to Leonardo da Vinci’s mirrored chambers and Yayoi Kusama’s infinity mirror rooms. These immersive experiences naturally lead to discussions of ongoing research on flat surfaces and billiards, opening deeper conversations about geometry, topology, and dynamical systems.

Gallery

In this work, we present a range of spaces rendered using our ray marching approach, including (thickened) translation surfaces, mirror rooms, (thickened) unfolded polyhedra, and translation prisms. Each example demonstrates how side identifications produce surprisingly immersive and repeated visual effects. We give the walls a subtle glass-like translucency so that viewers can recognize the sides of the base polygons while still seeing the repeating geometry beyond. This translucency creates a faint fog effect in the distance, visible in the rendered images.

We illustrate several representative examples. Figure 5 shows two translation surfaces: an L-shape and a double pentagon. The surfaces are constructed by identifying pairs of parallel sides through translations. Figure 6 shows a mirror room formed by reflecting rays across the walls of a triangle. Figure 7 shows an unfolded polyhedron arising from flattening the six square faces of a cube and re-gluing matching edges via translations and 90° rotations, making it the flat surface analogue of the sphere. Lastly, Appendix A extends this gallery with more visualizations of *mirror rooms*, *translation prisms*, and *thickened singularities*.

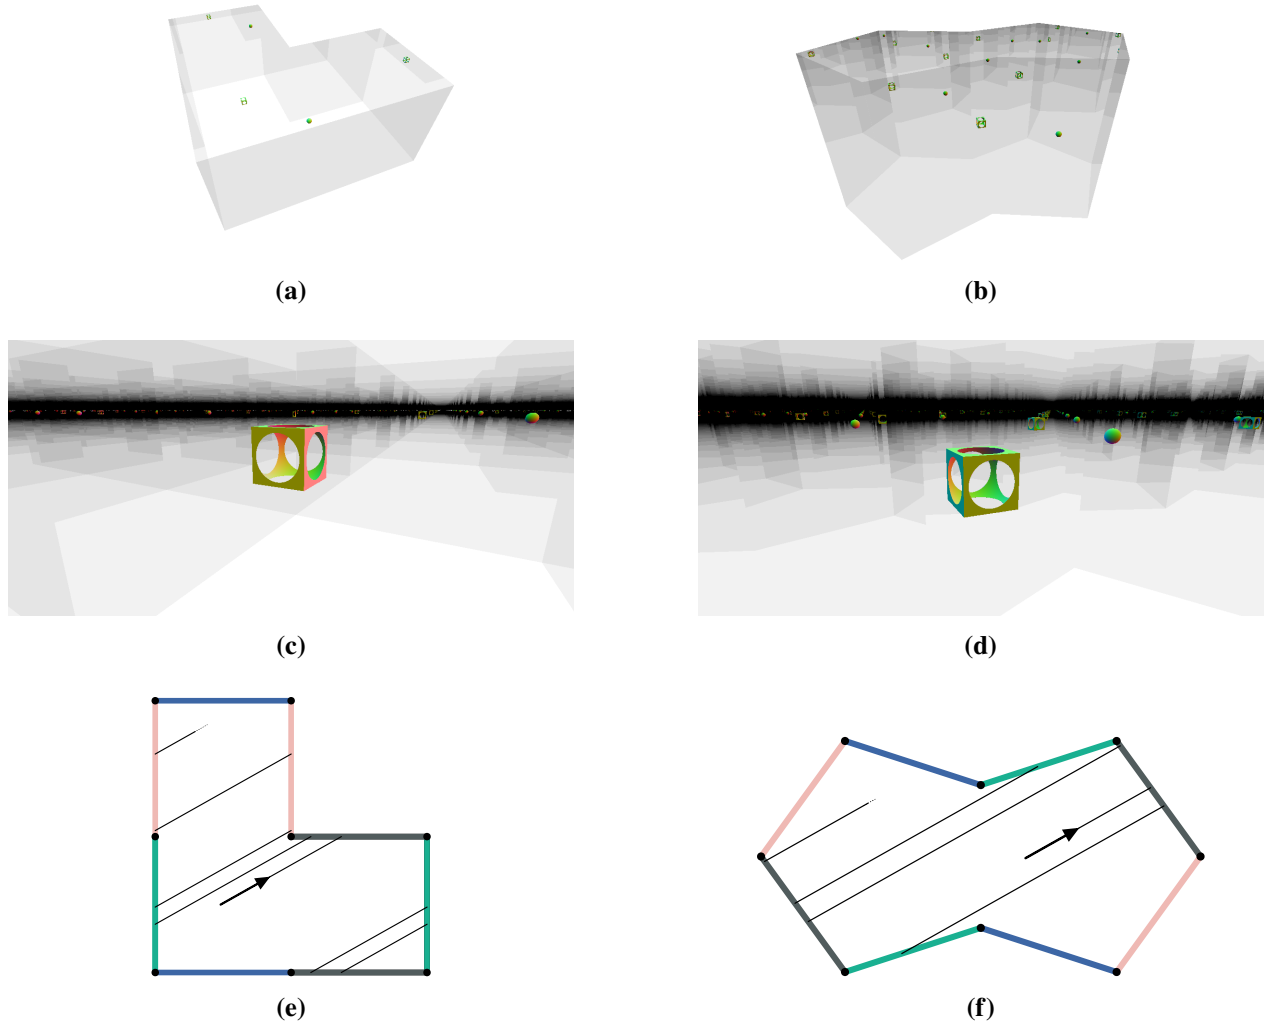


Figure 5: Visualization of translation surfaces. (a) Top-down view of a simple L-shape. (b) Top-down view of a double pentagon. (c) Inside view of a simple L-shape. (d) Inside view of a double pentagon. (e) Floor plan of the L-shape showing the path of a light ray. (f) Floor plan of the double pentagon showing the path of a light ray.




