

Image-Based Stereoscopic Painterly Rendering

E. Stavrakis and M. Gelautz

Interactive Media Systems Group, Vienna University of Technology, Austria

Abstract

We present a new image-based stereoscopic painterly algorithm that we use to automatically generate stereoscopic paintings. Our work is motivated by contemporary painters who have explored the aesthetic implications of painting stereo pairs of canvases. We base our method on two real images, acquired from spatially displaced cameras. We derive a depth map by utilizing computer vision depth-from-stereo techniques and use this information to plan and render stereo paintings. These paintings can be viewed stereoscopically, in which case the pictorial medium is perceptually extended by the viewer to better suggest the sense of distance.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation J.5 [Arts Humanities]: Fine Arts

1. Introduction

Stereoscopic painting is a special painting technique used by some painters. An artist creates two paintings of his composition, instead of one. The two paintings differ in that they are depicted from two horizontally displaced viewpoints. This canvas pair is to be viewed stereoscopically, so that one painting is viewed by each of the viewer's eyes. The viewer's brain is stimulated into fusing the paintings into a final composition, according to the principles of binocular vision [Wan95]. The major advantage of stereoscopic painting is that the artwork no longer remains flat and restricted to the two-dimensionality of the canvas itself. Once fused, elements of the stereoscopic composition appear to protrude in front of the display surface and others recede, making the painting more immersive, and in many cases, also more realistic.

The purpose of our paper is to present the idea of image-based stereoscopic painterly rendering. Our approach aims to eliminate the limitation of flatness of the 2D pictorial medium [Dur02] by re-introducing the binocular depth cues. We base our technique on images from real scenes that are content rich and enable us to produce more attractive and appealing paintings than those of 3D computer generated environments. We show how the existence of corresponding information between the two slightly different images can be used to enhance the appearance of the final stereo painting, as well as to speed up the overall rendering process.

We identify and discuss details that should be considered when devising painterly rendering algorithms which operate on images of real scenes.

Feature Correspondence. Stereo painting requires that depicted features correspond between the two canvases. [SK98] addressed the consistent editing of multiple views of the same scene in a plenoptic approach. In our case, painting elements, i.e. brush strokes, that cannot be matched in both views will inhibit stereo fusion and the viewers may experience discomfort. Also, large non-corresponding areas and deviations in paint color or style will produce similar undesired effects. Even though the brain is able to tolerate a small percentage of inconsistency, which varies from person to person, algorithms should strive on providing the best possible correspondence between the two paintings.

Randomness. A related issue to feature correspondence between views is the use of randomness. In some non-photorealistic rendering techniques [Sal97] [Her98] [Lit97], random numbers are used to inject irregularity into the process of abstraction or stylization. When randomness is used in stereoscopic painterly rendering, it must be as consistent as possible across a stereo image pair, so that irregularity can be equally modeled within both images.

Optimizations. Stereo pairs exhibit correlative information between their images that can be used to optimize a variety of algorithms [SGH*01] [ABC*91] (e.g. compression algorithms or rasterization of stereo pairs). Painterly algo-

gorithms can take advantage of stereo pair relationships (i.e. stereo disparity) to reduce the computational effort of planning or rendering of the paintings.

Depth from Stereo. As shown in various NPR works ([Mei96], [Shi00], [GCS02]), a depth map can be used to enhance the painting and its importance becomes critical in stereo painting, because it establishes a relationship between features of the two views. Depth of 3D computer generated scenes can be calculated fast and accurately. However, when dealing with images of real scenes, the computation of a corresponding depth map is much more complex. Depth maps derived by computer vision techniques, such as depth-from-stereo [SS02], are usually of considerably lower quality than those extracted from geometrically accurate 3D models. Subsequently, it is not practical to assume that the geometry can be reconstructed with accuracy. Painterly algorithms that use imperfect depth maps have to be designed to counterbalance the inaccurate geometric descriptions of the scene objects, without degrading the quality of the final painting.

