

EPSRC

Engineering and Physical Sciences
Research Council



InFoMM

Industrially Focused
Mathematical Modelling

EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling



Reflections from curved specular surfaces

Attila Kovacs



ARUP



Contents

1. Introduction	2
Background	2
Glossary.....	2
2. Image generation.....	3
3. Reconstruction	4
Known Environment	4
Unknown environment.....	5
Small disturbances.....	5
4. Results.....	5
5. Discussion, Conclusions & Recommendations	6
6. Potential Impact	7

1. Introduction

Background

Large glass sheets and reflective materials are often used as the external surfaces of modern buildings. One example can be seen on the title page, which shows a photo of part of a train station in Birmingham. These large reflective surfaces can exhibit complex behaviours, such as generating glare and hot spots due to the focusing of light. As glare and hot spots can be major annoyances in working environments, we wish to understand how the building shape generates them.

At Arup, the procedure for determining the position of hot spots and glare is to embed a CAD design of a building into ray tracing software and running the computationally expensive ray tracing algorithm on that design. However, this does not consider for the fact that geometrical changes can occur as the construction takes place. These changes include construction and fitting errors and bends, dents and other imperfections caused by the stress exerted on the surface. One example of a such a stress is the thermal stress that occurs because of the temperature difference between winter and summer in some surfaces, which can result in a small deformation of the building.

To carry out a meaningful analysis on the building, the geometrical model must closely match reality. This can be achieved by using various measurement techniques to build a computer representation of the constructed surface, a method which is called surface reconstruction. This reconstruction is frequently carried out using Lidar systems; however in the case of reflective surfaces this method fails, since only small amount of laser light reflects in the direction of the receiver, making the reconstruction near impossible.

Our aim is to formulate the problem of reconstruction of reflecting surfaces by using images of specular surfaces, so that Arup can offer better analysis for their customers. There are already multiple ideas about how to do the reconstruction, for example, by employing rotating and moving cameras, but these are not readily applicable to our case of large stationary reflective surfaces subjected to small disturbances.

Glossary

- **Canvas:** Plane in space on which the image appears.
- **Ray:** A model of light, where we assume that it can only propagate in a straight line. In principle, this has a direction originating from a light source and collected in the camera. However the mathematical model permits us to choose the direction as necessary.
- **Focus:** The point where all the rays originate from. The rays from the environment are coming towards the focus point, however we flip the directions for simplicity.
- **Specular surface:** A surface which has no intrinsic properties other than it reflects light from the surface, which makes it especially difficult to deal with it.
- **Forward problem:** We call the image generation problem a forward problem, since the generation of the image is done by explicit calculations.
- **Inverse problem:** The reconstruction of the surface is an inverse problem with respect to the image generation problem, which means we start from the result (the reflection) and trying to work back the cause (the specular surface).

Usual laser reconstruction of surfaces fails when presented with reflective surfaces, such as glossy metal. However, it is still possible to take images of the reflective surfaces

- **Pinhole camera:** A pinhole camera, as shown in Figure 1, is such that it has a focus point and the rays from it originate from this point. Its two main properties are the focal length and the field of view.
- **Orthogonal projection:** An orthogonal projection is a simplified pinhole camera, which has an infinitely large focal length. This means that all the rays are parallel to each other and orthogonal to the canvas. This results in much simpler calculations.




Figure 1: Schematics showing the operation of a pinhole camera (left) and orthogonal projection (right). Pinhole camera is harder to analyze, but has the advantage that it can recover depth, while the orthogonal projection has no depth information.

In order to address the issue of how to determine the shape of a reflecting surface, we will first look at how images are generated. We will then consider the problem of using images to reconstruct the reflecting surface.

