

Stylized multiresolution image representation

Mark Grundland*

Chris Gibbs

Neil A. Dodgson

University of Cambridge

Computer Laboratory

15 J. J. Thomson Avenue

Cambridge CB3 0FD, United Kingdom

E-mail: mark@eyemaginary.com

Abstract. We integrate stylized rendering with an efficient multiresolution image representation, enabling a user to control how compression affects the aesthetic appearance of an image. We adopt a point-based rendering approach to progressive image transmission and compression. We use a novel, adaptive farthest point sampling algorithm to represent the image at progressive levels of detail, balancing global coverage with local precision. 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 nonphotorealistic 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. We compare our results with conventional compression, and we discuss the implications of using nonphotorealistic representations for highly compressed imagery. © 2008 SPIE and IS&T.

[DOI: 10.1117/1.2898894]

Digital imaging does not attempt to present a complete reconstruction of an object. Instead, it samples the object at rapid rates, reproducing just enough to create the illusion of a complete representation, not unlike an impressionist painting.

—Richard O'Donnell¹

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. It allows 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 observa-

tion. An image can thus express more than what it records. Artists often seek to set the mood of a picture independently of its subject. Traditionally, artists use stylized rendering to set the stage for the reception of their 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. Deliberate stylized rendering shapes the viewer's impression of an image, allowing 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. Conventional image representations assume that images are only meant to inform, neglecting the fact that images also seek to impress. 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. Traditional art and electronic imaging value the economy of expression in visual representation, conveying the most information with the least effort. In image compression,¹ there is a need to reconcile efficiency with aesthetics.

These considerations motivate our framework for the stylized rendering of minimal data (Fig. 1). Our aim is to give the graphic designer control over the aesthetic appearance of a compressed image. We address the creative challenge 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 can support diverse styles of presentation, encompassing both photorealistic image reconstruction (PR) and nonphotorealistic image rendering (NPR). The task requires close cooperation between representation and stylization. Otherwise, if styling is treated as a mere afterthought naively applied on 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 ele-

*To obtain a color version of this paper as well as additional illustrations and animations, please visit:

<http://www.eyemaginary.com/Portfolio/Publications.html>

Paper 06149R received Aug. 26, 2006; revised manuscript received Aug. 28, 2007; accepted for publication Sep. 18, 2007; published online Apr. 2, 2008. This paper is a revision of a paper presented at the SPIE conference on Human Vision and Electronic Imaging X, January 2005, San Jose California. The paper presented there appears (unrefereed) in SPIE Proceedings Vol. 5666.

1017-9909/2008/17(1)/013009/17/\$25.00 © 2008 SPIE and IS&T.

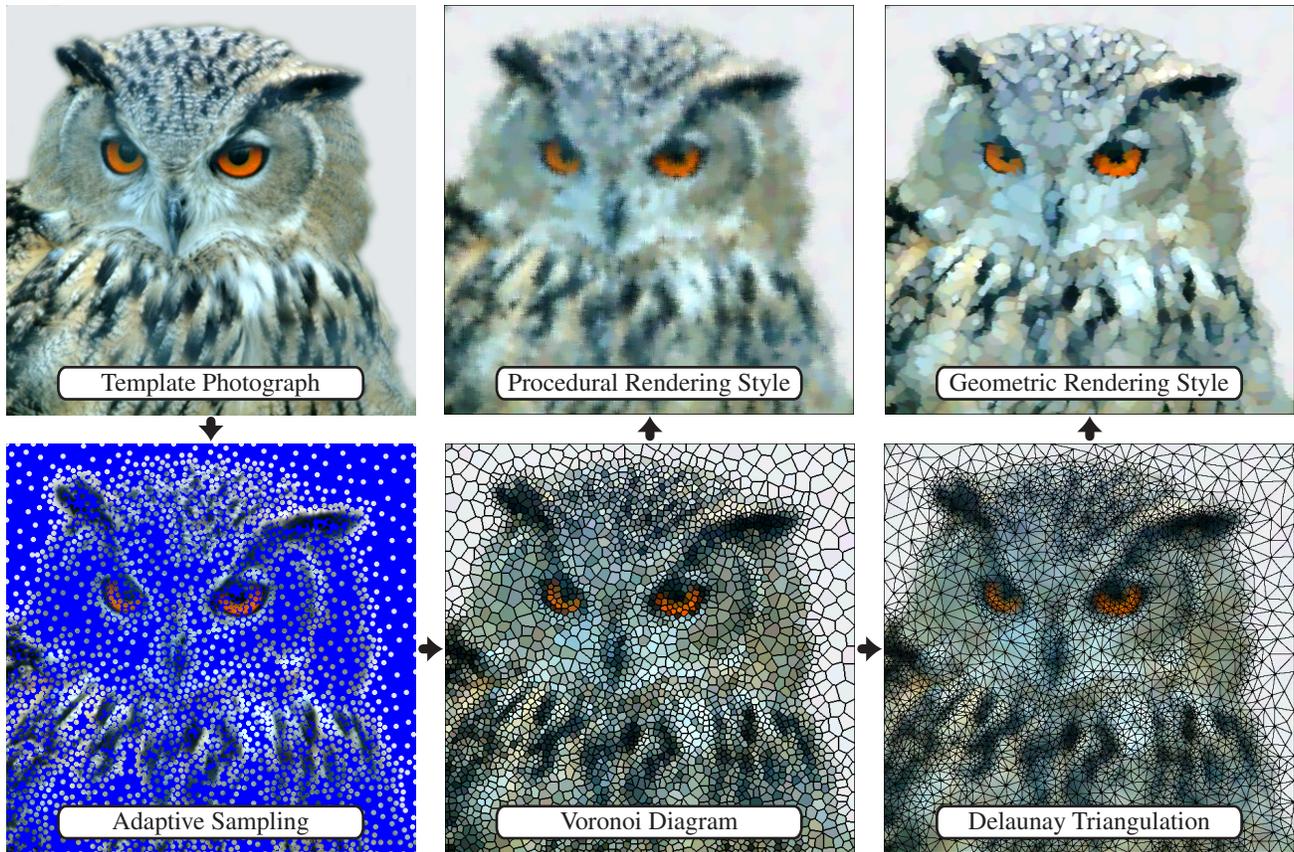


Fig. 1 Nonphotorealistic image rendering using our coverage adaptive sampling technique (3200 samples $\approx 2\%$).

ments of typical NPR techniques defy simple description. Indeed, a typical NPR rendition normally needs to be stored at the full resolution of the display device to avoid marring the artistic effect. We introduce an image representation suitable for stylized rendering and stylization techniques designed for compact encoding. Our approach thus supports 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. Typical applications include the progressive display of multimedia presentations, where images are transmitted over a narrow bandwidth network to display devices with variable resolutions, as well as the compact storage of picture collections, where the images are meant to be presented in a consistent artistic style. With such applications in mind, we propose an image representation that 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 reconstruc-

tion techniques and nonphotorealistic rendering styles.

- **Intentional:** Allows the artist to creatively formulate novel rendering styles.

This work gives an extended account of work presented at SPIE's Electronic Imaging 2005 conference.² We first briefly introduce our method (Sec. 2) and then outline its background and related work (Sec. 3). We describe our image representation and the way in which it supports a variety of progressive sampling mechanisms (Sec. 4) and a broad range of NPR styles (Sec. 5). Next, we show how a graphic designer can use a simple interface to design new NPR styles (Sec. 6). Finally, we compare our representation with conventional image compression (Sec. 7), showing how our results produced dramatically different visual artifacts, and we briefly discuss the implications of stylization for compact image representations (Sec. 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 shape, size, and placement of its region of influence on the rendition are all inferred 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 identify suitable locations for subsequent sample sites, according to the principle of sampling the center of the least known region of the image. Our novel adaptive farthest point sampling technique 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.

