Improving the creation process of catoptric anamorphosis using a reference pattern

Q. B. van Velthoven
Delft University of Technology

Catoptric anamorphosis is an image which seems to be distorted, but looks normal when observed from a specific point of view and indirectly via a reflective object. While some methods already exist that can create such images, they are limited in the amount of differently shaped reflective objects they support or they require you to obtain a 3D model of your object. We propose a new method, that can work with a reflective object of any shape and that requires minimal work from the user. After a suitable reference pattern is generated, the user can give a picture of both the reflective object and the pattern along with a picture of the reflection of the anamorphosis as input and the output is the corresponding anamorphosis. The results show that the method creates accurate catoptric anamorphoses for most differently shaped reflective objects. It is concluded that this method can help existing and new artists to create catoptric anamorphoses more easily while still offering the same features as older methods.

Introduction

While all the drawings of perspective art in the street may impose that this technique of creating optical illusions is fairly young, this is actually misleading. The art style, of which you can find an example in Figure 1a, has existed since the beginning of the fifteenth century, after linear perspective had been discovered. The term used for images as these is also not "perspective art" or "street art", but "anamorphoses" instead. An anamorphosis can be defined as a projection or a perspective of an image that requires viewers to see it from a specific position to see the image in its natural form. A special form of anamorphosis is catoptric anarmophosis, which means that a reflective object is required to obtain the image, as the natural looking image can only be obtained as a reflection. An example of a catoptric anamorphosis can be seen in Figure 1b.

(a) An example of an oblique anamorphosis.

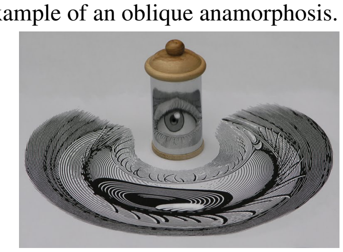

(b) An example of a catoptric anamorphosis.

Figure 1. Two examples of the two different kinds of anamorphosis/

After the earlier mentioned discovery of linear perspective in the beginning of the fifteenth century, it was Leanardo da Vinci who started understanding and exploring oblique anamorphosis (Collins, 1992). Catoptric anamorphosis was discovered later in the sixteenth century. The focus was then on cylindrical catoptric anamorphoses, which are catoptric anamorphoses which use a reflective cylinder and cylindrical anamorphoses thus are a subset of the catoptric anamorphoses.

Much research has been done on the subject of catoptric anamorphosis. In 1638 Jean-François Nicéron's La Perspective Curieuse came out, in which he described among others how to create several kinds of anamorphic art such as cylindrical, conical and pyramidal catoptric anamorphoses (Niceron, 1663; De Comité, 2010). According to Nicéron himself was his template to create cylindrical catoptric anamorphoses with not a very good approximation (Hunt et al., 2006). Nicéron also did not prove any of the techniques that he proposed in his book and many of them have later been proven to be not approximate and even incorrect. (Hunt et al., 2006).

Since La Perspective Curieuse was not completely correct or exact, the question had been where to find proper information on how to create catoptric anamorphoses. In 1999 Hunt, Nickel and Gigault described how to make plane, conical and cylindrical anamorphoses (Hunt, Nickel, & Gigault, 2000). For conical anamorphoses, which are anamorphoses that use a reflective cone, the equations necessary to draw them were given. Drawings of Nicéron were shown as well and it was verified that Nicéron's technique for drawing conical anamorphoses are a good approximation (Hunt et al., 2000). The same was done for cylindrical anamorphoses, however in this case it was shown that Nicéron's approximation was not that good (Hunt et al., 2000), just as Nicéron admitted. Hunt, Nickel and Gigault came up with the idea to use the inverse of the transformation function to create cylindrical anamorphoses (Hunt et al., 2000).

