UC Davis

IDAV Publications

Title

Non-Photorealistic Rendering using Watercolor Inspired Textures and Illumination

Permalink

<https://escholarship.org/uc/item/2g48f8kn>

Authors

Lum, Eric

Ma, Kwan-Liu

Publication Date

2001

Peer reviewed

Non-Photorealistic Rendering using Watercolor Inspired Textures and Illumination

Eric B. Lum Kwan-Liu Ma
Department of Computer Science
University of California, Davis
{lume,ma}@cs.ucdavis.edu

Abstract

In this paper, we present a watercolor inspired method for the rendering of surfaces. Our approach mimics the watercolor process by building up an illuminated scene through the compositing of several layers of semi-transparent paint. The key steps consist of creating textures for each layer using LIC of Perlin Noise, and then calculating the layer thickness distribution using an inverted subtractive lighting model. The resulting watercolor-style images have color coherence that results from the mixing of a limited palette of paints. The new lighting model helps to better convey large shape changes, while texture orientations give hints of less dominant features. The rendered images therefore possess perceptual clues to more effectively communicate shape and texture information.

1 Introduction

Through the careful use of a brush, a skilled watercolor artist is able to distribute semi-transparent pigments on a sheet of paper in a manner that creates a meaningful image. An artist can paint surfaces in a way that conveys the shape and texture of an object that might not be visible in reality. In addition, through the careful selecting and mixing of a limited palette of paints, an artist can present a sense of color harmony while at the same time adding a distinct mood to a scene. In this way, an artist can create a painting that may not be true to reality, but is rich with information.

Our work uses ideas from watercolor painting to communicate information about objects in a scene that might not be visible with more photo-realistic rendering methods. The goal of our work is not to produce realistic watercolor pictures, but rather it is to use watercolor techniques that an artist might employ to convey perceptual clues that could help illustrate shape and spatial relationships. Our work does not attempt to use painterly rendering for its own sake,

but rather as a means to present more information.

Our method builds up a scene by iteratively compositing layers of watercolor paint as shown in Figure 1. Each iteration requires the computation of a texture for that layer, and the determination of how the thickness of the paint in that layer should vary across the image. The texture for each layer is computed using Line Integral Convolution (LIC) [2] along one of the principal curvature directions. Then the layer thickness is determined using a modified Phong shading model. Since lighting is performed through the process of mixing pigments, a subtractive color model is used. This requires an inverted lighting model that indicates illumination not through the adding of color where an object is lit, but rather through the adding of paint based on an object's darkness.

Both of these techniques occurs in 3D object space, which stands in contrast to 2D painterly image filters that process images entirely in 2D projected image space. By avoiding image-space manipulation, frame-to-frame coherence is maintained, allowing for the creation of animation

Figure 1. An iterative non-photorealistic rendering process using a watercolor approach.

that avoids the "shower door effect" [1] or other jittering artifacts. In addition, by calculating textures in the 3D object space, it is possible to use painterly washes and textures that are derived from the objects themselves, rather than from a 2D projection of those objects. Our painterly illumination model allows for painterly variations of value and tone, which is derived from lighting information in the 3D object space. This would be difficult to derive from a projected 2D image.

2 Related Work

Several authors have described the use of principal curvature directions for the perceptually meaningful non-photorealistic rendering of surfaces. Interrante [9] describes how LIC along the principal curvature directions can be used to create brush stroke textures. Her work deals with the creation of individual brush strokes and describes issues involved in selecting their density, thickness, and color to produce images that convey information for scientific visualization applications. In a following work by Girshik and Interrante [5], lines are rendered on a surface by tracing along one of the principal curvature directions.