NPR research has predominantly taken the strategy of emulation, finding a rendering algorithm specific to a traditional medium or style. By contrast, we have developed algorithmic rendering styles that are distinctly digital, possibly appearing “painterly” but without any pretense to actual painting. 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 (Fig. 1) based on the Voronoi diagram and the Delaunay triangulation. The Voronoi diagram has already proven its utility in powerful photorealistic techniques of image reconstruction from scattered point samples, such as natural neighbor interpolation. We further explore its role in the development of novel NPR methods. Our approach takes advantage of the uncertainty inherent in interpolation to generate a variety of rendering styles. For a geometric style, the image mesh is refined through geometric subdivision, and the resulting tessellation of image marks 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 graphic designers with the tools to express their creative intentions by constructing novel procedural rendering styles through the process of interactive evolution by aesthetic selection. In this creative design process, the graphic designer contributes the indispensable aesthetic judgment to guide the evolution of his or her work in response to a selection of possible realizations presented by the program. This approach neatly combines the strengths of the two parties: the aesthetic judgement of the human and the computational power of the computer.

3 Background and Related Work

3.1 Nonphotorealistic Image Stylization

Our automated image rendering technique is based on painting with brush strokes. This popular NPR approach is the basis for numerous commercial artistic tools, such as Studio Artist,³ Corel Painter,⁴ and Piranesi.⁵ Pioneered by Haerberli⁶ and recently surveyed by Hertzmann,⁷ this approach represents a photographic image by a collection of brush strokes, basic rendering primitives parameterized by location, size, orientation, shape, color, texture, and opacity. The template image is sampled, the image samples deter-

mine the placement and properties of the brush strokes, and these brush strokes are composited to generate the stylized rendition. Many interactive,⁸ automated,⁹ and animated¹⁰ brush stroke rendering algorithms can produce a multiresolution image representation by painting a sequence of layered brush strokes. Although they can potentially be used for progressive image display, these systems were not designed for use in progressive image compression, since they generally assume that the full resolution template image is available throughout the rendering process. There are various alternative methods for image stylization, including halftoning,¹¹ texture transfer,¹² and image analogies,¹³ but they are likewise unsuitable for progressive image compression because they too rely on access to the full resolution template image. Our research into the graphic design of the compression and interpolation artifacts of color images is inspired by the artistic screens¹⁴ used to embed an expressive motif in the printed grain of a poster or a banknote.

We apply interactive evolution¹⁵ by aesthetic selection as a user interface for designing NPR styles. At each iteration of our interactive optimization technique, a genetic programming algorithm presents a selection of solutions to the user, who then subjectively evaluates their fitness. Interactive evolution has been successfully applied in the synthesis of graphical objects,¹⁶ especially decorative and abstract art.¹⁷ Normally in computer graphics, the evolutionary process directly transforms the contents of the image¹⁸ rather than producing a reusable image transformation. However, in medical imaging applications, interactive evolution has been used to construct transfer functions for volume rendering¹⁹ and coloring functions for image fusion.²⁰ A precedent for our approach can be found in Dalton’s work on NPR style design using automated optimization by neural networks and fuzzy logic,²¹ as well as interactive optimization by genetic algorithms,²² leading to the development²³ of Studio Artist.³ While Dalton originally applied genetic algorithms to tune the numeric parameters of image processing filters,²² we instead rely on genetic programming to construct the mathematical expressions that define our procedural rendering styles. In this way, we enable constructive design to augment combinatorial search. By reducing the development of NPR styles to the formulation of parametric procedural textures, we make it possible to create novel NPR styles by applying Sims’ genetic programming technique²⁴ for evolving procedural textures. Standard procedural coloring and texturing methods²⁵ have long proven valuable in multiresolution painting systems.²⁶ By relying on procedural rendering primitives, our system can render stylized images at any desired output resolution.

3.2 Nonphotorealistic Image Compression

Historically, image and video compression was the first successful, practical application of automated NPR systems. The bandwidth constraints of early imaging systems motivated the development of automated cartoon rendering sketched using edge detection. It enabled efficient coding of black-and-white²⁷ and grayscale²⁸ images for transmission at very low bit rates. For example, in 1985, cartoon rendering made possible a visual communication system for the deaf that worked in real time over telephone lines.²⁹

Image stylization in most of these early systems was limited to simple colorings of edge map regions. As such a paint-by-numbers rendering style can appear artificial, early image quantization³⁰ and compression³¹ systems sometimes added random noise to the picture to mask visual artifacts and restore a semblance of image grain. With increasing communication bandwidth and storage capacity, the visual simplicity of cartoon rendering became unnecessary for real-time compression performance, and stylized rendering was replaced by photorealistic reconstruction. As a result, the relevance of image stylization to image compression, the relationship of aesthetics and abstraction to efficiency and fidelity, has not been widely studied. Where painterly visual artifacts were observed in conventional image compression algorithms, such as fractal compression,³² morphological compression,³³ and edge coding,³⁴ they have usually been regarded as little more than a curiosity. Instead, we treat them as an opportunity for the development of novel approaches to the efficient representation of visual information. Our image representation allows the graphic designer to adapt the appearance of a picture to match its purpose. In the wider context of image reproduction, similar aims are pursued by color gamut mapping techniques that take into account the user's rendering intent.³⁵

The ability of stylized rendering to simplify visual representations is important for minimal graphics.³⁶ This fundamental open problem in computer graphics and vision involves the design of visually pleasing depictions that exhibit the minimal complexity necessary to convey a message. Schmidhuber,³⁷ motivated by both visual aesthetics and information theory, created an interactive drawing technique based on a self-similar grid that enables simple sketches to have a compact encoding. The graftal approach³⁸ offers a concise, multiresolution representation for an interactive NPR system, which combines various styles of hand-drawn illustration with the understanding of the scene geometry required to render the correct level of detail. In an automated NPR system, it is always possible to apply image stylization to a compressed or downsampled image, as suggested for scale-dependent pen-and-ink illustrations.³⁹ However, stylization cannot make up for the loss of relevant image detail, and it is prone to exaggerating any pre-existing compression or interpolation artifacts. Applying lossy compression to stylized images tends to unacceptably degrade the artistic effect, as the intricate surface textures of many NPR effects often prove difficult to compress efficiently, especially when they are produced by a pseudorandom process. Similarly, existing lossless image encodings are ill suited for stylized depictions. Experiments⁴⁰ show that stylized renditions are most compactly encoded as brush stroke sequences. Unfortunately, an explicit description of all the brush stroke properties is usually too complex to allow for efficient compression. Even reducing the number of brush strokes through time-consuming optimization is insufficient to make a brush stroke image representation as efficient as conventional image encoding.⁴¹ For automated color image compression, the NPR representations have so far proven considerably less compact than comparable photorealistic encodings. By contrast, our approach offers a concise, multiresolution image representation that is equally well suited to both stylized rendering and photorealistic reconstruction. Our

framework reduces a brush stroke image representation to just a sequence of pixel colors, which is carefully selected to give the best image approximation at each level of detail.

3.3 Voronoi Diagram and Delaunay Triangulation

Having deliberately limited ourselves to only working with 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.⁴² Each Voronoi polygon designates an image mark, the region of influence of a sample site on the rendition. A Voronoi diagram (Fig. 1, bottom center) is a proximity graph that subdivides the image plane by assigning each point in the plane to its closest sample site. A Voronoi polygon of a sample site is the region that is closest to its site. A Voronoi edge is equidistant to its two closest sites, while a Voronoi vertex is equidistant to three or more of its closest sites. For the Euclidean distance metric, the Voronoi polygons have convex shapes. If sample sites with adjacent Voronoi polygons are connected by edges, they form a dual graph, the Delaunay triangulation (Fig. 1, bottom right). These geometric data structures help our sampling and rendering algorithms to efficiently keep track of the sample sites, their spatial configuration, and their nearest neighbor relationships.

