Computers & Graphics 114 (2023) 210–218

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Special Section on SMI 2023

Bending the light: Next generation anamorphic sculptures

Louis Pratt, Andrew Johnston, Nico Pietroni *

Faculty of Engineering and Information Technology, University Of Technology Sydney, Broadway, Ultimo NSW 2007, Australia

ARTICLE INFO

Article history:
Received 17 May 2023
Accepted 30 May 2023
Available online 10 June 2023

Keywords:
Anamorphic
Art
Humanities

ABSTRACT

This paper presents a new method for generating artworks that extends the classical anamorphic archetype to use freeform reflective and refractive media and 3D surfaces instead of images. The methodology uses a mix of raytracing and surface deformation techniques to determine the proper deformation the object should undergo to be corrected by the optical tool once viewed by the observer in a specific location. Our approach also includes an optimization process that modifies the point of view and the location of the target image to avoid unwanted optical effects or occlusions. We successfully tested our technique on several practical examples and employed it to produce actual artworks.

1. Introduction

The word anamorphosis derives from the Greek ἀναμορφω, which translates as " transform ". A simple way to understand anamorphosis is to imagine your shadow cast from the sun and how this shadow might look when it is compressed or elongated at different times of the day, as depicted in Fig. 3. This projection of your shadow onto the ground exemplifies (oblique) anamorphosis. In this specific case, the viewer can correct the distorted shadows by moving to a specific point of view. Several artists used this principle to produce a target image through an optical distortion [1]. Arguably, one of the most famous examples is by Hans Holbein, the younger, with his painting "The Ambassador's" (1553), which hangs in the London National Gallery.

Anamorphic art can be classified into two distinct groups; Oblique and catoptric. In *catoptric anamorphosis*, the image is reconstituted via the reflective surface. Common examples include placing a conical or cylindrical mirror on top of the drawing or painting to transform a flat distorted image into an undistorted image. Contrary to oblique anamorphosis, the corrected image can be viewed from many different angles rather than a fixed perspective as in the case of oblique anamorphosis. An intriguing historical example of catoptric anamorphosis is its use to conceal an image like the 'secret' Catoptric anamorphic image of Prince Charles Edward Stuart painted on a wooden board to serve drinks (see Fig. 4). When a mirrored cylindrical drinking cup was placed correctly on the tray, the drinker could then toast to the exiled Prince and the Stuart family. The controversial image would be obscured without the mirrored cylinder, keeping the secret safe. In the context of art, oblique anamorphosis has been extensively explored both in painting, sculpture, and installation. Yung Hee Jo [2] and Jonty Hurwitz [3] have extended anamorphosis to sculpture rather than 2D images.

This paper presents a novel computational framework to generalize anamorphosis to create sculptures that interact with freeform mirrors or refractive objects such as those made of glass or water. Given a target sculpture and a set of refraction or reflective media, we simulate how the light interacts with the media using a raytracing technique to derive the target position of each vertex of the deformed sculpture such that the viewer will perceive the correct image from a given point of view. We implemented a framework that includes some interactive tools that allow the user to verify the result in real-time, optimize the sculpture's position automatically and modulate some parameters to fine-tune the installation to its final setup, ready for production. We initially verified the accuracy of the produced sculpture by simulating the final effect with available commercial software, then we fabricated tangible sculptures and exhibited them publicly in art prizes or at major art fair events. One artwork has been selected as a finalist for the prestigious Wynne Prize 2023 and exhibited at the art gallery of NSW [4]. Before then, our artwork was selected for exhibition at Sydney Contemporary 2022 [5], Fisher's Ghost Art Award 2022 [6], and North Sydney Art Prize [7]. One intriguing byproduct of our installation is that when we employed concave mirrors, the produced reflection is perceived in front of the mirror, creating a holographic effect.

2. Previous work

In the last decade, a significant amount of effort has been focused on using computational design and 3D printing to create

* Corresponding author.
E-mail addresses: art@louispratt.com (L. Pratt), andrew.johnston@uts.edu.au (A. Johnston), nico.pietroni@uts.edu.au (N. Pietroni).

https://doi.org/10.1016/j.cag.2023.05.023
0097-8493/

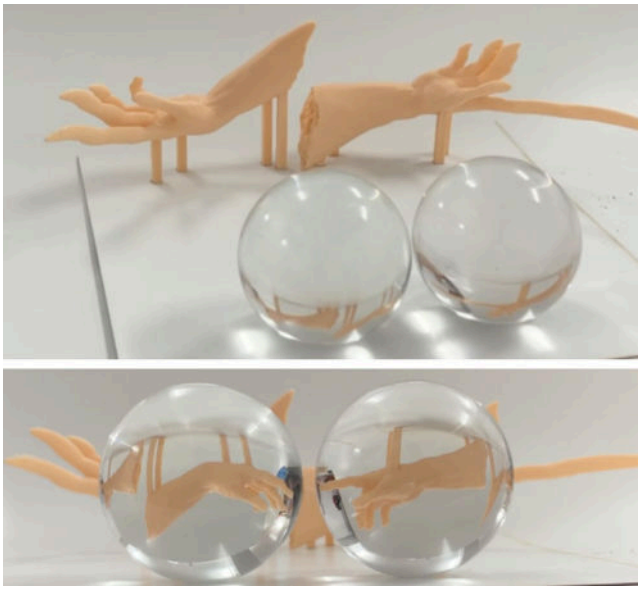

Fig. 1. An example of anamorphic sculpture using refractive objects.

alternative/artistic representations of 3D objects [8]. Some of these techniques include perspective distortions of 3D objects [9] or animations [10], including non-linear ray tracing [11]. Some artistic approaches composes multiple views [12]. Other technologies deform images to mimic the artist's style [13]. This section will first focus on technologies applied in the creation of anamorphic sculptures. Then we will also provide a quick overview of the technologies that exploit light refraction and shadow effects.