Creating catoptric anamorphoses with a reflective object of any shape became possible in 2010, as De Comité created a raytracer (De Comité, 2010). While this method technically has made it possible to create catoptric anamorphoses with more complex reflective objects, it is still not 'easy' for artists to create catoptric anamorphoses. If they want to work with complex reflective objects, they are obliged to recreate the complex objects as computer models or to scan the objects. The former is difficult, since the objects are complex and thus more troublesome to model. Scanning objects could be easier, but that process requires special hardware that is often not available.

The purpose of this study is to improve the creation process of catoptric anamorphoses, because anyone who wants to create catoptric anamorphoses is currently limited in the amount of reflective objects he can use or is required to recreate complex objects as computer models. This makes the creation process of catoptric anamorphoses both difficult and costly in time. We propose a new method that uses neither the earlier mentioned equations nor a raytracer solution, but only two different images: a picture of a 'reference pattern' and the reflective object together and an image that the user wants to obtain as the reflection of the anamorphosis. It is possible to read from the picture where points in the pattern are reflected on the reflective object, and since the pattern is created in such a way that points are clearly identified, we know from many points (pixels) the corresponding 'reflection points'. Now a mapping exists from the reflection pixels to the pattern pixels. If the other image is placed over the reflection, of many points in the new image it will be known where the corresponding points would be in the original. With this information and some extra smaller steps it is possible to create a catoptric anamorphosis. We hope to improve the current creation process of catoptric anamorphoses with the proposed method in a way that artists will be able to make catoptric anamorphoses with complex models which they could not create before, that the general public becomes more interested in creating catoptric anamorphoses and that catoptric anamorphoses become as prominent as oblique anamorphoses in everyday life.

First we have a look at the theory behind this method, to answer the question why the method should work from a theoretic standpoint. In that chapter is also discussed how the research question could be answered. Thereafter, the focus is on how the technique with 'reference patterns' can be incorporated into an algorithm. Design choices for this algorithm are mentioned and also the problems that we had to face with the design. After the method is explained, it is time to see how it works in practice and to explain any differences in the chapter 'Discussion'. Finally we answer the research question and point out directions for further research.

Theory

In the Introduction the current state of the art of catoptric anamorphoses creation methods was mentioned. On one hand there was the method of Hunt, Nickel and Gigault, which is exact but only works if reflective object has one of the three supported shapes (Hunt et al., 2000). On the other hand there is the not exact solution of De Comité which supports many different shapes, but requires users to come with a model of the reflective model themselves (De Comité, 2010). While not exact, De Comité's method is still based on a special technique to create catoptric anamorphoses that was written down in Diderot's Encyclopédie (De Comité, 2010). There seems to be no method yet that has the positive properties of both methods: a method that requires not too much work of the user and that supports many different shapes for the reflective object. We try to find out if the method that we propose can change that situation.

Our proposed method is not based on equations or such older techniques, but instead based on the way reflection works. To illustrate this, we imagine a scene with a reflective object, an object and a light source. A reflective object is an object that reflects an incoming ray of light to the observer. If the other object's reflection is on the reflective object, that means that rays coming from the light source are reflected diffusely and that some of those diffusely reflected rays reach the reflective object. Finally the reflective object reflects the rays in a specular fashion to the observer. In Figure 2 this process can be observed. If we now also focus on Figure 2 and imagine that over the current pattern an image is put over the place where the ray hits the pattern. How the reflection

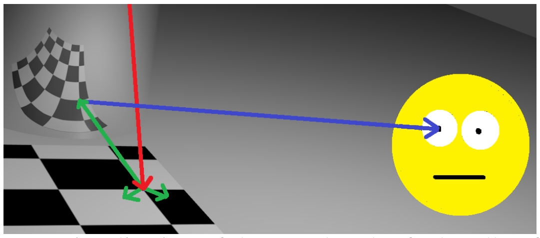

Figure 2. Visualisation of the travel path of a bundle of rays of light.

would then exactly look like is difficult to say, but it is clear that the reflection changes in the point where the diffused ray hits the reflection. The direction of rays does not depend on the colour of the object from which the rays reflect, thus if a point in the original pattern changes colour, then the corresponding point in the reflection changes of colour as well. It also works in the opposite direction: if a point in the reflection looks different from before, then the original has to look different as well in the corresponding point.