Adaptive image reconstruction is an important application for Voronoi diagrams.⁴³ Spatial partitions are used in a variety of methods for photorealistic image reconstruction from scattered point samples.⁴⁴ Delaunay triangulations support a hierarchical image representation offering antialiasing,⁴⁵ a technique readily applicable to our system. Fast linear interpolation by Gouraud shading can render Delaunay triangulations,⁴⁵ and it has been used for image compression.^{46,47} The more accurate natural neighbor interpolation relies on Voronoi diagrams,⁴⁸ and it too has proven useful for image compression.⁴⁹ In the context of NPR, Haerberli⁶ rendered images as geometric tilings with only a small number of optimally placed Voronoi tiles. Temporally coherent animations can be generated from a Voronoi diagram of a still frame⁵⁰ or a video sequence,⁵¹ and these approaches can be easily adapted to support our rendering techniques. Various kinds of Voronoi diagrams have been used for cubist stylization,^{51,52} stipple drawings,^{53,54} and ornamental mosaics.⁵⁵ Whereas these previous rendering methods have been very much application specific, we propose two general techniques for designing novel rendering styles.

Optimal incremental algorithms⁴² 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 rendering operations only involve the image pixels. Hence, we use a discrete Euclidean Voronoi diagram that maps each pixel to its closest site. To calculate an entire discrete Voronoi diagram, in time proportional to the area of the image, a fast distance propagation algorithm⁵⁶ visits most pixels only once. Danielsson's classic scan line algorithm,⁵⁷ a simpler technique that requires no extra storage, uses just six distance comparisons per pixel. Alternatively, a discrete Voronoi diagram can be constructed using the z -buffer on standard

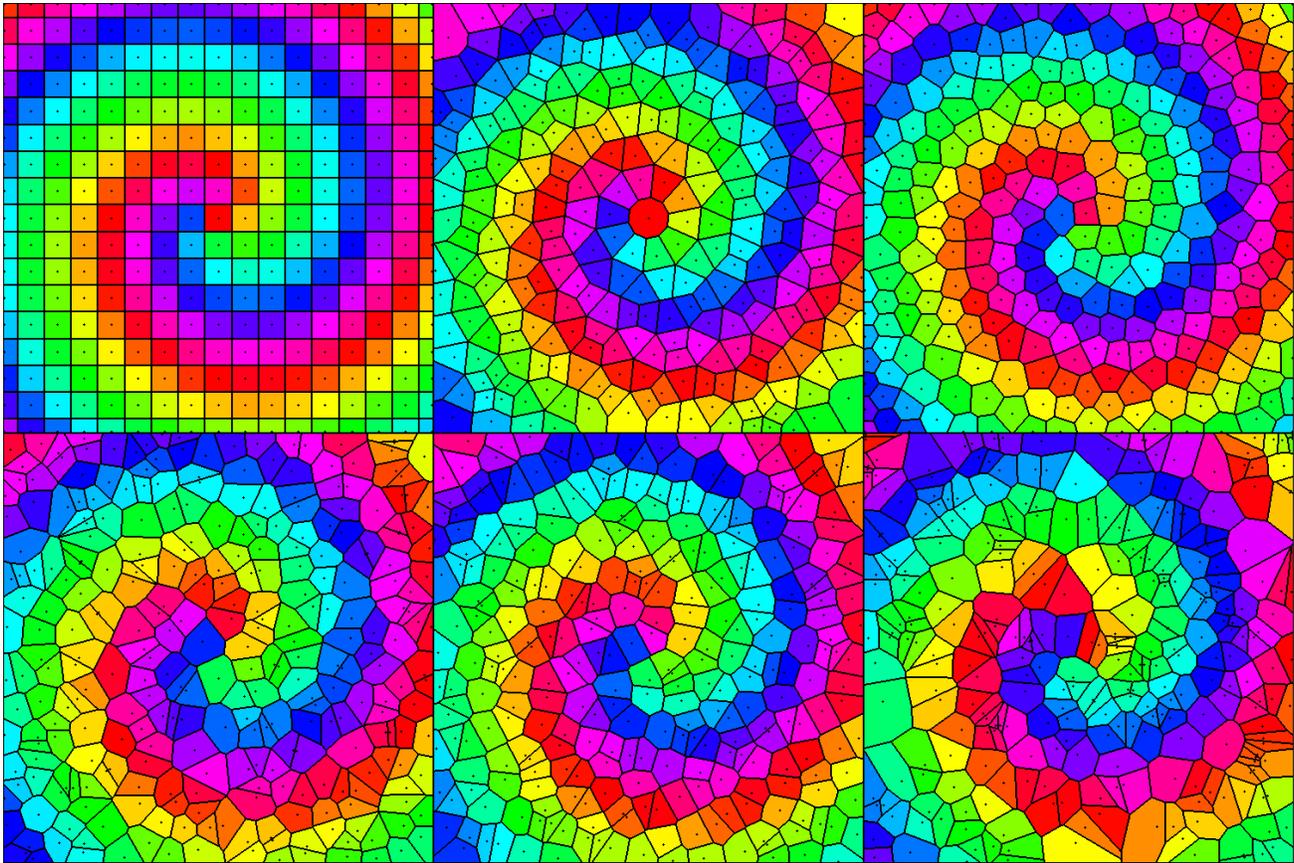


Fig. 2 Voronoi diagrams for nonadaptive sampling schemes. The top row, from left to right, shows periodic, nonperiodic, and farthest point sampling. The bottom row, from left to right, shows jittered, quasirandom, and random sampling.

3-D graphics hardware.⁵⁸ 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. The remaining approximation errors can be therefore safely ignored, 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 genuine Voronoi vertex from an arbitrarily short Voronoi edge. For our approximation to the Delaunay triangulation, it is sufficient to connect a pair of sample sites by an edge whenever we 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 duplicate site is placed there. The resulting strip of trapezoids, which frames the image rectangle, is straightforward to triangulate.

4 Sampling

We tested a range of sampling schemes for progressive selection of the sample sites. Nonadaptive sampling globally maintains a uniform resolution regardless of the image, importance sampling regionally adjusts its resolution to reflect the visual significance of image features, and adaptive sampling locally varies its resolution to capture the spatial distribution of image details. While nonadaptive and adaptive sampling are automatic, importance sampling enables the user to define regions of interest for the sampling process.

4.1 Nonadaptive Sampling

We start by reporting our investigations into the possibilities offered by nonadaptive progressive sampling schemes. These are important as building blocks for adaptive methods. At a given resolution, a nonadaptive sampling can be precomputed and stored as an array of pixel pointers. Either the sampling or its generating algorithm is assumed to be available to both the encoder and decoder, so it does not form part of the image representation itself. With no 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 number of sample sites placed in any region of the image should be proportional to its area, so

that the sampling density remains constant throughout the image. The distribution of sample sites should be uniform and isotropic while still allowing for a variety of different sample site configurations. Maintaining a minimum distance between sample sites serves to avoid clustering. Assuming that 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. We take all these considerations into account when choosing the nonadaptive sampling scheme that will provide site candidates for our adaptive sampling scheme.

We tested a number of nonadaptive sampling schemes⁵⁹ (Fig. 2). We used Voronoi diagrams to visualize the differences between them, and we studied their effect on rendered images. The various nonadaptive sampling schemes reflect different approaches to the inherent tradeoff between noise and aliasing. They offer different combinations of desirable properties, such as accurate reconstruction, progressive refinement, uniform coverage, isotropic distribution, blue noise spectrum, centroidal regions, and heterogeneous local configurations.

In regular periodic grids,⁵⁹ such as the familiar square lattice or the denser hexagonal lattice, the sample sites are spread out evenly. For multiresolution imaging, the order of sampling can be specified 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 nonperiodic grids are produced using the geometry of hierarchical substitution tilings⁶⁰ or the algebra of cut-and-project quasicrystals,⁶¹ which we used in our experiments. Nonperiodic tilings, such as Penrose tilings,⁶² Wang tilings,⁶³ and polyomino tilings,⁶⁴ have proven very useful in sampling. They guarantee a minimum distance between sample sites while exhibiting a less monotonous pattern of sample site configurations than periodic grids. The self-similar structure of some nonperiodic grids can be attractive to a graphic designer looking to endow the rendition with a sense of decorative symmetry (Fig. 6, top left).

