Isoluminant Color Picking and its Applications

Ankush Seth

School of Computer Science

McGill University
Montreal, Quebec, Canada
May 3, 2005

A thesis submitted to McGill University in partial fulfillment of the requirements
for the degree of Master of Science

Your file Votre référence
ISBN: 0-494-12538-1
Our file Notre référence
ISBN: 0-494-12538-1

ACKNOWLEDGMENTS

I would like to thank my advisor Dr. Allison Klein for the guidance, and advice she provided throughout the project on Isoluminant Color Picking for NPR, and during the course of this thesis. This thesis was born from the above mentioned project of which Dr. Klein, Trân-Quân Luong, Jason Lawrence, and I were a part. A lot of credit goes to Dr. Klein for the successful completion of the project as well as my thesis. Also, I would like to thank Soukia Savage for helping us out with the photo shoot. Finally, I would like to dedicate this thesis to my parents, without whose support, and encouragement none of this would have been possible.

TABLE OF CONTENTS

TABLE OF FIGURES

ABSTRACT

Color and luminance play important roles within the field of visual arts. The reason for this lies in the complexity of the human visual perception system. Our brains experience a perceptual tension when processing isoluminant fields (i.e. fields of equal luminance value) because the luminance and color processing pathways perceive the two fields differently. Skilled artists exploit this fact to great effect. This thesis looks at how isoluminance can be applied to the field of Non-Photorealistic Rendering. Specifically, this thesis makes a novel contribution to NPR by emphasizing the importance of isoluminance. It shows how isoluminant color picking can be used to improve existing NPR image filters, and to create new ones. It presents a geometric technique for isoluminant color picking and then applies it in a pointillist filter, a new Chuck Close inspired filter, and a unique image mosaic filter.

RÉSUMÉ

La couleur et la luminance jouent des rôles importants dans le champ des arts visuels. La raison se situe dans la complexité du système visuel humain de perception. Nos cerveaux éprouvent une tension perceptuelle en traitant les champs isoluminant (c.-à-d. champs de valeur égale de luminance) parce que les voies de traitement de luminance et de couleur perçoivent les deux champs différemment. Les artistes habiles savent exploiter ce fait pour produire des effets impressionnants. Cette thèse explore comment l'isoluminance peut être appliqué à la production d'images Non-Photorealistes. Plus spécifiquement, elle démontre comment la sélection de couleurs isoluminantes peut améliorer les filtres d'image de NPR existants, et en créer des neufs. Elle présente une technique géométrique pour la sélection de couleurs isoluminantes et l'applique ensuite dans un filtre pointilliste, un filtre inspiré par Chuck Close, et un filtre de mosaïque.

Chapter 1 Introduction

Humans see the world in two different ways. Our brain first processes the color and luminance information provided by the visual stimuli separately, and then integrates the two, giving us the final perceived image [Gol99]. It is in this regard that isoluminant colors have special properties.

When an object is isoluminant with respect to its background, it can be perceived by the color processing pathway (the 'What' system) but it is invisible to the luminance processing path (the 'Where' system)[Liv02]. This creates a perceptual tension which makes the object appear as if it is vibrating or in motion. Skilled artists exploit this idea to great effect. For example, the sun in Monet's painting *Impression Sunrise* appears to pulse or shimmer because the sun is isoluminant with the surrounding clouds [Liv02]. Similarly, in the painting *Saint-Tropez* by Paul Signac (figure 1.1(a)), the foreground trees appear to be more prominent than the surrounding background. This is because the colors used to paint the trees are isoluminant. Figure 1.1(b) clearly shows that the average luminance of the set of pixels that represent the foreground trees is the same as the average luminance of the background trees (figure 1.1(b)).

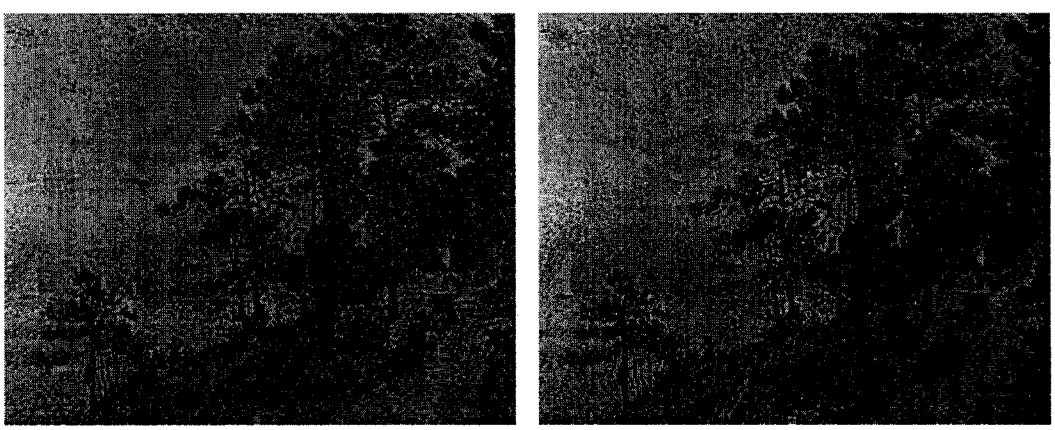

(a) Original version of Saint-Tropez (b) Luminance version

Figure 1.1: Saint-Tropez by Paul Signac. Notice that in (b) the trunks of the foreground trees are hard to distinguish from the trees in the background.

The majority of us find it hard to detect or choose isoluminant colors with differing hues (the attribute of color perception that enables one to distinguish the various parts of the color spectrum [Osb80]). Computers, on the other hand can choose isoluminant colors both quickly and easily. Surprisingly, not a lot of work has been done that incorporates isoluminance into NPR (Non-Photorealistic Rendering). Many painterly NPR filters use point or area sampling to pick colors. The chosen colors are then randomly perturbed to get the colors for the final image. These filters do not take any steps to maintain the input luminance values, and it is in this regard that this thesis makes its contribution.

This thesis presents a geometric technique for picking isoluminant colors and then looks at how such a color picking strategy can be used to both improve existing NPR filters and also to aid in the creation of new ones. The proposed technique is applied in a pointillist filter, a new Chuck Close inspired filter, and a novel type of image mosaic filter. The development of the Chuck Close and image mosaic filters also required edge analysis, texture generation, and image manipulation. Detailed information about these topics can be found in the paper *Isoluminant Color Picking for Non-Photorealistic Rendering*[1] and in the related thesis to be submitted by Trân-Quân Luong.

The rest of the thesis is organized as follows. Chapter 2 provides an in-depth discussion of relevant previous work. Chapter 3 describes the isoluminant color picking technique, and then elucidates how the technique can be applied in a pointillist filter, and in two new filters. It also looks at future work that can be done in order to improve the individual filters. Chapter 4 concludes the thesis by providing a brief summary, a short discussion of some other issues, and a section on higher-level future work.

---

[1] Accepted for publication in the 2005 *Graphics Interface* proceeding.

Chapter 2 Related Work

This chapter provides an in-depth review of related work. The color picking techniques and filters presented in this thesis mainly relate to the field of NPR, but because the filters also exploit the eye's natural ability to integrate colors within small spatial regions it is also important to look at offset color printing topics such as multi-color half-toning and artistic screening.

2.1 Non-Photorealistic Rendering

The field of Non-photorealistic rendering basically looks at ways of imitating artistic and other traditional styles of image production. In their book *Non-Photorealistic Rendering*, Gooch and Gooch write "Non-photorealistic rendering brings art and science together, concentrating less on the process and more on communicating the content of an image....Knowledge and techniques that have long been used by artists can be applied to computer graphics to emphasize specific features of a scene, expose subtle attributes, and omit extraneous information..." [GG01].

The NPR algorithm presented by Gooch, Coombe, and Shirley in [GCS02] generates painterly images using brush strokes. Their work mainly looks at how image segmentation based on the features contained within the image can be used to come up with a set of brush stroke positions. The color selection process for the strokes involves sampling the original image and then sometimes incorporating small random perturbations to come up with the new colors for the sampled region. (See figure 2.1(a)) Similarly, the algorithm by Hertzmann [Her98] can also be used to generate images that have a hand painted look to them. It allows the user to emulate various painterly styles like impressionist, pointillist and more. The output images in this case are rendered using spline brush strokes of varying sizes and styles. The color selection process is very simple and the system just tries to match the colors in the source image. Figure 2.1(b) shows a result image from [Her98]. The basic idea behind these and other image-based painterly filters is to convey the essence of an image using limited spatial and chromatic resolution [GCS02]. However the color selection schemes that these algorithms employ do not pay any special attention to isoluminance, and it is in this regard that the work presented in thesis is unique.

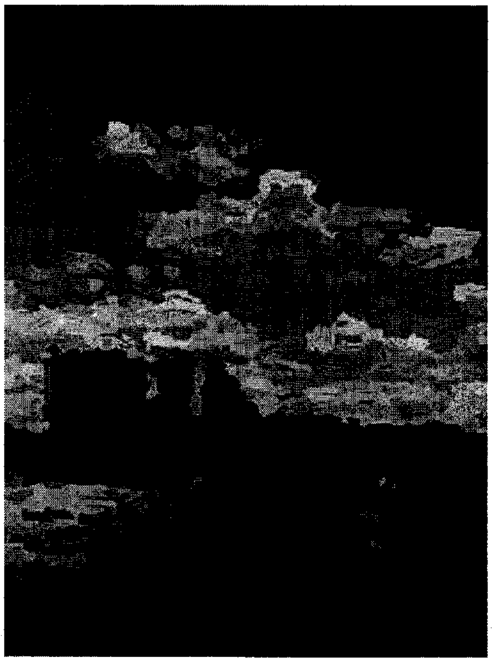

(a) A landscape painting of Hovenweep National Monument [GCS02].

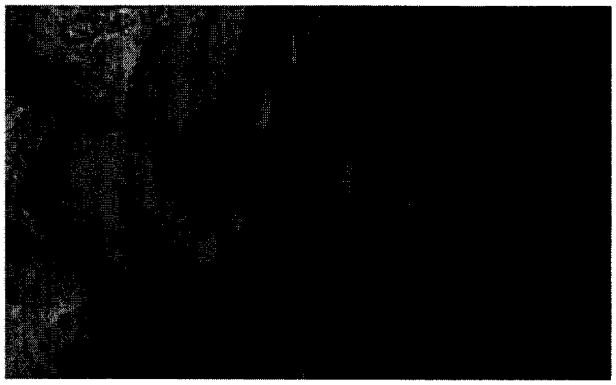

(b) Result image from [Her98].

Figure 2.1: Results from NPR painterly techniques that use point or area sampling during the color selection process.

A similar strategy is also followed by quite a few algorithms for automatic NPR renderings of 3D models. The color selection process that these algorithms follow is based on underlying material properties and some sort of stylized transformations of traditional shading models. For example, the algorithm presented in the paper *Cartoon-looking Rendering of 3D Scenes* by Philippe Decaudin [Dec96] uses a modified version of the traditional color formula to calculate the colors within a 3D scene, given the light sources, material properties, and shading information from the Phong shading model.

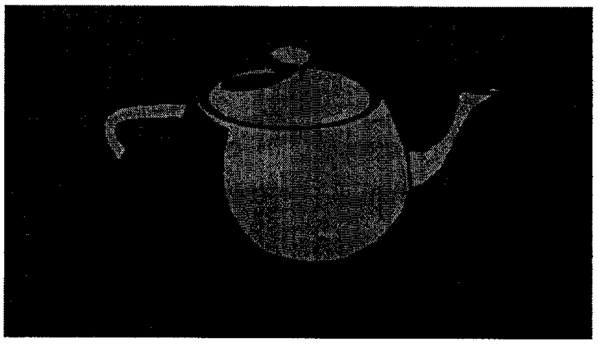

Figure 2.2: Result image from [Dec96].

The work presented in [GGSC98] and [GSG+99] uses the concepts of tone (the lightness or gradient of the color [Osb80]) and temperature (artist's understanding of whether a color is warm or cool [Hp]) in order to come up with shading models for technical illustrations. Even though the systems presented in [GGSC98] and [GSG+99] use changes in luminance to specify surface orientation, the visual prominence of edge lines and highlights are still the main focus of their algorithms, and therefore the inner coloring/shading is restricted to a fixed luminance range.

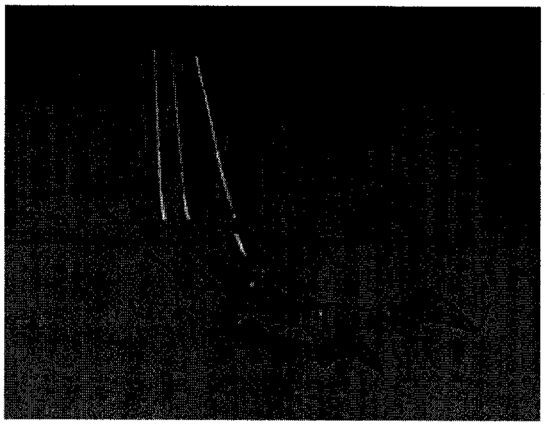

Figure 2.3: Approximating constant luminance tone rendering [GGSC98].

The fact that color luminance helps in portraying the shape or form of objects is a clear indication that an algorithm that not only takes tone and temperature into account but also looks at color luminance would benefit NPR applications. Figure 2.3 is an example image from [GGSC98] that shows how approximating constant luminance can help enhance subtle details within an image.

Finally, the work presented in the *The lit sphere: A model for capturing NPR shading from Art* by Sloan et al [SMGG01] looks at techniques for producing shading models based on a particular artist or artistic genre. An example image taken from the paper [SMGG01] is shown in Figure 2.4. The prime objective of the system is to enable non-artists to develop shading models that give the notion of light, depth, and material properties, something that artists are very good at portraying [SMGG01]. However, the technique presented in the paper does not pay attention to the benefits of constant luminance within its color selection process.

