Stylized rendering for multiresolution image representation

Mark Grundland, Chris Gibbs and Neil A. Dodgson

Computer Laboratory, University of Cambridge, 15 J. J. Thomson Ave., Cambridge CB3 0FD, UK

ABSTRACT

By integrating stylized rendering with an efficient multiresolution image representation, we enable the user to control how compression affects the aesthetic appearance of an image. Adopting a point-based rendering approach to progressive image transmission and compression, we represent an image by a sequence of color values. To best approximate the image at progressive levels of detail, a novel, adaptive farthest point sampling algorithm balances global coverage with local precision. Without storing any spatial information apart from the aspect ratio, the spatial position of each color value is inferred from the preceding members of the sampling sequence. Keeping track of the spatial influence of each sample on the rendition, a progressively generated discrete Voronoi diagram forms the common foundation for our sampling and rendering framework. This framework allows us to extend traditional photorealistic methods of image reconstruction by scattered data interpolation to encompass non-photorealistic rendering. It supports a wide variety of artistic rendering styles based on geometric subdivision or parametric procedural textures. Genetic programming enables the user to create original rendering styles through interactive evolution by aesthetic selection. Comparing our results with JPEG, we conclude with a brief overview of the implications of using non-photorealistic representations for highly compressed imagery.

Keywords: Non-photorealistic rendering, multiresolution image representation, progressive image compression, point based rendering, farthest point image sampling, Voronoi diagram, Delaunay triangulation, genetic programming, design of graphical styles, digital art

1. INTRODUCTION

Every image has a grain. Whether a brush stroke of paint on canvas or an artifact of interpolation and compression, it is the telltale mark of the image rendering process. By allowing the viewer to surmise the extent to which the image may be taken literally, this surface texture is the crucial visual cue that mediates the scale at which an image ceases to be informative. By leaving the details to the imagination, it is an invitation for interpretation to take the place of observation. An image can express more than what it records. An artist often seeks to set the mood of a picture independently of its subject. Traditionally, an artist uses stylized rendering to set the stage for the reception of his work. For instance, cinematography has the convention of using grainier film stock to distinguish a flashback from the central narrative, and the film grain may need to be synthesized for special effects to be believable. By shaping the viewer's impression of an image, deliberate stylized rendering allows for visual artifacts to play a constructive role in visual communication. Contrast the expression of artistic intent exhibited by the evocative or decorative motifs of traditional techniques, such as watercolors or engraving, with the computational expediency reflected in the blurring or blocking artifacts of imaging algorithms, such as interpolation and compression. Effective graphic design calls for the right balance of visual fidelity and visual style. Hence, an image representation should offer an integrated approach to both of these fundamental concerns of visual communication. In image compression, there is a need to reconcile efficiency with aesthetics.

These considerations motivate our framework for stylized rendering of minimal data. Our aim is to give the user creative control over the stylization of a compressed image. We address creative challenges faced by the graphic designer, the person responsible for the effective presentation of visual information. Through the choice of rendering style, the graphic designer can ensure that the image conveys an impression appropriate to its purpose and context. We investigate how an efficient, multiresolution image representation could be made to support diverse styles of presentation, encompassing both photorealistic image reconstruction (PR) and non-photorealistic image rendering (NPR). The task requires close cooperation of representation and stylization. Without it, if styling would be treated as a mere afterthought and naively applied on

Mark Grundland (Mark @ eyemaginary.com), Chris Gibbs (Gibbo @ cantab.net), and Neil A. Dodgson (Nad @ cl.cam.ac.uk).

A color version of this paper is available online: http://www.eyemaginary.com/Portfolio/Publications.html

Further author information: (Send correspondence to Mark Grundland)



Figure 1. Non-photorealistic image rendering using our adaptive sampling technique (3200 samples $\approx 2\%$).

top of normal image compression, the rendition would risk being inappropriately degraded by the information loss. Just as a painting can hardly be conveyed by describing the individual curves of its brush strokes, the rendering elements of conventional NPR techniques defy simple description and the entire rendition normally needs to be stored at the full resolution of the display device to avoid marring the artistic effect. By using an image representation suited to stylized rendering and rendering elements adapted to compact encoding, our approach provides for much more efficient storage of stylized image renditions. As an alternative to conventional means of compressing and rendering images, our method enables the design of novel image rendering styles that have the advantage of being fully compressible. A typical application would be the progressive display of multimedia presentations, where images are transmitted over a narrow bandwidth network to display devices with variable resolutions. With such an application in mind, we propose an image representation which is:

- Compact: Enables efficient lossless and lossy compression.
- Secure: Deters unauthorized access by scrambling the data.
- Progressive: Exhibits a smooth transition between multiple levels of detail, culminating in an exact reconstruction.
- Flexible: Supports diverse photorealistic reconstruction techniques and non-photorealistic rendering styles.
- Intentional: Allows the artist to creatively formulate novel rendering styles.

The paper is organized as follows. We first briefly introduce our method (Section 2) and then outline its background in related work (Section 3). We proceed to describe our image representation and the way in which it supports a variety of progressive sampling mechanisms (Section 4) as well as a broad range of NPR rendering styles (Section 5). Next, we show how an artist can use a simple interface to design new NPR styles (Section 6). Finally, we compare our representation with conventional image compression (Section 7), showing how our results produced dramatically different visual artifacts, and we briefly address the philosophical question of whether such representations are useful (Section 8).