The orderly appearance of any deterministic grid can always be concealed by randomly perturbing each site,⁵⁹ with the magnitude of the random displacement determining the minimum distance between sample sites. In our experiments, the resulting jittered grids produced visually inferior results compared to quasirandom sampling methods,⁵⁹ such as the Halton sequence, which seek to uniformly spread out sample sites without any apparent pattern. In both approaches, the sample sites are often placed far from the centroids of their Voronoi polygons, which is undesirable. They are still an improvement over random sampling from a uniform random distribution, which tends to cluster sample sites together, producing uneven image marks that give the rendition a grainy appearance.

Farthest point sampling⁶⁵ uses the Voronoi diagram to direct sample placement. 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 ver-

tex of the bounded Voronoi diagram. This is because, in a bounded Voronoi polygon, the vertices are the points farthest away from their closest sample site. 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 their closest sample sites. Farthest point sampling is guaranteed to produce a uniformly distributed sample set. The minimum distance between sites is provably at least half the maximum distance between any site and its closest neighbor. New sample sites are only placed at points equidistant to three or more of their closest sites. For this reason, farthest point sampling appears to naturally place sites close to the centroids of their Voronoi polygons. These sample sets are especially well suited for antialiasing, since they have been shown to have an isotropic power spectrum that mimics the ideal blue noise spectrum of the Poisson disk distribution.

In general, farthest point sampling (Fig. 3, top center) produced the best visual results of all the nonadaptive methods we tested. We recommend it for both its consistent performance and its intuitive appeal. The image marks produced by farthest point sampling tend to have similar shapes and sizes. If a greater variety of image mark shapes is desired, quasirandom Halton sampling (Fig. 3, top left) can simply be used instead. Once about 20% of the template image's pixels have been sampled, most of the remaining unsampled pixels will have a sampled pixel as one of their eight adjacent neighbors. Under these circumstances, we observed little perceptible difference between the different nonadaptive sampling schemes. For a progressive rendering that culminates in an exact reproduction, the remaining pixels can easily be sampled either in scan line or random order.

4.2 Importance Sampling

For selective emphasis, the user can specify a grayscale importance map for the template image to indicate the level of detail required. High importance is usually given to distinguishing features, boundaries between objects, and regions of inhomogeneous texture. In this way, a graphic designer can direct the system to capture foreground details with small image marks, while outlining background shading with large image marks. Normally, an importance map is interpreted by applying rejection sampling to filter site candidates generated by one of the nonadaptive sampling schemes. The gray level of each pixel of the importance map is equated with the probability of accepting a sample site at that location. In our framework, we extend farthest point sampling to support an importance map. At each turn, the algorithm inserts a new sample site at the Voronoi vertex that is the farthest scaled distance away from any existing sample site. We define the scaled distance between a Voronoi vertex and a sample site to be the squared Euclidean distance multiplied by the importance map value at the vertex. In effect, the importance map exerts a geometric distortion on the farthest point sampling process.

To cover the large-scale features of interest, it is sufficient to store a low-resolution importance map. In our examples, we stored the importance map as a 32×32 icon image with just 16 gray levels. Using such a low-resolution importance map (Fig. 3, top right) to guide nonadaptive sampling can dramatically improve the visual quality of the



Fig. 3 Sampling schemes rendered using “paint strokes” rendering style (6554 samples $\approx 2.5\%$). The top row, from left to right, shows nonadaptive sampling: quasirandom, farthest point, and importance-driven farthest point sampling. The bottom row, from left to right, shows adaptive sampling: bandwidth adaptive, coverage adaptive, and importance-driven coverage adaptive sampling.

rendition. Applying adaptive sampling in conjunction with an importance map (Fig. 3, bottom right) results in sharper edges and clearer details. In this approach, the user provides the aesthetic judgment and semantic understanding required to determine the global priorities for sampling, while the algorithm relies on geometric measurements and statistical analysis to determine the local placement of the sample sites. Without an importance map, adaptive sampling methods (Fig. 3, bottom left and bottom center) must rely entirely on local image properties, resulting in less focused renditions. While user-specified importance maps⁸ are a common feature of interactive NPR systems, automatic NPR systems have relied on computer-generated importance maps derived from color variance measures^{11,51} or perceptual salience models.⁵² As the encoding of an image representation is usually expected to proceed with minimal user intervention, similar automatic techniques for identifying regions of interest could be readily adapted for our purposes.

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. Its task is to summarize the picture by a point set. Taking advantage of the

local effect that a new sample site exerts on the rendition, a simple adaptive procedure can independently evaluate a sequence of site candidates provided by a nonadaptive sampling scheme. At each iteration, it selects the candidate that fosters the greatest improvement in the rendition. With each accepted color sample, it stores the number of preceding candidates that were skipped over. Alternatively, it is possible to adaptively place sample sites according to image features, such as edges,^{46,48} ridges and valleys,^{48,49} or least accurately approximated image regions.⁴⁷ However, in these techniques, storing the sample site positions reduces the space available for the sampled image values. As the sampling process progresses, the visual impact of optimal placement decreases, so that the extra effort and storage is unlikely to be justified.

Our adaptive sampling encodes an image solely through a sequence of its colors. To require no extra storage, a progressive adaptive sampling technique must base its choice of the next sample site entirely on the information contained in the preceding sites. To help decide where best to sample next, a Voronoi diagram keeps track of the spatial arrangement of the sample sites. Previous work⁴⁵ has suggested that new sample sites should be placed to randomly split either Delaunay edges exhibiting a large color difference or Voronoi polygons covering a large area, but no