(a) Still life by Cezanne

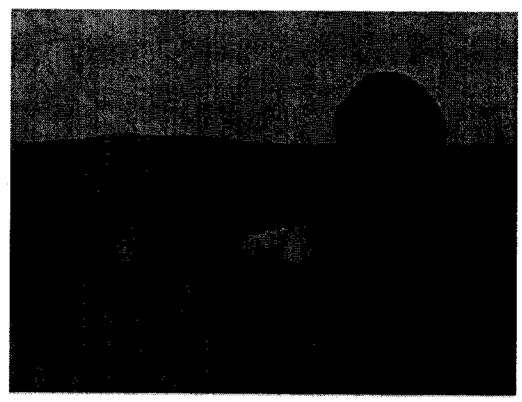

(b) Model illuminated using shading model captured from (a)

Figure 2.4: Result image from [SMGG01] produced by capturing shading model from artwork.

2.2 Half-toning and Artistic Screening

Like the isoluminant color picking technique presented in this thesis, both multi-color half-toning (the process of converting a continuous tone image into an image that can be printed with two or more color inks [Half]) and artistic screening (a technique that uses freely created artistic screen elements for generating halftones [OH95]) also depend upon the eye's natural ability to spatially integrate a multi-colored region, and therefore it is important to look at these offset color printing topics.

The work of Stollnitz, Ostromoukhov, and Salesin [SOS98] looks at issues within multi-color half-toning and how new algorithms can be used in conjunction with custom inks to accurately represent colors. They come up with new models and algorithms that help in addressing issues such as predicting the interaction between arbitrary inks when printed together, selecting the best set of inks to represent the colors within an image, and last but not the least, computing ink separations for three or more inks [SOS98]. The problem of finding the best set of custom inks to represent various colors is analogous to what the isoluminant color picking technique does. It resolves a given target color into several colors that satisfy the isoluminance and barycentric constraints, and depend upon the eye's ability to spatially integrate them.

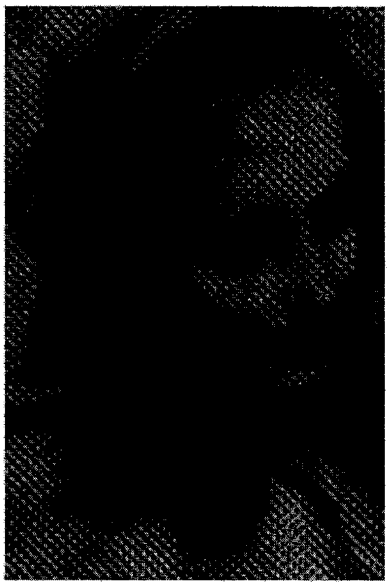

Figure 2.5: Color image reproduced using the Delaunay tetrahedrization [CH01].

Similarly, work done by Chosson, and Hersch [CH01] is also relevant. Their work looks at various ways of selecting the tetrahedrizations of colors that result in the least amount of visual artifacts within visible halftone patterns. An example from [CH01] can be seen in Figure 2.5. What makes the work presented in this thesis unique is the fact that in our case we want visible artifacts to be noticeable when the image is viewed from a close distance. This is contrary to the goals within offset color printing.

The "Chuck Close" and the image mosaic filters presented in this thesis use texture/image maps to control the distribution of color within the segmented regions and are therefore related to the work done by Ostromoukhov, and Hersch [OH95]. In their paper, Ostromoukhov, and Hersch come up with the unique half-toning technique of Artistic Screening. It allows the use of artistic screen elements for half-toning purposes so that image reproduction can be done in a way that conveys more artistic information. Similarly, the paper by Buchanan [Buc96] looks at ways of incorporating deliberate artifacts within images reproduced using half-toning techniques. Figure 2.6, a result from [Buc96] shows how artifacts can be used to simulate certain effects. In this case artifacts were introduced to simulate printing on coarse paper.

Figure 2.6: Result from [Buc96] that simulates printing on coarse paper.

Although the work presented in this thesis combines traditional half-toning techniques with more artistic goals, what makes it different is the fact that the color selection process takes place independently and locally for each region, whereas in the case of multi-color half-toning and artistic screening, the process takes place across the entire image. Moreover, with previous work the colors selected are both globally and locally close to the input image colors.

Finally, algorithms for generating image mosaics such as the one presented in the paper *Image Mosaics* by Finkelstein and Range, use a color selection process that is based on color proximity. The goal is to manipulate the "tile" image color histograms so that they resemble the color histograms of the corresponding region in the source image. (See figure 2.7) In contrast, the mosaic filter presented in chapter 3 transforms textures with colors that are chosen to be far from the target color, but get resolved to the desired colors when the image is viewed from a distance.

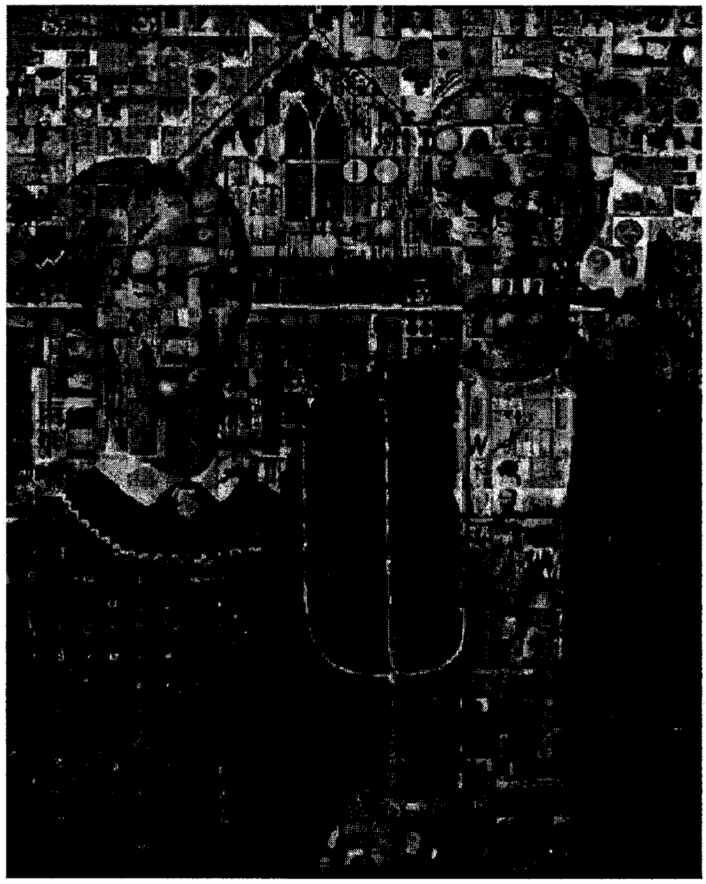

Figure 2.7: *American Gothic* composed of pictures from the Web [FR98].

Chapter 3 Isoluminant Color Picking

Chapter 3 first begins with a brief introduction to pointillism and then describes the target driven isoluminant color picking technique. It then looks at how isoluminant color picking can be applied in a pointillist filter, a 'Chuck Close' inspired filter, and a unique mosaic style filter.

The filters presented in this chapter are implemented in C++. The OpenGL library for C++ is used for drawing GL_POINTS in the case of the pointillist filter, and for drawing grid cells, and texture mapping in the case of the Chuck Close and image mosaic filters. Functions from the OpenCV (Open source Computer Vision) library are used during the edge detection and analysis phase of the Chuck Close and image mosaic filters.

3.1 Pointillism and Isoluminant Color Picking

Additive color theory is the fusion of colors obtained by combining two or more colors in the form of a mosaic too fine for the eye to resolve or by directing colored lights into the eye either simultaneously or in rapid succession [Osb80], and was one of the main reasons for the pointillist movement. Neo-Impressionists deemed that 'optical mixing' of colors obtained by using uniform touches of pure hues (called *points* in French, hence "pointillism") [CH99] produced brighter and more lively paintings compared to work based on subtractive color theory (color blending due to absorption or subtraction. It is usually accomplished by intermixing differently colored paints [Osb80]). Moreover, the colors selected by the pointillist artists not only averaged to the desired target color but also had a luminance value that was close to the luminance of the target color. At the time, "the novelty of these oddly luminescent, 'mechanical' paintings comprised of tiny dots of pure colors not only startled audiences, critics, but also the older Impressionists." [CH99]

Most pointillist filters that emulate the pointillist style pick colors by point sampling the input image and then perturbing the colors randomly without being concerned about the target luminance. The isoluminant color picking approach described below maintains isoluminance by generating clusters of points that average to the target color luminance, thereby producing results that are closer to the pointillist style.

Given a target color C, we want to find isoluminant colors that average to the target color but are still visually distinct. Mathematically, given a target color C in the RGB space, and a set of weights $\{w_1, w_2 ... w_n\}$ that form a barycentric combination (A weighted sum of points where the coefficients of the points sum up to one [FH05]), we seek colors $\{C_1, C_2 ... C_n\}$ such that:

$$C = \sum_{i=1}^{n} w_i C_i \quad (1)$$

and

$$Lum(C_i) = Lum(C) \quad (2)$$

Geometrically, we want to find an n-gon on the same luminance plane as the color C such that the barycentric co-ordinates of C with respect to this n-gon match the respective weights $\{w_i\}$. Note that a polygon satisfying the conditions above can be uniformly scaled or rotated within the luminance plane around C. The new polygon created as a result of the scaling or rotation will also satisfy the same constraints provided the polygon stays within the RGB cube. The further away the vertices of the polygon are from the center (target color) the more distinct the colors appear at a close range while still resolving to the target color when viewed from a distance. This is because we can use Euclidean distance in the RGB space as a reasonable approximation to perceptual distance. The fact that the luminance is kept constant is probably one of the reasons why the above approximation holds.

3.2 An Improved Pointillist Filter

This section looks at how the above isoluminant color picking technique leads to an improved pointillist filter. Rather than point sampling and perturbing the colors randomly, the technique from section 3.1 is applied during the color selection phase.

The filter starts off by dividing the input image *S* into a regular grid of small cells. For each grid cell the filter calculates its average color *C* and its luminance, *Lum*(*C*). It then picks *n* colors with corresponding weights such that they resolve to the grid cell's average (target) color. Since most pointillist artists use three or four colors to represent a target color, a value of four was used for *n*. This means that each color has a weight of 0.25. In the RGB cube the four colors correspond to the vertices of an isoluminant square centered at the target color *C*. The first vertex of the square is picked by selecting a point that is isoluminant with *C* and at a distance *d* from *C*. Given the first point, the remaining three can be calculated directly because a fixed size square is being constructed within the RGB cube. Once the square is constructed it is uniformly scaled until one or more vertices reach the boundary of the RGB cube. (See figures 3.1(a)-(c)) Figure 3.1(a) is a diagrammatic representation of the target color and its luminance plane inside the RGB cube. Figure 3.1(b) shows a scaled up depiction of the initial square that is constructed around the target color inside the RGB cube on the same luminance plane as the target color. The colored circles at the corners of the square in figure 3.1(b) not only average to the target color at the center but also have its luminance value. The outward pointing arrows represent the fact that the initial square is scaled in size. The larger the square the more distinct the final colors will be. Figure 3.1(c) is the luminance version of the same square. Note that the luminance value of the new colors is maintained during the expansion as well.

In most cases, the initial square will get scaled, thereby providing more color variation. However, for target colors near the sides of the RGB cube and depending upon the chosen value of *d*, it maybe the case that there is not enough room for the initial square to expand. In the worst case scenario the filter may not be able to even construct the initial square. In this case the filter uses the target color itself to color the corresponding region.

Regarding the selection of the first vertex of the isoluminant square, the filter can either make a random choice or it can choose the vertex to be at a fixed angle relative to *C*. The latter scheme produced results that we feel are visually more pleasing. This is because restricting the orientation of the square gives the sense of a fixed palette, and is therefore more coherent with what artists usually do.

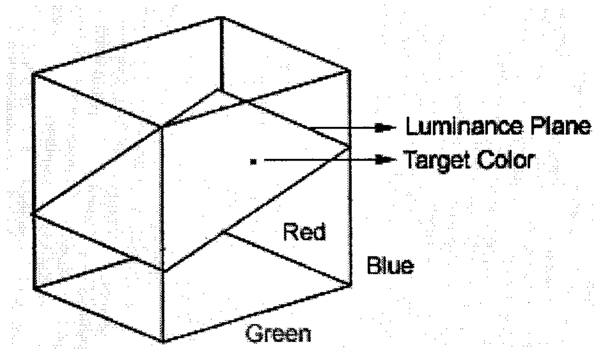

(a) Target color and its luminance plane inside the RGB cube

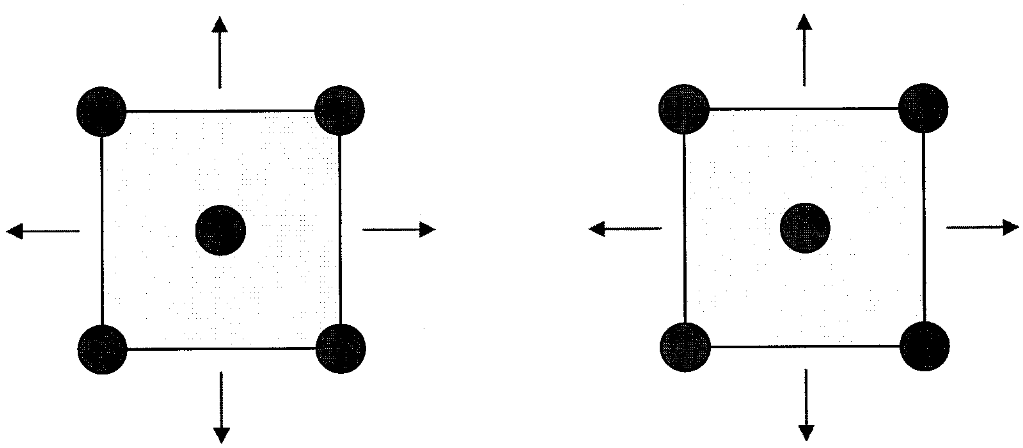

(b) Color version (c) Luminance version

Figure 3.1: Diagrammatic representation of the target color inside the RGB cube and the initial square constructed around it. Note that although all the vertices are distinctly colored in (b) they have the same luminance value as the target color at the center (c).