2. OVERVIEW

Our image representation simply consists of a sequence of colors sampled from the original template image. Only the color value of each sample site is stored explicitly. The placement and the size of its mark on the image is determined from the information carried by the preceding samples. A Voronoi spatial partition keeps track of the sample sites and their image marks, facilitating the efficient calculation of neighborhood relationships between the sites. This data structure can also help sampling methods to identify suitable locations for subsequent sample sites. Our novel adaptive farthest point sampling strategy balances the requirement to uniformly sample the image with the desire to accurately capture its variations. The result is a multiresolution image representation consisting of progressive levels of detail. Compression can be achieved by truncating the sequence and by efficiently encoding the differences between predicted and sampled colors. For security, a password can initialize the sampling sequence by determining the placement of the first few sample sites.

With regard to NPR styles, we take the approach of searching for general design principles rather than treating the construction of each NPR effect as a special case. Our framework acts as a template, translating the design of novel NPR styles into the application of more established techniques, such as parametric procedural textures. We present a unified sampling and rendering framework (Figure 1) based on the Voronoi diagram and the Delaunay triangulation. The Voronoi diagram has already been shown as a basis for powerful photorealistic techniques of image reconstruction from scattered samples, such as natural neighbor interpolation. We further explore its role in the development of novel NPR methods. For a geometric style, the tessellation of image marks is refined through geometric subdivision and the resulting mesh is rendered by various styles of shading. Alternatively, a procedural style resembles a procedural texture that is parameterized by the image mark's color sample and the spatial configuration of the nearby sample sites. Finally, we show how genetic programming provides the graphic designer with the tools to express his creative intentions by constructing novel procedural rendering styles through the process of interactive evolution by aesthetic selection.

3. BACKGROUND AND RELATED WORK

3.1. Non-photorealistic Image Rendering

In his pioneering work, Haeberli¹ explored the approximation of a photographic image by a collection of paint brush strokes, basic drawing primitives parameterized by location, size, orientation, shape, color, and texture. While such a rendition may require far fewer brush strokes than pixels, a full description of the brush strokes is still too complicated to be efficiently compressed. Reducing the number of brush strokes through time-consuming optimization may not prove sufficient to permit competitive compression.² We therefore deliberately limit ourselves to storing only the sampled colors, with the other brush stroke properties defined implicitly through the sampling technique and rendering style. Previously, progressive image compression was not regarded as a primary goal for other interactive,³ automated⁴ or animated⁵ brush stroke rendering systems. They assumed that the full template image is available throughout the rendering process. Our approach to the graphic design of the compression and aliasing artifacts of color images is partly inspired by the geometric⁶ and textured⁷ screens used in halftoning to embed an expressive motif in the printed grain of a poster or a banknote.

We have researched interactive evolution⁸ by aesthetic selection as an interface for designing NPR styles. While this approach has been successfully applied in the synthesis of graphical objects,⁹ it has been far less explored for the design of graphical styles for figurative depiction. Usually the evolutionary process operates directly on the picture representation¹⁰ rather than constructing a reusable transformation. The closest direct precedent for our approach is the evolution of coloring functions for medical image fusion.¹¹ By reducing the design of NPR styles to the formulation of parametric procedural textures, we can create novel NPR styles by applying Sims' technique¹² for evolving procedural textures. Standard procedural coloring and texturing methods¹³ have proven valuable in multiresolution painting.¹⁴ In pen-and-ink drawing, there has been wide use of parametric procedural stroke textures dependent on tone,¹⁵ orientation,¹⁶ and display resolution.¹⁷ For scale-dependent illustrations, Salisbury et al.¹⁷ propose that compression could be achieved by sending a rendering algorithm along with a stroke texture and a subsampled grayscale image to a PostScript display device. It is interesting to recall that the bandwidth constraints of early image compression were the original motivation for the development of cartoon rendering by edge detection,¹⁸ validating one of the first practical applications for NPR techniques.

3.2. Voronoi Diagram and Delaunay Triangulation

Having deliberately limited ourselves to using only a progressive sequence of point color samples, we need to make efficient use of this scarce resource. Hence, unlike the brush stroke methods that allow their strokes to overlap, our image

representation is based on the Voronoi spatial partition.¹⁹ A Voronoi diagram (Figure 1, bottom center) is a proximity graph that subdivides the image plane by assigning each point in the plane to its nearest sample site. Each Voronoi polygon designates an image mark, the region of influence of a sample site on the rendition. The Voronoi polygon is the portion of the image plane that is closer to that site than to any other site. For the Euclidean distance metric, the Voronoi polygons have convex shapes. When sites with adjacent Voronoi polygons are connected by edges they form a dual graph, the Delaunay triangulation (Figure 1, bottom right). These geometric data structures help our sampling and rendering algorithms to keep track of the sample sites, their spatial configuration, and their nearest neighbor relationships.

Adaptive image reconstruction from scattered point samples is an important application for Voronoi diagrams.²⁰ Delaunay triangulations support a hierarchical image representation offering antialiasing,²¹ a technique readily applicable to our system. For photorealistic imaging, fast Gouraud shading can render Delaunay triangulations²¹ while the more accurate natural neighbor interpolation relies on Voronoi diagrams.²² In the context of NPR, Haeberli¹ rendered images as geometric tilings with only a small number of optimally placed Voronoi tiles. Voronoi diagrams have been subsequently used for cubist animation,⁵ stipple drawings,^{23–25} inlay mosaics,²⁷ and other ornamental tilings.²⁶ Whereas these previous rendering methods have been very much application specific, we propose two general techniques for designing novel rendering styles. The approaches for generating temporally coherent animations from a Voronoi diagram of a still frame²⁶ or a video sequence⁵ are easily applicable to our rendering techniques.