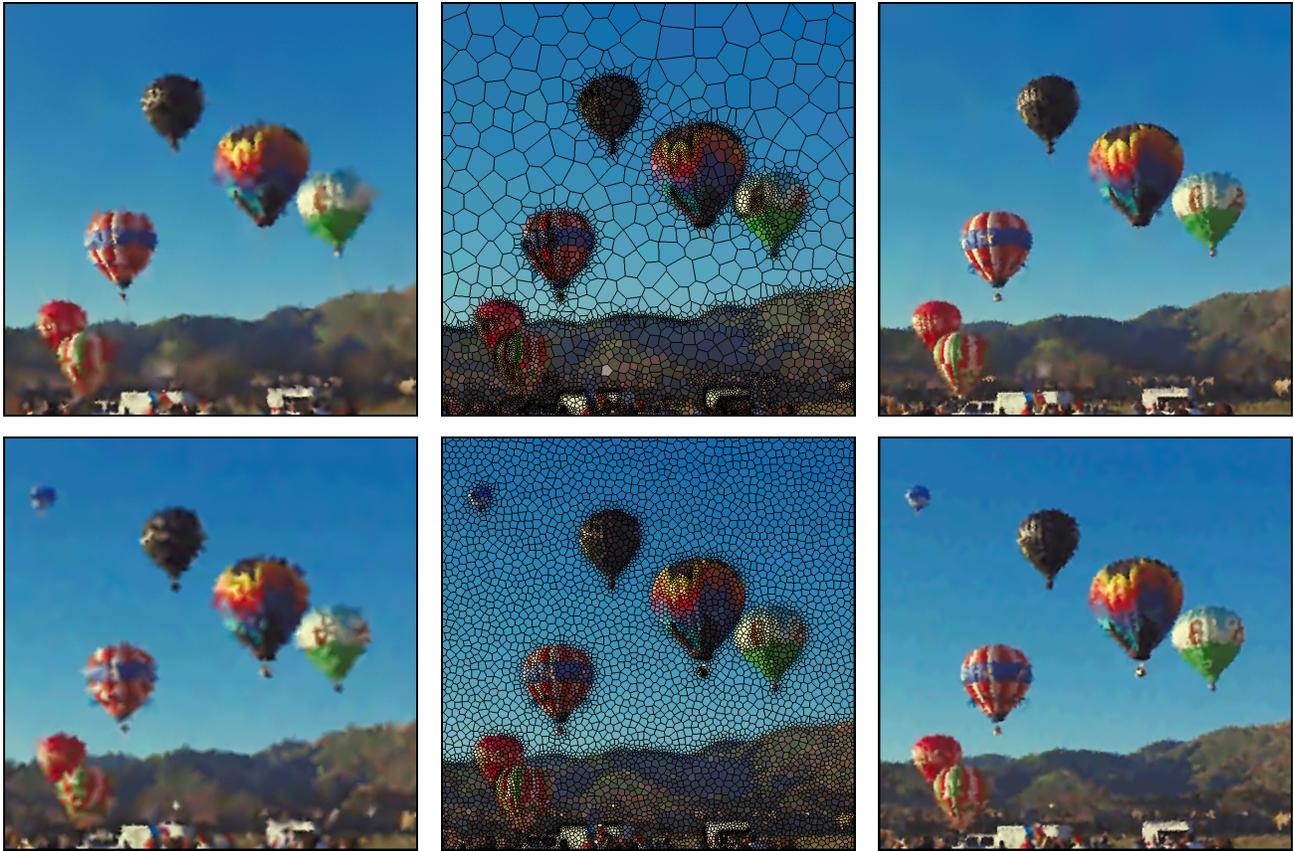


Fig. 4 Adaptive farthest point sampling techniques rendered with Voronoi and Gouraud shading. The left column displays a sparse sampling (2450 samples $\approx 2\%$), while the center and right columns display a denser sampling (7350 samples $\approx 6\%$). The top row shows bandwidth adaptive sampling (PSNR=24.31 and PSNR=26.58), while the bottom row shows our coverage adaptive sampling (PSNR=23.54 and PSNR=26.06). In the top row, note that the images are missing a balloon in the top left corner and a basket under the largest balloon.

explicit rule was given to help make the decision. Instead, adaptive farthest point sampling evaluates Voronoi vertices, which are natural candidates for locally uniform sampling, since they are the points farthest away from their closest sample sites. The bandwidth adaptive sampling scheme (Fig. 4, top row) originally proposed for farthest point sampling⁶⁵ has its drawbacks. In this scheme, Voronoi vertices are selected to maximize a rough estimate of their local bandwidth multiplied by their squared Euclidean distance to their closest sample sites, which is proportional to the unsampled circular area around them. In our experiments (Fig. 4), bandwidth adaptive sampling overly clusters sample sites in the vicinity of the high frequency details and high contrast contours it uncovers, yielding too little discernable refinement once the sampling has progressed sufficiently far. 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. When the rest of the image is seen to have sharply defined details, the viewer is apt to assume that the missing elements are absent from the picture.

Our coverage adaptive sampling (Fig. 4, bottom row) 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. To ensure global coverage, new sample sites should be placed in regions of low sampling density to uncover new image features. To ensure local precision, new sample sites should be placed in regions of high image frequency to refine previously uncovered image features. Our coverage adaptive sampling relies on basic robust statistics. Start by uniformly surveying the template image with nonadaptive farthest point sample sites O . Subsequently, to determine the best next sample site at each iteration (Fig. 5), 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 Euclidean distance r_i^2 to its closest sample sites and the luminance intensities l_n of its nearest neighboring sites $n \in N_i$. Scale its distance $w_i r_i^2$ by an optional importance map $w_i \geq 0$, which is assumed to be constant, $w_i = 1$, when not specified. Calculate the mean absolute luminance deviation d_i of its neighborhood:

$$d_i = \frac{1}{\|N_i\|} \sum_{n \in N_i} |l_n - \mu_{N_i}| \quad \text{for } \mu_{N_i} = \frac{1}{\|N_i\|} \sum_{n \in N_i} l_n. \quad (1)$$

For estimating variation, this dispersion measure is more resilient to outliers than statistical variance or Michelson contrast. Next, apply robust z -scores to standardize the two criteria \hat{r}_i^2 and \hat{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:

$$\hat{r}_i^2 = \frac{w_i r_i^2 - \mu_r}{\sigma_r} \quad \text{and} \quad \hat{d}_i = \frac{d_i - \mu_d}{\sigma_d}, \quad (2)$$

$$\sigma_r = \frac{1}{\|C\|} \sum_{j \in C} |w_j r_j^2 - \mu_r| \quad \text{and} \quad \mu_r = \frac{1}{\|C\|} \sum_{j \in C} w_j r_j^2, \quad (3)$$

$$\sigma_d = \frac{1}{\|C\|} \sum_{j \in C} |d_j - \mu_d| \quad \text{and} \quad \mu_d = \frac{1}{\|C\|} \sum_{j \in C} d_j. \quad (4)$$

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

$$e_i = \min(\hat{r}_i^2, \hat{d}_i) + \lambda \max(\hat{r}_i^2, \hat{d}_i). \quad (5)$$

This method of reconciling conflicting goals balances global coverage with local precision. 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. The top candidates must have either high z -scores for both criteria or an exceptionally high z -score for one of them. A single low z -score does not necessarily eliminate a candidate. Hence, the algorithm is able to select a candidate in a sparsely sampled area of the image, even when no local variation has yet been uncovered there. This strategy moderates the undersampling of low-frequency regions (high \hat{r}_i^2 and low \hat{d}_i) with the oversampling of high-frequency regions (low \hat{r}_i^2 and high \hat{d}_i). Our experiments use the parameters $\lambda=0.25$, $\|O\|=256$, $\|C\|=40$, and $\|N_i\|=6$. A lower tradeoff parameter would increase the sampling density in high contrast regions by decreasing it in low contrast regions.

When reconstruction error is measured by a sum of squared RGB color differences, our coverage adaptive sampling can exhibit slightly lower peak signal-to-noise ratio (PSNR) scores than the original bandwidth adaptive sampling method, even when we detect visually prominent features that the original method misses entirely (Fig. 4, top row, an entire balloon is missing). 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 uncovered. As the sampling progresses, the minimal local sampling density uniformly increases throughout the image, ensuring features with that resolution cannot elude discovery. This minimal resolution guarantee for progressive image display gives our coverage

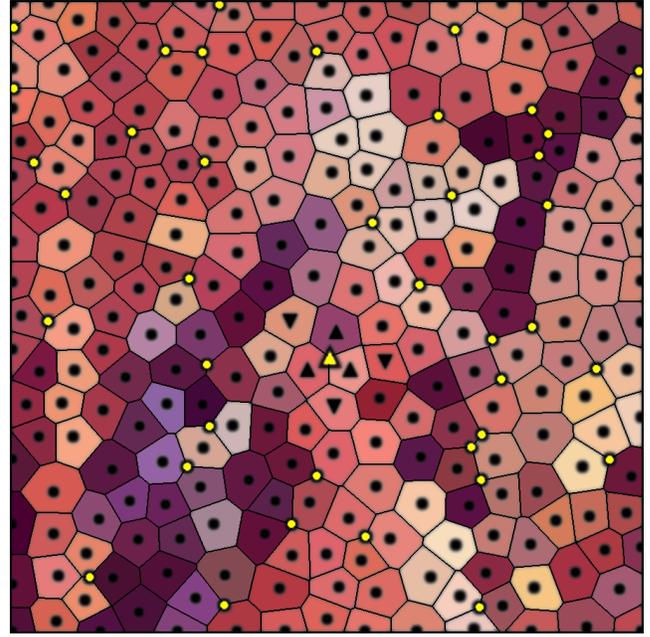


Fig. 5 Our coverage adaptive sampling and its Voronoi diagram. The adaptive sampling process starts from an initial set of sample sites O (●) obtained by farthest point sampling. At each iteration, site candidates C (○) are randomly selected from the vertices of the Voronoi diagram. Each candidate i (△) is evaluated according to robust z -scores, which provide relative measures of the spatial proximity \hat{r}_i^2 of its closest sample sites (▲) and the luminance variation \hat{d}_i of its nearest neighboring sample sites N_i (▲ and ▼).

adaptive sampling method a crucial advantage over the previous approach.