2.1. Anamorphic modeling

Since the first anamorphic illustration by Leonardo da Vinci in the Codex Atlanticus [14], anamorphism has raised significant interest in art. Generating an anamorphic drawing by hand is an arduous and error-prone process as it requires the artist to draw, thinking of the final optical effect from a different point of view. Hence, anamorphic art has remained confined for centuries to a few examples. Some papers formalized the anamorphosis from a mathematical point of view [15]. However, thanks to the advent of digital tools to manipulate images and 3D shapes, anamorphism has become more accessible. One of the first techniques to compute anamorphic distortion of images uses cylindrical mirrors [16]. Raytracing [17] is used to simulate the interaction of the light with simple mirrored geometric primitives and projects the color on each pixel of the final image. Anamorphic images have also been generated on large curved walls by placing modular elements using industrial robots [18]. Some studies [19] analyzed the perceptual problems relative to the viewer's position for Plane Geometric Anamorphosis.

Another class of methods employed anamorphism to deform 3D shapes such that the deformation vanishes when the shape is observed from a certain point of view [20]. A similar effect can be obtained by using continuous smooth deformation tools defined over the volume [21]. Appearance Mimicking surfaces [22] optimize the compression of surface normals to fit a certain volume such that they are perceived correctly from a given point of view. Anamorphism can also be used to create 3D rationalization of surfaces that generates Escher-like impossible figures when seen from a specific viewpoint [23]. To our knowledge, the first approach that included mirror and 3D shapes is the approach

Fig. 2. Examples of anamorphic sculptures using free-form mirrors: Left: The correct images; Right: the deformed sculpture as seen from a different point of view.

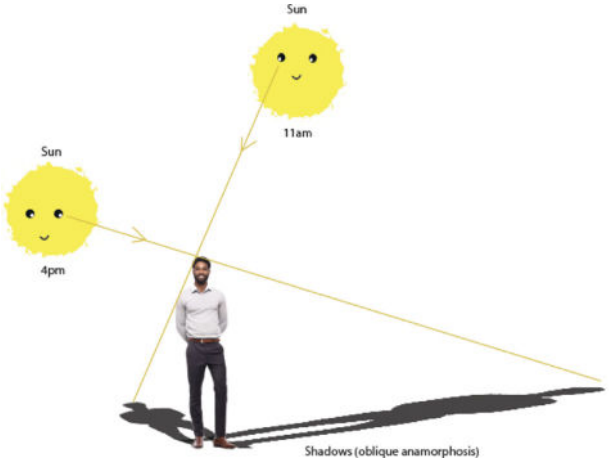

Fig. 3. A simple sketch to illustrate anamorphosis. The shape of the shadow changes with respect to the position of the light source.

proposed by Francesco de Comité in 2011 [24]. However, such an approach is limited to regular geometric primitives, such as spheres. The approach was extended in 2015 to work also with images [25]. A more recent method proposes a general computational framework for mirror cup and saucer art design [26]. An input shape is printed on a deformed surface so the user can observe the corrected reflection when a mirror cup is placed in a predefined spot.

To our knowledge, our framework is the first one that has been designed to work with free-form 3D shapes and also allows to use of both reflective and refractive materials. Furthermore, none of

Fig. 4. The 'secret' Catoptric anamorphic image of Prince Charles Edward Stuart.

the previous methods offers advanced control and functionality to the artist like the tool we present in this paper.

2.2. Playing with lights and shadows

Beyond anamorphosis, there is a long list of methods in computer graphics that employ light and shadows as the fabrication domains. Low reliefs are the most common example of this strategy put into practice. Low reliefs are height fields that, when lit, convince the viewer to perceive a complex 3D shape. They are an effective and durable method of depicting a 3D shape in a nearly 2D space. When it comes to manufacturing, low-reliefs offer the advantage of significantly reducing production costs, as they can be created using subtractive 2D milling techniques. High-reliefs are similar but relax the condition of having a height field keeping part of the original full 3D sculpting volume and allowing overhangs [27].

The current advancement in fabrication technologies allows for predicting and controlling how objects interact with lights. Additive manufacturing allows for precise and local control of the material employed in the fabrication process. Therefore it is possible to design optical effects by smoothly varying the optical properties within the volume. Refraction can be controlled by tessellating flat panels using a vocabulary of sticks made of acrylate resin [28]. The idea is that each stick will refract the light in a specific direction so that composing them in a specific manner makes it possible to control the final optical effect. A similar approach uses multi-material 3D printing to fabricate arbitrary surfaces with embedded optical fibres [29]. Each fibre is printed using materials with different refraction indexes, low refraction outside and high refraction inside. This technology forces the light to follow complex curved paths, transforming the object's surface into a custom-shaped display. The Magic Lens [30] facilitates the revelation of multiple images from different points of view. Another class of techniques optimizes the shape of transparent objects to produce a given image on a flat surface when the light hits them. These techniques require high precision milling processes that can be employed to fabricate either reflective micro facets [31] or to modify refractive objects [32–35] to concentrate light and control caustics projected on flat surfaces.

Shadows are another media that is employed to create evocative effects. Shadow Art [36] is a technique to retrieve a single 3D shape capable of projecting multiple silhouettes when illuminated from different directions. A more sophisticated approach allows the generation of perforated lampshades that focus the light along with predefined patterns and project target images on flat surfaces [37]. Recent works use cables and rods to provide interesting effects like levitating objects [38] or complex tensegrities [39]. However, these works rely on significantly different physical backgrounds to the proposed method.

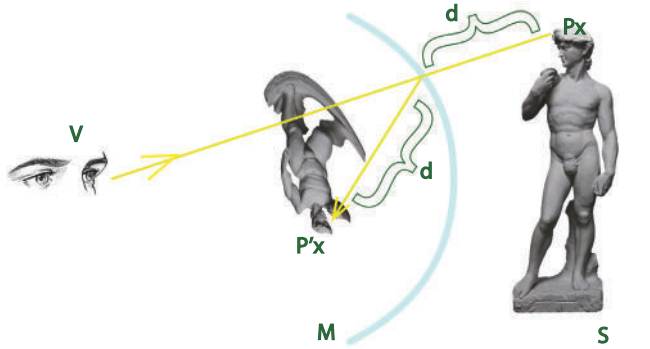

Fig. 5. A diagram illustrating our anamorphic modeling approach: the sculpture is deformed such that each point stays in the target position expressed by the reflected ray and the distance d.