3.2.1 Results

The results shown in figures 3.2 and 3.5 used four-by-four grid cells to calculate the average color. The initial distance d was set to $\sqrt{2}$. The radius of the colored points was set to 5 pixels. Figures 3.3 and 3.6 were also produced using the same values but the orientation of the initial square was picked at random rather than at a fixed angle relative to *C*. We feel the fixed angle approach produces results that are aesthetically more pleasing.

To ensure that there is neither significant overlap nor gaps, the pointillist filter also incorporates a minimum distance variable that determines the amount of overlap or gap between the colored points. For figures 3.2, 3.3, 3.5, and 3.6 the minimum distance variable was set to a value that resulted in full coverage i.e. the images generated contained no gaps between colored points. We feel that the results with full coverage look much more appealing when compared to the images with partial coverage. Figures 3.4 and 3.7 are examples of images with partial coverage. Note that in the case of images with partial coverage the introduction of gaps between colored points alters the final luminance value. This variation in luminance value is currently not taken into account by the filter. The source images for these results can be found in appendix A1.

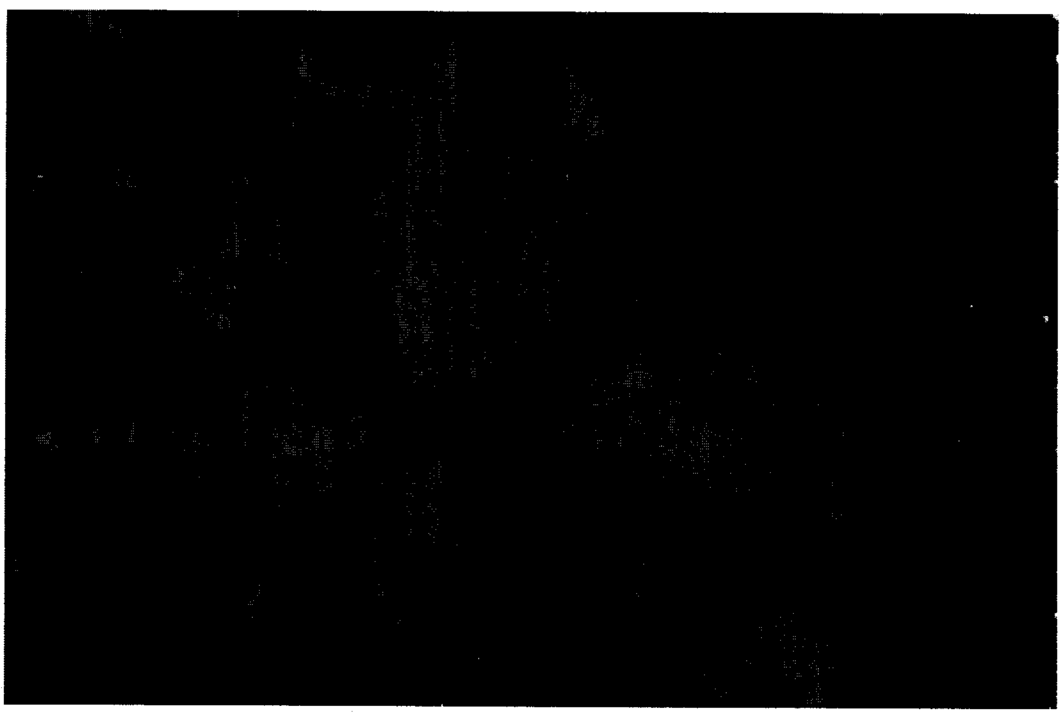

Figure 3.2: Pointillist image of a lily generated with fixed orientation, and full coverage.

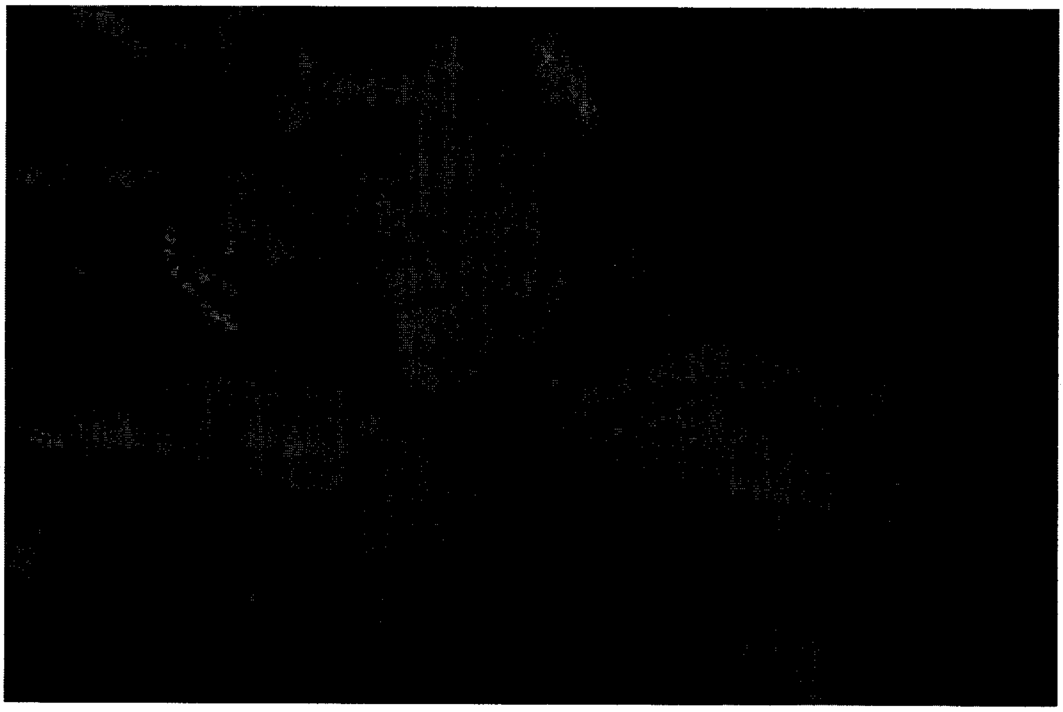

Figure 3.3: Pointillist image of a lily generated with random orientation, and full coverage. Note how randomly orienting the square results in bright green and red colored points within various sections of the image, thus making it less appealing.

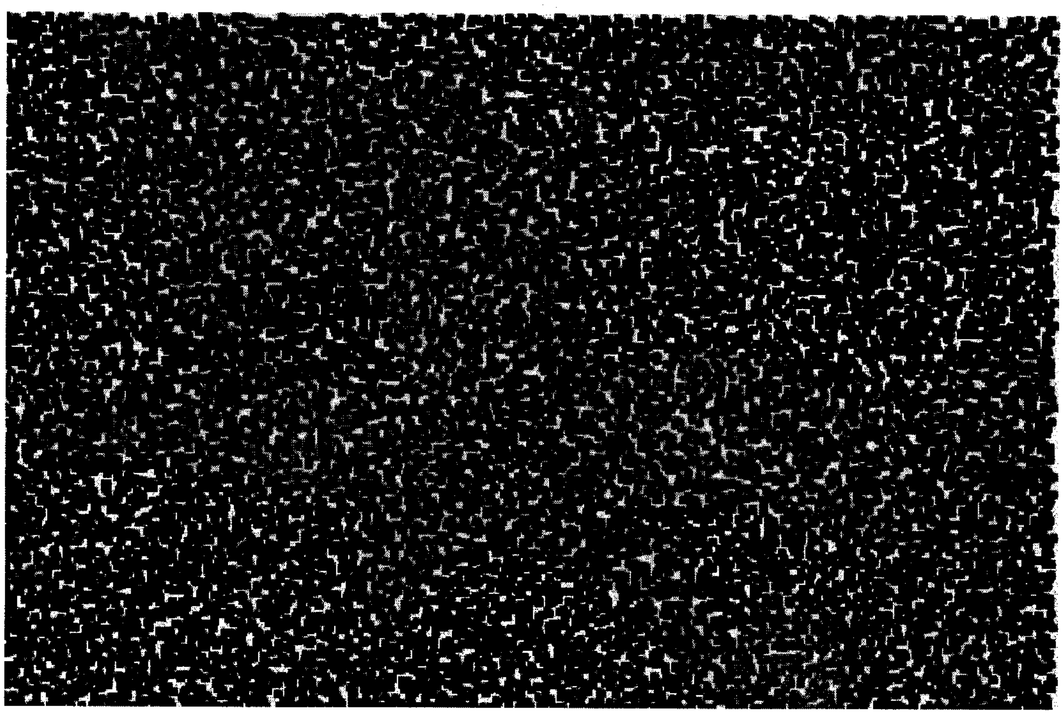

Figure 3.4: Pointillist image of a lily generated with fixed orientation, and partial coverage.

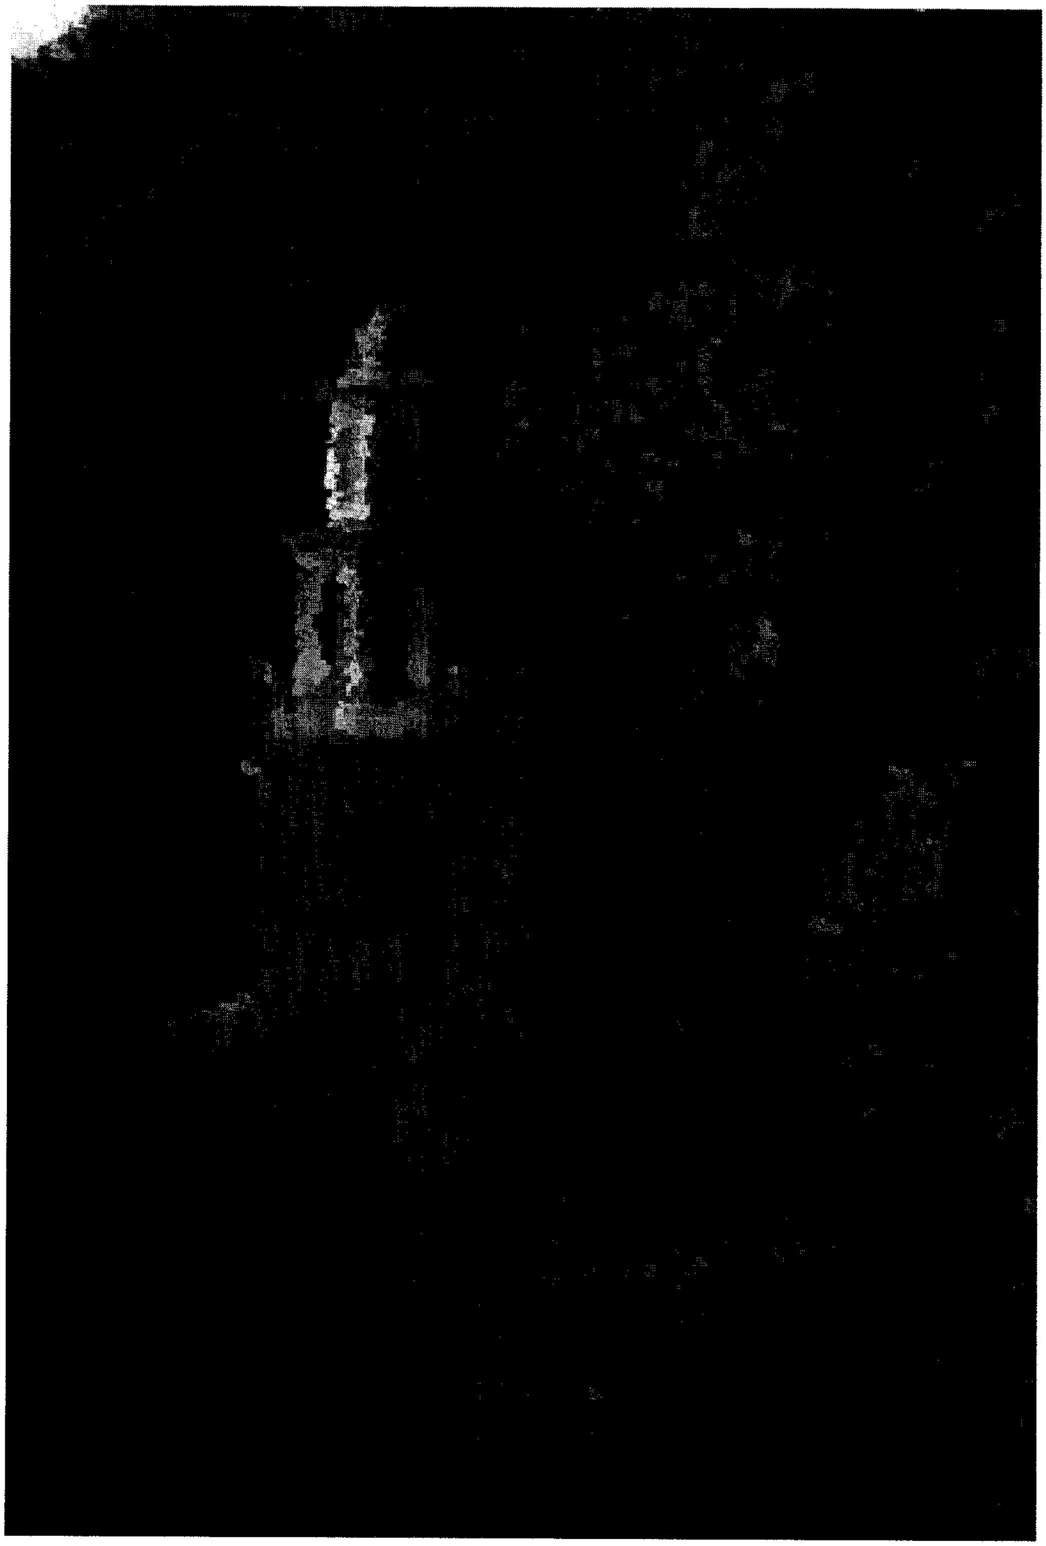

Figure 3.5: Pointillist image generated with fixed orientation, and full coverage.