5 Rendering

5.1 Geometric Rendering Styles

Our geometric rendering styles (Fig. 6, top) are based on the Delaunay triangulation,⁴² which is known to provide an optimal spatial partition for piece-wise linear interpolation. The Delaunay triangles serve as our basic rendering elements. Applying flat shading, by coloring the triangles according to the average color of their vertices,⁴⁵ has the generally undesirable effect of reducing the color contrast of the image. For a more accurate image approximation, we rely on Gouraud shading^{44–47} (Fig. 4, left and right), 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 patterns⁶⁶ by interpolating indices to a color table that defines the order and smoothness of the color changes. We experimented with replacing the usual linear interpolation, $f(p) = tf(a) + (1-t)f(b)$, between colors $f(a)$ and $f(b)$, by nonlinear interpolation, $f(p) = h(t)f(a) + h(1-t)f(b)$, using a symmetric power curve $h(t)$:

$$h(t) = \begin{cases} 2^{\alpha-1} t^\alpha & \text{when } 0 \leq t \leq \frac{1}{2} \\ 1 - 2^{\alpha-1} (1-t)^\alpha & \text{when } \frac{1}{2} \leq t \leq 1 \end{cases}. \quad (6)$$

Our “brush marks” style (Fig. 9) uses this symmetric power curve to interpolate the colors along the triangle edges. For

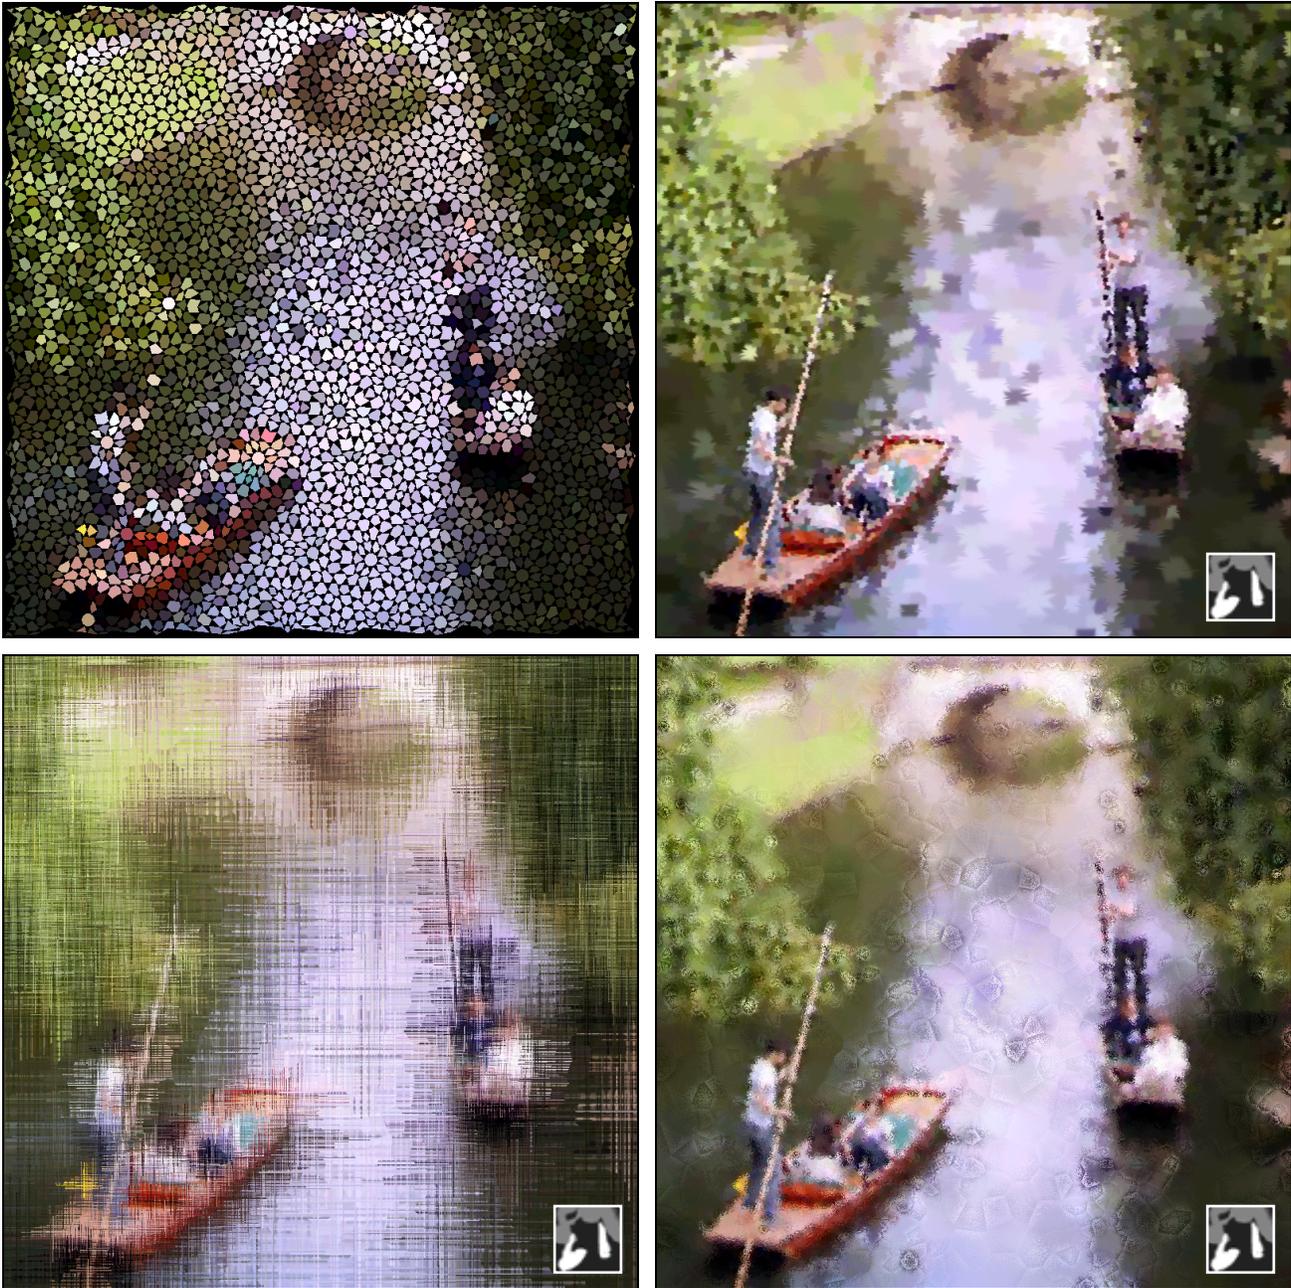


Fig. 6 Nonphotorealistic image rendering styles. The top row shows geometric rendering styles, “mosaic” (3200 samples $\approx 2\%$) on the left and “patchwork” (4800 samples $\approx 3\%$) on the right. The bottom row shows procedural rendering styles, “color hatching” (9600 samples $\approx 6\%$) on the left and “sponge painting” (4800 samples $\approx 3\%$) on the right. The mosaic style uses nonperiodic quasicrystal sampling, while the remaining styles rely on importance-driven coverage adaptive sampling.

a “patchwork” effect (Fig. 6, top right), it is also applied to nonlinearly interpolate the colors along each scan line.

In addition to shading, we also apply geometric subdivision to the Delaunay mesh to construct the tilings of our geometric rendering styles. Each Delaunay triangle can be subdivided either into four triangles by joining the midpoints of its sides or into six triangles by the intersection of its altitudes, its medians, or its angle bisectors. In our examples, the same geometric subdivision is performed once

for all the Delaunay triangles, and the resulting tiles are then shaded using a combination of flat, linear, and nonlinear shading. The design of a geometric rendering style consists of subdividing each Delaunay triangle, assigning colors either to 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 “mosaic” rendering style (Fig. 6, top left) 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 similar shape to create a pebble mosaic, while a self-similar quasicrystal sampling yields a decorative tiling with a small set of possible tile shapes.