3. Anamorphic deformation

We designed a system to create anamorphic sculptures in a general setup allowing the user to employ free-form reflecting surfaces, refractive objects or their combination. Our system requires as input all the elements involved in the optical effect, the point of view of the observer and the desired sculpture visualization. Then, it will output the position and the deformation of the sculpture that will produce the desired rendering when observed from the predefined point of view. Each of the involved entities can be either reflective or refractive.

To understand the main idea of our approach, let us refer to the 2D diagram in Fig. 5. That simple setup includes a sculpture S, a single curved reflective surface M and the observer position v. The main task of our deformation procedure is to derive for each point $p_x \in S$ the deformed 3D position p'_x such that p'_x will be reflected over M along the direction that connects v to the original position of p_x.

To produce the desired visual effect, It is sufficient to place p_x along the reflected direction. However, to obtain the correct shading of the visualized image, we should respect the local position of each point with the surrounding ones. A simple solution is to compute the distance d between a point on the target shape and its reflection point, then use this distance to place the deformed point over its reflected ray as shown in Fig. 5. There are infinite solutions that place the mesh's vertices along the reflected rays. To explore and choose one solution, we allow the user to control two parameters that we will explain in detail in Section 3.2.

We use a similar approach for refractive objects. In such a case, instead of computing the reflection line, we simulate the refraction using the well-known *Snells Law*. Snell's Law calculates the change in light direction as it moves between two media with varying refractive indexes. It is based on the difference in velocity of light propagation in the media. That effect is used in many commercial ray tracer engines. Snell's law states that, for a given pair of media, the ratio of the sines of the angle of incidence α_0 and angle of refraction α_1 is equal to the ratio between refractive indices (n_1/n_0) of the two media:

$$\frac{\sin\alpha_0}{\sin\alpha_1} = \frac{n_1}{n_0} \tag{1}$$

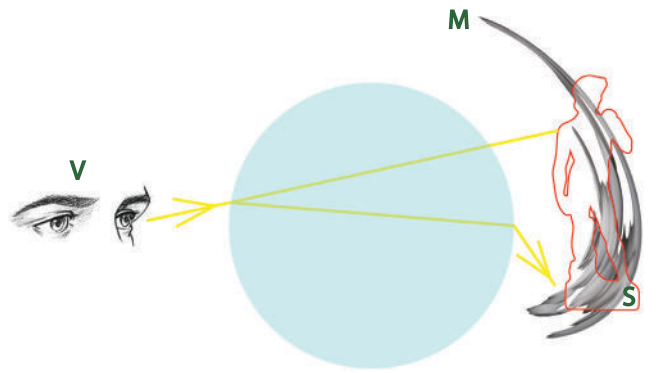

Fig. 6. A diagram illustrating our anamorphic modeling approach works with refractive media: Similarly to the case of the reflective surface each point is moved to the target position by simulating the light transportation. However in this case the sculpture will be placed beyond the refractive media.

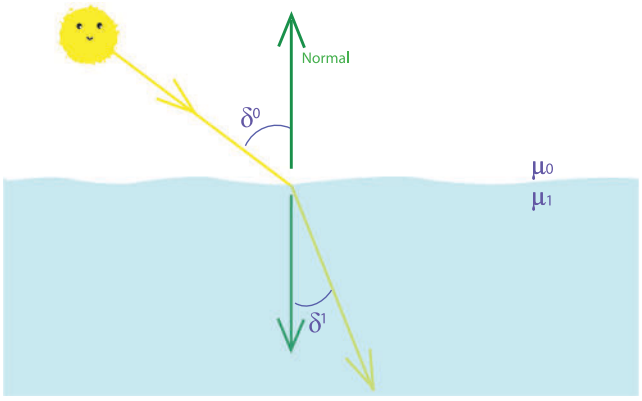

Fig. 7. A simple schema illustrating the Snell law modeling the propagation of the light through refractive media.

If we know the refractive indices n_0 and n_1, we can use Eq. (1) to compute how the light change direction when it enters and exit from the media. While Fig. 7 shows how the directions rotate in 2D, for the 3D case, we need also to define a rotation axis. In 3D, the rotation axis is orthogonal to the plane formed between the direction of the ray and the surface normal.

Similarly to the case of the reflective media, we compute the distance d from the first occurring intersection and the target position. We then position the target position along the last ray exiting from the refractive media at distance d (see Fig. 6).

While this strategy ensures, in theory, a correct visualization, the final physical optical effect can still suffer from some global artifacts like undesired reflections or self-occlusions. We will explain these effects in Sections 3.3–3.5. To cope with these artifacts, we provide two parameters to tune the distance of the sculpture from the optical media while keeping its local proportion correctly. Alternatively, these artifacts can be corrected by slightly changing the position of the viewer. We will explain these parameters in Section 3.2 and an automatic procedure to alleviate some of these artifacts in Section 3.6.1.

3.1. General setup

The proposed approach can be extended directly to a setup that combines multiple reflective or refractive objects. In that case, we need to trace a ray from the viewpoint, looking towards each vertex of the visualized sculpture. Then, we simulate how this ray of light moves between the optical elements of the scene utilizing the optical properties of the materials involved (reflection and refraction). Then move the vertex along the last reflected or refracted ray, using the distance from the first intersection to the target sculpture position.

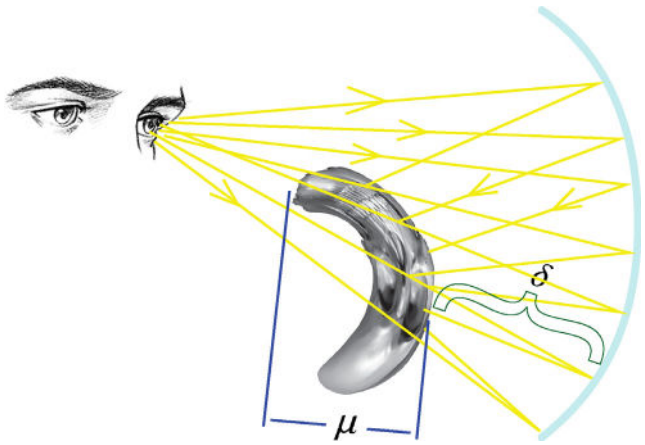

Fig. 8. The two parameters $\delta > 0$ and $\mu > 0$ allows to control the distance of the sculpture to the mirror and its size.