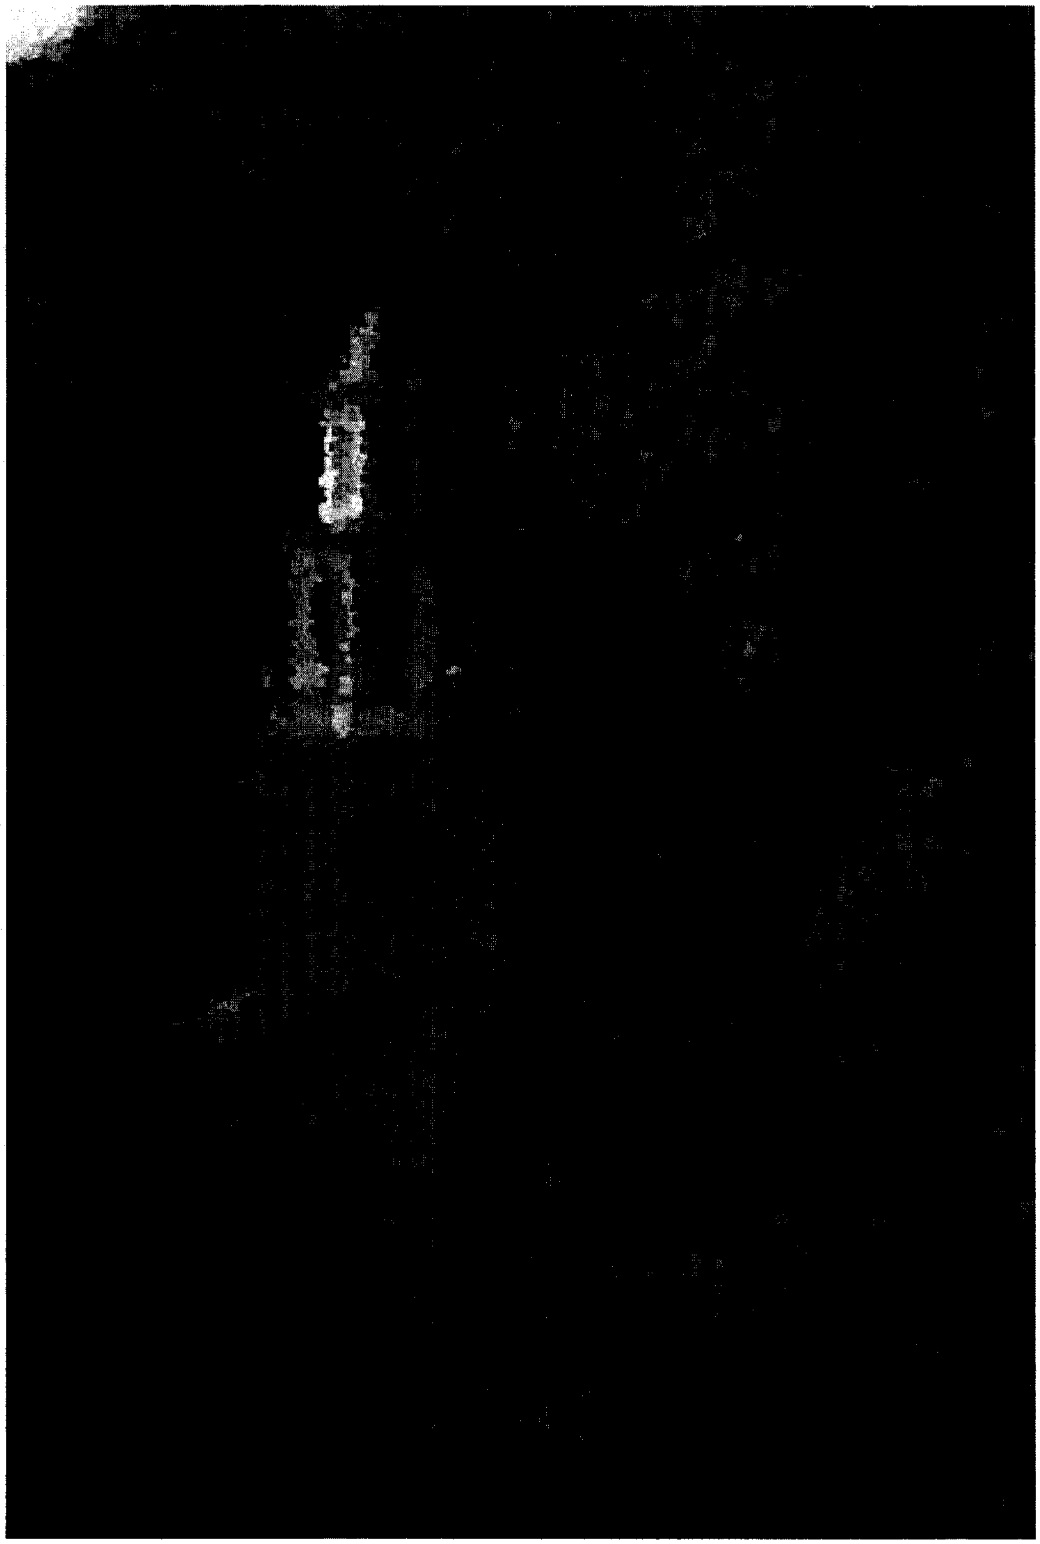

Figure 3.6: Pointillist image generated with random orientation, and full coverage. Note how randomly orienting the square results in bright red and pinkish colored points within various sections of the image, thus making it less appealing.

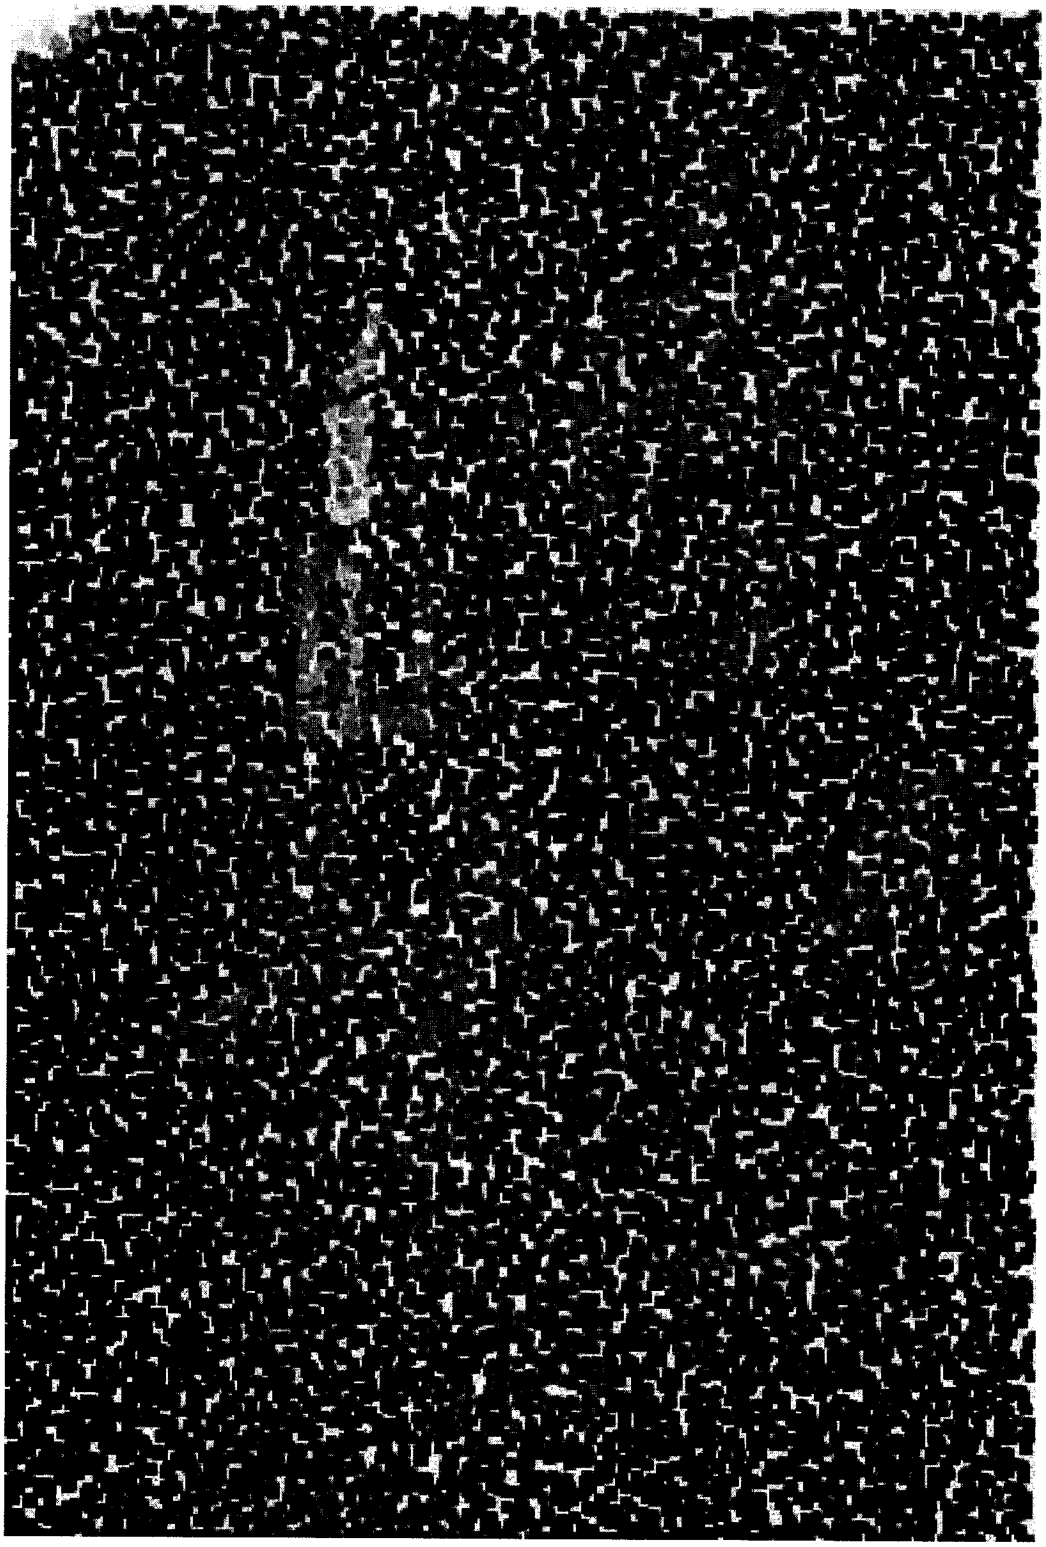

Figure 3.7: Pointillist image generated with fixed orientation, and partial coverage.

3.2.2 Observations

The level of detail visible within the output images is dependent on the grid size used during the average color calculation. The smaller the grid cells the more refined the resulting output. This is because as the grid size is reduced the probability of regions containing colors with varying hues decreases (i.e. the standard deviation between the color of a given pixel and the average color of the grid cell containing that pixel decreases), and therefore the average color becomes more representative of the input color.

In figure 3.8(a) we can observe that without luminance control the output image has a "grainy" quality to it. The image was generated using the perturbed sampling approach based on [Hae90]. On the other hand the result (figure 3.8(b)) based on the pointillist filter described in section 3.2 is able to achieve strong color variation as well as regional luminance control, thereby producing results that are closer to the pointillist approach.

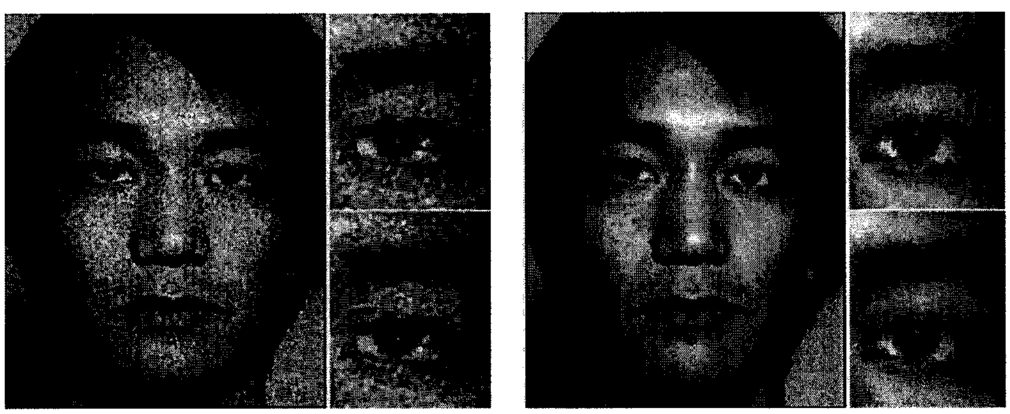

(a) Perturbed Sampling (b) Isoluminant Color Picking

Figure 3.8: Pointillist filter comparison. The inset (bottom right) in (a) clearly shows that isoluminance is not maintained, whereas in (b) the isoluminance is maintained even though the individual points are distinctly colored (top right).

As mentioned before, in areas of the input image where the target color is closer to the edges of the RGB cube, the isoluminant color picking technique produces points that have very little color variation. Although, this issue can be resolved by relaxing the luminance constraint, it is important to mention that pointillist artists sometimes prefer lower color variation in highly saturated regions. So it might be the case that the current approach is more consistent with the pointillist style.

The isoluminant color picking technique in itself is very flexible and can use greater or fewer colors so long as the assigned weights form a barycentric combination. Moreover, it can be extended to any linear color space i.e. color spaces where a linear combination of color values generates a target color, for example the YIQ and LEF color spaces.

