

Generating Pointillism Paintings Based on Seurat's Color Composition

Yi-Chian Wu¹, Yu-Ting Tsai², Wen-Chieh Lin^{†1}, and Wen-Hsin Li¹

¹Department of Computer Science, National Chiao Tung University, Taiwan

²Department of Computer Science and Engineering, Yuan Ze University, Taiwan

Abstract

This paper presents a novel example-based stippling technique that employs a simple and intuitive concept to convert a color image into a pointillism painting. Our method relies on analyzing and imitating the color distributions of Seurat's paintings to obtain a statistical color model. Then, this model can be easily combined with the modified multi-class blue noise sampling to stylize an input image with characteristics of color composition in Seurat's paintings. The blue noise property of the output image also ensures that the color points are randomly located but remain spatially uniform. In our experiments, the multivariate goodness-of-fit tests were adopted to quantitatively analyze the results of the proposed and previous methods, further confirming that the color composition of our results are more similar to Seurat's painting style than that of previous approaches. Additionally, we also conducted a user study participated by artists to qualitatively evaluate the synthesized images of the proposed method.

Categories and Subject Descriptors (according to ACM CCS): I.4.10 [Image Processing and Computer Vision]: Image Representation—Statistical; J.5 [Computer Applications]: Arts and Humanities—Fine arts

1. Introduction

In recent years, digital painting systems have become more and more popular. Pointillism paintings, especially artworks that mimic Seurat's style, are thus no longer restricted to using canvases and brushes, but can be generated by applying various non-photorealistic rendering algorithms. Traditional related methods, including stippling based on halftoning [Bay73, FS75, JN76, Ost01], texture synthesis [EF01, HJO*01, LSRY10, SWHS97, YCLL08], and Voronoi diagrams [DHvOS00, HHD03, Sec02], mainly focus on controlling the distributions of points with a single color or gray scale. There are also some previous articles [AML10, HE04, Her98, JH05, LSKL05, RP11, SY10, YY08] that instead investigate on synthesizing some special characteristics in Seurat's paintings, such as halos and complementary colors. However, these methods do not learn the color composition from Seurat's paintings. Their results thus tend to be perceptually different from real artworks.

In this paper, we propose an intuitive and simple approach to synthesize pointillism paintings that mimic Seurat's style.

Our method relies on analyzing Seurat's paintings to learn important color statistics, including when (or where) artists employ complementary colors and how artists blend a color using various color points. Specifically, we have observed that the color distributions of Seurat's paintings are spatially localized and can be approximated by multivariate Gaussians. The learned statistical color model thus can be applied to stylize an input image with characteristics of Seurat's paintings. We also perceptually compare the synthesized pointillism paintings of different approaches through a subjective user study. Some positive feedback of the proposed method was obtained, and comments from this user study also shed light on future research directions. In brief, this paper makes the following contributions:

- A statistical model of color distributions and composition learned from Seurat's paintings.
- An automatic image pointillism approach based on the learned statistical color model.
- A quantitative method and a qualitative user study to evaluate the quality of computer-synthesized pointillism paintings.

The system diagram of the proposed method is shown in Figure 1. Our system mainly consists of two stages: learning

[†] wclin@cs.nctu.edu.tw




Figure 1: System diagram of our method.

(top) and run-time (bottom). In the learning stage, we collect several Seurat's paintings, partition individual paintings into segments, and analyze their local color distributions using the proposed statistical color model (Section 3). In the run-time stage (Section 4), an input image is first segmented into several regions. Then, the learned statistical color model is applied to generate color points using multi-class sampling [Wei10], based on the color gradients of segmented regions in the input image. Our experimental results show that the output images can perceptually resemble the color distributions and composition of real artworks (Section 5). The quantitative comparison verifies that our results are closer to artist's paintings than the results of previous approaches. In addition, a subjective user study participated by artists is also presented to qualitatively evaluate the similarities and differences of the output images of different approaches with respect to Seurat's painting style.

2. Related Work

2.1. Example-based NPR

In recent years, example-based NPR approaches are commonly used for image stylization. Hertzmann et al. [HJO*01] transformed an input image into a painting of which style is learned from a pair of an image and its exemplar painting. Yan et al. [YCLL08] synthesized painterly art maps from stroke textures in paintings. These art maps then can be applied to render a 3D model in an artistic style similar to the paintings. Lee et al. [LSRY10] applied texture synthesis methods to transfer the style and texture of an example image based on the directional information of an input image. In this way, the result can aesthetically express the object shape of the input image.

Our work is also based on example-based NPR methods; however, we address this problem in a different way. Unlike previous methods that directly synthesize stroke textures extracted from an artist's painting, we construct a parametric statistical model from several paintings of an artist to capture the artist's style, and then apply this model to transform an arbitrary image into a painting with a similar style.

2.2. Stippling

Traditional stippling approaches often focus on adjusting positions, sizes, and shapes of points to convey the tone and texture of a gray-scale image. Some of them [DHvOS00, KSL*08, Sec02] rely on the centroidal Voronoi diagram to enable local controls of point distributions, while others instead employ texture synthesis [SWHS97], example-based methods [KMI*09, MALI10], graph theory [Mou07], etc.