3.2. Control parameters

We offer the artist several degrees of freedom to control the final result. First, we can tune the position and orientation of both the observer and the target geometry. However, if we utilize very curved mirrors or complex refractive objects, the final deformation might deliver large surfaces, and the fabrication cost can be prohibitive. Also, the position of the final sculpture might not fit into the predefined workspace. To produce a correct visualization, we need each vertex of the deformed sculpture (i) to be aligned with the computed final ray and (ii) to locally preserve the shape to ensure a proper rendering. Hence, we can move each vertex along their rays and preserve both (i) and (ii) if we are able to move them coherently. We implemented a simple solution that consists of two parameters: the first one, $\delta > 0$, controls the minimal distance of the sculpture to the ray emerging from last interaction with a mirror or a lens, and the second $\mu > 0$ scales linearly all the distances of the points along the rays (within the interval between the minimal and the maximal distance). This way, the user can control the final size of the sculpture to match the physical constraints involved in the fabrication process (see Fig. 8). We set δ as a fraction of the diagonal of the bounding box of the input sculpture. Values between 0.5 and 1 usually work well in practice, while we set μ to 1 by default.

3.3. Undesired reflections

While the system ensures that a vertex will appear precisely in the right spot when rendered from the specified point of view, the same vertex might also be visible in a different location. This undesired effect might interfere with the intended visual impact that an artist would like to construct. While this effect cannot occur in the presence of convex mirrors (currently employed in anamorphic art), it becomes a common artifact when we use concave mirrors or freeform refractive objects.

The effect is shown in Fig. 9. We can test if this artifact is present by raytracing with the target sculpture from the observer's point. To speed up the ray tracing process, for each optical media visible by the spectator, we derive a simple UV map as the screen coordinates of each vertex. We regularly sample the

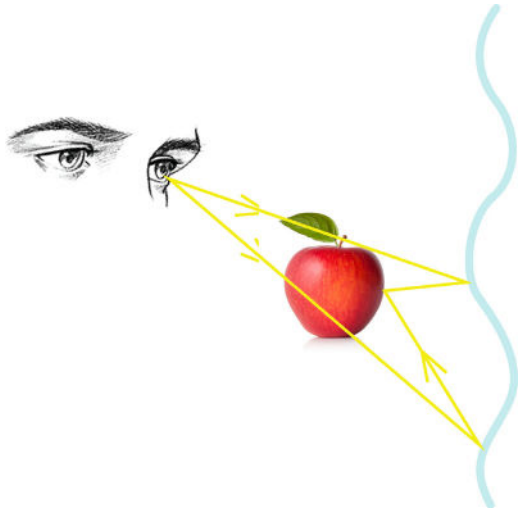

Fig. 9. Double rendering artifact generated by one point on the deformed sculpture that is reflected in two positions at the same time.

surface in screen space and we trace one ray for each sample, emanating from the view position and looking towards the sampled points. In this first step, we derive the target image by only considering the samples in correspondence with target image. Similarly, we produce the real-world image by physically-based ray-tracing the scene using all the samples in the entire surface of the optical medias. In this step we consider all the bounces produced by the light transportation. Then we can compare the two images to understand the accuracy of the created effect. Notice that to make the comparison effective and easy to compute, we store, for each pixel, the index of the visualized face. We then compare the index of faces produced by the two ray tracing operations (the one that must be visualized with the one that is produced by the final ray trace). This is a more robust and simple strategy than comparing directly the produced pixel colors.

3.4. Self occlusions

Even if the computed deformation produces the correct target image on the optical media, the final visual effect might not be visible to the observer since the deformed sculpture can impede it. This effect is visible in Fig. 10.a. This artifact might be pretty common when using mirrors. Indeed mirrors end up placing the deformed sculpture between the observer and the reflected image. In particular, if the mirror is convex, the final deformation will enlarge the reconstructed sculpture. These two factors, the positioning and the magnification, might frequently introduce these occlusion artifacts. Notice the traditional anamorphic sculptures employing cylindrical mirrors suffer from this artifact in most cases.

There is another possible way the sculpture can be not visible as a whole. The deformed sculpture is deformed in a way such that it occludes itself, making some part of its surface not visible. We illustrate this effect in Fig. 10.b. If we use the classical funny mirror switching from convex to concave regions, it is almost impossible to avoid this artifact. The sculpture can be placed very close to the mirror, but it will occlude the reflection generating the other occlusion artifacts we explained above. Generally, every mirror switching from concave to convex might induce this artifact if the target sculpture also spans across these different regions.

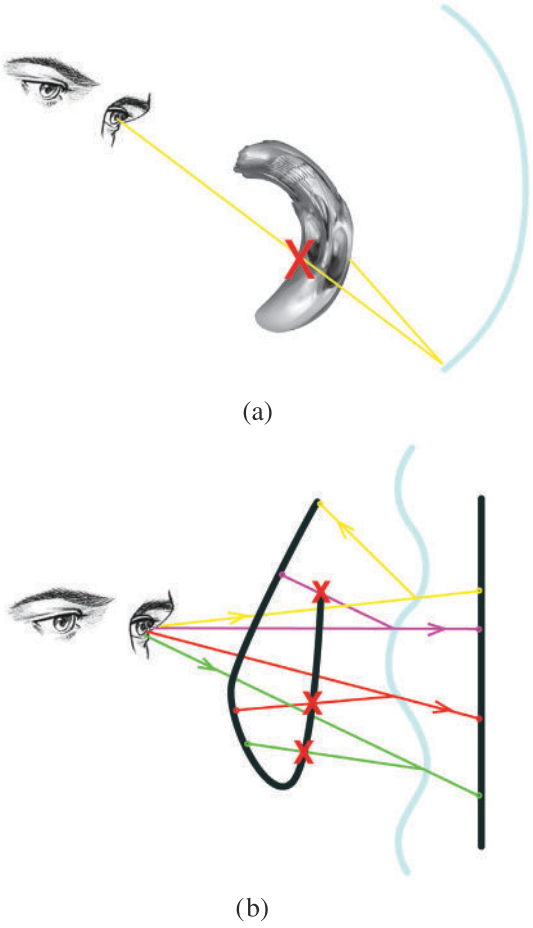

Fig. 10. (a) The sculpture occludes the point of view so the observer cannot see the effect on the mirror; (b) For complex mirror shape, with changes of concavity, the sculpture can self occlude.