In another example, our “paint strokes” rendering style (Fig. 3 and Fig. 8) uses the same subdivision as the “mosaic” style. For each Delaunay edge, its midpoint color is set to be the least neutral color of its two vertices, which is taken to be the color farthest from neutral gray. As in the “brush marks” style, nonlinear interpolation is used on the edges of the subdivided triangles while linear interpolation is used along the scan lines. The resulting style, reminiscent of the bold, angular strokes of a painting knife, assumes that plausible, saturated colors are preferable to colors that appear faded due to interpolation. Our framework offers plenty of scope to create different styles in search of a particular expression.

5.2 Procedural Rendering Styles

Our procedural rendering styles (Fig. 6, bottom), are based on the Voronoi diagram,⁴² which provides an efficient geometric data structure for keeping track of nearest neighbor relationships. They have the expressive power of parametric procedural textures²⁵ that adapt to local sampling properties such as color, density, and anisotropy. Hence, they can benefit from the rich library of functional components, such as multiresolution noise generators, and design methods, such as genetic programming, that already have been developed for procedural textures. Our approach is inspired by photorealistic image reconstruction through the use of local filters centered at the sample sites.⁵⁹ Our approach is also broadly related to the Shepard method for inverse distance weighted interpolation,⁴⁴ where the influence of nearby sample sites on an interpolated pixel decreases as their distance to the pixel increases. For each pixel p , we use the Voronoi diagram to find its closest sample site s_1 . To determine the pixel’s neighborhood, we approximate the pixel’s nearest neighbors by its closest sample site’s nearest neighbors. We select the K nearest neighboring sample sites s_k in order of increasing Euclidean distance $\Delta(s_1, s_k)$ from the pixel’s closest sample site s_1 . The pixel’s color $f(p)$ is calculated as a weighted sum of the colors $f(s_k)$ of its K neighboring sample sites s_k , with the weights determined by their local filter functions $\phi(p, s_k)$:

$$f(p) = \frac{\sum_{k=1}^K \phi(p, s_k) f(s_k)}{\sum_{k=1}^K \phi(p, s_k)}. \quad (7)$$

The neighborhood size K is chosen empirically to be large enough such that any further increase has negligible effect on the rendition, typically $10 \leq K \leq 40$; 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, non-negative filters that

are constrained to act locally within the neighborhood: as $\Delta(p, s_k) \rightarrow 0$, so $\phi(p, s_k) \rightarrow \infty$, while as $\Delta(p, s_k) \rightarrow \Delta(s_1, s_K)$, so $\phi(p, s_k) \rightarrow 0$. At each pixel, the filter function $\phi(p, s_k)$ usually puts much greater weight on sites relatively close to the pixel than on sites that are far away. Our nonlinear filter functions $\phi(p, s_k)$ 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 can be used to control luminance and chrominance. It is possible to further extend this approach by considering color properties fitted to the entire neighborhood, such as the color gradient.

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 closest sample site s_1 from the weighted sum $\phi(p, s_1) = 0$, thereby subdividing each Voronoi polygon into regions corresponding to its second closest sample sites. Two more examples show the richness of our framework. For the “sponge painting” style (Fig. 6, bottom right), the faceted glass style has been augmented by painting randomly chosen pixels with the color of their closest sample site. For the “color hatching” style (Fig. 6, bottom left), 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 sample site.

6 Evolution

An authentic artistic technique needs to offer the capacity for original expression. Graphic designers may not be satisfied with styles that come prepackaged 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 graphic designer should be in control of the rendering process without being required to grasp the complexity of how it works. For graphic designers who are not mathematicians, the development of algorithmic rendering styles must be a process of discovery rather than invention. Computer-assisted graphic design remains a creative task in that the requirements leave the form of the solution unspecified. We rely on interactive genetic programming as a means for original expression because of its ability to explore an open-ended parameter space.

Interactive evolution¹⁵ by aesthetic selection offers a user interface (Fig. 7, bottom) for exploring novel rendering styles. This “I-know-it-when-I-see-it” method of stochastic optimization employs the user’s artistic judgment to evaluate solutions proposed by the system. Our framework enables us to directly apply genetic programming to formulate the filter functions of our procedural rendering styles. A great variety of rendering styles (Fig. 7, top and right) become easily accessible. The details of our rudimentary implementation follow the classic work of Sims.²⁴ A symbolic expression tree defines the filter function of a procedural rendering style. Our expression trees have five categories of parameter leaf: *angle*, *Euclidean distance*, *Manhattan distance*, *constant*, and *pseudorandom*. They have 11 types of operator node: *cosine*, *add*, *subtract*, *divide*, *multiply*, *power*, *select*, *if*, *superellipse*, *Perlin noise*, and *texture map*. At each iteration, the algorithm randomly

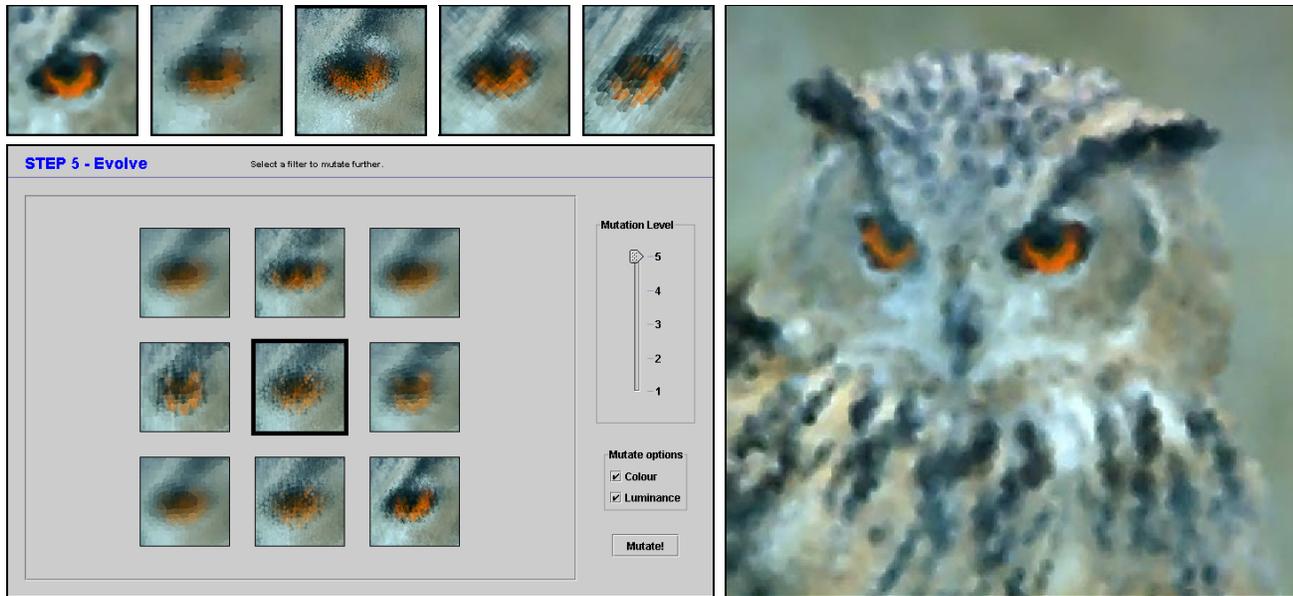


Fig. 7 Interactive evolution by aesthetic selection enables users to design the styles used in image rendering.

alters some of the parameters and operators of the expression tree. Different mutations can occur with different probabilities, dictating the prevalence of the various nodes. The resulting variations are then rendered and presented to the user, who selects which one should survive to produce a new generation of solutions for further refinement. To control the variability of the proposed solutions, the user can set the desired mutation rate. The process continues until the user is satisfied with the outcome. The challenge of interactive evolution is to rapidly 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 final result reflects a personal preference rather than an ideal solution. Compared to traditional genetic algorithms, this approach uses a small population of solutions that evolve for far fewer iterations in response to a far more intelligent fitness function embodied by the user's visual perception and aesthetic judgment.