As for color images, Hertzmann [Her98] proposed to paint an image with curved brush strokes of various sizes, so that pointillist images could be synthesized by using very short strokes. Hays and Essa [HE04] extracted appearance and motion information from images and videos to synthesize pointillism animations. Jang and Hong [JH05] integrated their color density model with Penrose tiling [Gla98, ODJ04] to determine point distributions. Luong et al. [LSKL05] improved pointillism filtering by utilizing isoluminant color points, as often exploited by skilled artists, to create perceptual tensions in an output image. From Seurat's paintings, Yang and Yang [YY08] analyzed elementary colors, complementary colors, point sizes, and halos. Seo and Yoon [SY10] further applied recursive Wang tiles [KC06] to produce point distributions based on the RYB color model that is effectively used in Seurat's paintings. Arroyo et al. [AML10] demonstrated that point positions can be stochastically determined by simulating the behaviors of artists with an expert system. Roguljic and Paćic [RP11] utilized a five-layer procedure to synthesize the pointillist style. They particularly performed color perturbations in the HSV color space to add more natural variations and human-like appearance in the results.

In this paper, we also focus on color images, but propose a novel statistical color model to simulate pointillism. Based on this model, quantitative comparisons among different pointillism algorithms thus can be easily achieved.

2.3. Halftoning

Halftoning is a process that simulates the continuous tone of a gray-scale image by controlling the sizes, shapes, and spacings of dots. The most representative work may be the famous Floyd-Steinberg error diffusion algorithm [FS75]. Ostromoukhov [Ost01] further improved the quality and performance of the Floyd-Steinberg algorithm to avoid visual artifacts in output images.

Nevertheless, halftoning approaches may not be suitable for color pointillism, since they compute output points in a given discrete sample space and tend to generate regular results. Besides, previous halftoning techniques, such as dithering and error diffusion, rely on a local neighborhood region with fixed size for sampling. Discretization artifacts may be obviously visible if the size of the neighborhood region is relatively much smaller than the resolution of the output image. Finally, although many halftoning methods

support color images, they can only synthesize a few colors. Therefore, the proposed pointillism stylization algorithm is not based on halftoning, but instead based on blue noise sampling to generate colorful and high-quality results.

2.4. Blue Noise Sampling

Blue noise sampling has been widely used in signal processing applications due to its spatially uniform and aliasing-free properties. Inspired by the local cone densities in the rhesus retina, Yellott [Yel83] suggested that Poisson disk distributions might be useful in artificial imaging systems. Lloyd's quantization [Llo82] and Cook's stochastic sampling [Coo86] introduced blue noise sampling algorithms that could achieve different trade-offs between quality and performance. After that, many sampling approaches based on Poisson disk distributions have been presented, including Wang tiles [CSHD03, KCODL06] and procedural object distributions [LD05]. Moreover, Ostromoukhov et al. [ODJ04] combined Penrose tiling [Gla98] and local importance analysis to efficiently generate sampling patterns with the blue noise property. This method was later applied by Jang and Hong [JH05] to distribute color points randomly, but still in accordance with color densities in the input image.

It is worth noting that most previous sampling techniques are single-class approaches. They only deal with a single class of samples at one time, but cannot maintain the desirable blue noise property among multi-class samples. Wei [Wei10] thus proposed multi-class Poisson disk sampling to fill this gap. Our pointillism algorithm also relies on multi-class sampling to synthesize evenly spaced stippling patterns with rich colors and identifiable structures.

3. Statistical Color Model

To realistically simulate the pointillist style, we collect several Seurat's paintings and analyze them using image processing and statistical methods. The learned statistical color model thus represents local color distributions and can be utilized to transform an input image into a pointillism painting. Figure 2 shows how the color composition is learned from Seurat's paintings. First, we partition the paintings into disjoint regions using image segmentation techniques (Section 3.1.1). For each region, clustering is further applied to extract a representative color and composite colors (Section 3.1.2). Finally, several covariance matrices are also computed for each region to form various color distributions within the region (Section 3.1.3).

The color statistics of each region, including the representative color, composite colors, and covariance matrices, thus statistically describe how Seurat locally blends a color using various color points in his paintings. Specifically, our color composition model is based on multivariate Gaussians, and includes two classes of color distributions: *representative*




Figure 2: Learning process of our statistical color composition model.

and *composite*. The representative color distribution $\text{Pr}_i(C)$ for the i -th region S_i is formulated as follows:

$$\text{Pr}_i(C) = \frac{\exp\left(-\frac{1}{2}(C - C_i^r)^T \mathbf{M}_i^{-1} (C - C_i^r)\right)}{(2\pi)^{3/2} |\mathbf{M}_i|^{1/2}}, \quad (1)$$

where $C_i^r \in \mathbb{R}^3$ and $\mathbf{M}_i \in \mathbb{R}^{3 \times 3}$ are the representative color and covariance matrix of S_i (in the RGB color space), respectively. This distribution is employed to generate new color distributions (Section 3.2).

Additionally, the composite color distribution $\text{Pr}_{i,j}(C)$ for the j -th composite color $C_{i,j}^c$ in S_i is formulated as follows:

$$\text{Pr}_{i,j}(C) = \frac{\exp\left(-\frac{1}{2}(C - C_{i,j}^c)^T \mathbf{M}_{i,j}^{-1} (C - C_{i,j}^c)\right)}{(2\pi)^{3/2} |\mathbf{M}_{i,j}|^{1/2}}, \quad (2)$$

where $\mathbf{M}_{i,j}$ is the composite covariance matrix of $C_{i,j}^c$. When converting an input image into a pointillism painting, the composite distributions will guide multi-class blue noise sampling to produce visually pleasant color points that mimic Seurat's style, which will be presented in Section 4.