Hertzmann and Zorin [8] describe an algorithm for the rendering of smooth surfaces in a line-art style. They create hatching patterns along the smooth curvature direction fields to communicate shape information. They also describe how C^2 subdivision surfaces and their principal curvature directions can be computed. Our work creates painterly textures for subdivision surfaces. Our algorithm first calculates raw principal curvature direction vectors using the method described by Hertzmann and Zorin [8]. These directions are then assigned a consistent orientation and smoothed. After this step, LIC along one of the principal curvature directions is performed like the work done by Interrante. Our work differs in that we apply LIC to Perlin Noise [13], allowing for a wide range of painterly textures. In addition our work uses several textures for each surface, with each texture applied to a different paint layer.

Gooch et al. [6] introduce a non-photorealistic lighting model that uses tone for illumination. By varying the color temperature along a surface from warm to cool they are able to convey visual clues as to the shape of objects while reserving extreme lights and darks for highlights and edge lines. Curtis et al. [3] describe a model for simulation of a wide range of watercolor techniques. They present an interactive watercolor paint system that allows a user to paint a number of realistic watercolor effects. They also show a process for the "watercolorization" of images and as an extension the "watercolorization" of projected 3D scenes. Nishita et al. [12] describe an algorithm for the display of individual brush strokes using Bezier Functions, while Way et al. [15] present a technique that uses of database of paint-

ing components that are applied to contours sketched by the user to create computer-generated Chinese paintings.

For our work we sought a lighting model that uses perceptually meaningful color temperature variations as used by Gooch et al., one that is produced in a manner analogous to how a real watercolor artist would apply paint a scene. This is accomplished through the compositing of several layers of paint of varying color temperature. Based on the lighting of a scene we calculate the paint thickness distribution for layers and then composite using the Kubelka-Munk (KM) color model [7] as used by Curtis et al. [3]. The KM model can be used to model pigments in terms of both their absorption and scattering characteristics. This makes possible more realistic calculation of colors which would result from the ordered application of successive layers of paint.

3 Watercolor

Watercolor is unique from other painting mediums in that semi-transparent pigments are used extensively, unlike the relatively opaque paints often used in oil or acrylic painting. These watered pigments can be applied using a wide range of brushes and techniques. One of the most common techniques used in watercolor is the application of thin washes of watered paint to regions of the paper. By compositing several consecutive washes on a sheet of paper, an artist can build up a scene. She can mix colors on the paper by applying a semi-transparent wash over washes that have already dried, allowing parts of the deeper wash to show through.

Unlike in oils painting, where an artist might mix white paint to form a lighter color, in watercolor, light colors are created by using less paint, permitting more of the underlying white paper to show through. White regions are often reserved for specular lighting, and can be achieved not through the application of white paint, but rather by avoiding painting in the highlight regions entirely. In some cases a masking fluid is applied that is later removed, preserving the underlying white color of the original paper.

An artist is able to convey a mood or atmosphere in a work through the careful selection of colors. In addition, by using a limited set of paints, and only using colors that can be created through the mixing of those paints, an artist is able to add a sense of "rightness" and color harmony [10] to a work. It is important to note that the watercolor pigments do not mix as light does. The mixing of red, green and blue paint does not yield white as it does with light, but instead might result in a bluish gray color.

An important part of lighting in watercolor involves not only the value or intensity of the applied paint, but also the color's temperature. Warm colors include yellow, red and orange, and can convey a sense of cheerfulness and excitement, while cool colors include blue, green and purple, and

Figure 2. The process of layer texture generation.

can give a refreshing, relaxing, or perhaps gloomy feel. Individual colors can be relatively cool or warm; for example, there are cool yellows and warm blues. Warm colors are often used to show surfaces that are directly lit by a light source, since light sources like the sun are usually warm in color. On the other hand, shadows are usually cooler in color since the light in shadowed areas is often the result of light reflected off other objects, or from a blue sky. In watercolor, shadows are rarely depicted as being simply darker, instead an artist will often look closely at a shadow in an attempt to identify any colors the shadow may possess. A shadow wash is often applied over an object that has already been painted with its diffuse color, making possible the underlying color of the object to show through.