Figure 6: Visualization of a mirror room with a nearly irrational triangle base. (a) First-person view inside the mirror room. (b) Floor plan showing the base triangle and part of a path of a light ray.

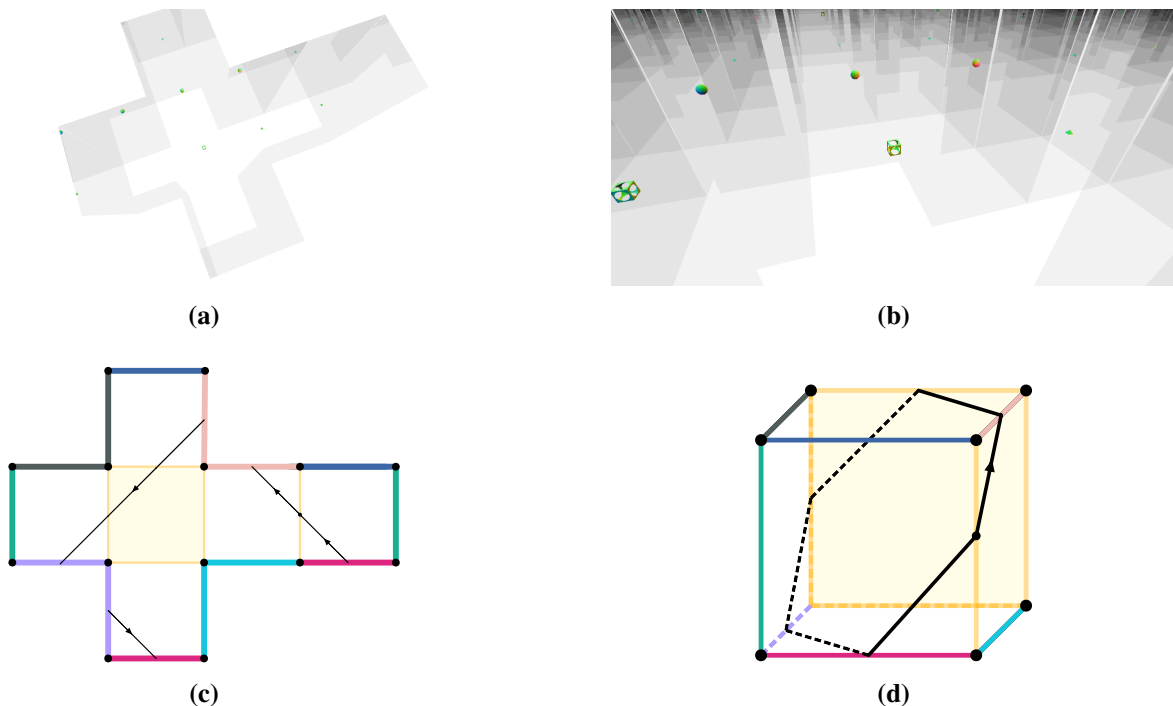


Figure 7: Visualization of an unfolded cube with an added Euclidean height component. (a) Top-down view of the net. (b) First-person perspective inside the unfolded net. (c) Floor plan of the cube net showing how edges are identified and the path of a light ray. (d) The same light ray path shown on the folded cube.

Summary and Conclusions

In this work, we use ray marching to visualize flat surfaces, providing an immersive, first-person experience of their intrinsic geometry. Future work could explore several directions: Artistically, real-time interactive installations within these surfaces could bring them to life as engaging public exhibits. Technically, implementing analytic solutions could enhance performance and rendering efficiency. Lastly, extending this work to affine or dilation surfaces would broaden the accessibility of mathematical spaces to the public through immersive visualization.

Acknowledgements

This work was carried out at the Heidelberg Experimental Geometry Lab at the University of Heidelberg and the Mathematics Lab at the Max Planck Institute for Mathematics in the Sciences. We thank and acknowledge Mara-Eliana Popescu for her contributions to an earlier version of this work, and Rémi Coulon for valuable discussions at the beginning of the project. Figure 2 was created by Ricardo Waibel in collaboration with the second author. Fabian Lander was partially supported by the Collaborative Research Center SFB/TRR 191 - 281071066 (Symplectic Structures in Geometry, Algebra and Dynamics) while working on this project. We sincerely thank the reviewers for their valuable comments and suggestions, which significantly improved the clarity and quality of this manuscript.