3.1. Analysis of Seurat's Paintings

3.1.1. Image Segmentation

Image segmentation plays an important role in our analysis process. We partition regions based on color information in Seurat's paintings, since the main purpose of segmentation in our algorithm is to find local color distributions for the proposed statistical model, not high-level object or pattern extraction. For each Seurat's painting, we first apply Gaussian filtering to blur sharp contours of color points, and then perform image segmentation based on mean shift clustering




Figure 3: An example of the representative color C_i^r and the composite colors $C_{i,j}^c$ of a segmented region S_i in a Seurat's painting. $K_{i,1}, K_{i,2}, \dots, K_{i,7}$ are the partitioned clusters of S_i .

[CM02]. In our experiments, parameters of mean-shift clustering were set as follows: spatial bandwidth: 30~80; range bandwidth: 8~20; minimum pixel count: 50~220. Note that currently the values of parameters are manually determined according to color variations in the input painting. Nevertheless, they can also be estimated using some data-dependent approaches, as suggested in [CM02]. In the second row of Figure 2, we show some examples of the segmented results of Seurat's paintings.

3.1.2. Representative and Composite Colors

Seurat believes that human eyes can locally mix color points, and perceive a different color when the points are viewed from a distance. Our method borrows this principle from pointillism, and analyzes what colors an artist would employ for a color that he really wants to express. Specifically, an image of pointillism painting is partitioned into disjoint regions after segmentation. We extract a representative color and several composite colors within each region to model the color composition that an artist applied. The representative color C_i^r of the region S_i is defined as the mean of pixel colors in S_i . For a region S_i , k-means clustering is utilized to classify pixel colors within S_i into k clusters. The mean of pixel colors of each cluster is defined as a composite color, e.g., $C_{i,j}^c$ denotes the j -th composite color of the j -th cluster $K_{i,j}$ of region S_i . More formally, C_i^r and $C_{i,j}^c$ are defined by

$$C_i^r = \frac{1}{n_i} \sum_{p \in S_i} C_p, \quad C_{i,j}^c = \frac{1}{n_{i,j}} \sum_{p \in K_{i,j}} C_p, \quad (3)$$

where n_i and $n_{i,j}$ are the numbers of pixels in S_i and $K_{i,j}$, respectively. Moreover, the composite density $d_{i,j}$ of the j -th cluster is also computed as $d_{i,j} = n_{i,j}/n_i$, in order to record the pixel ratio of the j -th cluster in region S_i .

Note that the clustering for extracting composite colors is different from the clustering for image segmentation in Section 3.1.1. The clustering here does not take the positions of pixels into account, but considers only pixel colors. Since the pixels in a segmented region are already localized and will be mixed by human eyes when viewed from a distance, there is no need to partition the region into several much smaller regions. Moreover, this clustering also can stochastically increase the color variations in the output image without substantially changing the perceived colors. In our experiments, the number of clusters in a region was set to 7, i.e., $k = 7$. This will keep the computational cost affordable, while preserving the representation efficiency. We show some examples of representative and composite colors learned from Seurat's paintings at the bottom row of Figure 2 and Figure 3(c).

3.1.3. Covariance Matrix Computation

The covariance matrix of a multivariate Gaussian summarizes the covariances of all pairs of observations/variables. To increase the robustness of our statistical color model, a representative covariance matrix and several composite covariance matrices are estimated for each region to account for the relations between each pair of color channels. The representative covariance matrix \mathbf{M}_i of the region S_i and the j -th composite covariance matrix $\mathbf{M}_{i,j}$ of the j -th cluster $K_{i,j}$ are respectively learned from pixels in S_i and $K_{i,j}$ by

$$\mathbf{M}_i = \frac{1}{n_i - 1} \sum_{p \in S_i} (C_p - C_i^r) (C_p - C_i^r)^T, \quad (4)$$

$$\mathbf{M}_{i,j} = \frac{1}{n_{i,j} - 1} \sum_{p \in K_{i,j}} (C_p - C_{i,j}^c) (C_p - C_{i,j}^c)^T. \quad (5)$$

3.2. Extending Color Model at Unsampled Space

A common problem of learning processes is how to extend the learned model when the available training data are limited. Currently, we have only referred five digital high-resolution Seurat's paintings as the training data. Thus, the number of representative and composite colors in these paintings may not be large enough to stylize an arbitrary input image. Figure 4(a) especially illustrates that the learned representative colors only cover a portion of the RGB color space. With such limited information, our statistical color model only produces results whose colors are similar to those of the reference paintings, but cannot be adaptive to dissimilar colors.

Rather than analyzing more Seurat's paintings, we adopt an extrapolation approach to solve this issue. For an arbitrary color in an input image, it is desirable that there is at least one representative color close enough to the input color. Thus, the RGB color space is first partitioned into a regular grid. If there are no representative colors in a grid cell, i.e., an empty cell, the extrapolation process is applied to generate new representative and composite color distributions for




Figure 4: Representative colors in the RGB color space before and after extrapolation.

the cell. In our experiments, a regular grid of size $8 \times 8 \times 8$ was utilized, since we have found that human eyes are not sensitive to the differences of colors within a cell of such size. The representative colors after extrapolation are shown in Figure 4(b).