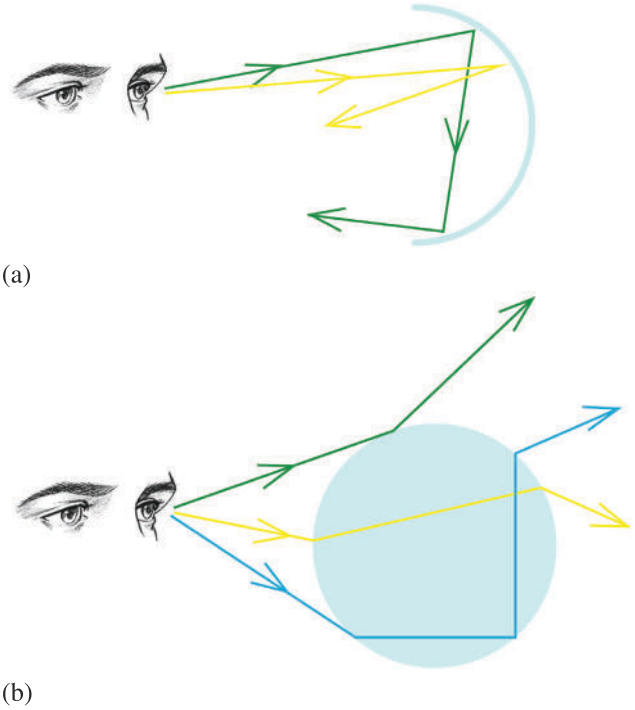

Fig. 11. (a) Different paths artifact on a mirrored object; (b) Different paths artifact on a refractive object.

3.5. Different propagation paths

It is possible that a pair of points that are close on the original surface will end up in very far apart places in deformed sculpture. This undesired effect might result in unexpected deformation in the final mesh, causing artifacts in the final rendering. A discontinuity in the ray-traced paths might occur both in the reflective and refractive case. For example, if we are using a concave mirror, then some paths can bounce multiple times on the same mirror before reaching the observer, while others bounce only once (see Fig. 11.a). If we use a refractive media, a traced ray might hit the surface at a critical angle; hence, the ray will bounce as reflected and scatter within the refractive media multiple times before it exits (see Fig. 11.b). This effect is inevitable for specific configurations in the real world using physically fabricated sculptures. However, we can detect the discontinuities and optimize the position of the sculpture to avoid them.

3.6. Optimization

We briefly introduced three sources of artifacts that might happen in our general anamorphic modeling pipeline, the *Undesired reflections*, *Self occlusions* and *Different propagation paths*. Unfortunately it is difficult to predict these events and sometimes even impossible to avoid in practice. For example, every mirror shape that exhibits changes in concavity will tend to produce unwanted reflections. Also, these effects are difficult to foresee due to a complex physical phenomenon that is difficult to differentiate and optimize directly. However, we can adjust the parameters iteratively to minimize these artifacts. To eliminate or mitigate these optical artifacts, we can explore the space of the possible degrees of freedom offered by our setup. The main scope of this optimization step is to minimize these artifacts while introducing as few as possible impacts in the final optical effect.

The most practical choice is to vary the parameters δ and μ since they keep the target effect intact without changing the final intent of the visualized figure. Unfortunately, we must limit these two parameters for practical reasons. Indeed large values of δ and μ can produce massive sculptures, which might significantly impact the production cost. Contrarily, a small δ can make the final sculpture too small and challenging to print at a given accuracy. Similarly, a small μ might introduce unwanted optical effects because we will reduce the variation of normals over the original surface. As a consequence, the rendering of the final sculpture will be perceived as significantly flatter than expected.

3.6.1. Automatic parameters exploration

In most simple cases, a quick manual tuning of these two parameters is usually sufficient to fix the issues (see Fig. 12). Unfortunately, we might need to adjust other parameters in the most complex cases. We can change the position and orientation of the viewpoint, as well we can apply all possible rigid transformations to the target sculpture and the optical media. Given the highly non-linear nature of the problem and the number of degrees of freedom involved, it is impossible to derive the globally optimal set of parameters in a closed-form. Hence we opted for an iterative gradient descent formulation.

For a given view position, we produce a ray-traced buffer containing the portion of the surface of the optical media that is visible by the spectator. Based on that rendering, we can count the number of pixels in the rendered image that corresponds to optical artifacts and write the energy $E_{artifacts}$ that we need to minimize as follows:

$$E_{artifacts} = w_0 R_{pix} + w_1 O_{pix} + w_2 P_{pix} \quad (2)$$

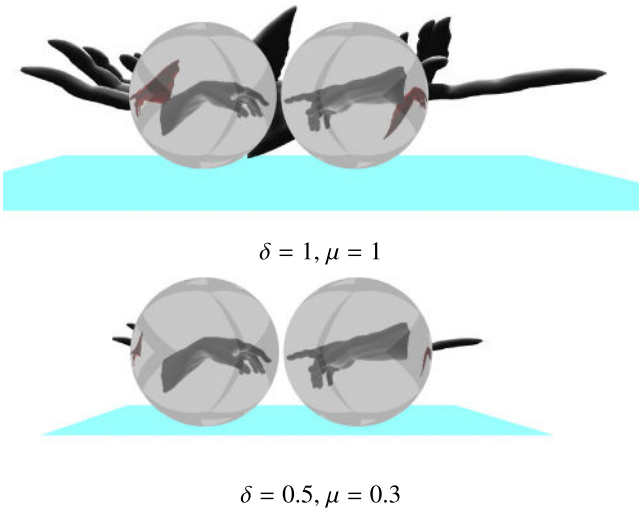

Fig. 12. Different δ and μ parameters can converge to different sculptures while the final optical effect remain the same.

Fig. 13. Auto optimization: the system slightly changed the position where the view of the observer was occluded by the sculpture (left) and obtained a valid configuration (right).