References

- [1] J. Applin. *Yayoi Kusama: Infinity Mirror Room – Phalli’s Field*. ser. One Work. Afterall Books, 2012.

- [2] J. S. Athreya, N. Bédaride, W. P. Hooper, and P. Hubert. “Linear Flows on Translation Prisms.” *arXiv preprint arXiv:2504.08955*, 2025.
- [3] R. Coulon, E. Matsumoto, H. Segerman, and S. Trettel. “Non-Euclidean Virtual Reality III: Nil.” *Proceedings of Bridges 2020: Mathematics, Art, Music, Architecture, Education, Culture*. C. Yackel, R. Bosch, E. Torrence, and K. Fenyvesi, Eds. Phoenix, Arizona: Tessellations Publishing, 2020. pp. 153–160. <http://archive.bridgesmathart.org/2020/bridges2020-153.html>.
- [4] R. Coulon, E. Matsumoto, H. Segerman, and S. Trettel. “Non-Euclidean Virtual Reality IV: Sol.” *Proceedings of Bridges 2020: Mathematics, Art, Music, Architecture, Education, Culture*. C. Yackel, R. Bosch, E. Torrence, and K. Fenyvesi, Eds. Phoenix, Arizona: Tessellations Publishing, 2020. pp. 161–168. <http://archive.bridgesmathart.org/2020/bridges2020-161.html>.
- [5] R. Coulon, E. A. Matsumoto, H. Segerman, and S. J. Trettel. “Ray-Marching Thurston Geometries.” *Experimental Mathematics*, vol. 31, no. 4, 2022, pp. 1197–1277.
- [6] R. A. Drebin, L. Carpenter, and P. Hanrahan. “Volume rendering.” *SIGGRAPH Comput. Graph.*, vol. 22, no. 4, Jun. 1988, p. 65–74.
- [7] J. D. Foley. *Computer Graphics: Principles and Practice*. Addison-Wesley Professional, 1996. vol. 12110.
- [8] J. C. Hart. “Sphere Tracing: A Geometric Method for the Antialiased Ray Tracing of Implicit Surfaces.” *The Visual Computer*, vol. 12, no. 10, 1996, pp. 527–545.
- [9] V. Hart, A. Hawksley, E. Matsumoto, and H. Segerman. “Non-Euclidean Virtual Reality I: Explorations of H^3 .” *Proceedings of Bridges 2017: Mathematics, Art, Music, Architecture, Education, Culture*. D. Swart, C. H. Séquin, and K. Fenyvesi, Eds. Phoenix, Arizona: Tessellations Publishing, 2017. pp. 33–40. <http://archive.bridgesmathart.org/2017/bridges2017-33.html>.
- [10] V. Hart, A. Hawksley, E. Matsumoto, and H. Segerman. “Non-Euclidean Virtual Reality II: Explorations of $H^2 \times E$.” *Proceedings of Bridges 2017: Mathematics, Art, Music, Architecture, Education, Culture*. D. Swart, C. H. Séquin, and K. Fenyvesi, Eds. Phoenix, Arizona: Tessellations Publishing, 2017. pp. 41–48. <http://archive.bridgesmathart.org/2017/bridges2017-41.html>.
- [11] F. Lander and D. Taha. “Ray Marching Flat Surfaces Demo.” <https://fabianlander.github.io/apps/raymarchingflatsurfacesapp/>. 2025.
- [12] F. Lander and D. Taha. “RayMarchingFlatSurfaces.” <https://github.com/FabianLander/RayMarchingFlatSurfaces.git>. 2025.
- [13] R. Nelson and H. Segerman. “Visualizing Hyperbolic Honeycombs.” *Journal of Mathematics and the Arts*, vol. 11, no. 1, 2017, pp. 4–39.
- [14] K. Perlin and E. M. Hoffert. “Hypertexture.” *Proceedings of the 16th annual conference on Computer graphics and interactive techniques*. 1989. pp. 253–262.
- [15] I. Quilez. “Rendering Worlds with Two Triangles.” *Proceedings of the 2008 NVSCENE Demoparty Technical Seminar*. 2008. <https://www.iquilezles.org/www/material/nvscene2008/rwwtt.pdf>.
- [16] I. Quílez. “Distance Functions.” <https://iquilezles.org/articles/distfunctions/>. 2015.
- [17] Z. Wang. “3D Representation Methods: A Survey.” *arXiv preprint arXiv:2410.06475*, 2024.
- [18] J. Weeks. “Real-Time Rendering in Curved Spaces.” *IEEE Computer Graphics and Applications*, vol. 22, no. 6, 2002, pp. 90–99.
- [19] J. R. Weeks. “Real-Time Animation in Hyperbolic, Spherical, and Product Geometries.” *Non-Euclidean Geometries: János Bolyai Memorial Volume*, 2006, pp. 287–305.
- [20] A. Zorich. “Flat Surfaces.” *arXiv preprint math/0609392*, 2006.