3.2.3 Pointillist Filter Future Work

As mentioned earlier the initial starting point for the isoluminant square is either picked at a fixed angle relative to the target color C or randomly. This situation can probably be handled in a more sophisticated manner and is something that should be part of future work in this area. One possible solution would be to use the average tone of each grid cell as a parameter in deciding what orientation should be picked for the initial square for that particular grid cell. This could be regarded as an attempt at maintaining 'harmony' in terms of color within the local regions.

Another possible improvement would be to add the ability to control the d parameter directly. This feature could be useful for specifying the visual divergence of the isoluminant color sets at near viewing distances.

3.3 Chuck Close Filter

This section looks at an image filter that emulates the mosaic-like painted portraits of Chuck Close. The main focus is to show how the isoluminant color picking technique can be used to create a new filter.

Born in 1940, Chuck Close is more than a painter, printmaker and photographer. He is an artist that builds "painting experiences for the viewer." [CC] Considered as one of the most influential artists of the modern era, Chuck Close has produced artworks that have a striking visual effect.

Chuck Close creates his portraits by first tiling an enlarged picture of the subject and the destination canvas into diamond shaped grid cells. Chuck Close refers to them as *marks*. He then reproduces each *mark* in the source picture as a collection of concentric colored blobs, such that from a distance, the canvas appears to be a seamless portrait, but as we come closer, Close's use of color, isoluminance and texture clutter the visual field [Liv02], invoking "a competition between the face and its constituent blocks to engage our perception of shape from shading."[Pel99] (See figure 3.9)

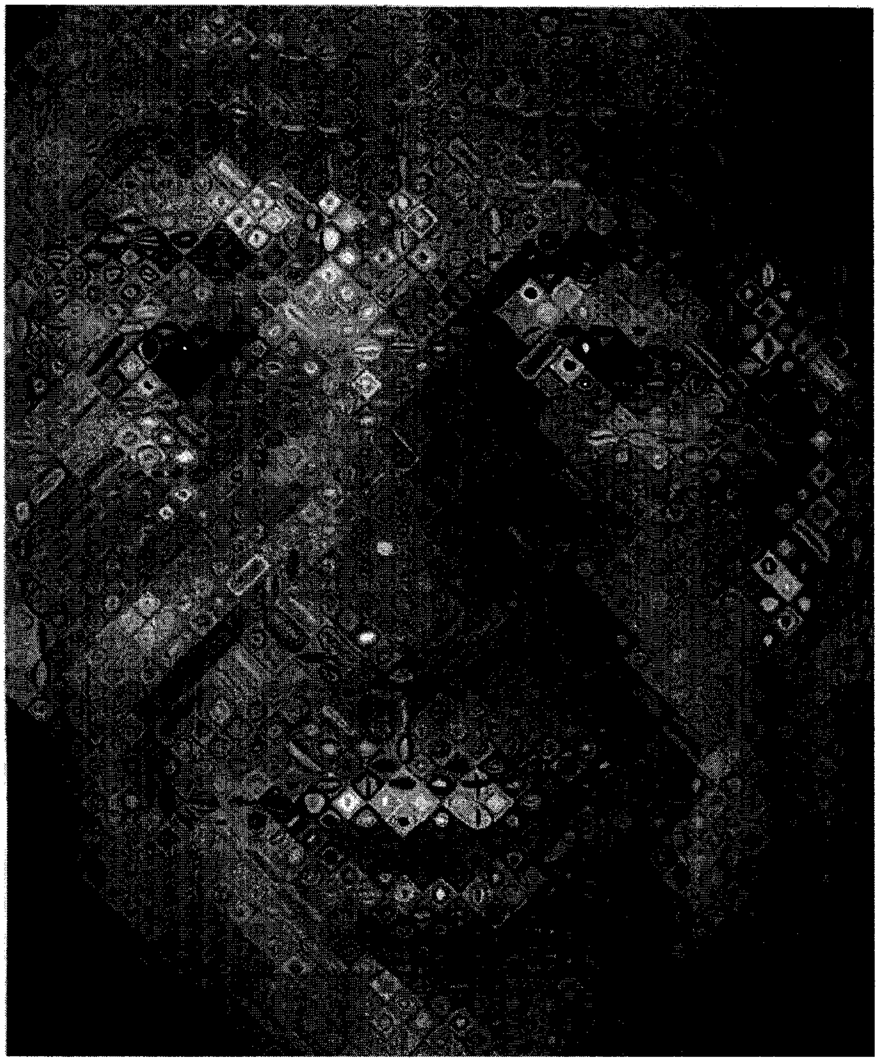

Figure 3.9: Robert 1997, by Chuck Close.

In order to make the output of the filter representative of his style of painting, one must keep in mind some of the apparent conventions that Chuck Close follows while creating his portraits. Given below is a list that enumerates these conventions:

1. Each *mark* contains one or two nested quads and several concentric circles or ellipses that don't have a perfect shape. The decision about which shape should be used inside the *mark* is based on the edge information contained within the source grid cell. In general the *marks* fall into one of the following four categories:
   a. Circular
   b. Vertical
   c. Horizontal
   d. Diagonal
2. Chuck Close often colors the first quad of each *mark* with a random color, and then adds successive layers so as to move closer to the target color [Sto98]. Quite often these layers alternate in luminance values. However, each *mark* in general contains fewer luminance levels as compared to the number of colors used to paint the mark. This can be observed in the figure below. Figure 3.11(a) is a close up of a section from figure 3.9 and figure 3.11(b) is the corresponding grey-scale version.

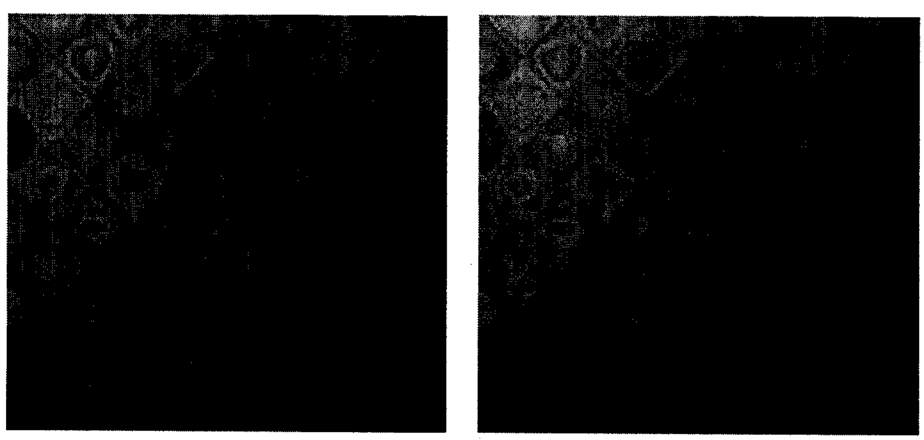

(a) Color version (b) Grey-scale version

Figure 3.10: Number of Colors vs. number of Luminance levels.

3. Adjoining *marks* that are comprised of similarly orientated diagonal ellipses, and have similar colors maybe merged together to represent the underlying diagonal edge. Two to three *marks* may be merged at a time to form the larger diagonal (figure 3.11) or the rare V-shaped *mark*. However, it is important to note that *marks* containing horizontal or vertical ellipses are never connected. Also, once the marks are merged they are regarded as single marks during the color selection process.

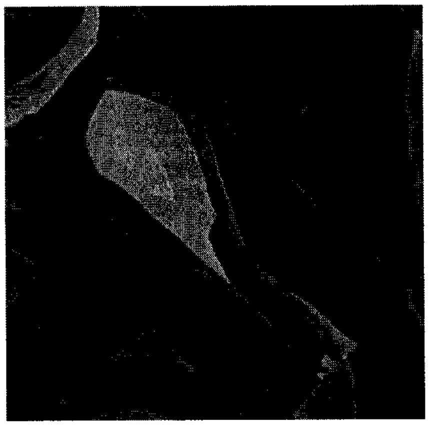

Figure 3.11: Example image of a merged grid cell.

4. Large regions that have little or no edge information often contain diagonal *marks* that have been merged together so as to disrupt the uniformity within the visual field.

3.3.1 Shape Selection

The shape selection process is based on the underlying edge information in the source image. The edge information is extracted from the input image using Canny edge detection (see figure 3.12). This information is further processed so as to decide which shape needs to be assigned to the marks in the output image. More information about how the shape selection process works can be found in the related thesis to be submitted by Trân-Quân Luong.

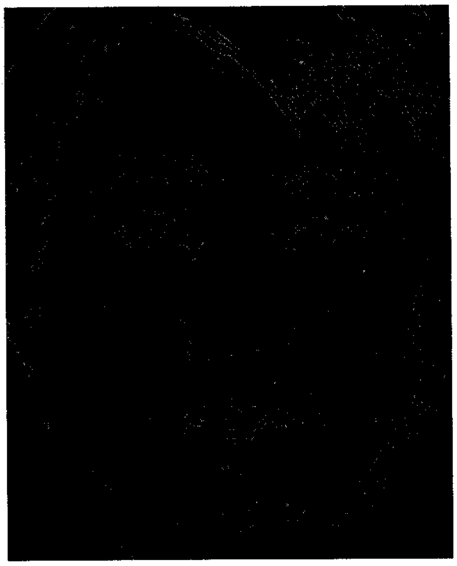

Figure 3.12: Edge information of the input image

3.3.2 Color Selection

This section provides an in-depth look at the color picking process followed during the color selection phase of the filter. Because there are five layers to each mark, colors are picked so that layers one and three, and layers two and four are isoluminant. Moreover, because the layers vary in size we calculate the relative areas of each layer in the following manner:

1. Each layer L within a given mark is rendered in white, and all other layers are rendered in black.
2. At this point, we sum up the intensities of all pixels within the mark. This sum divided by the total area of the mark in pixels yields the percent covered by layer L. This value corresponds to L's weight in the color picking approach.

Figures 3.13(a)-(e) are a series of snapshots that show the above procedure in action. Figure 3.13(a) shows the outermost layer of every mark rendered in white. Similarly, figure 3.13(b) shows the next layer, and so forth up to figure 3.13(e) that shows the innermost layer of every mark rendered in white.

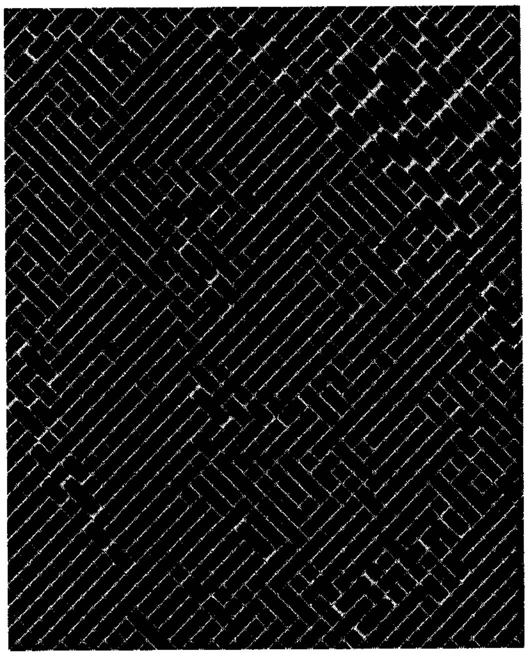

(a) Layer 1 (Outermost Quadilateral)

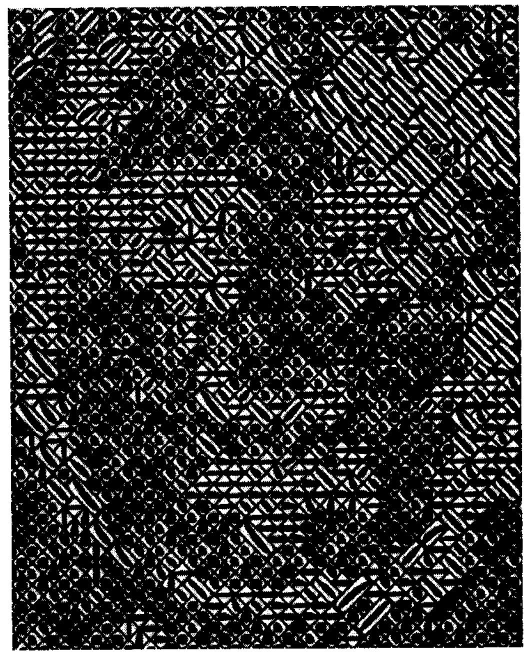

(b) Layer 2 (2nd level Quadilateral)

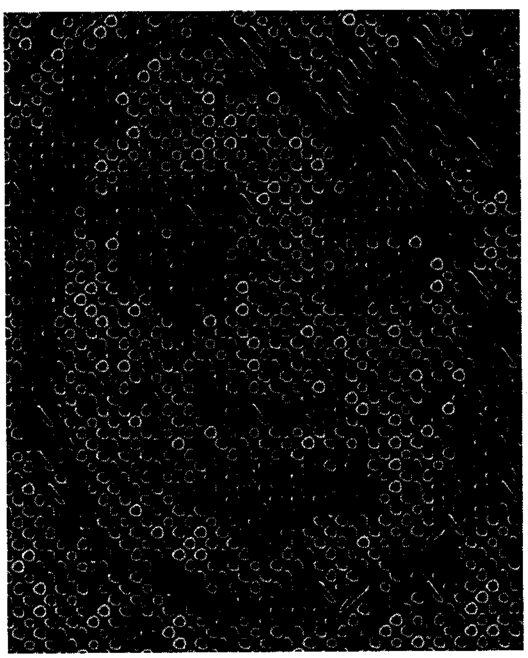

(c) Layer 3 (1st level of nested blobs)

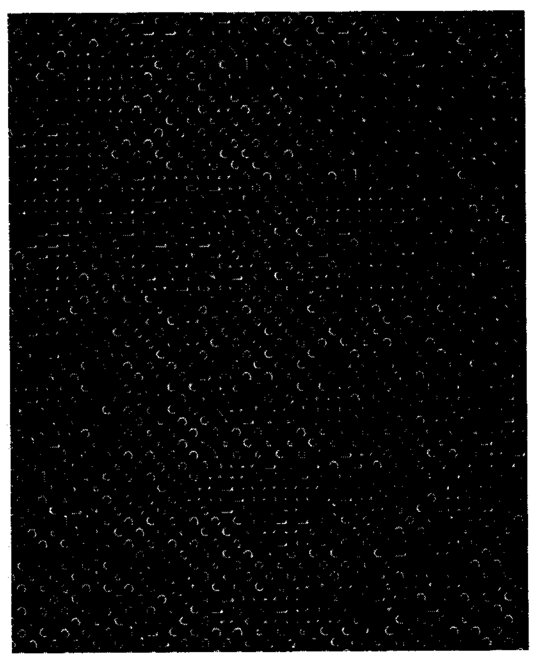

(d) Layer 4 (2nd level of nested blobs)

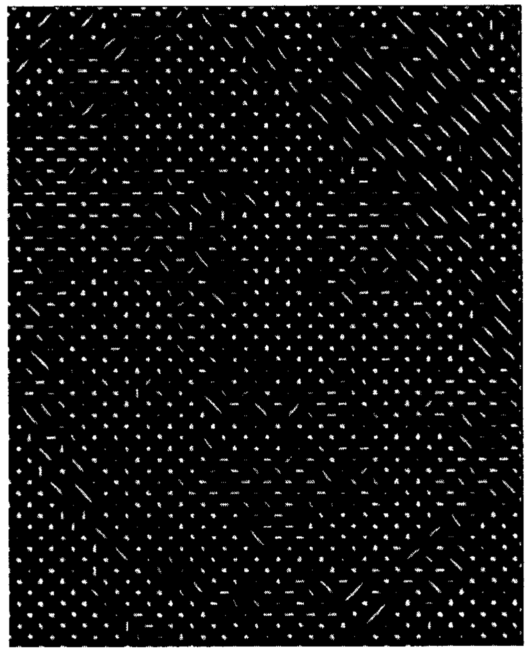

(e) Layer 5 (3rd level of nested blobs)

Figure 3.13: Snapshots from the relative area computation.