Given that the user provided an image that contains both the reflective object and a pattern, he needs to find as many points as possible in both the original and the reflection and match the corresponding points. In that case the user can change either the reflection or the original and create the new version of the other properly, because a higher amount of matched points results in a higher amount of points of which the user knows how they should change in colour. Here the user should change the reflection by placing an image over the pattern's reflection, since the idea of catoptric anamorphoses is that the reflection looks normal. Then the original can change to match the reflection and an anamorphosis is created, since the anamorphosis' reflection only looks like the image that the user provided as reflection from the position that the big image is taken. The original by that reflection is then by definition an anamorphosis.

This is the technique that makes our method work in general, but our method has some extra steps that are required to ensure that even in worse conditions our program can give a fairly accurate result. All the steps in our method are described in the next chapter.

Method

When the main theory behind our method was discussed, it may have seemed that the technique was able to solve the problem easily: the user simply had to give an image with both the reflective object and the pattern and an image that the user wanted the anamorphosis' reflection to be and then only points needed to be matched. However, there are multiple decisions yet to be made on how to deal with several situations that the theory alone cannot deal with. For instance, maybe even the biggest question is how points can be found and matched with each other. Another problem that can and probably will come up is that the maximum amount of points that can be found in the reflection, i.e. the amount of pixels that cover the reflection area, is less than the maximum amount of points that the original/pattern contains. These problems need to be solved and can often be solved in different ways, thus in this chapter the question is answered which problems need to be solved and explain why we chose for certain solutions.

Our proposed method exists of multiple stages and the difficulties and design choices of these stages are discussed in order. The stages are as follows:

1. Create a reference pattern.
2. Indicate where the anamorphosis' reflection should be.
3. Find points in the reflection and the original and match corresponding points.
4. Generate triangles with the points in the reflection.
5. Colour all points in the original that correspond to points in the new reflection.
6. Colour empty pixels in the original.
7. Rectify the new original to obtain the anamorphosis.

Reference pattern and point matching

Before the user gives an image containing both the reflective object and the object/pattern on which the anamorphosis is going to be drawn as input, it is useful to put a special pattern on the object. Whether the user has to indicate the points manually or the points are found automatically, it is easier to find points when there is a clear difference with neighbour points (read pixels) in several directions (Moravec, 1980). This is clearer when we take a look at Figure 2 again. It is possible to indicate points at the intersection between two white squares and two black squares, because in two different directions there is a big change in colour. A chequer board pattern has many of such intersection and can thus result in a large amount of distinctive points, which makes it a good idea to continue such a pattern. Another reason to choose for a chequer board pattern is that the points are spread evenly over the pattern, or at least in the original and often fairly well in the reflection too. This allows the user to place the image where he/she wants without having to worry about a low point density, which lowers the quality of the created anamorphosis.

While it has become clear why a chequer board pattern is useful in this case, we still do not know how to colour the squares to our advantage. Only two colours are required to create points and thus find points, as can be seen in Figure 2. However, using many different colours may make it easier to match corresponding points in the reflection and the original later on: if a colour is only used once in the chequer board, then matching the point with that colour in the reflection and the original should be easy. Unfortunately, this solution is not reliable. First of all, there is often a slight colour difference between points in the reflection and in the original based on the position of the light source and other circumstances such as dust on the reflective object. Secondly, if the colour of the light is not white but, for instance, more yellow, then generally the more white squares will become more yellow (white objects reflect rays of all colours and absorbs none), the yellow squares stay yellow and the other squares will become more black, as they absorbs rays of all colour except the colour of that square. White light is a bundle of rays of all colours, yellow light of only yellow rays, thus that square's colour is not present in the bundle of rays thus the square will not reflect any rays and appears black. Even when there is a slight colour pitch and there are still some bundles of white light, many colours will change colour and matching on basis of colour becomes then impossible. One thing that is clear for a pattern with any amount of colours is that it is beneficial to only put squares next to each other with distinct colours: if the pattern is seen in an image from a distance and the pattern is not evenly lit, even then points ought to be found.