The only remaining issue is how to obtain new representative and composite color distributions. To allow high-quality output results, we extrapolate the parameters of distributions, including the means and the covariance matrices. The extrapolation process thus consists of four steps:

1. Search for an empty cell G with the most number of non-empty neighboring cells.
2. Obtain the representative color C_{new}^r and the covariance matrix \mathbf{M}_{new} for G .
3. Based on C_{new}^r and \mathbf{M}_{new} , generate composite color distributions for G .
4. Repeat Steps 1.~3. until there are no empty cells.

In our experiments, the neighboring cells of G are defined as the 26 neighbors in the adjacent $3 \times 3 \times 3$ cells of G . For Step 2., we simply set C_{new}^r as the color value of the center of G , while \mathbf{M}_{new} is computed by

$$\mathbf{M}_{new} = \frac{\sum_{i \in N(G)} w_i \mathbf{M}_i^r}{\sum_{i \in N(G)} w_i}, \quad (6)$$

where $N(G)$ denotes the index set of representative colors that are in the non-empty neighboring cells of G , and the extrapolation weight w_i is inversely proportional to the distance from C_i^r to C_{new}^r

$$w_i = \frac{1}{\|C_i^r - C_{new}^r\|}. \quad (7)$$

Note that \mathbf{M}_{new} is derived only from the parameters of representative color distributions in the neighboring cells of G in this step. As for Step 3., we first sample many new colors, say 2000 colors in our experiments, from the color distribution defined by C_{new}^r and \mathbf{M}_{new} , and then apply the analysis methods in Sections 3.1.2 and 3.1.3 to compute the composite colors and covariance matrices from these color samples. Figure 5 compares the results of our algorithm with/without extending the statistical color model. In this figure, the extrapolation process successfully generates richer colors in the stylized result that expresses more detailed features and color variations of the input image.




Figure 5: Comparison of our results with/without extending the statistical color model.

4. Image Pointillism

After learning the parameters of the proposed statistical color model, an input image then can be transformed into Seurat's style. The input image is first segmented into disjoint regions using mean shift clustering [CM02] with the following parameter settings: spatial bandwidth: 20~80; range bandwidth: 2~5; minimum pixel count: 20~50. For each segmented region, we search for a learned representative color $C_{i^*}^r$ that is closest to the mean color of the region. Given $C_{i^*}^r$, our approach generates composite colors by sampling from the learned statistical color model ($\{C_{i^*,j}^c\}_{j=1}^k$ and $\{\mathbf{M}_{i^*,j}\}_{j=1}^k$) using Equation 2. Each sampled composite color is treated as a class of samples in the modified multi-class blue noise sampling process in which color points (samples) of different composite colors (different classes) are placed in the segmented region based on the blue noise property (Section 4.1). The proposed framework also allows changing the shape of a point into an ellipse (Section 4.2), which further improves the quality of our results.

4.1. Point Sampling

In Seurat's paintings, points with similar colors tend to be distributed with the blue noise property in a local region. Several previous articles [JH05, KC06, LD05] also produced visually pleasant results by applying blue noise sampling methods to pointillism stylization. Nevertheless, they are single-class methods. This means that all points are placed regardless of their colors. To solve this problem, we apply a modified blue noise sampling approach based on the multi-class hard disk technique [Wei10]. Please refer to [Wei10] for detailed discussion on the differences of single-class and multi-class blue noise sampling.

The proposed sampling process generates points of each segmented region separately. After finding the closest representative color $C_{i^*}^r$ for a segmented region in the input image, we compute the R -matrix of a region for multi-class

sampling to determine the positions of points. The R -matrix, which is a generalization of the distance constraint in single-class Poisson disk sampling, specifies the desired minimal distance between each pair of the composite colors in the region. Thus, the R -matrix can be utilized to randomly place points that belong to different classes, while keeping the distances among them well constrained according to our statistical color model. Finally, the colors of points are drawn from the composite color distributions of the region.

In our algorithm, the R -matrix of a segmented region is computed according to the suggestion in [Wei10], but slightly modified. Specifically, the diagonal terms of the R -matrix need to be manually specified in [Wei10], while these terms are computed automatically in our algorithm:

$$(\mathbf{R}_{i^*})_{j,j} = \frac{1}{d_{i^*,j}} \frac{a_l + a_s}{3}, \quad (8)$$

where $d_{i^*,n}$ is the density of the composite color n obtained from Seurat's paintings (section 3.1.2). a_l and a_s are the user-defined transverse and conjugate diameters of an elliptic point used to fill the region. $(\mathbf{R}_{i^*})_{m,n}$ is the entry of the R -matrix in the m -th row and the n -th column. $(\mathbf{R}_{i^*})_{m,n}$ denotes the minimal distance between class m ($C_{i^*,m}^c$) and class n ($C_{i^*,n}^c$) when color points of these two classes are placed (sampled) to a region. Hence, Equation 8 guides the sampling process to distribute color points according to the distribution in Seurat's paintings since the minimal distance between two color points is inversely proportional to the composite density. Additionally, the second term is multiplied to account for the size of an elliptic color point. This avoids too much overlapping between two color points. Once the diagonal terms of the R -matrix is computed, the off-diagonal terms can be computed automatically using Wei's method [Wei10].