4 Texture Generation

In our work, watercolor style textures are created in the 3D object space rather than in the final image plane. This can be contrasted with how a real artist creates a painting where the paint is manipulated exclusively on the image plane, specifically on a white sheet of paper. This deviation from the real watercolor creation process is necessary to create animation that avoids temporal artifacts. Our work does not focus on the synthesis of individual brush strokes, but instead focuses on the creation of wash-like textures.

Figure 2 summarizes the process of layer texture generation. Textures are created by performing LIC of Perlin Noise along one of the principal curvature directions of a surface.

Although principal curvature direction field lines are not

visible in reality, they can be used to approximate one way an artist might use a brush to convey the shape of a surface. Raw principal curvature directions are calculated for each vertex of a triangular subdivision-surface using the technique described by Hertzmann and Zornin [8]. Their optimization technique for smoothing the principal curvature directions is not employed in order to avoid discarding features that might be eliminated in the optimization process.

Each principal direction is not a vector since the maximum and minimum curvatures at a point on a surface are the same forward as backward. However, in order to perform proper linear interpolation between vertex samples on the surface, a consistent vector orientation is defined. It is desirable to define an orientation of these vectors such that the vectors are smooth with respect to their neighbors. Rather than performing a global optimization of these orientations, a greedy *breadth-first* search approximation is utilized. An initial vertex seed is chosen at random. Then, in a breadth first search manner, the orientations of the neighboring vertices is defined such that smoothness is maximized with respect to those vertices that have had their orientation already defined. This process is repeated several times with different seeds, keeping the set of orientations that resulted in the highest amount of smoothness. After this step, vector averaging can be performed between each vertex and its neighbors to smooth the curvature vector fields further.

Once the smoothed, vectorized principal curvature directions are calculated, they are used to perform LIC of Perlin Noise. LIC is a technique that low-pass filters a volume along a direction and has been used extensively for vector field visualization. Unlike the work of Interrante [9] that uses LIC along principal curvature directions to create brushstrokes, we apply LIC to Perlin Noise to create textures that more closely resemble a watercolor wash. We have found that the smoothing of Perlin Noise can be used to create a wide range of textures that can roughly approximate the appearance of watercolor. By applying a *gain* and *bias* [4] to the noise prior to LIC it is also possible to create textures that range from smooth and wash-like (Figure 3) to distinct and brush-like (Figure 4). It is also possible to further soften the texture by applying a gain and bias after the LIC step. Convolution is done by tracing a path along the surface in 3D space while accumulating the volumetric noise function. This smooths the noise along one of the principal curvature directions, resulting in an appearance reminiscent of which a wet brush would create. Although the smoothing of Perlin Noise along a principal curvature direction does not have any physical relationship with the distribution of watered pigments on paper, we have found that the resulting textures can be perceptually meaningful and visually pleasing.

Figure 3. Smooth and wash-like appearance. A wash-like texture does not show brushstrokes, but rather has an appearance reminiscent of one that might be created while distributing watered pigment with a large brush.

Figure 4. Distinct and brush-like appearance. By sharpening the Perlin Noise prior to LIC a texture is created that is much more distinct, with a more brush-like appearance.

5 Pigment Based Lighting Model

Through the compositing of a limited palette of semi-transparent paints, a watercolor artist is able to create lighting effects that not only provide clues to the shapes of objects in a scene, but also present a desired mood or atmosphere. We present a pigment-based illumination method to mimic the way light might be painted in watercolor.

Combining colors from a limited palette is an important part of watercolor painting. Like the work done by Curtis et al. [3], the KM model is used to simulate the compositing of pigment layers of watercolor paints. If a simple additive color model were used instead, a different set of colors would result from the mixing of a limited palette. These colors would not give the same sense of color coherence that would occur with real watercolor paints.