2. Image generation

We consider the problem of determining the image of a known environment reflected by a given surface. The geometry of this problem can be seen in Figure 2. The key mathematics we use are the geometrical properties of surfaces and curves, such as the surface normal. Working through the calculations, it is possible to calculate the normal and the mapping between the inbound and outbound rays, where inbound refers to rays that originate from the camera focus and outbound means rays that are reflected from the surface and directed at the environment. To generate the image, we cast rays from every point of the canvas by using one of the camera models and, knowing the surface, these rays can be followed from the camera to the model of the environment. An example of generating an image is shown in Figure 3. Here we follow one of the rays from a pinhole camera which has its focal point at the origin, which reflects on the surface and propagates toward an infinitely far environment.

In everyday life, we have two dimensional surfaces in three spatial dimensions and the images taken by a camera are also two dimensional. To further simplify our model, we will reduce the dimension and consider one dimensional curved surfaces in two spatial directions and similarly we will have images that are one dimensional and can be represented as a function.

First step is to have a simple enough algorithm for generating images from a surface and environment




Figure 2: Schematic shows the main parts of the problem. (O) denotes the position of the camera facing in direction (D) towards the surface (f). One of the light rays (R) is tracked through the scene, hitting a part of the environment (E).




Figure 3: The environment is the colored dome, and the rectangle is the image of the environment that the observer (black dot) facing the reflective surface (blue curve) sees. Additionally, the inbound (facing the camera) and outbound rays (facing the environment) are also shown.

The main idea for reconstruction is to define a measure of similarity between images and use this to compare the reflection of the environment for an unknown surface to our generated images to find the shape of the surface

3. Reconstruction

Now that we can calculate images, if we are presented with an image of the environment seen in some unknown surface, it should be possible to reconstruct the reflective surface. However, given that in general we do not know the environment, there is an inherent uncertainty because different combinations of surface and environment can result in the same image. If we have additional information, for example, either we know the environment, or we know that the surface is only slightly disturbed, we can make progress easily. However, we can also make progress in the completely general case by increasing the amount of information that we use by using more than one picture.

Known Environment

One piece of additional information that we can use to enable the reconstruction is to assume that we know the model of the environment that generated the image. We then use the image generation method described in Section 2 to generate images for arbitrary surfaces. We then want to compare these images to the real one by defining a suitable distance between images, which is zero when the two images are the same. We build an iterative search on top of this where, in each iteration, we change our guess to get closer to

the truth. This iterative comparison between real and generated images is an optimization problem.

Unknown environment

Our optimization method only works when the environment is known. However, we usually do not know the model of the environment, because it is hard to get and in arbitrary real-life cases the assumption that the environment is far away is not always valid. An example for this is when there are objects in the reflected image, which are close to the surface or the camera. This motivates using other approaches, for example, incorporating multiple images with different cameras, or using a camera that moves along a given path and observing how the reflection changes, as this contains a lot of information about the surface.

Small disturbances

We also consider the case of small disturbances on a given surface, for example some small amplitude oscillations on a planar mirror. In this case, the reconstruction can be cast into a much simpler form and can be solved explicitly. The advantage of this small disturbance reconstruction method is that it does not need expensive optimization for the reconstruction. However, the reconstruction of the surface is only accurate when the disturbance is much smaller than the distance between the camera and the closest point on the surface. Furthermore, if the disturbance is too small, the inherent noise (caused by neglected effects such as camera nonlinearities, noisy image capture etc) means that the reconstruction will fail.

For small disturbances, simple reconstruction is possible in the simplest cases

4. Results

To evaluate our reconstruction methods, we generate random images and environments, calculate a reflected image, and then try to reconstruct the true surface by using only the environment and the reflection. In the 2D case shown in Figure 4, the reconstruction is fast and accurate and the optimisation method gives a result which is indistinguishable from the true surface. We also see that the disturbance-based method gives an excellent fit for small angles.




Figure 4: Reconstruction of the reduced dimensionality problem. The true surface is shown in black. The dashed light blue line is the reconstructed surface calculated using the optimization method and the green line is the reconstructed surface using the disturbance-based method. The horizontal axis represents angle.