Optimal incremental algorithms¹⁹ generate a Voronoi diagram for N sites in $O(N \log N)$ worst-case running time. However, the overhead of maintaining the complex data structures required by these computational geometry algorithms may not be justified since our operations only involve the image pixels. Trading some flexibility for ease of implementation, we use a discrete Euclidean Voronoi diagram that maps each pixel to its nearest site. To calculate an entire discrete Voronoi diagram at once in time proportional to the area of the image, a fast distance propagation algorithm²⁸ visits most pixels only once. A simpler technique that requires no extra storage, Danielsson's classic scan line algorithm²⁹ uses just six distance comparisons per pixel. Alternatively, a discrete Voronoi diagram can be constructed using the z-buffer on standard 3D graphics hardware.³⁰ In our system, the samples are received sequentially. To incrementally insert a new sample site into the discrete Voronoi diagram, we trace the perimeter of the new Voronoi polygon. The number of distance comparisons required is proportional to its perimeter while the number of pixel updates is proportional to its area.

The discrete Voronoi diagram has some minor pitfalls that require careful implementation. We impose the assumption that each discretized Voronoi polygon is simply connected. We can thus safely ignore the remaining approximation errors since they are confined to isolated pixels. On rare occasions, it may not be possible to exactly recover a Delaunay triangulation because of the difficulty of distinguishing a Voronoi vertex from an arbitrarily short Voronoi edge. For our approximation to the Delaunay triangulation, it is sufficient to connect a pair of sites by an edge whenever we can locally detect a shared edge between their discretized Voronoi polygons. Since the Delaunay triangulation only covers the convex hull of its sites, we need to extend the triangulation to the whole of the image rectangle. For each site whose Voronoi polygon intersects an image boundary, we project the site perpendicularly onto the boundary and a dummy site is placed there. Hence, it becomes straightforward to triangulate the resulting strip of trapezoids that frames the image rectangle.

4. SAMPLING

We tested a range of sampling strategies for progressive selection of the sample sites. We classify these as non-adaptive sampling, importance sampling, and adaptive sampling.

4.1. Non-adaptive Sampling

We start by investigating the possibilities of non-adaptive progressive sampling schemes. They are important as building blocks for adaptive methods. At a given resolution, a non-adaptive sampling can be pre-computed and stored as an array of pixel pointers. Either the sampling or its generating algorithm is assumed to be present in both the encoder and decoder so it does not form part of the image representation itself. Without preconceptions about the distribution of visually salient features in the template image, the same amount of information should be devoted to representing each part of the image. Hence, the sampling density should be uniform and isotropic. The sample sites should be spread evenly throughout the image, avoiding clustering by maintaining a minimum distance between samples. Assuming the correlation between pixels decreases with distance, for a sample to be most representative of its image mark, it should lie close to the centroid of its Voronoi polygon. This property can be enforced by Lloyd's optimization algorithm but it is difficult to apply directly in

incremental sampling.³¹ We used these criteria to select a non-adaptive sampling strategy to provide the candidate sites for our adaptive sampling scheme. We implemented a number of standard non-adaptive sampling schemes³² and we studied their effect on rendered images. They reflect different approaches to the inherent tradeoff between noise and aliasing.

Regular periodic grids,³² such as the familiar square lattice or the denser hexagonal lattice, spread samples out evenly. For multiresolution imaging, the order of sampling can be designated by a quad-tree or a recursively defined space-filling curve. Regular grids suffer from repetitive aliasing artifacts, made more distracting by being aligned along straight lines.

Self-similar non-periodic grids, implemented through the geometry of hierarchical substitution tilings³³ or the algebra of cut-and-project quasicrystals,³⁴ guarantee a minimum distance between samples while exhibiting a less monotonous pattern of recurrent aliasing artifacts. Their self-similar structure may be attractive to an artist looking to endow the rendition with a sense of decorative symmetry (Figure 3, top right).

The orderly appearance of any deterministic grid can always be concealed by pseudo-randomly perturbing each site,³² with the magnitude of the random displacement determining the minimum distance between samples. In our experiments, the resulting jittered grids performed similarly to quasirandom sampling methods,³² such as the Halton sequence, which seeks to uniformly spread out samples without any apparent pattern. However, in both approaches, the sample sites are often placed far from the centroids of their Voronoi polygons, which is undesirable. They are still a significant improvement over pseudo-random sampling from a uniform random distribution, which tends to cluster samples together, producing uneven image marks that give the rendition a grainy appearance.

The Voronoi diagram is instrumental in farthest point sampling.³⁵ It is the iterative strategy of sampling at the point of least information, which is taken to be the point farthest from all the previous sites. One starts by sampling the corners of the image rectangle and a few randomly chosen internal sites. When the intersections of the image rectangle with the edges of the Voronoi diagram are included as vertices, the farthest point is necessarily a vertex of the bounded Voronoi diagram. This is because a vertex in the Voronoi diagram is equidistant from the surrounding sites and no other site may be closer to it. It is possible to incrementally sample N sites in $O(N \log N)$ time by maintaining a balanced binary tree of Voronoi vertices ordered by their distance from the nearest site. Farthest point sampling is guaranteed to produce a uniformly distributed sample set. The minimum distance between sites is at least half of the maximum distance between any site and its nearest neighbor. For this reason, farthest point sampling appears to naturally place sites close to the centroids of their Voronoi polygons. These sample sets have been shown to have an isotropic power spectrum that mimics the blue noise spectrum of the Poisson disk distribution, making them ideally suited for antialiasing.

Farthest point sampling produced the best visual results of all the non-adaptive methods we tested. We recommend it as much for its deterministic performance as for its intuitive appeal. In addition, we observed that once about 20% of the template image's pixels have been sampled, there is little perceptible difference between the different non-adaptive sampling strategies. For a progressive rendering that culminates in an exact reproduction, the remaining pixels can easily be sampled either in lexicographic or pseudo-random order.