Illumination through painting requires the inversion of traditional lighting equations since instead of adding light to a scene, one composites colored paints onto white paper. In traditional computer graphics, objects are often thought of as having an ambient color, to which diffuse and specular lighting is added. Since our method uses a subtractive pigment color model, this process must be reversed. Instead of adding color where a surface is lit, we apply additional pigment where a surface is dark. This occurs through the compositing of several wash-like layers of paint not only to indicate diffuse lighting, but more importantly to add paint to unlit regions.

First, the specular or highlight regions of the scene

are identified. These bright regions correspond to areas on a painting where the watercolor pigment should not be applied or should be applied in a very thin layer. A layer of paint is then added for diffuse lighting. Objects are typically painted entirely in their diffuse color except for those regions that have been identified as having a strong specular component, in which case paint is either not applied, or is applied thinly. This can be defined by:

$$\text{diffuse thickness} = (1 - (V \cdot R)^{n_{\text{specular}}})^{k_{\text{diffuse}}}$$

where V is the view vector and R is the reflected light direction.

Then a layer of paint is composited to indicate shading. This "unlit" layer indicates to what extent an object does not receive diffuse lighting and can be defined by:

$$\text{unlit layer thickness} = (1 - (L \cdot N))^{k_{\text{ambient}}}$$

where N is the surface normal and L is the light source direction. In other words, those regions that have no diffuse lighting as calculated by a Phong shading model receive the thickest layer of paint. Paint thickness is reduced in those areas that receive the most lighting. Figure 5 displays the diffuse-light layer (top), the unlit layer (middle), and their composited result (bottom) for a two cups model.

Finally, a brush-like layer of brush shaped paint can be added. This corresponds to dry brush stroke techniques a watercolor artist might use after the final wash has dried. This layer is created in the same manner as the wash-like

Layer 1: Diffuse lighting

Layer 1 + Layer 2

Layer 2: Unlit wash

Layer 3: Brushed shaped paint

Composited result (Layer 1 + Layer 2)

Final Composited result (Layer 1 + Layer 2 + Layer 3)

Figure 5. Layers of paint are added for diffuse and unlit color contributions.

Figure 6. A third brush-like layer of paint is added to produce the final image shown on the bottom.

layers, except with Perlin Noise that has had a gain and bias applied to it prior to LIC. Figure 6 shows the result of adding the brush-like layer to produce the final image.

Watercolor paintings rarely have smooth Phong shading. Not only is it very difficult for a human to paint smooth Phong shadings with a brush, but by making the lighting transitions more abrupt it is possible to make the large shape changes more distinct. In order to reduce the appearance of extremely smooth computer generated Phong shading, a gain and bias is applied to the unlit paint layer thickness. Figure 7 shows this difference. The left image has smooth Phong shading while the right one displays rough transitions as a result of mimicking the watercolor process.

An artist often fades the edges of a surface to give an object a softer appearance. This effect can be simulated by thinning a layer of paint, depending on the degree that a polygon faces the viewpoint. The paint thickness is varied depending on the dot product between the viewing direction, and the surface normal. This is similar to techniques that have been used for silhouette detection [11], except that instead of drawing a dark silhouette, the edges are softened. The resulting images have a softer appearance as desired. Figure 8 shows images before and after softening. By attenuating the ambient layer of paint in the silhouette regions, the cup has a much softer appearance. In particular the left side of the cup fades into the background and the dark line on the right side of the cup has been removed.

Our pigment based illumination model allows shading to occur using a discrete palette of paints that are mixed to produce a set of colors like those found in watercolor. This would not be the case if a simple additive color-mixing model were used. Through the careful selection of paints, it is possible to create a variety of atmospheres that have a sense of coherence made possible by the mixing of pigments from a limited palette. In Figure 9, the left image uses cool colors and has a gloomy appearance while the right image was generated by using a bright palette of colors to achieve a much more cheerful atmosphere. Notice in the left image the cool blue used for the unlit layer and the relatively cool red for directly illuminated regions. In the right image, a very warm orange is used for the diffuse layer with a warm yellowish-green used for shading. The unlit layer still has a cooler color than the diffuse layer. This change in temperature as well as intensity helps give the shading a more watercolor styled appearance.