The reconstruction with disturbance method gets better when the amplitude of disturbance decreases, however arbitrary small reconstruction is not possible due to noise

For real images, the problem is harder and the calculation is much slower, but in certain cases where the surface disturbance is not too large the reconstruction can be achieved. An example of surface reconstruction via the optimization method for 2D surfaces using a real test picture as the unknown environment can be seen in Figures 5 and 6. The procedure is as follows. We take the standard Lena image and we project it onto our environment. We

then use the true surface shown in Figure 5 (left) to generate an image as shown in Figure 6 (left). We then pick a surface shape and use the projection of Lena in the environment to generate an image, and we then iterate the surface shape until this image agrees with the one generated using the true surface. The final modelled surface is shown in Figure 5 (right) and we see that it is identical to the true surface. In Figure 6 (right) we also show the image generated assuming that there are no disturbances to the surface and comparing the two pictures in Figure 6 we see that the small disturbances only cause small changes to the image.




Figure 5: Contour plot of the true surface (left) and the reconstructed surface (right)

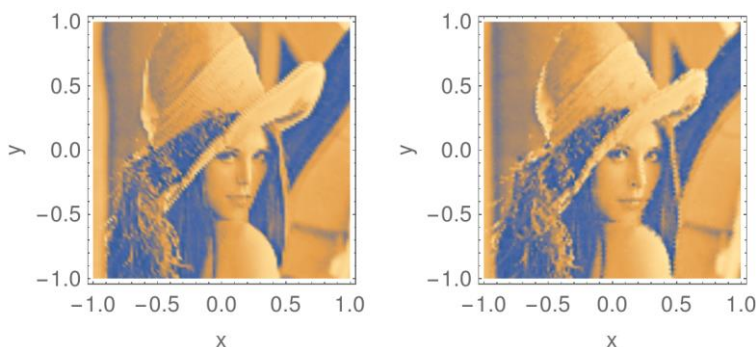



Figure 6: Reflection of the environment using the true surface shown in Figure 5 (left) and the reflection of the environment in a plane mirror (right). Here we use the standard Lena image projected to a 3D sphere.

5. Discussion, Conclusions & Recommendations

We have developed a mathematical framework for reconstructing the shape of reflecting surfaces. We formulated the forward problem of generating images, given the shape of a reflecting surface, for real 2D images and, for simplicity, for a reduced dimensional version as well. We found that we were able to generate reflected images of an environment and, by comparing these images to pictures where we do not know the surface, we used an iterative method find the best fitting surface, which results in good reconstruction for synthetic problems.

For the special case when we are dealing with small disturbances on a known surface, the reconstruction can be done explicitly in 2D and the synthetic results show that the



disturbance-based method can be used as proper reconstruction procedure, but this method is not trivial to extend to 3D, thus further work may be necessary.

Finally, because the model of the environment is not always available, we also formulated an optimization-based reconstruction method which uses multiple cameras and does not require us to capture the environment. Instead, in this method we compare images from different cameras directly.

6. Potential Impact

In the short-term, Arup will attempt to apply our methodology and code to reconstruct simple real world and ray-traced images. Our work has excited interest from the Façade Engineering Team who, in turn, have interest from a leading glass manufacturer in investigating and developing these ideas further for quality checking of glass sheets. They will investigate the possibility of undertaking further research in this area. Several novel and potentially powerful innovations could follow in longer term (e.g. fluid free surface reconstruction).

Steve Walker, Associate Director of Advanced Technology and Research commented, *“We are very pleased with the progress made on this issue and the process that we can experiment with ourselves. It was sufficiently encouraging for the façade engineering group and a major manufacturer of glass as well to look at developing this research direction further, especially the multiple cameras and the unknown environment part, which could be interesting areas of further research. Arup, together with another manufacturer, will investigate funding and scope of such a research project.”*