4.2. Importance Sampling

For selective emphasis, the user can provide a grayscale importance map over the template image to indicate the level of detail required.³ High importance should be assigned to distinguishing features, the boundaries between objects, and regions of inhomogeneous texture. We interpret the importance map by applying rejection sampling to filter candidate sites generated by one of the non-adaptive sampling strategies. The gray level of each pixel of the importance map is equated with the probability of accepting a sample site at that position. To cover the large-scale features of interest, it is sufficient to store a low-resolution importance map. For instance (Figure 3, bottom left), given a 400×400 template image, the importance map can be stored efficiently as a 32×32 icon image with just 16 gray levels.

4.3. Adaptive Sampling

Adaptive progressive sampling seeks the sample site that maximizes the perceptual similarity between the given template image and the emerging rendition. Taking advantage of the local style-dependent effect of a new sample site on the rendition, a simple adaptive procedure independently evaluates a sequence of alternative sites chosen according to a non-adaptive sampling strategy. At each iteration, it keeps the sample site that fosters the greatest improvement in the rendition. With each accepted color sample, it stores the number of preceding candidates that were skipped over. However, the extra effort and storage is unlikely to be justified once the progressive sampling has grown sufficiently dense.



Figure 2. Adaptive farthest point sampling techniques rendered with Voronoi and Gouraud shading (10485 samples $\approx 4\%$). On the left is the original adaptive farthest point sampling (PSNR = 27.11). On the right is our new adaptive sampling (PSNR = 26.70).

To require no extra storage, a progressive adaptive sampling strategy must base its choice of the next sample site entirely on the information contained in the preceding sites. By keeping track of the spatial distribution of sample sites, a Voronoi diagram helps in deciding where best to sample next. Darsa and Costa²¹ suggest that new samples should be placed to pseudo-randomly split either Delaunay edges exhibiting a large color difference or Voronoi polygons covering a large area. They do not describe a way to balance these criteria during the sampling process. An alternative is to consider Voronoi vertices, which are equidistant from their nearest sites and thus are natural candidates for locally uniform, adaptive sampling. The adaptive sampling scheme originally proposed for farthest point sampling³⁵ has its drawbacks (Figure 2, left). In this scheme, Voronoi vertices are selected to maximize the area of their local empty circle (i.e. the squared distance to the nearest sample site) weighted by a rough estimate of their local bandwidth. In our experiments, this strategy overly clusters samples in the vicinity of the high frequency details and sharp contours it uncovers, yielding too little discernable refinement once the sampling sufficiently progresses. Also, since a minimum local sampling density is not upheld, significant features may elude discovery for as long as their surroundings are deemed to have low bandwidth.

Our improved adaptive sampling criterion is designed to balance the need to uniformly sample the smooth tones of the template image with the desire to accurately capture its edges and details. Our stochastic approach (Figure 2, right) relies on basic robust statistics. Start by uniformly surveying the template image with non-adaptive farthest point sample sites, O. Subsequently, to determine the best next sample site at each iteration, first randomly select a small set of site candidates, C, from the vertices of the bounded Voronoi diagram of the preceding sites. The properties of this random subset are taken to be representative of the statistics of the entire population of Voronoi vertices. For each candidate, $i \in C$, find the squared distance, r_i^2 , to its nearest sample site and the luminance intensities, l_n , of its nearest neighboring sites, $n \in N_i$. Calculate the mean absolute luminance deviation of its neighborhood:

$$d_{i} = \frac{1}{\|N_{i}\|} \sum_{n \in N_{i}} |l_{n} - \mu_{l}| \quad \text{for} \quad \mu_{l} = \frac{1}{\|N_{i}\|} \sum_{n \in N_{i}} l_{n} \,. \tag{1}$$

As an estimate of variation, this dispersion measure is more resilient to outliers than variance or contrast. Next, apply robust z-scores to standardize the two criteria, r_i^2 and d_i , so that they may be sensibly compared. Scaling each criterion by its mean absolute deviation, robust z-scores give a relative measure of how far each criterion deviates from its mean:

$$\hat{r}_i^2 = \frac{r_i^2 - \mu_r}{\sigma_r} \quad \text{for} \quad \sigma_r = \frac{1}{\|C\|} \sum_{j \in C} |r_j^2 - \mu_r| \quad \text{and} \quad \mu_r = \frac{1}{\|C\|} \sum_{j \in C} r_j^2 \,, \tag{2}$$

$$\hat{d}_i = \frac{d_i - \mu_d}{\sigma_d} \quad \text{for} \quad \sigma_d = \frac{1}{\|C\|} \sum_{j \in C} |d_j - \mu_d| \quad \text{and} \quad \mu_d = \frac{1}{\|C\|} \sum_{j \in C} d_j .$$
 (3)

Finally, evaluate the combined scores, e_i , and select the candidate with the top combined score:

$$e_i = \min(\hat{r}_i^2, \hat{d}_i) + \lambda \max(\hat{r}_i^2, \hat{d}_i)$$
 (4)

This method of reconciling conflicting goals balances global coverage with local precision. If an importance map is available, it can be used to multiplicatively weigh the final scores. Given a z-score for each criterion, the final score gives greater weight to the lower z-score over the higher z-score according to the tradeoff parameter, $0 < \lambda < 1$. Hence, the tradeoff parameter dictates how high the higher z-score must be to dominate the influence of the lower z-score. Exceptionally high z-scores are bound to get noticed. This strategy moderates the undersampling of low frequency regions (high r_i^2 and low d_i) with the oversampling of high frequency regions (low r_i^2 and high d_i). Our experiments used the parameters, $\lambda = 0.25$, ||O|| = 256, ||C|| = 40, and $||N_i|| = 6$. A lower tradeoff parameter would result in crisper edges but coarser gradients.