Given the various percentages, the colors for the various layers can be calculated in the following manner:

1. Given a target color C, and a set of five weights $\{w_i\}$, we find two colors C_1 and C_2 such that they satisfy the following condition:

   $$C = w_0L_0 + (w_1 + w_3)C_1 + (w_2 + w_4)C_2,$$

   where L_0 is an arbitrarily chosen color for the outermost layer of the *mark*.

2. In this step the color picking algorithm is applied in a manner that allows isoluminant pairs of colors for unequal weights to be generated. So for the two colors C_1 and C_2 that were generated in step one we produce a set of layer colors $\{L_i\}$ that satisfy the following conditions:

   $$C_1 = (w_1L_1 + w_3L_3) / (w_1 + w_3)$$
   $$C_2 = (w_2L_2 + w_4L_4) / (w_2 + w_4)$$
   $$Lum(L_1) = Lum(L_3)$$
   $$Lum(L_2) = Lum(L_4),$$

   where the distance between the pairs of isoluminant colors is maximal.

In terms of implementing the above steps, we first start off by randomly choosing L_0 such that the result of $(C - w_0L_0)(1 - w_0)$ is a valid color. The process is repeated until such a color is found. Similarly, a new color C_1 is chosen such that $C - w_0L_0 - (w_1 + w_3)C_1$ results in a valid color. For each C_1, we then iterate over all possible L_1s. Because each L_1 uniquely determines L_3, we need to find a pair that is as far apart as possible from each other while being contained within the RGB cube. Once this is done we proceed in a similar fashion to find layer colors for C_2. Note that C_2 is uniquely determined by each C_1. Moreover, we need to iterate over all the plausible C_1s until we find a valid set of colors.

3.3.3 Overall Algorithm

This section consolidates the various processes discussed in the previous sections. Given below is the overall algorithm for the "Chuck Close" filter.

1. Segment the image into diamond shaped grids.
2. For each *mark*, calculate the average color, and figure out the underlying prominent edge orientation by averaging the edge values determined by a per-pixel edge filter.
3. If adjoining *marks* contain an edge with the same diagonal edge orientation, and also have a similar average color, then these *marks* are merged together.
4. At this point, each *mark* is assigned a set of nested blob textures of the relevant shape. The relative weight of each layer within the *mark* being considered is computed. Each layer is then assigned a color using the color selection process mentioned in section 3.3.2. Finally, the *mark* is rendered with the appropriate colors.

3.3.4 Results

Figures 3.14 and 3.15 are some of the results produced using this filter. The source images for figures 3.14 and 3.15 can be found in appendix A1. The results presented contain 24 by 29 *marks* (grid cells). These are the default dimensions used by Chuck Close himself. Since the number of *marks* produced is a parameter in the filter, it can therefore be set to any user defined value.

Figure 3.14: 'Chuck Close' style image of Quân. A close up of some of the marks can be seen in the inset (bottom left).

Figure 3.15: 'Chuck Close' style image of Ankush. A close up of some of the marks can be seen in the inset (bottom left).

3.3.5 Observations

The "Chuck Close" filter takes about two seconds to produce the final output on a Pentium 4 class machine, and therefore for all practical purposes the filter is relatively fast. The color picking technique only maximizes the distance between individual isoluminant pairs and not over all the color pair distances. This is because the latter idea is computationally more expensive. Moreover, the results produced by it were visually similar to the results produced by the initial technique.

Currently, the "Chuck Close" filter does not modify the shapes of the marks around the eyes or the nostrils so as to make them more representative of the facial feature. This is something that Chuck Close pays attention to while creating portraits.

3.3.6 Chuck Close Filter Future work

The color selection approach used in the "Chuck Close" filter does not restrict the choice of colors to any particular palette and therefore we can sometimes notice colors that do not seem to be used by Chuck Close. Finding an optimal solution to this problem is something that can be looked into in the future. One possible solution would be to restrict the selected colors to a distribution of colors based on all of Chuck Close's artworks. We tested this approach by using a distribution based on a single Chuck Close portrait. The output generated had an increased semblance of authenticity, and this was because it contained colors that were a better approximation to what Chuck Close would probably have used. Figure 3.16 is the image that was generated with this additional constraint, and figure 3.17 was produced without that constraint. A comparison between the two results shows that this simple idea works.

The problem with this technique however, is that, in a lot of cases given an input color the color picking technique simply cannot generate a set of colors that satisfies all the constraints. This is because a distribution of colors based on Chuck Close's work is not broad enough that any given color can be resolved into five new colors that satisfy the luminance constraints and also stay within the distribution. Even if the distribution is based on all of Chuck Close's work, we probably will still face the same problem. Moreover, the additional constraint of restricting the colors to a distribution drastically increases the time it takes for the filter to produce the final image. Instead of the average time of five seconds, the modified filter takes approximately six minutes to generate the output on a Pentium 4 class machine.

A better approach would be to combine this idea with some sort of machine learning technique so that a new mapping can be predicated based on previous knowledge.

Figure 3.16: Output image with colors restricted to a distribution based on a single Chuck Close portrait.

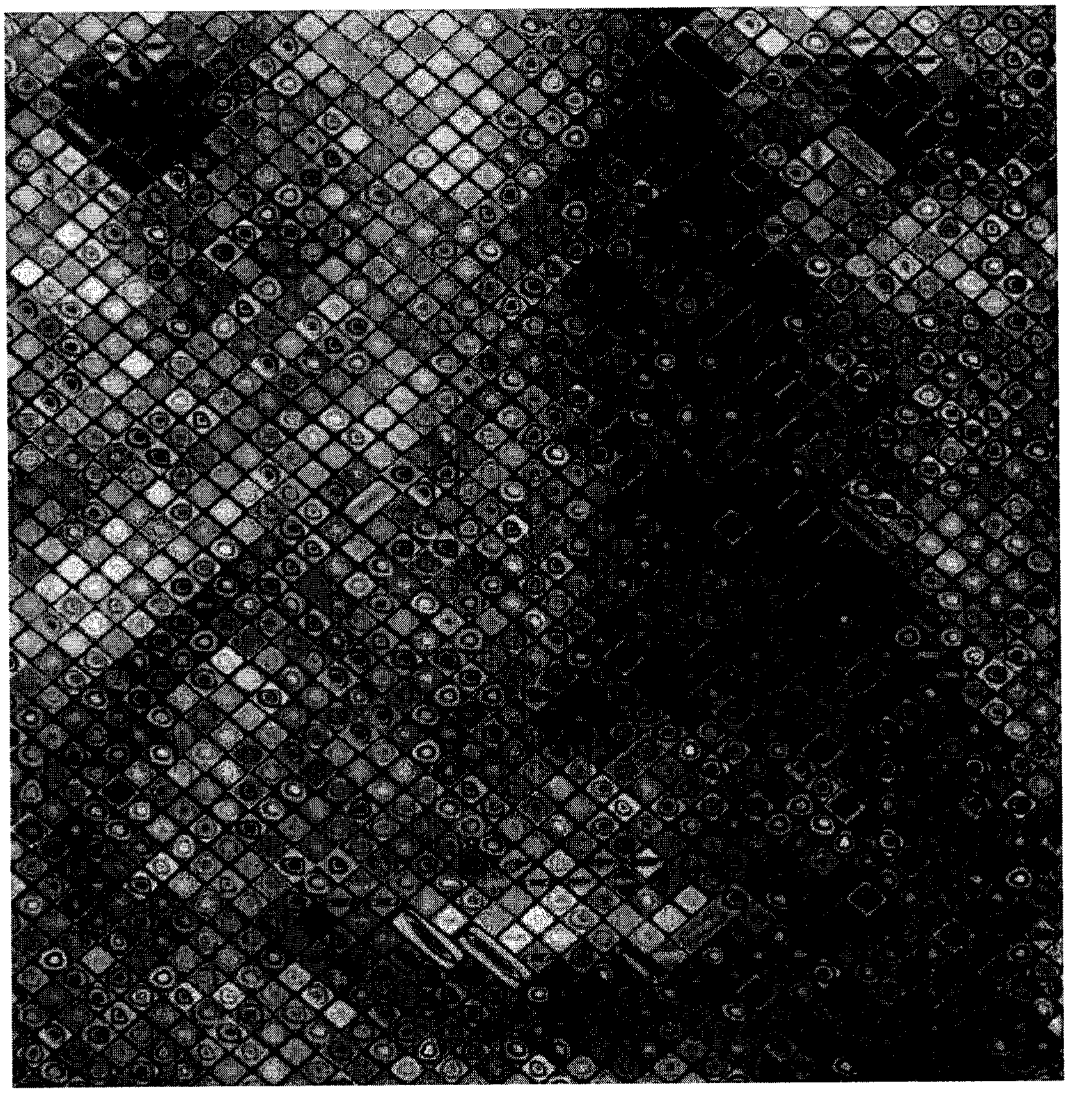

Figure 3.17: Standard Chuck Close filter output.

3.4 Image Mosaic Filter

The image mosaic filter presented in this section is the outcome of a unique blend of ideas from the *Photomosaics* paper [SH97], and the "Chuck Close" filter presented in section 3.3. The following procedure is followed in order to generate an image mosaic:

1. The small images are first categorized based on shape. This is done by applying the shape selection rules from the "Chuck Close" filter.

2. The images are then segmented into several layers. In our case the number of layers is five.
3. Finally, the color picking technique is used to paint the various layers with colors, based on their relative areas. The chosen colors can either all be isoluminant as in the case of the pointillist filter or they can be selected in pairs as in the case of the Chuck Close filter. The images presented in the results section use the latter approach.

For detailed information on how the images were categorized and segmented please refer to the related thesis to be submitted by Trân-Quân Luong.

3.4.1 Results

Figures 3.18, 3.19, and 3.20 were produced using the modified image mosaic filter. Figure 3.18 is comprised of small images of different flowers. The small images were categorized based on shape, and then segmented into several layers so that colors could be assigned to them at runtime. Figure 3.19 was created in a similar fashion, using images of fish, and other marine life forms. Figure 3.20 is based on images from Andy Warhol's collection.

3.4.2 Observations

The filter presented above needs very few source images for it to generate the image mosaic, and therefore it eliminates the task of creating a large library of source imagery. Also, both the image selection and the color picking processes are fast. Therefore the filter can generate visually captivating results in under five seconds on a Pentium 4 class machine.