Paint Spilling. Stereo painting cannot be considered as a process of painting two flat surfaces, it is a 3D process and varying depth levels exist within the final fused stereo painting. When painting the two canvases, surfaces at different distances from the viewer must be depicted with the appropriate colors. We call *paint spilling* the result of inadequately extending painting elements across surfaces at different depths. Intersurface paint spilling does not necessarily degrade the quality of a single painting, however in stereo painting it is a more noticeable and mostly unpleasant effect. The viewer may be confused since separation of the depth layers becomes more difficult.

Our technique is based on the acquisition of two still images (we will refer to them as the *left* and *right*) from spatially displaced cameras, which simulate the viewpoints from which an artist would depict his envisioned artwork. We correct the two images for lens distortions and bring them into epipolar geometry. We use automated stereo matching algorithms to calculate a disparity map in the geometry of the left image, as well as an occlusion map that encodes the locations of occluded scene points. We fix small artifacts of the disparity map by simple hole-filling. We then paint the left image with a modified version of Hertzmann's [Her98] painting algorithm that takes into consideration depth discontinuities. Using the occlusion map we paint occluded areas of the right view. We warp the left painting into the geometry of the right, based on the disparity map, to complete the right painting. As a final step, we display the stereo painting and use a method to automatically adjust stereo pair separation in order to reduce eye strain.

2. Background

The work presented in this paper is based on knowledge drawn from both the art and science communities. We make

a short summary of depth perception and its relevance to visual arts. Stereoscopic painting is a relatively new artistic technique that is not widely known. We give pointers to related work and literature in stereoscopic painting and finally we review computer graphics research closely related to our approach.

2.1. Depth perception in arts

In fine arts, human perception plays a central role in establishing a channel between the artist and his audience over which emotions, feelings and ideas may be communicated. Artists who have been able to use, and sometimes abuse, principles of human perception in creating and presenting their work, have attracted large audiences. In painting, some of the *monocular* depth cues (light and shade, relative size, interposition, textural gradient, aerial perspective, motion parallax, linear perspective) have been vastly exploited and exaggerated to compensate for the absence of *binocular* depth cues (binocular disparity and convergence). Binocular depth cues are provided by the two retinal images perceived, from our left and right eyes. In the presence of binocular depth cues, the human visual system is able to evaluate and appreciate depth information [Wan95]. When stereo pairs of images are created and presented to each eye, care must be taken to properly reproduce these cues, otherwise the viewer will experience discomfort [MSJ96].

2.2. Stereoscopic painting

Stereoscopy itself was first described by Wheatstone [Whe79], however, it has not been linked directly with painting, but has been mostly utilized in photography. In his manifesto Ferragallo [Fer74] calls the fine arts community to reintegrate the vast knowledge of painting into a framework of stereoscopic aesthetics. Girling [Gir90] and Futo [Fut91] record their experiences in anaglyph drawing and describe the technicalities and subtleties of laying down stereoscopic artwork.

Rene Magritte's 'Man with Newspaper' (1928) is an early painting that can be viewed stereoscopically. Michael Kupka creates stunning stereo oil paintings and drawings from stereo pairs of photographs. Roger Ferragallo and Oskar Fischinger have also been painting stereoscopic canvases, mostly by using geometric shapes to do abstract compositions. Salvador Dali [Mau89] has painted several stereo pairs of canvases. His mastery allowed him to blend real and surreal elements in stereo paintings with great fidelity and technical expertise. He called this art form 'Metaphysical Hyperrealism'.

The technical details of creating stereo pairs are inherent to the stereoscopic painting process. The artist has the labor intensive and cumbersome task of reproducing the same composition twice from slightly different viewpoints, while preserving these viewpoints horizontally aligned. To



Figure 1: Source left image.



Figure 2: Source right image.

overcome the technicalities involved in stereo pair creation, artists used stereo photography to base their compositions on, or they have been restricted to geometric forms. We believe that the technical requirements and excess effort in creating stereoscopic paintings, as well as presenting and viewing them, were the main reasons why the technique has not widely spread and remains still a mostly unexplored technique in painting.

2.3. Non-photorealistic image-based painterly rendering

Generation of aesthetically pleasing images is being actively researched and a wide range of techniques have already been introduced [GSS*99]. In our work, we focus on stroke-based rendering [Her03] and algorithms that use depth information to enhance the painting.