When reconstruction error is measured by a sum of squared pixel differences, our method can exhibit slightly lower PSNR scores than the original adaptive farthest point sampling method, even when we detect visually prominent features that the original method misses entirely (for example, Figure 2, left, is missing a nostril). This is because our method purposefully devotes sample sites to maintaining a minimal sampling density even in regions where local variation has yet to be discovered. As the sampling progresses, the minimal local sampling density becomes uniformly denser throughout the image, and features at that resolution cannot elude discovery. Because of this minimal resolution guarantee for progressive image display, our adaptive sampling method represents an important improvement over the previous approach.

5. RENDERING

5.1. Geometric Rendering Styles

Delaunay triangles serve as the basic element of the geometric rendering styles (Figure 3, top). It is possible to apply flat shading, such as coloring according to the average color of the vertices,²¹ but this reduces the color contrast of the image. For a more accurate image approximation, apply Gouraud shading, which first linearly interpolates the colors along each edge and then linearly interpolates between edges across each horizontal scan line. As well as directly interpolating the colors, Gouraud shading can be used to generate decorative patterns³⁶ by interpolating indices to a color table that defines the order and smoothness of the color changes. We experimented with nonlinear interpolation between colors:

$$h(t) = 2^{\alpha - 1} t^{\alpha}$$
 when $0 \le t \le \frac{1}{2}$ and $h(t) = 1 - 2^{\alpha - 1} (1 - t)^{\alpha}$ when $\frac{1}{2} \le t \le 1$. (5)

Our "brush marks" style (Figure 6) uses this symmetric power curve to interpolate the colors along the triangle edges. For a "patchwork" effect (Figure 3, top left), it is also applied to nonlinearly interpolate the colors along each scan line.

In addition to shading, we can also apply geometric subdivision to the Delaunay mesh to construct the tilings of our geometric rendering styles. The simplest cases are to subdivide the triangle into four by joining the midpoints of its sides or into six by the intersection of its altitudes, its medians, or its angle bisectors. The design of a geometric rendering style consists of subdividing each Delaunay triangle, assigning colors to either the newly created vertices of the subdivided tiling or to the newly created tiles themselves, and choosing a shading method for each newly created tile. For instance, our "inlay mosaic" rendering style (Figure 3, top right) is formed by joining the midpoints of the edges of each triangle to make three outer triangles and one inner triangle. The outer triangles are flat colored with the original colors of the associated samples while the central triangle is colored black. In this way, each sample site gives rise to a star-shaped polygon, while the black central triangles serve as grout between the tiles. The mosaic tiles appear as tightly packed as possible and their layout reflects the structure of the sampling. Farthest point sampling produces tiles of uniform size and individual shape to create a pebble mosaic while a self-similar grid yields a decorative tiling with a small set of possible tile shapes.

In another example, our "paint strokes" rendering style (Figure 5) uses the same subdivision as the "inlay mosaic". For each Delaunay edge, its midpoint color is set to be the least neutral color (farthest from neutral gray) of its two vertices. Like the "brush marks" style, nonlinear interpolation is used on the edges of the subdivided triangles while linear interpolation is used along the scan lines. Reminiscent of the bold, angular, layered strokes of a painting knife, the resulting style assumes that saturated, plausible colors are preferable to colors that appear faded due to interpolation. Different people will have different reactions to these styles and there is plenty of scope to create different styles in search of a particular expression.

5.2. Procedural Rendering Styles

Based on the Voronoi diagram, our procedural rendering styles (Figure 3, bottom) are inspired by photorealistic image reconstruction through the use of local filters centered at the sample sites.³² They have the expressive power of parametric procedural textures.¹³ The color, f(p), of pixel p is expressed as a weighted sum of the colors, $f(s_k)$, of its first K nearest neighboring sample sites, s_k , ordered by increasing Euclidean distance, $\Delta(s_k, s_1)$, from its nearest sample site s_1 :

$$f(p) = \frac{\sum_{k=1}^{K} \phi(p, s_k) f(s_k)}{\sum_{k=1}^{K} \phi(p, s_k)} \,. \tag{6}$$



Figure 3. Non-photorealistic image rendering styles. The top row shows geometric rendering styles, "patchwork" (6400 samples $\approx 4\%$) on the left and "inlay mosaic" (3200 samples $\approx 2\%$) on the right. The bottom row shows procedural rendering styles, "sponge painting" (6400 samples $\approx 4\%$) on the left and "color hatching" (12800 samples $\approx 8\%$) on the right. The mosaic style uses quasicrystal sampling, while the remaining styles rely on importance map sampling.

The neighborhood size, K, is chosen empirically to be large enough such that any further increase has negligible effect on the rendition, typically $10 \le K \le 40$; when K is too small, spurious discontinuities may appear along the edges of some Voronoi polygons. The design of a procedural rendering style is encapsulated by its filter function, $\phi(p, s_k)$. Usually, we use spatially invariant nonnegative filters that are constrained to act locally within the neighborhood: $\phi(p, s_k) \to \infty$ as $\Delta(p, s_k) \to 0$ while $\phi(p, s_k) \to 0$ as $\Delta(p, s_k) \to \Delta(s_1, s_K)$. The filter function $\phi(p, s_k)$ can depend on the distance and angle between the pixel and the site, on the site's index, and on its sampled color. Separate filters may be applied for luminance and chrominance. It is possible to extend this concept by considering color properties fitted to the entire neighborhood, such as the color gradient, but this is beyond the scope of this paper.

As with classical procedural textures, it is easy to build up diverse rendering styles out of simple functional components. For instance, weighting samples by inverse distance $\phi(p, s_k) = \Delta(p, s_k)^{-\alpha}$ renders Voronoi polygons with soft edges. For an effect akin to looking at an image through faceted glass, omit the nearest neighbor s_1 from the weighted sum, $\phi(p, s_1) = 0$, thereby subdividing each Voronoi polygon into regions corresponding to the second nearest neighbors of its sample site. Two more examples show the richness of this framework. For the "sponge painting" style (Figure 3, bottom left), the facetted glass style has been augmented by painting randomly chosen pixels with the color of their nearest sample site. For the "color hatching" style (Figure 3, bottom right), orientation is used. Its filter function is inversely related to the absolute difference between the Euclidean and Manhattan distances between the pixel and its neighboring site.