In 1976, Appel and Haken proved that every planar map can be coloured with at most four colours (Appel & Haken, 1976). A planar map is topologically equivalent to a plane graph, which is a drawing of a planar graph and a drawing in which edges of the graph do not cross each other (Trudeau, 2015). A chequer board pattern is thus a plane graph and that means that with at most four colours there is no square next to a square with the same colour. A way to colour the squares can be found at Figure 3. We see in that figure that the colouring with four colours is fairly trivial for a chequer board pattern, which means that no algorithm such as that of Chaitin is required to colour the squares (Chaitin, 1982). To make the colours as distinct as possible, we used the colours that each use one colour channel (we used the RGB format to define colours) and one colour that uses two channels.

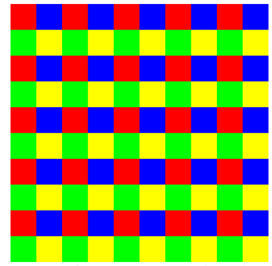

Figure 3. An example of a pattern with four different colours and with no squares of the same colour next to each other.

We need to emphasise that the use of four different colours for the pattern is only actually required if an algorithm is used to match the points in the pattern and the reflection. This is something that we want to achieve, as the amount of matched points has a positive contribution to the quality of the anamorphosis that will be created. If the user is required to manually match all points, the workload increases for the user. If there are also many points in a small area, it may become very difficult for the user to click the right one and match them. Unfortunately, it seems to be the case that manually selecting and matching points is the only reliable solution at the moment. There are many different methods in the field of feature point detection and matching that are capable of finding points, such as the earlier cited method of Moravec and the method of Harris and Stephens (Moravec, 1980; Harris & Stephens, 1988). These methods only work with grey images, but if we want to use the different colours in the pattern, then that is possible as well (Gevers, Van De Weijer, & Stokman, 2006). The problem occurs when we try to match points. In general, we cannot say anything about the transformation function that turns original points into reflection points, because in general we know nothing about the shape of the reflective object, since we allow the user to use any reflective object. If the reflective object was, for instance, a regular mirror, we could speak of an affine transformation since parallel lines in the original pattern are still parallel in the reflection (Berger, 2010). Li, X. Huang, J. Huang and Zhang (2014) found that "existing feature matching methods support either a specific or a small set of transformation models" (p.2407). This means that there is no solution yet to match the points when the transformation function is not known. The alternative is requiring users to match the points themselves. While this task may look enormous, we have found ways to reduce the workload heavily.

Position of the anamorphosis' reflection

One method to limit the work that the user has to do is offering the user to indicate where on the reflection the anamorphosis' reflection should be. Only points that are in the indicated region and around the border of the region are actually used to create the anamorphosis. Instead of having the user find and match every point in the pattern and its reflection, he ought to only collect points that are in the mentioned area. Once all the necessary points are collected, it is time for the computer to create the anamorphosis.

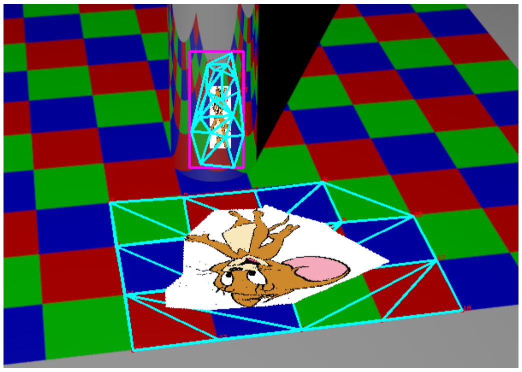

Figure 4. Delaunay triangulation used to interpolate.