In [Shi00] a depth map, computed by ray-tracing, is used to adaptively control the parameters of their image moment-based painterly rendering algorithm. Specifically, they adjust the size of their stroke textures so that distant objects are approximated by larger strokes, whilst objects closer to the viewer are rendered with smaller strokes. They also use depth to introduce a mask that restricts each stroke to approximate only one object within the window that encompasses their stroke. Objects of the same color that lie at different depths are preserved as separate objects on the final painted image.

Gooch et al. [GCS02] use segmentation and morphological operations to automatically create a painting by planning and rendering a resolution independent set of strokes. In the same work, an extension of the algorithm uses depth maps, computed from 3D scenes, to assist the segmentation process via which better results are reported on the final painted images. A discussion about depth-from-stereo highlights the problems arising from such depth maps.

Hertzmann [Her98] introduced the use of curved brush strokes of multiple sizes that are planned and rendered in a variety of styles. The technique uses a coarse-to-fine concept, where the painting is progressively refined with differently sized spline brush strokes, similarly to how painters work.

2.4. Overview

The next section gives a description of our stereo painting algorithm and an extension for stereo painting presentation. In section 4, we present some results and discuss our experiences with users of our system. Finally, in section 5 we summarize this work and discuss future directions of stereo painterly rendering.

3. Stereoscopic painterly rendering

The proposed stereoscopic painterly rendering algorithm can be divided into two steps. In the first, we establish a relationship between the two views. In the subsequent step we use this relationship to paint the two images. As an optional step, we show how this relationship can be used when presenting the final paintings on digital displays.

3.1. Stereo painting algorithm inputs

Our algorithm uses an image pair that is acquired from two spatially displaced cameras. Since we are using cameras that do not have the properties of the ideal pinhole camera model [Fau93], our images are subjected to lens distortions. We use a standard calibration technique [Zha00] to calibrate each of our cameras and bring the stereo rig into epipolar geometry, so that scanlines are horizontally correspondent [Fau93]. We will refer to the corrected images of the left and right cameras as I_L and I_R respectively. A scene point visible from




Figure 3: Depth map calculated by stereo matching in the geometry of the right view. Missing disparity values (e.g. due to occlusions) are marked as black.

both cameras is mapped to a pixel $P_L(x_L, y_L)$ in the left image and pixel $P_R(x_R, y_R)$ in the right. The vector describing the pixel distance between these pixels is the disparity vector $\vec{d}(x_R - x_L, y_R - y_L)$. Since we have a horizontal correspondence of features between the two images in epipolar geometry ($y_L = y_R$), the disparity vector will be $\vec{d}(x_R - x_L, 0)$. A crucial step in stereo analysis is the computation of a disparity map $I_{disparity}$, using automated stereo matching techniques [SS02]. We note that disparity is inversely proportional to the distance of each point from the viewpoint and hence a depth map I_{depth} is inferred from the disparity map by using

$$I_{depth} = s \cdot \frac{1}{I_{disparity}}$$

where s is a scaling factor. The disparity map can be used to transform the first stereo view into the geometry of the second one, and vice versa. Stereo-derived depth values can normally only be computed for those pixels that are visible



Figure 4: Depth map in the geometry of the left view. Missing values reconstructed from neighboring values.

in both the left and right image. This means that some pixels may not be assigned a depth value. In our implementation, we estimate and reconstruct the missing depth values by using reliable neighboring depth values. We store the location of those pixels, usually a small portion of the whole image, in a binary map, which we refer to as occlusion map I_{occ} . The occlusion map can be computed for both the left and the right view. The pixels encoded in I_{occ} require special consideration during the painting process.

3.2. Stereo painting algorithm

Our stereo painting algorithm is a modified version of Hertzmann's [Her98] single image painting algorithm. For completeness, we briefly describe the original algorithm together with the modifications and a discussion on the motivation behind each modification. Readers are encouraged to consult the original paper for a thorough discussion of all properties of the painting algorithm.

