Stereo Painting: Pleasing the Third Eye*

Efstathios Stavrakis – Margrit Gelautz
Technical University of Vienna

Painting and the 3rd Dimension

Painting, as we are accustomed to understand it, involves an artist's creation on a two dimensional surface. The ultimate goal of many art movements has been to present reality through an artistic "sensor" that captures and transforms it into a world of aesthetics. The style of paintings varies among artists and chronological periods, but paintings share the limitation of being bound to the flatness of the pictorial medium[1]. Linear perspective, texture, shading and lighting have been vastly used and sometimes abused to achieve certain perceptual effects and have been further used to enhance the compositional ingredients that make up the artwork. Unsurprisingly enough, one such effect is to capture the essence of depth and accurately represent distances. Spatial reality through the artist's eye and hand is transformed and distorted in unpredictable ways, but remains comprehensible, even though true depth and distance cannot be provided by a single painting, despite the mastery of its creator.

Stereoscopic Painting

An almost unknown, yet approached, painting form that lifts this barrier of flatness is stereoscopic painting[2]. A stereoscopic painting is composed of two canvases, each simulating the retinal image of each of the artist's eyes. This two-viewpoint depiction can be thought of as a pair of projections of an illusionary 3D composition space onto the two canvases. The two final painted surfaces can perfectly stand as artistic pieces of work by themselves. However, there is a distinct advantage; viewing the painting stereoscopically, with the 'third eye', the spectator can see inside the artist's composition space and how he placed the paint media on this peculiar three dimensional world. The binocular perception of a stereoscopic painting brings into play the third dimension of the painted piece, which artists have otherwise tried to depict in single canvas paintings by using monocular 'perceptual depth' cues. The stereo painting on the other hand detaches from the 2D surface and either protrudes into it or recedes from it, allowing the viewer to see the depth of the scene and the distances between objects to their full effect. In addition, while the viewer is looking at the two components of a stereo painting, the real world itself defocuses and is superseded by the stereoscopically fused painting. This secondary effect of absorbing the viewers' attention within the artwork is an effect many visual artists would be pleased with, if achieved under any circumstances.

Painters and Stereoscopy

Artists, being creative individuals, tend to experiment with different tools and apply, combine and exploit various techniques in their work. Stereoscopy has not escaped them either. Stereoscopic drawings can be found within most early technical articles about stereoscopy, but mainly for the purpose of illustration and idea communication rather than artistic composition. In the middle of the 20th century, painters technically aware of the effects of binocular vision in the human visual system have undertaken the tedious task of painting two stereoscopic canvases instead of one. René Magritte (1898-1967) with his "L'Homme au Journal" (1928) appears to have executed a quasi-stereo painting, but did not further exploit the possibilities. At the beginning of the '70s, Salvador Dalí (1904-89), Roger Ferragallo, Oscar Fischinger (1900-1967), Michael Kupka and Heinz Günther Leitner, almost concurrently, created several stereo paintings proving the fact that stereoscopy is applicable to painting. Their works can be generally categorized as photographically assisted and non-photographically assisted. The former, mainly exploited by Dalí and Kupka, is executed by the artist, who sets the foundation of the piece based on stereo photographs. The latter, usually, presents more abstract subjects (e.g. geometric formations) that are easier to manipulate, but allow artists to communicate imaginary spaces to their audiences. An example painting can be seen in figure 1.

* This work was supported by the Austrian Science Fund (FWF) under project P15663.

Figure 1. Stereoscopic painting "Eye" (1985) by Heinz Günther Leitner.
(Used by permission from VESCON Systemtechnik GmbH, Gleisdorf, Austria)

The Artists' Constraints

Looking at the handmade stereo paintings, it appears that the technicalities of preserving the viewpoints and colour consistency across the canvas pairs had an impact on the results. While most are meticulously executed, there seem to be some inconsistencies that do not always provide a comfortable stereo viewing experience. While Ferragallo states that he had been free-viewing the paintings while executing them, a rather cumbersome technique, there is little known as to the processes used by the other artists. Whether they have been painting simultaneously on both canvases or executed the two paintings sequentially one after the other, is not known. Observing some of these paintings, one can easily identify areas where the colour does not correlate; however fusing the stereo painting gives a blended appearance, due mostly to tonal variation and the effects of optical colour mixture. Nevertheless, the excessive labour and the technical expertise required are not the main reasons that stereo painting remains vastly unexplored until today. The major challenge for the stereo painters is the audience's inability to engage readily in stereoscopic free-viewing. Whether technology can provide a permanent and non-intrusive solution to this remains to be explored.