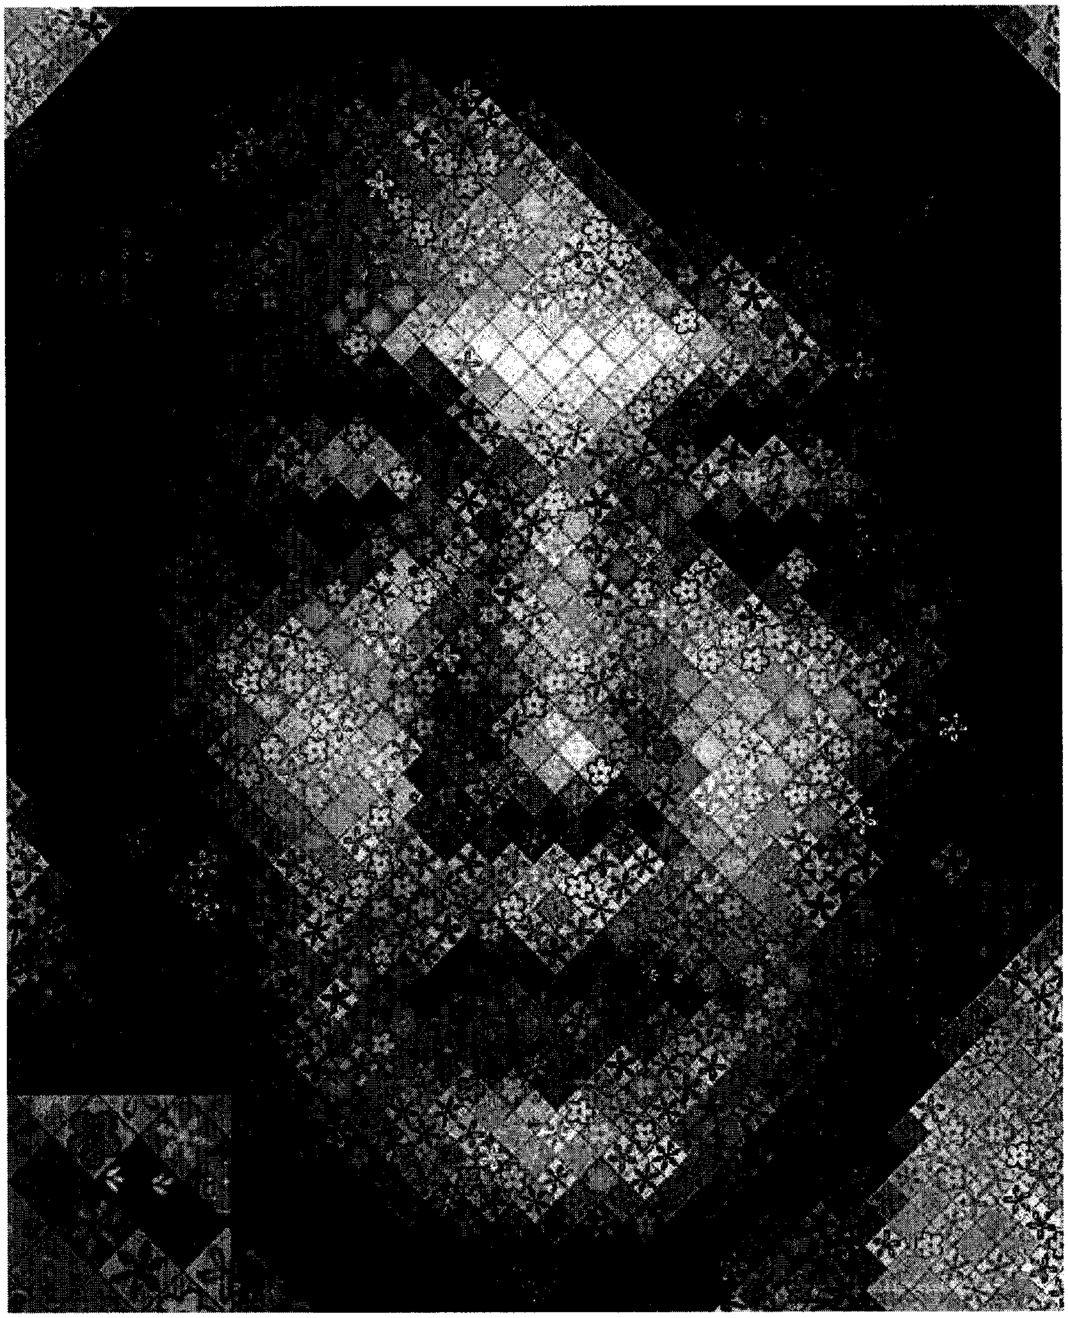

Figure 3.18: Image mosaic created using images of flowers. A close up of some of the tile images can be seen in the inset (bottom left).

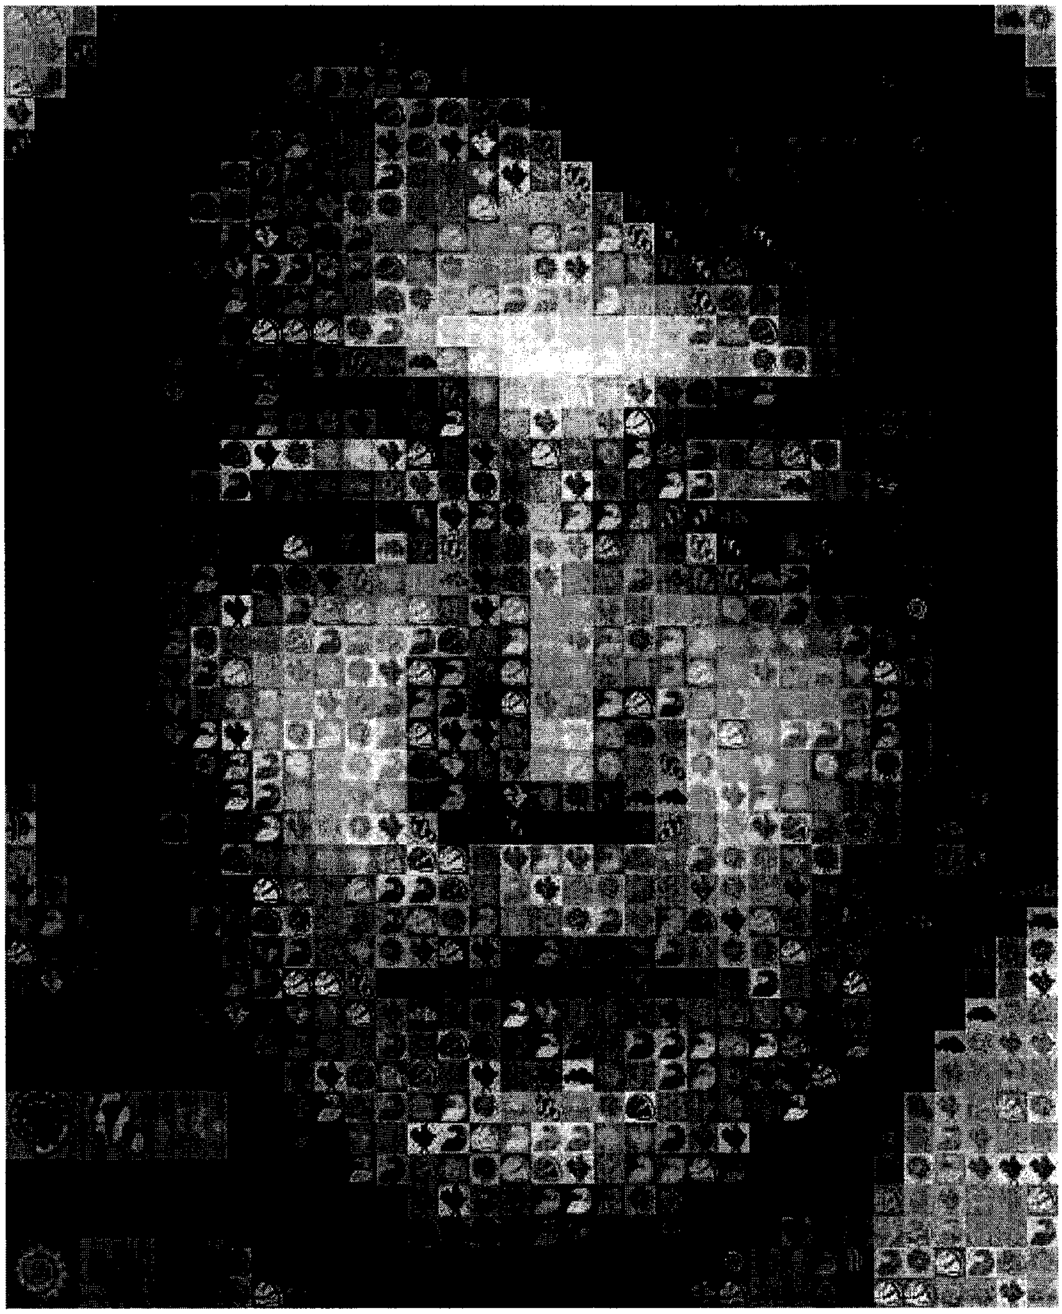

Figure 3.19: Image mosaic created using images of marine life forms. A close up of some of the tile images can be seen in the inset (bottom left).

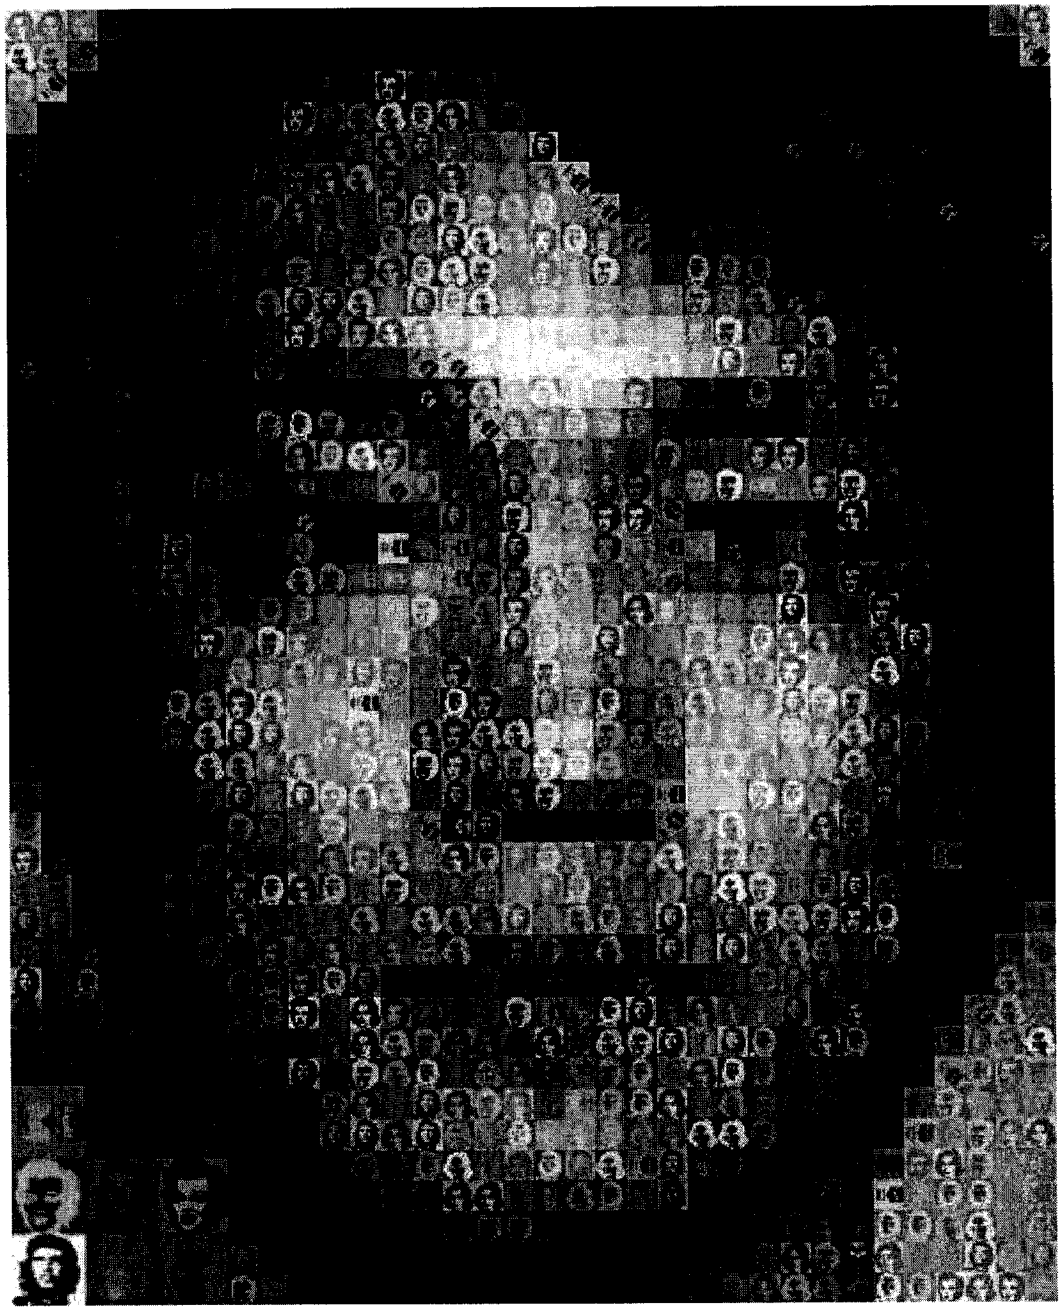

Figure 3.20: Image mosaic created using Andy Warhol images. A close up of some of the tile images can be seen in the inset (bottom left).

Chapter 4 Conclusion

This thesis presented a fast, geometric technique for isoluminant color picking in additive color spaces. It looked at how the color picking technique could be used to both improve existing NPR image filters and to create new ones. The color picking approach was used to come up with an improved pointillist filter, a new Chuck Close inspired NPR image filter, and a novel image mosaic filter.

One of the main goals of this thesis was to bring out the importance of isoluminance within the field of visual arts, and the benefits of incorporating the idea into NPR filters. To this end a generic isoluminant color picking technique, along with three image filters were presented as part of the thesis. However, abstract questions such as "should the isoluminance constraint be applied globally throughout the image or should it be applied only within certain regions? And if applied locally, which regions within a given image should be painted with isoluminant colors?" were not addressed by this thesis, as they were beyond the scope of this thesis topic. Formulizing definitive answers to such questions would probably be impossible because part of answer is subjective and would vary from one person to the other, but basic guidelines would definitely help in designing more sophisticated algorithms and filters.

4.1 Future Work

Regarding future work, the filters presented in this thesis used isoluminance to create layered imagery that has visual appeal at various viewing distances. Using isoluminant color picking to enhance other techniques in painterly rendering would be something interesting to look at in the future. For example, instead of the usual method of sampling the source image during the color selection process, isoluminance could be used to direct attention towards or away from specific regions within a painterly image. Some of the other potential applications of isoluminance are using it to increase or decrease the perceived delineation between text and its background, and intentional exaggeration or minimization of the perceived depth of objects within an image.

Bibliography

[Buc96] JohnW. Buchanan. Special effects with half-toning. *Computer Graphics Forum*, 15(3):97–108, August 1996. ISSN 1067-7055.