A list of brush stroke sizes of radii $R_1 \dots R_n$ is used to paint progressively a series of layers, from coarse-to-fine. Each layer is created by filtering the original image I_{source} (either I_L or I_R) with a median filter, to obtain a reference image I_{ref} . The median filtered image is computed by convolving I_{source} with a pixel-centered kernel of size $K \times K$, where $K = 2 \cdot f_g R_i + 1$, with f_g as the blurring factor. The median filter's response gives rise to intensity edges that are helpful in distinguishing objects in stereo viewing. We use a jittered grid to place strokes as described in the original algorithm. The jittered grid is an equally spaced grid with cell size $f_g R_i$, where f_g is a grid scaling factor. At a neighborhood $M \in I_{ref}$ of size $(f_g R_i) \times (f_g R_i)$ around each grid point, we find the pixel in M that has the maximum Euclidean distance, in RGB color space, from its corresponding pixel in the image I_{camas} we progressively paint on. If this distance is above a threshold T we plan a stroke by defining the first



Figure 5: Depth map in the geometry of the right view. Missing values reconstructed from neighboring values.




Figure 6: Initial occlusion map and resulting map after dilation.




Figure 7: Painting occluded areas of the right view, by using the occlusion maps of figure 6. Painting of occluded areas by using the undilated occlusion map (shown left) is less accurate than the one provided by the dilated occlusion map (shown right).



Figure 8: Final painting of right view. On the left, painting generated without dilation of the occlusion map (visible seam due to depth map inaccuracies). On the right, painting generated after dilation of the occlusion map (no seams visible, depth inaccuracies compensated).

control point (x_0, y_0) of the spline brush stroke at that location.

To order the strokes that paint the respective canvas I_{canvas} for I_{source} , we generate a random list of numbers that is equal to the maximum value of the depth map's histogram. The maximum value of this histogram can be thought of as the maximum amount of brush strokes that may be painted on the canvas at a particular depth value. From the largest to the

smallest depth value, we order and render all strokes at each depth value by using the random list, effectively avoiding regularity that may occur in areas where successive strokes may have the same depth (i.e. a surface perpendicular to the viewing direction).

To create the stereo painting, we first use the process described above on I_L to generate the left painting. We apply paint on all pixels of the left image. We then process I_R , but we now take into consideration the pixel locations of the right view that could not be seen from the left, which are encoded within I_{occ} . We allow strokes to be initiated only from the image locations of occluded pixels. The occlusion map remains constant between refining layers and the strokes are allowed to expand beyond the borders of the occluded regions. If we restricted brush strokes to be planned only within the occluded regions, the brush stroke termination criteria, discussed later, would prohibit the placement of most of the larger strokes, impeding the refining process of the algorithm. Once the painting of occluded areas is completed, the disparity map is used to transform the left painting atop of the current painting of I_R .

As described earlier, the disparity maps we are using are not perfect and certain errors are inherent in the painting process. One of them is that near depth discontinuities, where usually occlusions are present, the disparity maps may slightly be expanded or contracted. The effect of these inaccuracies is that 'seams' between the generated and the propagated painting are visible after the warping operation. To overcome this problem, prior to painting the right image, we expand the occlusion map by applying a dilation morphological operator to I_{occ} . Depending on the input images used, a few iterations of dilation may be required. We found that a number of iterations equal to $\max\{R_1, \dots, R_n\}$ produced good results in most cases. Another problem arising when warping the painting is that even though large occluded areas will be fixed in the above step, areas formed only by a few pixels will still remain as empty holes. To fix them we copy the pixel colors of I_{ref} .

Our results, discussed in section 4, show that our strategies for correcting the artifacts, produced by the low-quality depth maps, are producing good results.

3.3. Stroke creation

The spline stroke planning procedure places control points (x_i, y_i) normal to the direction of the gradient of the Sobel-filtered luminance of I_{ref} , at a distance of R_i . The color of the brush stroke is set to the color of I_{ref} at the location of the first control point (x_0, y_0) . Note that a stroke can only be initiated from a location at which an occluded pixel is present, when painting the right image.

We continue to place control points until one of the following criteria is met:

(a) a predefined maximum number of control points has been reached,

(b) the difference between the color of I_{ref} at (x_{i-1}, y_{i-1}) and the constant brush stroke color is larger than the difference between the color of I_{ref} and the color of I_{camas} at (x_{i-1}, y_{i-1}) ,

(c) the magnitude of the gradient of the Sobel-filtered luminance of I_{ref} is a vanishing gradient,

(d) any pixel P within a distance from (x_i, y_i) smaller than R_i has an absolute difference of depth value from the depth value at (x_i, y_i) which is greater than a predefined depth threshold T_{depth} .