6. EVOLUTION

An authentic artistic technique needs to offer the capacity for original expression. Artists may not be satisfied with styles that come pre-packaged and ready to use, but rather require tools to create their own personal styles. Just as the primary concern of a painter is not the chemistry of paint, the artist should be in control of the rendering process without being required to grasp the complexity of how it works. NPR research has predominantly taken the strategy of emulation, finding a rendering algorithm specific to a traditional medium or style. Instead, we have developed algorithmic rendering styles that are distinctly digital, possibly appearing "painterly" but without any pretense to actual "painting".

Interactive evolution⁸ by aesthetic selection offers a viable user interface (Figure 4, bottom) for exploring novel rendering styles. For artists who are not mathematicians, the development of algorithmic rendering styles must be a process of discovery rather than invention. This "I-know-it-when-Isee-it" method of stochastic optimization employs the user's artistic judgment to evaluate solutions proposed by the system. Computer-assisted graphic design remains a creative task in that the requirements leave the form of the solution unspecified. Constructive design needs to augment combinatorial search. We chose genetic programming as a means for original expression because of its ability to explore an open-ended parameter space. A key advantage of our approach is that we can apply genetic programming directly to formulate the filter functions of our procedural rendering styles. A great variety of rendering styles (Figure 4, top) becomes easily accessible. The details of our rudimentary implementation follow the classic work of Sims.¹² A symbolic expression tree defines the filter function. Our expression trees have the following parameter leaves: angle, Euclidean distance, Manhattan distance,



Figure 4. Interactive evolution of image rendering styles.

constant, and pseudo-random. They have the following operator nodes: *cosine, add, divide, multiply, power, select, if, superellipse, Perlin noise, and texture map.* At each iteration, the algorithm randomly alters the parameters and operators of the expression tree. The resulting variations are then rendered and presented to the user who selects which one should survive to produce new mutations for further refinement. The user can also specify the desired rate of mutation. The process continues until the user is satisfied with the outcome. The challenge of interactive evolution is to converge on the user's intentions by continuously offering a variety of relevant alternatives. This is a creative feedback loop, as the rendition can affect the vision that shapes it. The approach raises some open research questions. Can the algorithm's suggestions be optimized by analyzing the artist's past selections? Can an algorithm adjust the rendering styles along perceptual dimensions, such as vibrant, rough, or curly? How should multiple rendering styles be combined in a single rendition?

7. COMPRESSION

The progressive nature of our sampling methods allows these NPR styles to be used for progressively rendering an image as it is received over a limited-bandwidth network. The nature of the representation also makes it a candidate for image



Figure 5. Progressive image display using the "paint strokes" rendering style with our adaptive sampling (1%, 2%, 4%, 8%, 16%, 32%).

compression. Sending a truncated sequence of samples is itself a form of image compression. In addition, our sampling methods are amenable to any lossy or lossless compression algorithm for a stream of scattered color samples. To establish a baseline for future improvement, we experimented with a simple scheme based on the Delaunay triangulation. Apart from the aspect ratio and an optional importance map, we only store the colors of our sample sites. For security purposes,²⁰ a password can be translated into the seed for the random number generator that helps to determine the first few sample site locations. The rendering style definition could also be embedded, if it is not already available to the renderer.

To construct the encoder, we start with an invertible predictive mapping. With the initial sample sites evenly surveying the template image, we keep the literal values of the first ||O|| = 256 color samples. By seeding the rendition with a few representative colors, we minimize the color distortion of later lossy compression. Regardless of the sampling method, we predict the color of the new sample from the three preceding sites that form its surrounding Delaunay triangle. We store the difference between the color sampled from the template image and the color predicted by linear interpolation. Color information is initially quantized as 5-bit color components in the perceptually uniform CIE-LAB color space, assuming the template image starts out with 16 bits of RGB color per pixel. We truncate and quantize the differences between the color component difference to the nearest power of 2. Finally, to perform symbol encoding we apply static Huffman compression. In this way, the number of bits required to encode each successive color difference is proportional to the frequency of its occurrence. We chose this encoding since the scattering of the sample sites tends to remove any correlation between them.

A lossless multiresolution image representation (Figure 5), a progressive sampling of the entire template image, requires exact storage of the color differences to preserve the color of each sample site when rendering its pixel. This is achieved by omitting the truncation and quantization step. The compressed image data can then be rendered with different styles.

At high compression ratios, our method compares well with the alternative strategy of storing the template image using conventional lossy compression and then either displaying the decompressed image directly or applying NPR to the decompressed image. We compare (Figure 6) our compression scheme to standard JPEG. With the image compressed to roughly the same size, we see that our encoding produces a perceptually more attractive rendering with less contouring and color loss. However, if our rendering styles are applied to the JPEG compressed image, effectively rendering the image twice, the result (Figure 6, top left) appears clearly degraded by the JPEG artifacts which are absent from our encoding.



Figure 6. Image compression using the "brush marks" rendering style (10485 samples $\approx 4\%$). In the top row, from left to right, the template image (512×512, 16 bit color) is compressed using standard JPEG (7.0K at 73:1) and then the JPEG image is rendered using quasirandom Halton sampling (without additional compression). The bottom row shows our combined sampling and compression algorithms rendering the template image with the same number of samples: first lossless quasirandom Halton sampling (12.8K at 40:1), next lossy quasirandom Halton sampling (6.1K at 84:1), and finally lossy adaptive sampling using our method (7.2K at 71:1).