Generate triangles and interpolation

Once all points are found in both the reflection and the original, it is time to recolour the original to match the new reflection. Unfortunately, it is very unlikely that the user has manually found and matched all pixels that are covered by the image that serves as the anamorphosis' reflection. Not only is it difficult if not practically impossible to generate a pattern of which the reflection gives one point per pixel consistently, it also requires much patience and pixel perfect clicking of the user. Fortunately, it is possible to fill these gaps of knowledge with interpolation. There is a high chance that the reflection of the resulting anamorphosis is not a perfect replica of the image that the user wanted to have as the anamorphosis' reflection, because interpolation assumes close points can tell something about a missing point. Nevertheless, since the mentioned alternative is practically impossible to do, we decided to settle for interpolation.

There are several ways to spatially interpolate, but some techniques only work when the points are spread evenly, such as bilinear and bicubic interpolation. The points in both the reflection and the original are not guaranteed to be uniformly spaced, which means that another technique must be found that does not require a regular grid of points. We decided to generate Delaunay triangles and to interpolate using Barycentric coordinates. By generating triangles between the points, smaller areas are created over which can be interpolated. The choice for Delaunay triangulation can be justified: it is not only a well-known method for which multiple implementations are created, it also has the property to maximise the minimal angle within triangles (Delaunay, 1934). This is positive for us, since that favours triangles of which the points are not too far away from each other, as a far point would result in a small angle. The smaller the triangles are, the more accurate the anamorphosis will be. In Figure 4 a Delaunay triangulation can be spotted.

After the creation of the smaller areas, interpolation can take place. First all the points in the reflection (and therefore also the triangles) are scaled up so that the indicated reflection area has the same height and width in pixels as the image that acts as the anamorphosis' reflection. Then for every pixel in the just mentioned image the colour is read, the triangle in which the pixel is situated is looked up and the Barycentric coordinates of that pixel position in that triangle are saved. The corresponding not scaled up points in the original are found for the three scaled up points that make up the triangle. The pixel to colour in the original is calculated with the Barycentric coordinates and the three original points. The colour of that pixel becomes the earlier read colour.

The reason to scale up the reflection points and not to minimise the anamorphosis' reflection instead is that minimising the image makes the image quality worse: there are simply fewer pixels thus not all information can be saved. If we then read colours from the minimised image, the program will be much faster since the method spends most of the time on going through all the pixels. On the other hand, there will be large gaps between the drawn pixels in the original since the draw function is called less often. It is also not known if the pixels that are drawn actually have an accurate colour value, since interpolation was used by minimising the image. It may take much more time to compute the anamorphosis, but scaling up the points and thus not losing any details is in our eyes better than to minimise the image and obtain an anamorphosis of which the reflection does not resemble the image. It is thus advised to use a relatively high resolution image to use as the anamorphosis' reflection, with a total number of pixels higher than the amount of pixels the anamorphosis is expected to cover. This way there will very likely be no gaps in the anamorphosis.

Fortunately, our method can also deal with gaps by filling them up. Once again Delaunay triangles can be generated with the drawn pixels as the input points and interpolation can be used to fill these gaps, or the gaps are simply filled by copying the colour of a neighbour. If a low resolution image is used, the gaps are big and the first choice will be better. If the resolution of the anamorphosis' reflection is high, the gaps will be small and the latter will be fast and give a fairly accurate result.

Perspective fix

Technically no anamorphosis has been drawn yet, but instead how the anamorphosis would look like in the big image. To actually obtain the anamorphosis so that the user can use it in reality or as a texture in a raytracer, it is required to remove the perspective element from the image. As we said earlier, parallel lines remain parallel after an affine transformation (Berger, 2010). The pattern in the big image can be considered the result of an affine transformation, thus it is possible to remove the perspective from the pattern by reversing the transformation so that the user looks right upon the pattern. The user has to indicate the four corners of the pattern by clicking on it and provide an image that is taken from right above the pattern and an algorithm will reverse the affine transformation. Since the anamorphosis was drawn right on the pattern, the anamorphosis has been corrected for the perspective as well. This was the last step and the anamorphosis is now ready to be used.