4.2. Point Shapes

From our observation, the point shapes in Seurat's paintings are certainly not circles, but are similar to ellipses whose main axes roughly align with the contours of objects. This not only emphasizes sharp features of image contents, but also makes object shapes more distinguishable. Thus, our algorithm further allows changing the shape of a generated point into an ellipse with user-defined transverse and conjugate diameters. We also compute the gradients of luminance of the input image, and rotate each elliptic point to align its minor axis with the direction of the gradient at the same position in the image. To further emphasize the contours of objects, color points are filled in a order from gradients of larger magnitude to gradients of lower magnitude.

5. Results

In previous work, the quality of computer-synthesized pointillism paintings is judged mostly by human eyes, but visual judgement can often be too *subjective* for digital artworks. In this paper, we present a quantitative approach to

compare the image quality based on the learned statistical color model. We believe that comparing the color composition and distributions between synthesized and Seurat's paintings can provide more objective and persuadable judgement. In addition to the quantitative approach, we also conducted a user study to evaluate the perceptual quality of synthesized results. Finally, some limitations of the proposed method are discussed.

5.1. Quantitative Comparison

To measure the similarity between our results and Seurat's paintings, we apply the multivariate goodness-of-fit tests [AZ05] for quantitative comparison. A synthesized result is first partitioned into disjoint regions using image segmentation. For each segmented region S_i , we search for a corresponding region S_j in Seurat's paintings such that the mean color of S_j is the closest to that of S_i . The multivariate goodness-of-fit tests, with the logarithmic distance function [AZ05], are then employed to compare the color distributions of S_i and S_j . To compare with other methods, we also measure the similarity between the output images of previous methods [HE04, Her98, LSCL05, YY08] and Seurat's paintings (shown in Figures 6-9).

Figure 6(e) plots the results of goodness-of-fit tests for the input image "pink flower". The horizontal axis represents the index of segments in the image, while the vertical axis represents the energy of the goodness-of-fit test. Lower energy means higher similarity of the distribution and composition of the color points between the image and Seurat's paintings. From Figure 6(e), one can find that the energy of our approach is generally smaller than the energy of other methods. Table 1 further lists the statistics of goodness-of-fit tests on the synthesized results of different input images. To evaluate if the distributions of the goodness-of-fit tests across segments in other results and ours is statistically significant, we compute the p -values of each method against ours. The p -values in Table 1 especially show that most pairwise comparisons are statistically significant. This means that the color composition of our output images is the most similar to that of Seurat's paintings and has statistical significance when compared to other methods.

From Figures 6-9, it appears that the approach proposed in [HE04] tends to produce a pointillism painting with more light colors. The output image of [Her98] is often visually mottled in color, while the method of Luong et al. [LSCL05] may synthesize too blurry result (shown in Figure 8(b)). The contours of points cannot be clearly observed in the results of [YY08], which also do not exhibit rich colors. By contrast, our algorithm can clearly render the contours of points, and simulate strokes with colors that are indeed employed in pointillism paintings.

Other results. Figure 10 shows the input of Figure 4(a) and our result. More synthesized results are shown in Figure




Figure 6: Comparison of synthesized results (pink flower). (e) is the energy of goodness-of-fit tests in different segmented regions. The computed p -values verify that our result is closest to the Seurat's painting, which is statistically significant.

Method	Ours	[YY08]	[HE04]	[Her98]	[LSKL05]
pink flower	avg=4.62 std=0.27 $p=3.0e-3$	avg=4.66 std=0.40 $p=3.78e-79$	avg=5.00 std=0.38 $p=0.0$	n/a n/a n/a	n/a n/a n/a
cow	avg=4.66 std=0.37	avg=4.94 std=0.61 $p=0.0$	n/a n/a n/a	avg=4.66 std=0.37 $p=0.18$	n/a n/a n/a
columns	avg=4.62 std=0.40	avg=4.85 std=0.55 $p=0.0$	n/a n/a n/a	n/a std=0.45 $p=7.68e-11$	avg=4.58 std=0.45 $p=0.0$
geisha	avg=4.70 std=0.37	avg=4.79 std=0.54 $p=1.41e-14$	avg=4.80 std=0.38 $p=9.0e-18$	n/a n/a n/a	n/a n/a n/a

Table 1: Average energy of goodness-of-fit tests and its standard deviation across different segments. The p -values of each method against ours are also shown in this table.




Figure 7: Comparison of synthesized results (cow).

14. Additionally, we compared our synthesized result with one of Seurat's paintings in Figure 12. This figure demonstrates that the color composition and point distributions of our result are close to those in the Seurat's painting. Please see also the supplemental materials for more results.

The learning time for our statistical color model varies with the digitized size of a painting. For 5 Seurat's paintings used in our experiments, the average learning time is about 4020 seconds (average size: 1000×787). Table 2 lists the timings and the traverse/conjugate diameters of the ellipses used for rendering our results. The timings were measured on a PC with Intel® CPU i7-720QM and 4 GB RAM.




Figure 8: Comparison of synthesized results (columns).




Figure 9: Comparison of synthesized results (geisha).




Figure 10: The input of Figure 4(a) and our result (bridge).