In addition, by selecting paints for unlit and illuminated layers based on temperature, objects can have warm diffuse lighting and have cool shadows, as is often found in traditional watercolor. Figure 10 shows a rendering of the Utah Teapot using a standard Phong lighting model, while Figure 11 demonstrates the use of varying temperatures.

Usually the diffuse layer of pigment covers an object in its entirety. Thus, after the unlit layer of paint is added,

the darker regions still have parts of their diffuse color present. However, it is sometimes desirable to minimize the amount of mixing that occurs between the diffuse and unlit colors of an object, in which case the diffuse layer of paint can fade as the amount of diffuse lighting on an object is reduced. The corresponding equation would look like:

$$\text{diffuse thickness} = (1 - (V \cdot R)^{n_{\text{specular}}})(L \cdot N)k_{\text{ambient}}$$

This can be particularly effective in changing the temperature of an object from one extreme to the other as seen in Figure 12, but also causes less of the object’s underlying color to be shown in the shaded regions.

The rendering framework is implemented using ray tracing with the layer textures generated procedurally on a per pixel basis. For better performance a more efficient implementation could use pre-calculated textures generated using Fast LIC [14]. The current ray tracing implementation does not render at real time rates at higher resolutions. Nevertheless, low-resolution 128x128 images of the scene shown in Figure 9 can be rendered in under 1 second with an AMD Athlon 1.4 GHz processor which makes possible interactive tuning of texture and viewing parameters.

6 Conclusions

We have presented a method for non-photorealistic rendering using watercolor inspired textures and illumination. The method has produced encouraging results that provide perceptual clues about an object’s shape that might not be visible with more photorealistic rendering methods. The illumination technique used allows the rendered images to have a sense of color coherence. In addition color temperature can be varied as well as color value to indicate an object’s shading. The combination of these techniques allow us to create images that can be more meaningful than a photorealistic rendering of the same scene. Since both of these methods are applied in 3D object space, animation sequences can be generated without temporal artifacts that could detract from the shape information being presented.

Future work includes investigating how even more meaning could be conveyed using watercolor methods. Watercolor artists have a vast array of techniques at their disposal. This paper has only covered a few of them. By using additional techniques even more meaningful images could be created. However, some watercolor techniques are not well suited for indicating the properties of a given shape. For example, the fading of paint at the edges of an object might mimic a particular watercolor style, and can show softness in an object, but the perceptual clues provided from having sharper edges can also be lost. Future work on improving user interfaces to allow intuitive selection of a variety of painting styles could permit the user to quickly select

Figure 7. The left image has smooth Phong shading that would be difficult for a watercolor artist to create with a brush. By applying a gain and bias to the thickness of the unlit layer the user can make the shading much more abrupt as shown in the right image.

Figure 8. Before (left) and after (right) edge softening. Notice that the left side of the cup fades into the background and the dark line on the right side of the cup has been removed.

Figure 9. The left image uses cool colors and has a gloomy appearance. Through the use of a warmer palette of colors a much more cheerful atmosphere can be achieved as shown in the right image.

Figure 10. The Utah Teapot rendered using the standard Phong Shading model.

Figure 12. A result of minimizing the amount of mixing between the diffuse and unlit paint.

Figure 11. Utah Teapot rendered using our technique. The selection of paints for each layer permits temperature to vary based on the lighting in the scene.

Figure 13. The wash-like textures help indicate the shape of the surface. In particular notice the curved regions where strokes follow the shape of the surface.

rendering parameters appropriate for the presentation of a given object.