Results and Discussion

An experiment is conducted to help us answer whether the proposed method can improve the creation process for catoptric anamorphoses in practice. There are many choices that an artist can make when he wants to create a catoptric anamorphosis: for example, the look and position of the pattern, the shape and position of the reflective object, the position where the reflection of the anamorphosis should be, the reflection of the anamorphosis itself and of course the amount of points with their locations are all decided by the user. There are still even some more technical choices, such as the resolution of the images. Therefore, we have chosen to not do a full analysis with all kinds of different benchmarks, since creating enough anamorphoses to be able to confidently say something about relations between different variables would take too much time. Even if we had created so many different anamorphoses, there is the problem that there is no way that all information could be put in a table and still be able to be read easily. Instead a different approach is taken to test the method in practice.

We have taken images of a few different scenes and tried to put the anamorphosis' reflection in spots such that the anamorphosis became more and more likely to be drawn wrongly. An example of such a situation is when the reflection of the anamorphosis is placed in an area with a low point density. Meanwhile other variables are meant to be constant, so that changes in the quality can be contributed to a factor more easily. All of the images are renders made with the use of the open source raytracing program Blender. The main reason to not use photographs is that different renders can be made with the camera in the same exact position, which reduces the error when the anamorphosis' reflection that the user wanted and the anamorphosis' actual reflection are compared on pixel basis. Otherwise it could be very attractive to change the viewpoint a little in order to get the exact reflection that we should obtain with the generated anamorphosis. Another reason is that in a program such as Blender complex models can be created to act as reflective objects, while it is often more difficult in reality to obtain complex reflective objects. Using a render does not give unfair results in this case, since we collected and matched points manually and that task should not be more difficult in a photograph than in a rendered image. While the images of the scene differ, the chosen anamorphosis' reflection will stay the same. If this image is kept constant, then the influence of other factors becomes more noticeable.

In the first test a reflective object that vaguely resembles a cylinder was used and put next to the pattern, as can be seen in Figure 5a. A big curve was present in the reflection which is the main difficulty in this scene. The method then also handled this case very well, as you can see in Figure 5b. The actual reflection of the anamorphosis is slightly longer and a bit wider in the waist area than the resized image. The head is also slightly more turned. Overall they look a lot like each other and these small differences are probably caused by a lack of points in the area. More points between the current found points would have resulted in smaller triangles and interpolation would then give more accurate results. One problem that is not fixed by the method due to time constraints is that the method assumes that the anamorphosis and the reflection have exactly the same colour. In reality lightning makes either the original or the reflection (slightly) lighter than the other. The problem is clearly visible in Figure 5b as the actual reflection of the anamorphosis is much darker than the user intended. Due to time constraints we were unable to incorporate a 'colour correction' function in our method, thus in every test the anamorphosis' reflection is darker or lighter than the intended reflection.

In the second test a much thinner reflective cylinder was used, but this time it was placed on the pattern itself, which resulted in the reflection having more curves than in the previous test. Further was not the complete pattern visible and was the reflection of the anamorphosis much closer to the anamorphosis itself. The main difficulty in the test was that in the reflection many squares were visible in only a small area of the reflection. We wanted to see how this affected the anamorphosis and decided to put the reflection in the mentioned area. The result was worse than expected, as we expected that the anamorphosis would look too much like the reflection itself, because a lack of points should mean more interpolation and thus that the anamorphosis would look more as the original. However, the opposite was actually true as can be observed in Figure 6a. Drawn together with the triangles it became clear that the lack of points resulted in the strange sharp angle in the anamorphosis. This is a great example of a problem that could be solved if there had been a method to automatically find and match points. The resolution was too low for a human to pick points reliably, because a pixel off in the reflection meant here multiple pixels off in the original. With an automatic way to find points and an increase of squares on the pattern it should have been possible to match

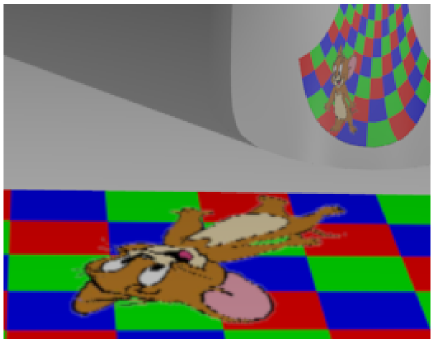

(a) Image of an accurate catoptric anamorphosis and its reflection.