Termination criteria (a) through (c) are part of the original algorithm, we introduced criterion (d) to handle *paint spilling* as explained next.

As we discussed in the introduction, *paint spilling* over neighboring surfaces at different depths may be uncomfortable to see in a stereo painting. Termination criterion (d) of the stroke creation algorithm prevents *paint spilling*. The problem arises because, as the spline stroke is being planned, control points may be located inadequately close to a depth discontinuity, which will cause paint to spill beyond the object's boundaries. This happens because the stroke has a width of R_i and when painted, it will extend further from where the actual spline has been planned. Calculating the difference between the depth value at the current control point location and each pixel's depth value within a window of size $(2 \cdot R_i + 1) \times (2 \cdot R_i + 1)$ around it, ensures that no depth discontinuities larger than T_{depth} are in proximity and the control point can be added to the stroke. In addition, the pixels tested for depth difference include all the possible locations for the next candidate control point, which should be placed at a distance R_i .

This test would not prevent paint spilling, if the distance between the control points were larger than R_i . In this case, an additional test of the absolute depth difference must be performed between subsequent control points and T_{depth} .

3.4. Stereo painting presentation and viewing

The final output of our algorithm is two paintings with correlative information that can be viewed stereoscopically. The generated stereo painting is a full color stereo pair, which can be presented and viewed with any of the standard methods (e.g. printed on paper and viewed with a mirror stereoscope, viewed on a digital display in a side-by-side image arrangement, etc.). We have chosen to present our paintings on both conventional digital displays and large stereo immersive tables. The stereo paintings are presented by using standard page-flipping stereo and viewers use active shutter glasses.

A stereo pair presented on the flat surface, e.g. a digital display, can be adjusted by altering the stereo pair separation in order to increase or decrease the effect of depth per-




Figure 9: Adjusting the plane of zero parallax by horizontal image translation. On the left, the plane of zero parallax coincides with the background, while on the right the plane of zero parallax lies on the body of the subject.

ception. Habitually in the real world, when looking at an object the human eyes converge and accommodate (focus) causing the object to be seen singly. However, when viewing stereo image pairs on a digital display, accommodation remains fixed on the display surface, while the eyes can rotate, so that the lines of sight converge at a particular distance [HF93]. When two corresponding image points as seen from the left and the right eyes coincide, then the fused point lies on the display surface and has zero parallax. Any points that have negative parallax will appear to be between the viewer and the display surface, while those points with positive parallax will be perceived as being behind the display. By adjusting the parallax between particular points in the scene, it is possible to set different scene points to fall on the plane of zero parallax, effectively bringing out or pushing behind the display different scene objects.

Parallax in our case may be changed by horizontally translating the two paintings. We employ a simple method for automatic horizontal image translation based on the disparity map. We translate the two paintings horizontally, bringing the corresponding image points with minimum disparity on the plane of zero parallax. This guarantees that none of the scene objects will appear to be in front of the displaying surface. When fused objects appear to be in front of the surrounding of the display, this information comes into conflict with the monocular depth cue of interposition, which signals the user's visual system that the frame of the display surface is in front of everything displayed within it [MSJ96]. Thus to avoid this problem, which may cause eye strain or breakdown of stereopsis, we initially set the whole scene behind the display's image plane and we allow the user to interactively adjust which scene points will be on the plane of zero parallax, via mouse interaction. We track the mouse movement and read the corresponding disparity value at the mouse position and smoothly translate the images so that the corresponding points, to this disparity value, coincide.

4. Tests and results

We have presented stereo paintings generated with our algorithm to several users. While most of the users were not aware of stereoscopic painting, they found the resulting paintings appealing and particularly pointed out the sense of distance.