This demonstrates the need for the image encoding scheme to take account of the image rendering method.

It is difficult to give any quantitative difference measures comparing the standard compression algorithms with our NPR method. Non-photorealistic renditions appeal to highly non-linear aspects of human vision. To assess the perceptual quality of our results, we cannot rely on standard measures, such as PSNR. For example, a rendering which uses the "sponge painting" style introduces pseudo-random noise into the final image. This is clearly going to decrease the signal to noise ratio but may well improve the human perception of the image. Applying models of human visual processing to evaluate non-photorealistic image representations is a research challenge that has only begun to be addressed.³⁷

8. DISCUSSION

When NPR is applied to image compression, it changes the criteria for deciding what to keep and what to discard. NPR has the potential to make the artifacts of partial sampling and color quantization appear more plausible in highly compressed images. For example, consider the way a portrait painter depicts hair with only a few broad brush strokes. Similarly, in our technique, regions of homogeneous texture with a high frequency component may not require a proportionally high sampling rate to be rendered believably. Since NPR appears to reduce the complexity of representation,³⁸ its potential for image compression has certainly been mentioned before, but without a practical automated implementation for color photographic illustrations.

Conventional image compression has been conceived as a fully automatic process, concerned with photorealistic reproduction quality as an objective property of the image alone, without individual consideration given to the purpose or context of its presentation. Existing image compression algorithms typically optimize the encoding of an image with respect to a predetermined rendering technique. The graphic designer can only control the rate of data loss but not its visual consequences. Our approach provides an efficient image representation designed to support diverse styles of presentation. Moreover, our image representation does not presuppose the style of its presentation, photorealistic or not. Analogous to the way word processors distinguish text from font, our image representation explicitly separates the description of image content from the specification of image style so that each one can be saved, changed and reused independently.

Directing the viewer's attention through stylized presentation allows for selective emphasis. Deliberate omission of detail can stimulate the viewer's interest. As reviewed by Strothotte and Strothotte,³⁹ psychological studies have found that ambiguity invites scrutiny thereby improving memory performance. Abstraction can serve to engage the imagination. By filling in the empty spaces between the brush strokes with projections of his own expectations, the viewer is drawn into the picture, possibly becoming more inclined to identify with its message. Rendering styles that encourage the viewer to complete the picture could be considered a powerful form of compression.

Every image has a resolution at which stylized depiction becomes inevitable. Our work is based on the idea that, when visible distortion due to compression, or aliasing due to interpolation cannot be avoided, its appearance should be determined by the designer and not the algorithm. Wherever imperfection cannot be hidden, the user should be given the option of putting it to good use. An image may be more likely to receive the benefit of the doubt when its appearance clearly manifests an intentional choice. Intentionality can make the difference between a visible artifact being regarded as an accidental mechanical flaw or an essential part of a picture's unique character. We give the graphic designer the tools required to express his creative intentions by developing personalized rendering styles that are specially suited for displaying compressed imagery. By exercising control over the compression artifacts, the designer may be able to improve the perceived visual quality of a highly compressed image by shifting the viewer's expectation from a photographic reproduction to an artistic expression.

9. CONCLUSION

Our technique gives the graphic designer the freedom to choose the rendering style of a compressed image. We present a straightforward approach to automated stylized rendering for use with progressive image compression, where a wide range of expressive image rendering styles may be generated from a common multiresolution image representation designed to support a compact, secure encoding. We have developed a novel adaptive farthest-point sampling algorithm and a novel framework for NPR styles. Clearly, these methods are not aimed at compression applications where objective visual fidelity is all that matters. In practice, they are most appropriate in contexts where images communicate ideas or illustrate narratives. NPR abandons the conventional goal of exact reproduction in pursuit of the evocative capacity for visual communication. As an alternative technique of image compression, it has the advantage of making abstraction and simplification appear legitimate.

ACKNOWLEDGMENTS

We wish to thank Malcolm Sabin, Carsten Moenning, Alan Blackwell, Peter Robinson, and Victor Ostromoukhov for their time, insight, and advice. The images are taken from FreeFoto.com and the Waterloo Brag Zone.

REFERENCES

- 1. P. Haeberli, "Paint by Numbers: Abstract Image Representations," in *Proceedings of SIGGRAPH*, pp. 207-214, 1990.
- 2. T. Sziranyi and Z. Toth, "Random Paintbrush Transformation," in *Proceedings of 15th International Conference on Pattern Recognition*, **3**, pp. 151-154, 2000.
- 3. F. Durand, V. Ostromoukhov, M. Miller, F. Duranleau, and J. Dorsey, "Decoupling Strokes and High-Level Attributes for Interactive Traditional Drawing," in *Proceedings of Eurographics Rendering Workshop*, pp. 71-82, 2001.
- 4. A. Hertzmann, "Painterly Rendering with Curved Brush Strokes of Multiple Sizes," in *Proceedings of SIGGRAPH*, pp. 453-460, 1998.
- A. Klein, P.-P. Sloan, A. Colburn, A. Finkelstein, and M. F. Cohen, "Video Cubism," in *Microsoft Research Technical* Report MSR-TR-2001-45, 2001.
- 6. V. Ostromoukhov and R. D. Hersch, "Artistic Screening," in Proceedings of SIGGRAPH, pp. 219-228, 1995.
- 7. O. Veryovka and J. Buchanan, "Texture-Based Dither Matrices," Computer Graphics Forum, 19, 1, pp. 51-64, 2000.
- 8. H. Takagi, "Interactive Evolutionary Computation: Fusion of the Capabilities of EC Optimization and Human Evaluation," *Proceedings of the IEEE*, **89**, 9, pp. 1275-1296, 2001.
- 9. P. Bentley, Evolutionary Design by Computers, Morgan Kaufmann, San Francisco, 1999.