(b) Comparison of the wanted reflection and the anamorphosis' actual reflection.

Figure 5. Results of the first test.

these points without a problem. This problem was now solved by rendering the image with a higher resolution, so that slightly misplacing a point in the reflection would have fewer consequences.

Another problem of manually collecting points is exposed in Figure 7. While a large part of the anamorphosis seems to be in line with the expectation, given the location of its reflection, something seemed to have gone wrong in the bottom left. If we again have a look at the triangles, we see that the triangulation with the reflection points went well, as there are no crossing edges. However, in the original clearly some crossing edges can be seen. This is caused if a point in the original is (slightly) misplaced and placed before another point. An edge between two reflection points results in an edge between the two original points that correspond with those reflection points, thus if another point is placed before one of those, edges will cross. This may seem not too harmful, but crossing edges means overlapping triangles, and then it becomes unknown which triangle will be used in the interpolation process. If the larger triangle is used, then

(a) Result of the second test with a lower resolution scene image.

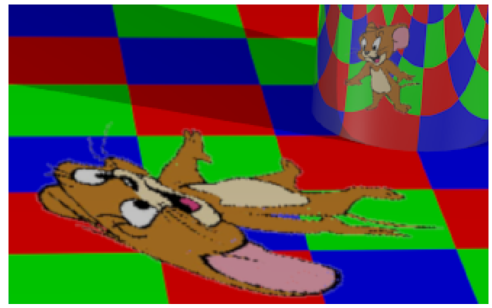

(b) Result of the second test with a higher resolution scene image.

Figure 6. The results of the second test.

the result of the interpolation will be less accurate and thus result in a partly wrong anamorphosis.

After having seen both positive and negative results, it was difficult to answer whether the proposed method is reliable enough to be used in practice. Under certain circumstances, such as high resolution image of the pattern and the reflective object when there are a lot of points in the reflection close together and enough squares and the right amount of squares to secure enough points for enough triangles but not to mistakenly click the wrong point, the method gives fairly accurate results. In theory it is also the case that the method should be able to give a pixel perfect anamorphosis when a point can be found at every pixel. Unfortunately, as long as finding and matching points automatically is out of the question, there will be at least small inaccuracies in the anamorphosis. Nonetheless, even with the current amount of tasks that the user has to do to secure a fairly accurate anamorphosis (generating the pattern with the right amount of squares, obtaining high resolution images, placing and matching many points) we conclude that the method is reliable. The workload could drastically be lowered with automatic point finding and matching, but the method generates fairly accurate catoptric anamorphoses as long as the user takes enough time to place the points correctly. Even then could be argued that the method is an alternative for the other solutions, since making a complex model or using the equations takes some time as well and not necessarily less than this method.

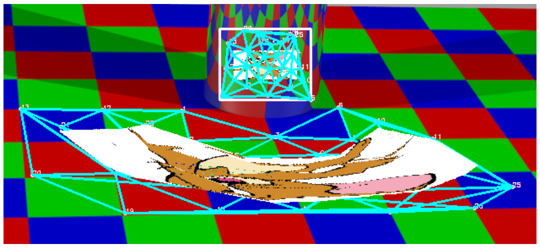

Figure 7. Crossing edges ruin the anamorphosis.

Conclusions and Future Work

With this research we tried to find a new method to create catoptric anamorphoses with, since the existing methods limited users in the amount of reflective shapes they could use or required the user to model the objects themselves. We have come up with a new technique to generate anamorphoses that works by providing an image with both a pattern and the reflection of that pattern with a mapping of points in the reflection and the original that correspond. Our proposed method should be able to work without these limits, should not require too much work of the user and should result in catoptric anamorphoses that give reflections that differ at most very little from the user's chosen reflection. The question that we try to answer is whether the method is indeed capable of delivering a reliable and better process for the creation of catoptric anamorphoses.