We have started preliminary work on how this technique can be used for scientific visualization applications. Figure 13 shows how our approach can create images that could be helpful to scientists and engineers in their work. Future work includes how our technique could be used for a wider variety of scientific datasets.

Other further work includes research into how images could be produced that would resemble more closely real watercolor paintings. Since real watercolor paintings are created entirely in the 2D image plane, it appears a number of watercolor effects would require manipulation in the image plane rather than object space. As such, creating realistic watercolors, while preserving frame-to-frame coherence for animation, offers a number of challenges.

Acknowledgements

This work has been sponsored by the National Science Foundation under contracts ACI 9983641 (PECASE Award) and ACI 9982251 (LSSDSV program). Thanks to Cyberware for the rocker arm dataset.

References

- [1] BARBARA, J. M. Painterly rendering for animation. In *SIGGRAPH '96 Conference Proceedings* (August 1996), pp. 477–484.
- [2] CABRAL, B., AND LEEDOM, L. Imaging vector fields using line integral convolution. In *SIGGRAPH '93 Conference Proceedings* (August 1993), pp. 263–270.
- [3] CURTIS, C. J., ANDERSON, S. E., SEIMS, J. E., FLEISCHER, K. W., AND SALESIN, D. H. Computer-generated watercolor. In *SIGGRAPH '97 Conference Proceedings* (August 1997), pp. 421–430.
- [4] EBERT, D., MUSGRAVE, F. K., PEACHEY, D., PERLIN, K., AND WORLEY, S. *Texturing and Modeling: A Procedural Approach*. Academic Press Limited, 1994, pp. 194–196.
- [5] GIRSHICK, A., INTERRANTE, V., HAKER, S., AND LEMOINE, T. Line direction matters: An argument for the use of principal directions in 3d line drawings. In *Proceedings of the First International Symposium on Non-Photorealistic Animation and Rendering (NPAR 2000)* (June 2000), pp. 43–52.
- [6] GOOCH, A., GOOCH, B., SHIRLEY, P., AND COHEN, E. A non-photorealistic lighting model for automatic technical illustration. In *SIGGRAPH '98 Conference Proceedings* (July 1998), pp. 447–452.
- [7] HAASE, C. S., AND MEYER, G. W. Modeling pigmented materials for realistic image synthesis. *ACM Transactions on Graphics* 11, 4 (October 1992), 305–335.
- [8] HERTZMANN, A., AND ZORIN, D. Illustrating smooth surfaces. In *SIGGRAPH 2000 Conference Proceedings* (August 2000), pp. 517–526.
- [9] INTERRANTE, V. Illustrating surface shape in volume data via principal direction-driven 3d line integral convolution. In *SIGGRAPH '97 Conference Proceedings* (August 1997), pp. 109–116.
- [10] MACKENZIE, G. *The Watercolorist's Essential Notebook*. North Light Books, 1999.
- [11] MARKOSIAN, L., KOWALSKI, M. A., TRYCHIN, S. J., BOURDEV, L. D., GOLDSTEIN, D., AND HUGHES, J. F. Real-time nonphotorealistic rendering. In *SIGGRAPH '97 Conference Proceedings* (August 1997), pp. 415–420.
- [12] NISHITA, T., TAKITA, S., NAKAMAE, E. A display algorithm of brush strokes using Bezier functions. In *CG International '93 Proceedings* (June 1993), pp. 244–257.
- [13] PERLIN, K. An image synthesizer. In *SIGGRAPH '85 Conference Proceedings* (July 1985), pp. 287–296.
- [14] STALLING, D., HEGE, H. C. Fast and Resolution independent line integral convolution. In *SIGGRAPH '95 Conference Proceedings* (August 1995), pp. 249–256.
- [15] WAY, D.L., HSU, C.W., CHIU, H.Y., SHIH, Z.C. Computer-generated Chinese painting for landscapes and portraits. In *WSCG 2001 Conference Proceedings* (February 2001).