- 10. E. Baker and M. Seltzer, "Evolving Line Drawings," in Proceedings of Graphics Interface, pp. 91-100, 1994.
- 11. R. Poli and S. Cagnoni, "Genetic Programming with User-Driven Selection: Experiments on the Evolution of Algorithms for Image Enhancement," in *Proceedings of Genetic Programming*, pp. 269-277, 1997.
- 12. K. Sims, "Interactive Evolution of Equations for Procedural Models," Visual Computer, 9, 8, pp. 466-476, 1993.
- D. S. Ebert, F. K. Musgrave, D. Peachey, K. Perlin, and S. Worley, *Texturing & Modeling*, 2 ed, AP Professional, San Diego, 1998.
- 14. K. Perlin and L. Velho, "Live Paint: Painting with Procedural Multiscale Textures," in *Proceedings of SIGGRAPH*, pp. 153-160, 1995.
- M. P. Salisbury, S. E. Anderson, R. Barzel, and D. H. Salesin, "Interactive Pen-and-Ink Illustration," in *Proceedings* of SIGGRAPH, pp. 101-108, 1994.
- M. P. Salisbury, M. T. Wong, J. F. Hughes, and D. H. Salesin, "Orientable Textures for Image-Based Pen-and-Ink Illustration," in *Proceedings of SIGGRAPH*, pp. 401-406, 1997.
- M. Salisbury, C. Anderson, D. Lischinski, and D. H. Salesin, "Scale-Dependent Reproduction of Pen-and-Ink Illustrations," in *Proceedings of SIGGRAPH*, Computer Graphics, pp. 461-468, 1996.
- 18. D. E. Pearson and J. A. Robinson, "Visual Communication at Very Low Data Rates," *Proceedings of the IEEE*, **73**, 4, pp. 795-812, 1985.
- 19. F. Aurenhammer, "Voronoi Diagrams: a Survey of a Fundamental Geometric Data Structure," ACM Computing Surveys, 23, 3, pp. 345-405, 1991.
- N. Ahuja, B. An, and B. Schachter, "Image Representation Using Voronoi Tessellation," *Computer Vision, Graphics, & Image Processing*, 29, 3, pp. 286-295, 1985.
- L. Darsa and B. Costa, "Multiresolution Representation and Reconstruction of Adaptively Sampled Images," in *Proceedings of SIBGRAPI*, pp. 321–328, 1996.
- F. Anton, D. Mioc, and A. Fournier, "Reconstructing 2D Images with Natural Neighbour Interpolation," *Visual Computer*, 17, 3, pp. 134-146, 2001.
- O. Deussen, S. Hiller, C. van Overveld, and T. Strothotte, "Floating Points: A Method for Computing Stipple Drawings," *Computer Graphics Forum*, 19, 3, pp. C41-C50, 2000.
- 24. A. Secord, "Weighted Voronoi Stippling," in *Proceedings of the Second International Symposium on Non-photorealistic Animation and Rendering*, pp. 37-43, 2002.
- 25. S. Hiller, H. Hellwig, and O. Deussen, "Beyond Stippling: Methods for Distributing Objects on the Plane," in *Computer Graphics Forum*, **22**, 3, pp. 515-22, 2003.
- C. S. Kaplan, "Voronoi Diagrams and Ornamental Design," in *The First Annual Symposium of the International Society for the Arts, Mathematics, and Architecture*, pp. 277-283, 1999.
- 27. A. Hausner, "Simulating Decorative Mosaics," in *Proceedings of SIGGRAPH*, pp. 573 580, 2001.
- 28. I. Ragnemalm, "Neighborhoods for Distance Transformations Using Ordered Propagation," CVGIP: Image Understanding, 56, 3, pp. 399-409, 1992.
- 29. P. E. Danielsson, "Euclidean Distance Mapping," Computer Graphics & Image Processing, 14, 3, pp. 227-248, 1980.
- 30. K. E. Hoff, II, T. Culver, J. Keyser, L. Ming, and D. Manocha, "Fast Computation of Generalized Voronoi Diagrams Using Graphics Hardware," in *Proceedings of SIGGRAPH*, pp. 277-286, 1999.
- 31. S. Hiller, O. Deussen, and A. Keller, "Tiled Blue Noise Samples," in *Proceedings of Vision, Modeling and Visualization*, pp. 265-271, 2001.
- 32. A. S. Glassner, Principles of Digital Image Synthesis, 1, Morgan Kaufmann, San Francisco, 1995.
- C. Goodman-Strauss, "Aperiodic Hierarchical Tilings," in *Foams, Emulsions, and Cellular Materials*, Proceedings of NATO-ASI, E-354, pp. 481-496, Kluwer, 1999.
- 34. J. Patera, "Non-Crystallographic Root Systems and Quasicrystals," in *The Mathematics of Long-Range Aperiodic Order*, Proceedings of NATO-ASI, C-489, pp. 443-465, Kluwer, 1997.
- 35. Y. Eldar, M. Lindenbaum, M. Porat, and Y. Y. Zeevi, "The Farthest Point Strategy for Progressive Image Sampling," *IEEE Transactions of Image Processing*, **6**, 9, pp. 1305-1315, 1997.
- 36. H. Zhang, "Pattern Generation with Color Map Gouraud Shading," Computers & Graphics, 20, 1, pp. 157-162, 1996.
- 37. M. Cadik, "Human Perception and Computer Graphics," in *Czech Technical University Postgraduate Study Report* DC-PSR-2004-06, 2004.
- 38. I. Herman and D. Duke, "Minimal Graphics," *IEEE Computer Graphics & Applications*, **21**, 6, pp. 18-21, 2001.
- 39. C. Strothotte and T. Strothotte, Seeing between the Pixels: Pictures in Interactive Systems, Springer, Berlin, 1997.