Figures 1 and 2 show the original stereo image pair. The depth maps for the left and right view are computed by automated stereo analysis and can be seen in figures 4 and 5 respectively. Figure 6 gives the occlusion map for the right painting before and after dilation. The results of painting brush strokes emanating from the occluded regions using the respective occlusion maps of figure 6 are shown in figure 7. The final painting in the geometry of the right view as presented in figure 8, shows the difference between using an undilated and a dilated occlusion map. In figure 9, the left and right painted images are plotted on top of each other to illustrate the effect of parallax adjustment. Stereoscopic paintings generated by our algorithm are presented in figures 10 through 12, along with the original image pairs.

We have presented the stereo paintings on a large stereo display, where users have reported better immersion than on the smaller conventional computer monitors. We adjusted the painting parameters to present both rough and refined paintings to the users. To our surprise many of them preferred the coarser paintings, where the painterly effect was more recognizable.

In one of these tests we created two versions of the same stereo painting, one rendered with our stereo painterly algorithm and one by painting the two images individually. Some of the users reported difficulties in fusing the individually painted images and some users with technical knowledge of stereo immediately pointed out the areas where non-correlating features existed. The difficulty in fusing the individually painted images increased proportionally to the coarseness of the painting.

On the technical aspect of our algorithm, we achieved a speed increase over the individual painting of the two images, since we avoided painting most of the second image. Table 1 gives an idea of the performance increase we achieved by painting with our stereo algorithm, in comparison to painting the two views individually. This computational enhancement is significant when creating high resolution stereo paintings.

5. Summary and future work

We have presented an algorithm for generating stereo paintings of real images. We used depth-from-stereo to establish a relationship of corresponding features between the two images. We painted only the one image and used this relationship to propagate paint to the second unpainted image. We tackled the problem of inter-surface *paint spilling* and

Image Size	Improvement
640x480px	1%
1280x1024px	9%
1600x1200px	13%

Table 1: Measurements of performance improvement of the stereo painting algorithm over individually painted views.

provided a solution to painting occluded areas of the second view where paint could not be propagated. The stereo paintings produced by our algorithm give the impression of a painted world, in which the brush strokes are attached to objects.

Stereoscopic painting as an art form exhibits divergence from single-view painting; in creating, presenting and viewing it. Stereopsis apart from being a visual mechanism can be used to affect perception of aesthetically pleasing pictorial information.

The duality of a stereoscopic painting, allowing active participation of both the artist and the viewer, needs to be further explored. We believe that stereoscopic fusion can be mixed together with other mechanisms that enhance depth perception (e.g. contrast adjustment, stylization of shapes) in new ways. This may allow greater degrees of artistic freedom in stereo painting. For example, it may be obvious that encoding the same emphasis on both images will give rise to a targeted artistic expression or effect, it is not so obvious what will be the aesthetic implications of encoding different percentages of emphasis in each of the stereo images.

As discussed in [Sim03], stereoscopic projection systems can be a useful tool for artists in exploring space of stereoscopic artwork. However, apart from stereo-viewing devices, new interfaces need to be defined for interactive stereoscopic painting of real images in a computer, where the artist is given more control over the painting style and spatial relations of his compositions. For instance, as described in this work, depth-from-stereo can be used to paint scene objects at their relative positions in space. An interface that allows the artist to directly manipulate the depth map in order to exaggerate the aesthetic value of scene objects may be useful. Furthermore, depth maps from different stereo sources could be composed into a single stereo painting to produce interesting spatial effects.