Where R_{pix} is the number of pixels producing *double reflections*, O_{pix} is the number of pixels having *self occlusions* and P_{pix} is the number of pixels producing a different path from the most common one. For practical purposes, we tolerate double reflections more than self-occlusion and different paths since they still allow the complete visualization of the target image. Hence it is sufficient to set up a weight w_0 much smaller than w_1 and w_2. Our optimization procedure performs an iterative search in the space of parameters and updates the current state with the best-explored solution. We first start by exploring nearby values of δ and μ since they do not modify the target visualized image. Every time we test a parameter, we ray-trace the scene and derive the test value of $E_{artifacts}$. If we find a solution that reduces the energy, we update δ and μ with the new values and continue iterating; otherwise, we start exploring the solutions that move the camera or the target object. We then sample nearby translation and rotations along the three axis and we iterate with our optimization procedure until it is not possible to further reduce the energy. The effect of this optimization step is visible in Fig. 13.

Notice that we explicitly omitted freeform deformation on the optical media as this will require the technology for accurately fabricating glass or mirrored objects. Despite theoretically possible, it requires expensive machining, i.e. 6DOF degree CNC milling. We decided first to fabricate the media and then 3D scan it to get its digital representation. This way, we can accurately represent the physical optical objects used in the scene.

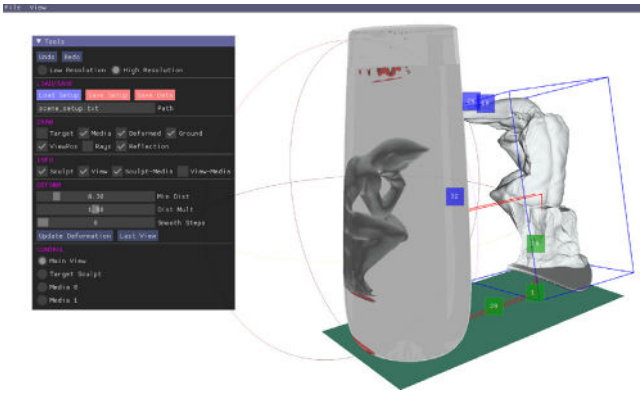

Fig. 14. A screenshot of our interactive editor.

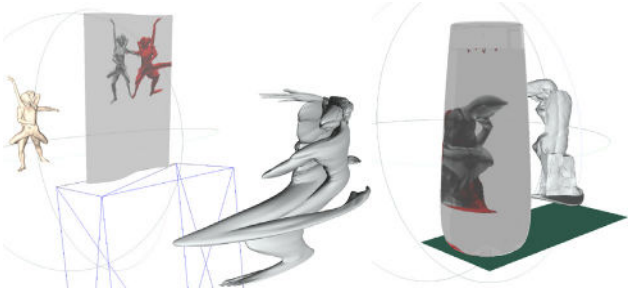

Fig. 15. Two example of undesired reflection artifacts in the final optical effect. To capture the attention of the artist, we report the undesired artifacts in red.

4. Implementation

We integrated our methods into an interactive application that allows the artist to preview the final effect of the produced sculpture, tune the parameters and adjust the result until they achieve the desired effect. We implemented our method in c++ using the *vcg* library (the same used to implement *meshlab* [40]) and *LibIgl* [41]. The system enables users to manipulate the view position and movement of the sculpture in real-time (see Fig. 14). Additionally, users have control over all parameters governing the final deformations, such as $\delta > 0$ and $\mu > 0$, and can quickly visualize the resulting changes. The system also outputs the final size of the sculpture (using a bounding box) and the local arrangement of all the entities in the installation. This feature allows the artist to perceive the space and physical feasibility of the final installation.

Every time a modification occurs, the system uses our raytracing procedure to update the deformed mesh and the produced reflected image in the optical media. We report any artifacts arising in the final optical effect (undesired reflections, self-occlusions or different propagation paths) using a different color in the rendered image to trigger the artist's attention (see Fig. 15).

Ray tracer. As previously discussed in Section 3, to compute the final deformation, we need to displace the mesh vertices along the direction of the last reflected ray. We interpolate the normals over the surface from vertices using barycentric coordinates. This simple strategy provides us with smooth varying bouncing directions over the surface. Unfortunately, since we compute the distance d (used to displace the vertex of the final deformed surface) by intersecting rays with the triangulated surfaces, then the distance values will variate linearly within each triangle and might have abrupt changes across different triangles. To cope with this artifact, we took inspiration from the approach used by Garland and colleagues [42] to approximate a surface for quadric edge collapse decimation. We first initialize a quadric for each vertex. Each quadric expresses the distance to a plane passing on the vertex and oriented along its normal. Quadrics can be summed and averaged to simultaneously represent the distance to multiple planes. Given an intersection of a ray with a triangular face, we express the intersection point using barycentric coordinates. Then we use such weights to average the three quadrics on the face's vertices. This strategy provides a function representing a weighted distance to a set of planes. Then, as [42], we compute the point that minimizes this distance function and use it as the new intersection point. This simple strategy makes the system more invariant to different input triangulation.

Fig. 16. Our hydro-forming tool.

4.1. Fabrication

One of the main difficulties in realizing anamorphic sculptures is fabricating accurate freeform mirrored surfaces. We used lost wax casting and manually polished freeform mirrored surfaces, like the saddle surface shown in Fig. 2. To create large reflecting surfaces, we used Hydroforming, a popular technique employed in the 1940s to manufacture smooth metal surfaces. We developed a hydroforming tool to inflate stainless steel into a dome. Stainless steel was an ideal material for this project because it is durable and can be mirror polished. Fig. 16 shows a picture of our first hydroforming tool capable of generating curved mirrors up to 1 meter in diameter. After fabricating our mirrors, we acquire a digital model for our experiments using standard 3D scanning technology.