5.2. User Study

Computer-synthesized paintings are digital artworks that sometimes cannot be simply judged by a quantitative method. We thus performed a subjective user study to qualitatively compare the synthesized images of different pointillism algorithms. Another important purpose of this study is to further understand the perceptual quality, similarities, and differences of our results with respect to Seurat's style, especially from the point of view of an artist, not a scientist.

Input image	Segmentation	Sampling	Drawing	Traverse	Conjugate
pink flower	950s	1687s	1891s	5	2
cow	1409s	1353s	2369s	5	2
columns	1262s	1435s	2101s	5	2
geisha	709s	1369s	2543s	4	2
bridge	199s	270s	2282s	5	2
beach	339s	1475s	1293s	5	2.5
cat	502s	641s	973s	4	2
peacock	274s	430s	2062s	5	2
tulip	183s	285s	867s	5	2
sunset	742s	1570s	1021s	5	2

Table 2: Timings (in seconds) and traverse/conjugate diameters (in pixels) of the ellipses used for rendering our results. The drawing step includes retrieving point color from learned model and plotting ellipses.




Figure 11: Comparison of our synthesized results using different parameters for image segmentation.




Figure 12: Color composition and point distribution of our result is close to those in the Seurat's painting. (a) and (b) correspond to the blue box region in (c) and (d), respectively.

Three artists, whose painting experience ranges from 6 to 30 years, participated in our user study. One of them particularly has the experience of pointillism paintings. There were two tests conducted in our user study. Each artist was shown the same image set and could evaluate images at will without time limit. In the first test, the artists were asked to first examine four input images and their synthesized results of different methods (shown in Figures 6–9), and then to give detailed aesthetic comments on the results. In the second test, the artists were asked to directly give comments on five syn-



Figure 13: Synthesized result of our method with manually-edited halos (cat).




Figure 14: Additional synthesized results of our method.

thesized results of our method which were synthesized from different input images (shown in Figures 10 and 14). In both tests, no information or hints about which result was generated by which approach were provided to the artists. Each artist examined images independently and the output images of different approaches were shown in random order.

Major comments of the artists are summarized as follows, where the source of each comment is shown in parentheses:

1. Seurat often employs dark or gray tones in his paintings; however, color tones do not resemble Seurat's style in many synthesized results, including results by ours and other approaches. (from both tests)
2. Composition and layout in some synthesized results, particularly Figure 14(c), may not be similar to Seurat's style. (from test 2)
3. Brush strokes in most synthesized results, except for ours, are neither directional nor aligned with object silhouettes. The directional property, however, is very important since it presents a *visual rhythm* that is an essential element in Seurat's paintings to enliven subjects. (from test 1)
4. In most results, except for ours, the shape and size of a point almost do not vary with the contents of an input image. (from test 1)
5. There is too much empty space in both results of Hays and Essa [HE04], i.e., Figures 6(b) and 9(b). Since no grounding colors are presented, their results are also too bright when compared to other methods. (from test 1)
6. The results of Yang and Yang [YY08] are usually more similar to cut-and-paste artworks or mosaic paintings, not pointillism paintings. (from test 1)
7. Compared to real pointillism paintings, colors in some synthesized results, except for ours, often do not change gradually, especially in Figure 7(c). (from test 1)

Comments 1. and 2. are actually irrelevant to synthesis methods, since composition/layout and color tones in each

result just follow its input image, no matter which approach was employed. In fact, it is a dilemma whether the output painting should be faithful to its input image or Seurat's painting style when their styles in composition/layout and color tones conflict. We believe that these two issues can be better and easily handled by carefully choosing input images from the aesthetic point of view. This also leaves a user more freedom for creation. In comments 3. and 4., the artists especially indicate that the orientation, shape, and size of a point are visually important elements in Seurat's paintings. In this aspect, our method usually performs better than other approaches, since the proposed method indeed changes the orientation and shape of a point according to the contents of an input image (Section 4.2). However, the size of a point is still constant in a synthesized result in the current implementation. We seek to investigate this issue in the future.

Furthermore, some positive comments on our results from the artists are summarized as follows:

1. The color composition and illumination effects of Figure 9(d) are natural and clear. Its overall style is particularly similar to Seurat's famous painting: "A Sunday Afternoon on the Island of La Grande Jatte". (from test 1)
2. In Figure 14(a), the gradual color changes are as gentle and soft as those in Seurat's paintings, especially for peacock feathers. (from test 2)
3. In Figure 14(c), the color composition and natural appearance of the sky are appreciated. The employed colors are rich and expressive, but not peculiar or artificial at all. (from test 2)

5.3. Discussions

In our framework, image segmentation of the input image affects the fine-scale features in the synthesized pointillism paintings. Figure 11 compares the synthesized results of our method with different parameters for image segmentation. Figure 11(b) were synthesized from Figure 11(a) with large segmented regions, while Figure 11(d) from Figure 11(c) with small segmented regions. One can find that detailed features are lost in Figure 11(b). Currently, three parameters of the applied image segmentation algorithm are manually tuned. This provides users more flexibility to control the output. If a fully automatic approach is desired, some parameter estimation techniques can be applied to solve this problem. We consider this as a future research direction.