6. Acknowledgement

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

References

- [ABC*91] ADELSON S. J., BENTLEY J. B., CHONG I. S., HODGES L. F., WINOGRAD J.: Simultaneous generation of stereoscopic views. *Computer Graphics Forum* 10, 1 (1991), 3–10. 1
- [Dur02] 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. 1
- [Fau93] FAUGERAS O.: *Three-Dimensional Computer Vision: a Geometric Viewpoint*. The MIT Press, 1993. 3
- [Fer74] FERRAGALLO R.: On stereoscopic painting. *Leonardo* 7, 2 (Spring 1974), 97–104. 2
- [Fut91] FUTU L. P.: *Anaglyph Painting*. Zurich: Phänomen-Art Verlag Zurich, 1991. 2
- [GCS02] GOOCH B., COOMBE G., SHIRLEY P.: Artistic vision: Painterly rendering using computer vision techniques. In *Proc. NPAR '02* (2002), ACM Press, pp. 83–90. 2, 3
- [Gir90] GIRLING A. N.: *Stereoscopic Drawing: A Theory of 3-D Vision and Its Application to Stereoscopic Drawing*. AN Girling, 1990. 2
- [GSS*99] GREEN S., SALESIN D., SCHOFIELD S., HERTZMANN A., LITWINOWICZ P., GOOCH A., CURTIS C., GOOCH B.: Non-photorealistic rendering. *Course Notes for SIGGRAPH 1999* (August 1999). 3
- [Her98] HERTZMANN A.: Painterly rendering with curved brush strokes of multiple sizes. In *Proc. SIGGRAPH '98* (1998), ACM Press, pp. 453–460. 1, 2, 3, 4
- [Her03] HERTZMANN A.: A survey of stroke-based rendering. *IEEE Computer Graphics & Applications, Special Issue on Non-Photorealistic Rendering* 23, 4 (July/August 2003), pp. 70–81. 3
- [HF93] HIRUMA N., FUKUDA T.: Accommodation response to binocular stereoscopic tv images and their viewing conditions. *Society of Motion Picture and Television Engineers* 102, 12 (1993), 1137–1144. 6
- [Lit97] LITWINOWICZ P.: Processing images and video for an impressionist effect. In *Proc. SIGGRAPH '97* (1997), vol. 31, pp. 407–414. 1
- [Mau98] MAUR K. V.: *Salvador Dali, 1904-1989*. Verlag Gerd Hatje, 1989. 2
- [Mei96] MEIER B. J.: Painterly rendering for animation. In *Proc. SIGGRAPH '96* (1996), vol. 30, pp. 477–484. 2
- [MSJ96] MCVEIGH J. S., SIEGEL M. W., JORDAN A. G.: Algorithm for automated eye-strain reduction in real stereoscopic images and sequences. In *Proc. SPIE '96, Human Vision and Electronic Imaging* (1996), vol. 2657, pp. 307–316. 2, 6
- [Sal97] SALISBURY M. P.: *Image-Based Pen-and-Ink Illustration*. PhD thesis, Department of Computer Science and Engineering, University of Washington, 1997. 1
- [SGH*01] SAWHNEY H. S., GUO Y., HANNA K., KUMAR R., ADKINS S., ZHOU S.: Hybrid stereo camera: an IBR approach for synthesis of very high resolution stereoscopic image sequences. In *Proc. SIGGRAPH '01* (2001), ACM Press, pp. 451–460. 1
- [Shi00] SHIRAIISHI M.: Adaptive parameter control for image moment-based painterly rendering. In *Proc. 9th International Conference on Geometry and Graphics* (2000). 2, 3
- [Sim03] SIMÓ A.: A horizontal stereoscopic projection system for working at the artist studio & mimesis, the function that made the organ. In *Sketches & Applications, SIGGRAPH '03* (2003). 7
- [SK98] SEITZ S. M., KUTULAKOS K. N.: Plenoptic image editing. In *Proceedings of the Sixth International Conference on Computer Vision* (1998), IEEE Computer Society, p. 17. 1
- [SS02] SCHARSTEIN D., SZELISKI R.: A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *International Journal of Computer Vision* 22, 1/2/3 (April-June 2002), 7–47. 2, 4
- [Wan95] WANDELL B. A.: *Foundations of Vision*. Sinauer Associates, Inc, 1995. 1, 2
- [Whe79] WHEATSTONE S. C.: Contributions to the Physiology of Vision. Part the First; On Some Remarkable, and Hitherto Unobserved, Phenomena of Binocular Vision. *Philosophical Transactions of the Royal Society* (1879), 371–394. 2
- [Zha00] ZHANG Z.: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22, 11 (2000), 1330–1334. 3



Figure 10: On the left, a source stereo pair. On the right, the resulting painted stereo pair.



Figure 11: On the left, source image pair from Middlebury College, with provided ground truth disparity map. On the right, stereo pair generated by our algorithm.



Figure 12: On the left, another source image pair is shown. On the right, the painted stereo pair generated by our algorithm.