5. Results

We used our method to produce several anamorphic sculptures as shown in Figs. 2, 18, 17 and 1. The final rendering of the sculpture in the mirror looks as intended, while the final sculpture is highly deformed and almost unrecognizable (see Fig. 2). One of the most delicate and challenging tasks was to polish the mirror without introducing artifacts over the surface. Although we did our best, the final surface was not perfect, and

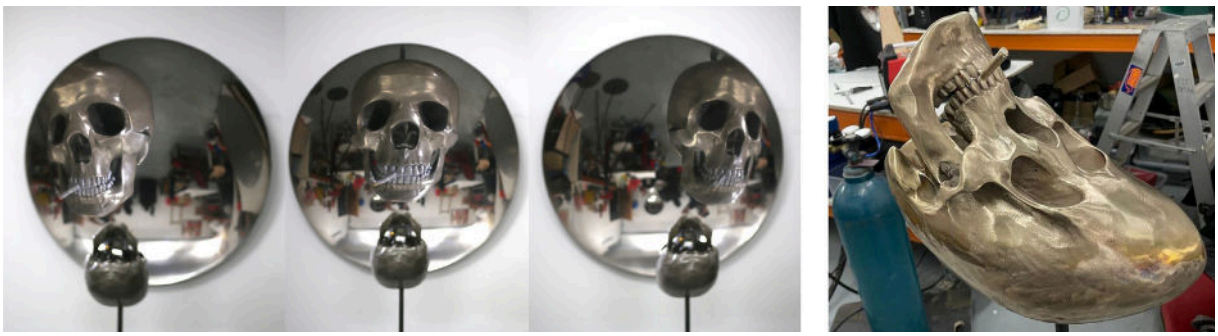

Fig. 17. The final installation at the art Gallery of New South Wales (left) and the deformed skull before the final assembly (right) realized in bronze.

Fig. 18. A distorted sculpture (left) that is correctly visualized through a jar of water (right). Notice the high distortion in the area of the head caused by the high refraction effect happening on the side of the jar. This effect is correctly rendered in the final sculpture.

we are planning perhaps to improve our finishing skills for the future.

In Fig. 1 we experimented with our technique using two glass spheres. Interestingly the hands point outwards where they are not reflected in the spheres, creating a nice final effect. Fig. 18 shows a sculpture that uses a glass vase filled with water. Finally, in Fig. 17 we show an installation with a mirrored concave dome of 1000 mm and a bronze sculpture of a skull to be reflected in it. Surprisingly the concave mirror creates an *holographic* effect that is quite impressive to be visualized on such a large scale: the observer perceives the reflected image as positioned in front of the mirror.

6. Conclusions

We proposed a new technology to extend anamorphosis in the art to exploit the recent advancements in ray tracing and additive manufacturing. While previous technologies could only utilize simple geometric primitives (cylinders or cones), our technology can work with many complex surfaces. We also showed how our technology could utilize either reflective or refractive objects. To our knowledge, the use of refractive media in anamorphosis is entirely new.

We employed our technology to fabricate innovative sculptures and exhibited them at prestigious international art events. We analyzed the intrinsic limitations depending on the physics of light transport in the paper. Unfortunately, we cannot avoid such restrictions entirely. Still, we can partially mitigate them by optimizing the parameters that define the entire scene, such as the position of the viewer, or of the sculpture, or of other factors determining the resulting deformation. One exciting direction to overcome some of the current limitations is to optimize the reflecting/refracting media's shape. However, this will require the ability to accurately fabricate reflective or refractive shapes from a digital shape representation. Another interesting future direction to explore is discovering new ways to maximize the deformation of the final sculpture, eventually including details in non-visible areas. This might increase the surprise of the viewer to observe the difference between the manufactured object and its reflection in the mirror. An additional area for improvement is to broaden the scope of the simulation to account for the stereoscopic impact of an individual's perspective. Currently, our computations rely solely on a monocular viewpoint, which is a useful estimation but not entirely precise. So, although the produced effect is correct, the computation could be more accurate. Furthermore, we observed that the viewer must be at a specific distance from the optical media to achieve proper focus, whether the effects are refractive or reflective. A more comprehensive examination of the optimal focal distance may require more precise modeling of the viewer.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: We exposed artworks using the technology presented in this paper

Data availability

The data that has been used is confidential.

Acknowledgments

The authors thank Paolo Cignoni, Riccardo Scateni and Davide Fara for the initial brainstorming and initial prototype implementation, Simon Bethune for the help with the hydroforming tool. This work was supported in part by the Australian Government Research Training Program Stipend.

References

[1] Baltrusaitis J. Anamorphic art. Cambridge; 1977.
[2] Jo YH. About looking. 2014, Savina Museum of Contemporary Art.
[3] Hurwitz J. The art of Jonty Hurwitz. 2023, https://jontyhurwitz.com.
[4] Wynne prize, art gallery of NSW. 2023, https://www.artgallery.nsw.gov.au/prizes/wynne/.
[5] Sydney contemporary, CarriageWorks Sydney. 2022, https://sydneycontemporary.com.au.
[6] Fisher's ghost art award 2022, campbelltown arts centre. 2022, https://c-a-c.com.au/fgaa-2022/.
[7] North Sydney art prize. 2022, https://northsydneyliving.com.au/the-north-sydney-art-prize-an-exhibition-no-one-will-want-to-miss/.