It also can be found that some colors in the synthesized result are somewhat dissimilar to the input image. This is due to the fact that the learned and extrapolated representative colors are not densely distributed over the entire RGB color space. To solve this issue, we can employ a grid of larger size for the extrapolation process in Section 3.2, for example, increasing the grid size from $8 \times 8 \times 8$ to $16 \times 16 \times 16$, so that more extrapolated representative colors are generated and distributed in the RGB color space. Nevertheless, this will

result in more computational time for the run-time stage, since the cost for finding the closest representative color (Section 4.1) will be definitely increased when the number of total representative colors increases.

Furthermore, previous articles often consider that halos in pointillism paintings are likely to be applied at sharp edges by artists [SY10, YY08]. Nevertheless, Seurat sometimes added halos to smooth regions in some paintings, where the color variations of halos usually cannot be estimated solely based on the luminance. An obvious example is Seurat's famous painting: "A Sunday Afternoon on the Island of La Grande Jatte". In the bottom left of this painting, the hat edges of the man with a tobacco pipe are rather smooth with respect to the dark-color suit. If halos were not added around the hat, its edges would be easily ignored. Similar examples also can be found in other Seurat's paintings, such as "The Circus" and "The Models". In fact, where to apply the halos may be an artist's subjective decision. Therefore, our approach currently does not simulate the halos in Seurat's paintings. However, our approach can still assist a user to simulate halos by generating pointillism paintings from input images with manually-edited halos. For example, Figure 13 shows the synthesized result of an input image with halos around the cat's back. This actually provides a more flexible way for users to control where halos should be added. In the future, we would like to further investigate an automatic halo synthesis method in pointillism paintings.

As described in [Wei10], multi-class blue noise sampling is required to distribute color points in pointillism paintings by regarding each color channel as a separate class. In this paper, we instead treat a learned composite color in a segmented region as a separate class. The R -matrices acquired from Seurat's paintings also reflect that multi-class blue noise sampling is needed as most of the off-diagonal terms of the R -matrices are nonzero.

6. Conclusions

This paper presents a simple and intuitive algorithm to synthesize a pointillism painting from an arbitrary input image. In the learning stage, we statistically analyze the color composition in real paintings by artists, especially by Seurat, to learn a color model. This model thus can be applied to convert the input image into Seurat's style using the modified multi-class blue noise sampling method. Moreover, the shape of a point is allowed to be changed into an ellipse, and then oriented according to the gradients in the input image. Our user study also particularly shows that the proposed method can synthesize pointillism paintings whose color composition is perceptually similar to Seurat's style.

In the future, we are interested in a more accurate and robust technique for analyzing the color composition of real pointillism paintings. The color distributions of Seurat's paintings may be imitated even better based on a color space

other than RGB, such as CIELAB or $l\text{-}\alpha\text{-}\beta$. The influence of the selected color space needs to be further studied. Moreover, learning a shape and orientation model for color points in Seurat's paintings is a possible research direction. We also would like to investigate automatic approaches for estimating image segmentation parameters in our framework.

Acknowledgments

The authors would like to thank the anonymous reviewers for profound suggestions, and Prof. James Hays for providing some of his results. This work was supported in part by the Taiwan National Science Council under Grant Number: NSC101-2628-E-009-021-MY3, and the UST-UCSD International Center of Excellence in Advanced Bioengineering sponsored by the Taiwan National Science Council IRICE Program under Grant Number: NSC-101-2911-I-009-101.