[CC] *Chuck Close: Process and Collaboration.* Available: http://www.chuckclose.coe.uh.edu/life/index.html

[CH01] S.M. Chosson and Roger D. Hersch. Visually-based color space tetrahedrizations for printing with custom inks. In *Symposium Electronic Imaging 2001, Color Imaging:Device-Independent Color, Color Hardcopy, and Graphic Arts VI*, pages 81–92, 2001.

[CH99] Russell T. Clement and Annick Houzé. *Neo-Impressionist Painters: A Sourcebook on Georges Seurat, Camille Pissarro, Paul Signac, Theo Van Rysselberghe, Henri Edmond Cross, Charles Angrand, Maximilien Luce, and Albert Dubois-Pillet.* Greenwood Press, 1999.

[Dec96] Phillippe Decaudin. Cartoon-looking rendering of 3d-scenes. Technical Report INRIA 2919, Universite de Technologie de Compiegne, France, June 1996.

[FH05] Gerald Farin and Dianne Hansford. *Practical Linear Algebra – A Geometry Toolbox.* AK Peters Ltd., 2005.

[FR98] Adam Finkelstein and Marisa Range. Image mosaics. *Lecture Notes in Computer Science*, 1375, 1998.

[Gol99] E. Bruce Goldstein. *Sensation and Perception.* Brooks/Cole Publishing, 1999.

[GCS02] Bruce Gooch, Greg Coombe, and Peter Shirley. Artistic vision: painterly rendering using computer vision techniques. In *Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering*, pages 83–ff. ACM Press, 2002.

[GG01] Bruce Gooch, Amy Gooch. *Non-Photorealistic Rendering*. AK Peters, Ltd., July 2001.

[GGSC98] Amy Gooch, Bruce Gooch, Peter Shirley, and Elaine Cohen. A non-photorealistic lighting model for automatic technical illustration. *Proceedings of SIGGRAPH 98*, pages 447–452, July 1998.

[GSG+99] Bruce Gooch, Peter-Pike Sloan, Amy Gooch, Peter Shirley, and Richard Riesenfeld. Interactive Technical Illustration. *Interactive 3D Conference Proceedings*, April1999.

[Hae90] Paul Haeberli. Paint by numbers: abstract image representations. In *Proceedings of SIGGRAPH 90*, pages 207–214. ACM Press, 1990.

[Half] *Halftoning*. Available: http://fhctech.org/fhc/imaging/halftone.htm.

[Her90] Aaron Hertzmann. Painterly rendering with curved brush strokes of multiple sizes. In *Proceedings of SIGGRAPH 98*, pages 453–460. ACM Press, 1998.

[Hp] *Handprint: an artist's color theory*. Available: http://www.handprint.com/HP/WCL/color9.html#top

[KB96] Peter K. Kaiser and Robert M. Boynton. *Human Color Vision*. Optical Society of America, 1996.

[KM+02] Robert D. Kalnins, Lee Markosian, Barbara J. Meier, Michael A. Kowalski, Joseph C. Lee, Philip L. Davidson, Matthew Webb, John F. Hughes, and Adam Finkelstein. Wysiwyg npr: drawing strokes directly on 3d models. In *Proceedings of SIGGRAPH 02*, pages 755–762. ACM Press, 2002.

[KP02] Junhwan Kim and Fabio Pellacini. Jigsaw image mosaics. In *Proceedings of SIGGRAPH 02*, pages 657–664. ACM Press, 2002.

[Liv02] Margaret Livingstone. *Vision and Art: The Biology of Seeing*. Harry N. Abrams, 2002.

[Osb80] Roy Osborne. *Lights and Pigments: Colour Principles for Artists*. Harper & Row, Publishers, Inc., 1980.

[OH95] Victor Ostromoukhov and Roger D. Hersch. Artistic screening. In *Proceedings of SIGGRAPH 95*, pages 219–228. ACM Press, 1995.

[OH99] Victor Ostromoukhov and Roger D. Hersch. Multi-color and artistic dithering. In *Proceedings of SIGGRAPH 99*, pages 425–432. ACM Press/Addison-Wesley Publishing Co., 1999.

[Pel99] Denis G. Pelli. Close encounters – an artist shows that size affects shape. *Science*, 285:844–846, August 1999.

[Sto98] Robert Storr. *Chuck Close*. Museum of Modern Art, 1998.

[SH97] Robert Silvers and Michael Hawley. *Photomosaics*. Henry Holt and Co., Inc., 1997.

[SMGG01] Peter-Pike Sloan, William Martin, Amy Gooch, and Bruce Gooch. The lit sphere: A model for capturing npr shading from art. *Graphics Interface 2001*, pages 143–150, June 2001. ISBN 0-96888-080-0.

[SOS98] Eric J. Stollnitz, Victor Ostromoukhov, and David H. Salesin. Reproducing color images using custom inks. In *Proceedings of SIGGRAPH 98*, pages 267–274. ACM Press, 1998.

Appendix A1: Source Images

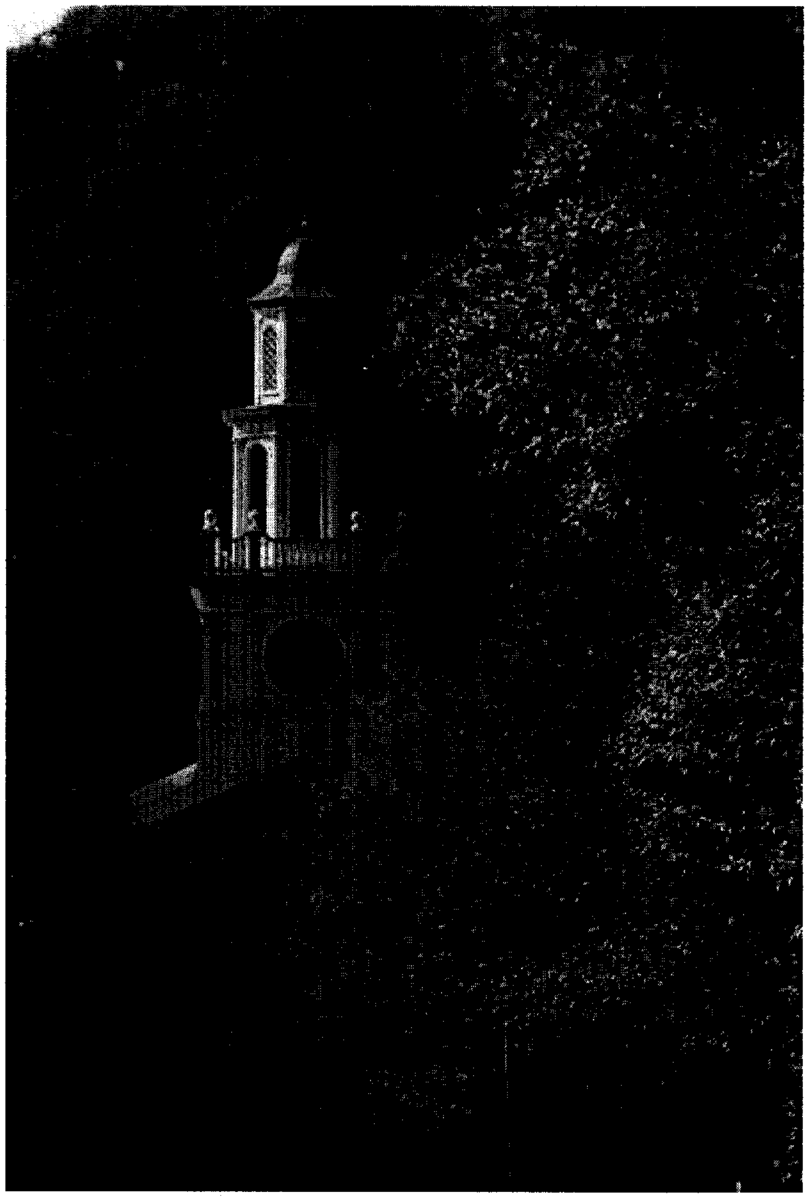

Source: www.freefoto.com.

Available: http://www.freefoto.com/preview.jsp?id=19-03-70&k=Autumn+color+in+Vermont

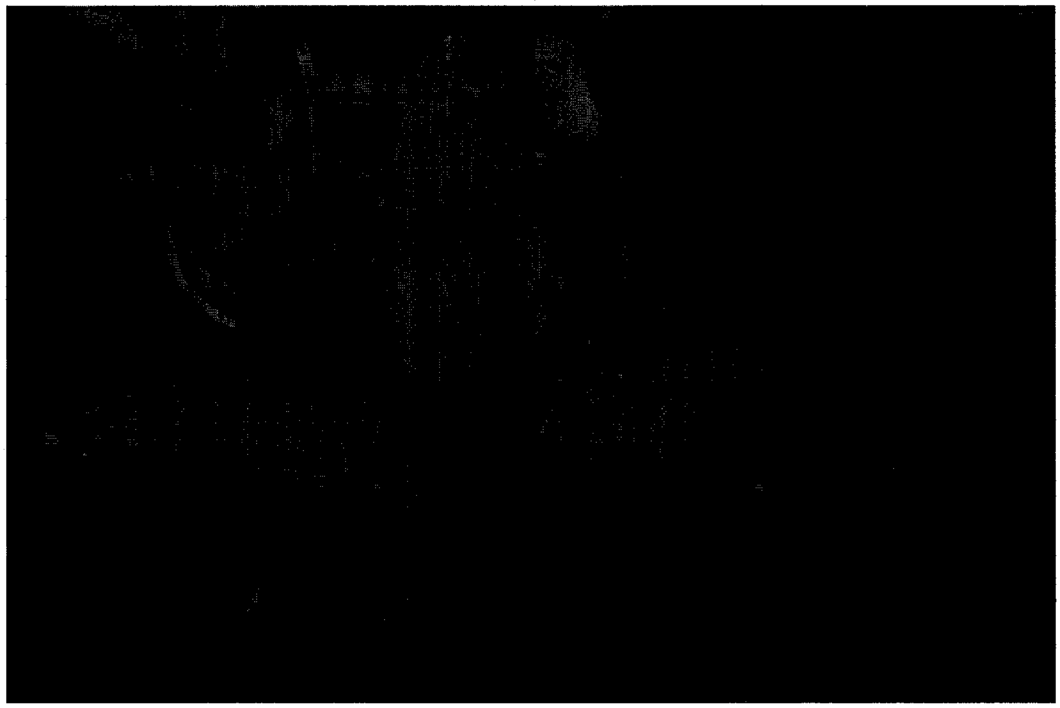

Source: www.freefoto.com.

Available: http://www.freefoto.com/preview.jsp?id=12-12-12&k=Lily

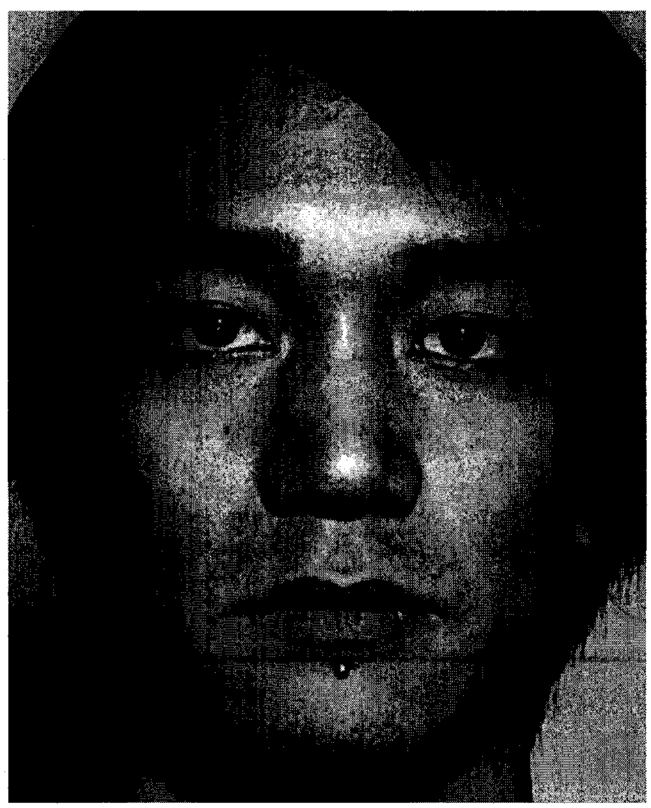

Source: Picture taken by Soukia Savage.

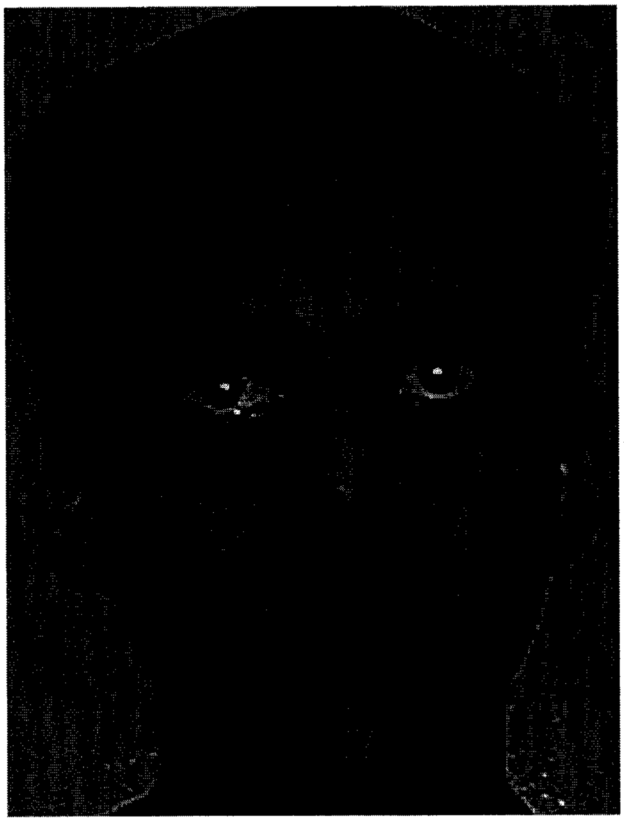

Source: Picture taken by me.