[8] Pietroni N, Bickel B, Malomo L, Cignoni P. State of the art on stylized fabrication. In: SIGGRAPH Asia. ACM; 2019, p. 118:1.
[9] Glaeser G. On nonlinear perspectives in science, art and nature. In: Neumann L, Sbert M, Gooch B, Purgathofer W, editors. Computational aesthetics in graphics, visualization and imaging. The Eurographics Association; 2005, http://dx.doi.org/10.2312/COMPAESTH/COMPAESTH05/123-132.
[10] Coleman P, Singh K. Ryan: Rendering your animation nonlinearly projected. In: Proceedings of the 3rd international symposium on non-photorealistic animation and rendering. New York, NY, USA: Association for Computing Machinery; 2004, p. 129–56. http://dx.doi.org/10.1145/987657.987678.
[11] Gröller E. Nonlinear ray-tracing: Visualizing strange worlds. Vis Comput 1995;11(5):263–74.
[12] Agrawala M, Zorin D, Munzner T. Artistic multiprojection rendering. In: Peroche B, Rushmeier HE, editors. Proceedings of the eurographics workshop on rendering techniques 2000. Eurographics, Springer; 2000, p. 125–36.
[13] Collomosse JP, Hall PM. Cubist style rendering from photographs. IEEE Trans Vis Comput Graph 2003;9(4):443–53, URL http://doi.ieeecomputersociety.org/10.1109/TVCG.2003.1260739.
[14] Navoni M, Buzzi F. Leonardo Da Vinci and the secrets of the codex atlanticus. White Star Publishers; 2019, URL https://books.google.com.au/books?id=EeNzuwEACAAJ.
[15] De Rosa A, Bortot A. Anamorphosis: Between perspective and catoptrics. In: Sriraman B, editor. Handbook of the mathematics of the arts and sciences. Cham: Springer International Publishing; 2021, p. 243–89. http://dx.doi.org/10.1007/978-3-319-57072-3_38, https://doi.org/10.1007/978-3-319-570723_38.
[16] de Comite F. A general procedure for the construction of mirror anamorphoses. In: Bridges 2010: Mathematics, music, art, architecture, culture. Pecs, Hungary; 2010, p. 231–8, URL https://hal.archives-ouvertes.fr/hal-00861388.
[17] Pharr M, Jakob W, Humphreys G. Physically based rendering: From theory to implementation. 3rd ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.; 2016.
[18] Jovanovic M. Generating an anamorphic image on a curved surface utilizing robotic fabrication process. In: Proceedings of the 34th international conference on education and research in computer aided architectural design in Europe, vol. 1. 2016.
[19] Stojaković V, Tepavcevic B. Distortion minimization: A framework for the design of plane geometric anamorphosis. Nexus Netw J 2016;18. http://dx.doi.org/10.1007/s00004-016-0302-z.
[20] Hansford D, Collins DL. Anamorphic 3D geometry. Computing 2007;79(2–4):211–23.
[21] Sánchez-Reyes J, Chacón JM. Anamorphic free-form deformation. Comput Aided Geom Design 2016;46:30–42. http://dx.doi.org/10.1016/j.cagd.2016.06.002.
[22] Schüller C, Panozzo D, Sorkine-Hornung O. Appearance-mimicking surfaces. ACM Trans Graph 2014;33(6). http://dx.doi.org/10.1145/2661229.2661267.
[23] Sánchez-Reyes J, Chacón J. How to make impossible objects possible: Anamorphic deformation of textured NURBS. Comput Aided Geom Design 2020;78:101826. http://dx.doi.org/10.1016/j.cagd.2020.101826.
[24] de Comite F. A new kind of three-dimensional anamorphosis. In: Bridges 2011: Mathematics, music, art, architecture, culture. Coimbra, Portugal; 2011, p. 33–8, URL https://hal.archives-ouvertes.fr/hal-00861392.
[25] Comité FD, Grisoni L. Numerical anamorphosis: an artistic exploration. In: SIGGRAPH Asia art papers. ACM; 2015, p. 1:1–7.
[26] Wu K, Chen R, Fu X-M, Liu L. Computational mirror cup and saucer art. ACM Trans Graph 2022;41(5). http://dx.doi.org/10.1145/3517120.
[27] Kerber J, Wang M, Chang J, Zhang JJ, Belyaev A, Seidel H-P. Computer assisted relief generation—A survey. Comput Graph Forum 2012;31(8):2363–77. http://dx.doi.org/10.1111/j.1467-8659.2012.03185.x.
[28] Yue Y, Iwasaki K, Chen B-Y, Dobashi Y, Nishita T. Pixel art with refracted light by rearrangeable sticks. Comput Graph Forum 2012;31(2pt3):575–82.
[29] Pereira T, Rusinkiewicz S, Matusik W. Computational light routing: 3D printed optical fibers for sensing and display. ACM Trans Graph 2014;33(3):24:1–24:13.
[30] Papas M, Houit T, Nowrouzezahrai D, Gross MH, Jarosz W. The magic lens: refractive steganography. ACM Trans Graph 2012;31(6):186:1–186:10.
[31] Weyrich T, Peers P, Matusik W, Rusinkiewicz S. Fabricating microgeometry for custom surface reflectance. ACM Trans Graph 2009;28(3):32.
[32] Papas M, Jarosz W, Jakob W, Rusinkiewicz S, Matusik W, Weyrich T. Goal-based caustics. Comput Graph Forum 2011;30(2):503–11.
[33] Kiser T, Eigensatz M, Nguyen MM, Bompas P, Pauly M. Architectural caustics - controlling light with geometry. In: Hesselgren L, Sharma S, 0001 JW, Baldassini N, Bompas P, Raynaud J, editors. AAG. Springer; 2012, p. 91–106.
[34] Yue Y, Iwasaki K, Chen B-Y, Dobashi Y, Nishita T. Poisson-based continuous surface generation for goal-based caustics. ACM Trans Graph 2014;33(3):31:1–7.
[35] Schwartzburg Y, Testuz R, Tagliasacchi A, Pauly M. High-contrast computational caustic design. ACM Trans Graph 2014;33(4):74:1–74:11.
[36] Mitra NJ, Pauly M. Shadow art. ACM Trans Graph 2009;28(5):156.
[37] Zhao H, Lu L, Wei Y, Lischinski D, Sharf A, Cohen-Or D, et al. Printed perforated lampshades for continuous projective images. ACM Trans Graph 2016;35(5):154:1–154:11.
[38] Kushner S, Ulinski R, Singh K, Levin DI, Jacobson A. Levitating rigid objects with hidden rods and wires. Comput Graph Forum 2021.
[39] Pietroni N, Tarini M, Vaxman A, Panozzo D, Cignoni P. Position-based tensegrity design. ACM Trans Graph 2017;36(6). http://dx.doi.org/10.1145/3130800.3130809, https://doi.org/10.1145/3130800.3130809.
[40] Cignoni P, Callieri M, Corsini M, Dellepiane M, Ganovelli F, Ranzuglia G. MeshLab: an open-source mesh processing tool. In: Scarano V, Chiara RD, Erra U, editors. Eurographics Italian chapter conference. The Eurographics Association; 2008, http://dx.doi.org/10.2312/LocalChapterEvents/ItalChap/ItalianChapConf2008/129-136.
[41] Jacobson A, Panozzo D, et al. libigl: A simple C++ geometry processing library. 2018, https://libigl.github.io/.
[42] Garland M, Heckbert PS. Surface simplification using quadric error metrics. In: Proceedings of the 24th annual conference on computer graphics and interactive techniques. USA: ACM Press/Addison-Wesley Publishing Co.; 1997, p. 209–16. http://dx.doi.org/10.1145/258734.258849.