References

- [AML10] ARROYO G., MARTÍN D., LUZÓN M. V.: A stochastic approach to simulate artists behaviour for automatic felt-tipped stippling. In *Proc. of IEEE CEC '10* (2010), pp. 1–8. [1](#), [2](#)
- [AZ05] ASLAN B., ZECH G.: Statistical energy as a tool for binning-free, multivariate goodness-of-fit tests, two-sample comparison and unfolding. *Nucl. Instrum. Methods Phys. Res. A* 537, 3 (2005), 1–8. [6](#)
- [Bay73] BAYER B. E.: An optimum method for two level rendition of continuous tone pictures. In *Proc. of IEEE ICC '73* (1973), pp. 2611–2615. [1](#)
- [CM02] COMANICIU D., MEER P.: Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 24, 5 (2002), 603–619. [4](#), [5](#)
- [Coo86] COOK R. L.: Stochastic sampling in computer graphics. *ACM Trans. Graph.* 5, 1 (1986), 51–72. [3](#)
- [CSHD03] COHEN M. F., SHADE J., HILLER S., DEUSSEN O.: Wang tiles for image and texture generation. *ACM Trans. Graph.* 22, 3 (2003), 287–294. [3](#)
- [DHvOS00] DEUSSEN O., HILLER S., VAN OVERVELD C. W. A. M., STROTHOTTE T.: Floating points: A method for computing stipple drawings. *Comput. Graph. Forum* 19, 3 (2000). [1](#), [2](#)
- [EF01] EFROS A. A., FREEMAN W. T.: Image quilting for texture synthesis and transfer. In *Proc. of ACM SIGGRAPH '01* (2001), pp. 341–346. [1](#)
- [FS75] FLOYD R. W., STEINBERG L.: An adaptive algorithm for spatial gray scale. In *Proc. of SID Dig. Tech. Papers '75* (1975), pp. 36–37. [1](#), [2](#)
- [Gla98] GLASSNER A. S.: Penrose tiling. *IEEE Comput. Graph. Appl.* 18, 4 (1998), 78–86. [2](#), [3](#)
- [HE04] HAYS J., ESSA I. A.: Image and video based painterly animation. In *Proc. of NPAR '04* (2004), pp. 113–120. [1](#), [2](#), [6](#), [7](#), [8](#)
- [Her98] HERTZMANN A.: Painterly rendering with curved brush strokes of multiple sizes. In *Proc. of ACM SIGGRAPH '98* (1998), pp. 453–460. [1](#), [2](#), [6](#), [7](#)
- [HHD03] HILLER S., HELLWIG H., DEUSSEN O.: Beyond stippling - methods for distributing objects on the plane. *Comput. Graph. Forum* 22, 3 (2003), 515–522. [1](#)
- [HJO*01] HERTZMANN A., JACOBS C. E., OLIVER N., CURRESS B., SALESIN D.: Image analogies. In *Proc. of ACM SIGGRAPH '01* (2001), pp. 327–340. [1](#), [2](#)
- [JH05] JANG S., HONG H.-K.: Stippling technique based on color analysis. In *Proc. of PCM '05, Part 2* (2005), pp. 782–793. [1](#), [2](#), [3](#), [5](#)
- [JJN76] JARVIS J. F., JUDICE C. N., NINKE W. H.: A survey of techniques for the display of continuous tone pictures on bilevel displays. *Comput. Graph. Image Process.* 5, 1 (1976), 13–40. [1](#)
- [KCODL06] KOPF J., COHEN-OR D., DEUSSEN O., LISCHINSKI D.: Recursive wang tiles for real-time blue noise. *ACM Trans. Graph.* 25, 3 (2006), 509–518. [2](#), [3](#), [5](#)
- [KMI*09] KIM S., MACIEJEWSKI R., ISENBERG T., ANDREWS W. M., CHEN W., SOUSA M. C., EBERT D. S.: Stippling by example. In *Proc. of NPAR '09* (2009), pp. 41–50. [2](#)
- [KSL*08] KIM D., SON M., LEE Y., KANG H., LEE S.: Feature-guided image stippling. *Comput. Graph. Forum* 27, 4 (2008), 1209–1216. [2](#)
- [LD05] LAGAE A., DUTRÉ P.: A procedural object distribution function. *ACM Trans. Graph.* 24, 4 (2005), 1442–1461. [3](#), [5](#)
- [Llo82] LLOYD S. P.: Least squares quantization in pcm. *IEEE Trans. Inf. Theory* 28, 2 (1982), 129–136. [3](#)
- [LSKL05] LUONG T.-Q., SETH A., KLEIN A., LAWRENCE J.: Isoluminant color picking for non-photorealistic rendering. In *Proc. of GI '05* (2005), pp. 233–240. [1](#), [2](#), [6](#), [7](#)
- [LSRY10] LEE H., SEO S., RYOO S., YOON K.: Directional texture transfer. In *Proc. of NPAR '10* (2010), pp. 43–48. [1](#), [2](#)
- [MALI10] MARTÍN D., ARROYO G., LUZÓN M. V., ISENBERG T.: Example-based stippling using a scale-dependent grayscale process. In *Proc. of NPAR '10* (2010), pp. 51–61. [2](#)
- [Mou07] MOULD D.: Stipple placement using distance in a weighted graph. In *Proc. of CAE '07* (2007), pp. 45–52. [2](#)
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph.* 23, 3 (2004), 488–495. [2](#), [3](#)
- [Ost01] OSTROMOUKHOV V.: A simple and efficient error-diffusion algorithm. In *Proc. of ACM SIGGRAPH '01* (2001), pp. 567–572. [1](#), [2](#)
- [RP11] ROGULJIC D., PAPIC V.: Procedure for non-photorealistic rendering in pointillist style. In *Proc. of ELMAR '11* (2011), pp. 111–114. [1](#), [2](#)
- [Sec02] SECORD A.: Weighted voronoi stippling. In *Proc. of NPAR '02* (2002), pp. 37–43. [1](#), [2](#)
- [SWHS97] SALISBURY M., WONG M. T., HUGHES J. F., SALESIN D.: Orientable textures for image-based pen-and-ink illustration. In *Proc. of ACM SIGGRAPH '97* (1997), pp. 401–406. [1](#), [2](#)
- [SY10] SEO S., YOON K.: Color juxtaposition for pointillism based on an artistic color model and a statistical analysis. *Visual Comput.* 26, 6–8 (2010), 421–431. [1](#), [2](#), [9](#)
- [Wei10] WEI L.-Y.: Multi-class blue noise sampling. *ACM Trans. Graph.* 29, 4 (2010). [2](#), [3](#), [5](#), [6](#), [9](#)
- [YCLL08] YAN C.-R., CHI M.-T., LEE T.-Y., LIN W.-C.: Stylized rendering using samples of a painted image. *IEEE Trans. Vis. Comput. Graph.* 14, 2 (2008), 468–480. [1](#), [2](#)
- [Yel83] YELLOTT J. I.: Spectral consequences of photoreceptor sampling in the rhesus retina. *Science* 221, 4608 (1983), 382–385. [3](#)
- [YY08] YANG C.-K., YANG H.-L.: Realization of seurat's pointillism via non-photorealistic rendering. *Visual Comput.* 24, 5 (2008), 303–322. [1](#), [2](#), [6](#), [7](#), [8](#), [9](#)