Technology & Stereo Painting

Since photographs can be a good base for preserving the viewpoint, they have sometimes been used by stereo painters. In our time, where computational and imaging devices come in all shapes, sizes and different capabilities, an interesting question arises: Can this more advanced technology be put to use in order to assist in the artistic process of stereo painting? The immediate answer is positive, but not a straightforward one. Nowadays, an artist can use a computer to paint directly in a digital form. The simplest method is to digitize a stereo pair of photos or to use stereo devices to feed his digital studio directly with a basis for developing his artwork, much like the photographs of Marc Lacroix that Dalí appears to have used. Unfortunately, conventional input devices (i.e. mouse and keyboard) are not convenient for the artist to interact with his working area. Therefore, less conventional input devices and computational mechanisms that translate human creative intention into a digital form have been invented over the past few decades, to assist in the artistic process (i.e. varying pressure pens on digital tablets, haptic input devices, etc). Furthermore, software to simulate painting media, brushes and surfaces has been in widespread use for some time now. Unfortunately, some problems impairing stereo painting creation still remain. The artist is bound to have to 'paint', with patience, almost the same composition twice, with no less effort than doing it by hand on a real canvas. Since the viewpoint can be accurately preserved, and since the basic images are fixed and the colours can be more or less managed across the two paintings, the ideal case for the artist would be to paint on one image, and automatically a *deus ex machina* helper process would fill in the other painting; the artist would then go over and correct any mistakes or add tonal effects.

Fortunately, such a helping agent exists, although it is not yet fully utilised. In computer vision, and particularly computer stereo vision, an open research topic is the automatic extraction of correspondence between features of the stereo pair. The idea governing this process, called *stereo matching*, is to search and extract from both images those scene points that are visible from both viewpoints. Even though this task may sound trivial, particular aspects of the stereo matching process form a rather difficult problem for the computer vision community which is beyond the scope of this article. The results of automatic stereo matching depend on many factors, such as the scene structure, texture, lighting and shading, to name but a few. However, assuming that feature correspondence can be established between the two components of a stereo image pair, referred to as the *disparity map* (see fig. 2(c)), it is possible to use this map to transfer paint from one image to the other. When a painting element (e.g. a brush stroke) is positioned on top of the one image by the artist, the disparity map encodes the corresponding position of it in the other image. It is obvious that the disparity map cannot provide any information for those points in the scene that are occluded in one of the views, as marked in fig. 2(d). For these, the artist must retouch the painting on the second image to complete the stereo painting, or else those regions must be reconstructed by performing more advanced image analysis.

Computer Generated Stereo Painting

As part of ongoing scientific research, we have formulated a semi automatic computational process[3] that can generate stereo paintings in a few seconds (see figure 2.). This system allows the user to provide high-level parameters (brush stroke parameters to use, and the method to establish scene point correspondences); it then extracts, automatically, a disparity map that is used in the generation of the second view of the stereo painting. Briefly, the painting algorithm[4] slices the image into a grid and positions brush strokes at grid points. The brush strokes follow the image colour and stop where large colour changes take place. The brush strokes are not guaranteed to paint a single object, since the algorithm operates in 2D space. To overcome this limitation we can restrict the brush stroke creation so that it does not extend between surfaces, an effect we call *paint spilling.* We can estimate the object interposition by using the disparity map, much like the human visual system extracts distance information from the two retinal images, and we do this by finding large changes in the disparity map; if two neighbouring disparity values have a significant difference, then the two points described lie on different surfaces. Another challenge is the filling in of the painting on the occluded areas. The artist, in our case the computer itself, cannot see the surface in the working image that needs to be filled in with paint in the second image. This is tackled by identifying these occluded points on the second image and painting directly on these areas on the second image. These newly placed paint elements are then overlaid with the brush strokes the artist painted originally on the working image, automatically, by using the disparity map.

(a) Source Left | (b) Source Right | (c) Disparity Map | (d) Occlusion Map

Left Painting | Right Painting

Figure 2. An automatically generated stereo painting.

Future Directions

As described earlier, technological advances, especially artistic content creation algorithms, allow many 2D software packages (e.g. Adobe Photoshop, Corel Painter) to include painting simulation processes amongst their many features. Nevertheless, if these effects are applied individually on the two input stereo images, the result is not guaranteed to be stereoscopically consistent. By contrast, for software packages directly working with three-dimensional models it is much easier to create stereoscopic paintings, since the geometrical structure of the scene is fully known and scene point correspondence can be mathematically calculated. Unfortunately, three-dimensional models are not trivially constructed and require either advanced 3D scanning equipment or professionals who will manually build them over time. The advantage of using image analysis is that it requires less effort and the complexity of the scenes can be much larger than that composed by 3D models. While the algorithm which we have developed works fairly well in generating stereo paintings, and new techniques will soon tackle sufficiently the technicalities of stereo matching, it is unlikely that the human artist can be replaced. The most interesting direction for our work is a hybrid of human-computer stereo painting system, where the user provides the human input and the computer takes up the repetitive labour intensive tasks.

References

[1] Durand F.: Limitations of the medium and pictorial techniques. *Perceptual and Artistic Principles for Effective Computer Depiction*, Course Notes for ACM SIGGRAPH 2002 (July 2002), pp. 27–45.

[2] Ferragallo R.: *On stereoscopic painting*, Leonardo 7, 2 (Spring 1974), pp. 97–104.

[3] Stavrakis E. and Gelautz M., *Image-based stereoscopic painterly rendering*, in Eurographics Symposium on Rendering (June 2004), pp. 53–60.

[4] Hertzmann A., *Painterly rendering with curved brush strokes of multiple sizes*, SIGGRAPH ’98, pp. 453–460.