Results have shown that this method is capable of producing high quality anamorphoses. When the user found enough points in and around the area where the anamorphosis' reflection should be, an anamorphosis would be created of which the reflection would differ slightly from the image that had to act as the anamorphosis' reflection. However, the results also showed some fields in which our method did not perform very well. For instance, the method is very prone to human error and a misplaced point could ruin the anamorphosis. Nonetheless, the method is reliable as long as the user makes no mistakes.

The research question can then also be answered positively, since the method is capable of producing fairly accurate anamorphoses, does not require the user to make complex models and can work with a reflective object with any shape. Therefore can be concluded that the creation process has been improved with the coming of the proposed method and that was the research question that we wanted to answer.

For the future, we would recommend to look into a way to make finding and matching points automatically possible. Current methods have not been proven to support point matching when the reflection cannot be reduced to an affine transformation of the original. A robust method that succeeds in matching points automatically should result in more found points and thus a higher point density in and around the area of the anamorphosis' reflection, which on its turn will result in anamorphoses of higher quality and of course a lower workload for the user. Furthermore, the current implementation does not use hardware acceleration. Since much time is spent checking in which triangle a point is situated and since this is often a task of a graphics processing unit (GPU), a huge speedup may be achieved by passing work on to the GPU. Finally, correcting the anamorphosis for the difference in colour between the reflection and the original is a task we could not fulfill due to time constraints, but something that needs to done to improve the catoptric anamorphoses.

References

Appel, K., & Haken, W. (1976). Every planar map is four colorable. *Bulletin of the American mathematical Society*, *82*(5), 711–712.

Berger, M. (2010). *Geometry revealed: a jacobs ladder to modern higher geometry*. Springer.

Chaitin, G. J. (1982). Register allocation & spilling via graph coloring. In *Acm sigplan notices* (Vol. 17, pp. 98–105).

Collins, D. L. (1992). Anamorphosis and the eccentric observer: Inverted perspective and construction of the gaze. *Leonardo*, 73–82.

De Comité, F. (2010). A general procedure for the construction of mirror anamorphoses. In *Bridges 2010: Mathematics, music, art, architecture, culture* (pp. 231–238).

Delaunay, B. (1934). Sur la sphere vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, *7*(793-800), 1–2.

Gevers, T., Van De Weijer, J., & Stokman, H. (2006). *Color feature detection.* CRC press.

Harris, C., & Stephens, M. (1988). A combined corner and edge detector. In *Alvey vision conference* (Vol. 15, pp. 10–5244).

Hunt, J. L., Nickel, B. G., & Gigault, C. (2000). Anamorphic images. *American Journal of Physics*, *68*(3), 232–237.

Hunt, J. L., et al. (2006). J-f. niceron's la perspective curieuse revisited. *Bridges London: Mathematics, Music, Art, Architecture, Culture, eds. Reza Sarhangi and John Sharp*, *459464*.

Li, H., Huang, X., Huang, J., & Zhang, S. (2014). Feature matching with affine-function transformation models. *IEEE transactions on pattern analysis and machine intelligence*, *36*(12), 2407–2422.

Moravec, H. P. (1980). *Obstacle avoidance and navigation in the real world by a seeing robot rover.* (Tech. Rep.). STANFORD UNIV CA DEPT OF COMPUTER SCIENCE.

Niceron, J. F. (1663). *La perspective curieuse*. JD Puis.

Orosz, I. (2014). *Catoptric anamorphosis of an eye*. Retrieved from `https://3.bp.blogspot.com/_lCs_XEfXbGM/TIQAqni298I/AAAAAAAAEGU/qaYEjnUFIX8/s640/eyed.jpg`

Trudeau, R. J. (2015). *Introduction to graph theory*. Dover Publications.

Zinn, D. (2012). *Chalk art*. Retrieved from http://www.streetartutopia.com/wp-content/uploads/2012/10/Calk-Art-by-David-Zinn-14.jpg