In future research, we wish to make the rendering styles proposed by our interactive evolutionary algorithm more relevant to the user's aesthetic preferences. The user's past choices could be taken into account when generating a new selection of styles. The algorithm should not suggest the same style again if its visual effect has been previously rejected by the user. Before proposing a new selection of styles, the algorithm should also compare the styles with each other to ensure that they visually differ by approximately the same amount, as stipulated by the mutation rate, so that each new style actually offers a distinct rendering possibility. Styles would need to be compared according to their rendered images rather than their expression trees, because structurally distinct expression trees can yield visually indistinguishable results. In practice, an approximate comparison would be performed by rendering just a small subset of the image pixels. In another approach, the user

could take a proactive role in directing the evolutionary algorithm's creative priorities, rather than being confined to the reactive role of judging its creative results. The user would designate the perceptual criteria most relevant for success, such as vibrant, rough, or curly. Various image analysis techniques would automatically score style candidates according to the desired perceptual criteria, and only the best candidates would be presented to the user for further consideration. A final approach could allow the user to apply different styles to different regions of the picture. Instead of designing a single style for the entire image, the user would be free to vary styles according to image content to emphasize the visual composition of the picture. Interactive image segmentation would assist in this process. To generate a family of styles, a user parameter would be either embedded in the expression tree of a single master style or applied to control a linear combination of two distinct styles, such as a foreground style and a background style. Hence, the style parameter of each sample site would need to be encoded in the image representation.

7 Compression

The progressive nature of our sampling methods allows our NPR styles to be used for progressively rendering an image as it is received over a narrow bandwidth network. The nature of the representation also makes it a candidate for image 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,⁴³ a password can be



Fig. 8 Progressive image rendering using the “paint strokes” style with coverage adaptive sampling: 1%, 2%, 4%, 8%, 16%, and 32%.

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. For the initial sample sites 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 rendering style, 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 *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 actual and predicted color values. First we clear the two least significant bits and then round the magnitude of 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 color difference is proportional to the frequency of its occurrence. This encoding is appropriate for our use, since the scattering of the sample sites tends to remove any correlation between them.

A lossless multiresolution image representation (Fig. 8), 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 stylized rendering to the decompressed image. We compare (Fig. 9) our compression scheme to standard JPEG. With the image compressed to roughly the same size, we see that, in this example, 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 (Fig. 9, top right) appears clearly degraded by the JPEG artifacts, which are absent from our encoding. This demonstrates the need for the image encoding scheme to take account of the image rendering method. Previous research on photorealistic image representations^{46,47} has shown that, for high compression ratios, linear interpolation of Delaunay triangulations obtained through adaptive sampling can yield visually superior images compared to classical transform encodings, such as JPEG’s discrete cosine

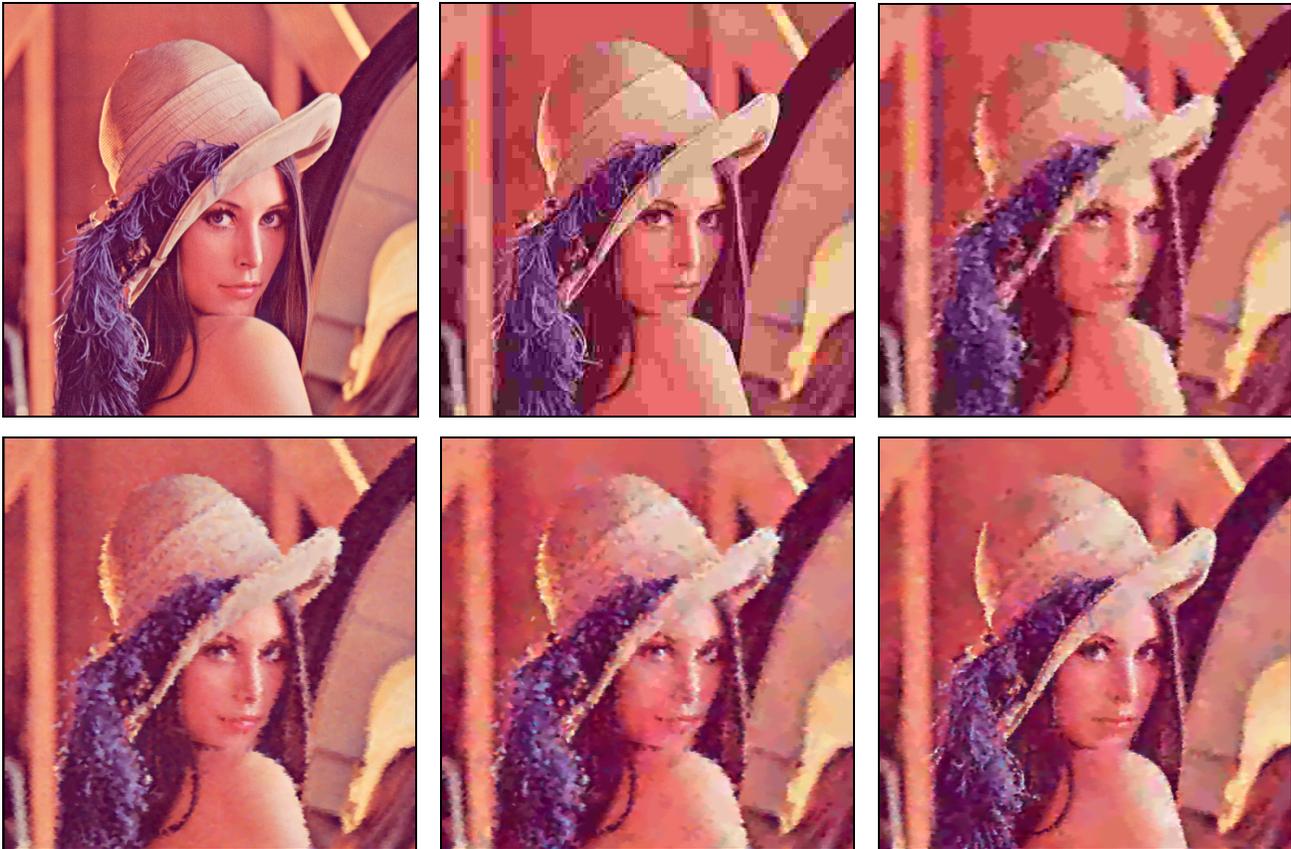


Fig. 9 Image compression using the “brush marks” rendering style (10485 samples $\approx 4\%$). In the top row, from left to right, the template image (512×512 , 16 bit color) is compressed using standard JPEG (7.0 K at 73:1), and then the JPEG image is rendered using quasirandom Halton sampling (without 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 (12.8 K at 40:1), next lossy quasirandom Halton sampling (6.1 K at 84:1), and finally lossy coverage adaptive sampling (7.2 K at 71:1).

transform. Our image representation extends these findings to stylized rendering. However, closer integration of sampling and encoding may be needed to compete with the storage efficiency of more recent transform encodings, such as JPEG2000’s discrete wavelet transform.

It is difficult to give any quantitative difference measures comparing the standard compression algorithms with our NPR method. NPR renditions appeal to highly nonlinear 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 that uses the “sponge painting” style introduces random noise into the picture. While such a use of random noise is likely to decrease the signal-to-noise ratio, it may well improve the human perception of the rendered image, as demonstrated by early approaches to image quantization³⁰ and compression.³¹ Applying models of human visual processing to evaluate NPR image representations is a research challenge that has only begun to be addressed.⁶⁷

8 Discussion

A concise visual representation demands that every element be essential to imparting its message. However, a physi-

cally accurate portrayal of a scene, specific enough to make every detail explicit, results in highly complex models. While such descriptions may be encoded as efficiently as their entropy allows, their intrinsic complexity cannot be perceptibly reduced, even when it is superfluous to the purpose of the image. Upholding the suspension of disbelief needed for the viewer to equate representation with reality requires a uniform resolution of detail, regardless of its relevance for visual communication. This places a fundamental constraint on the efficiency of any photorealistic image representation. How many needles should it take to draw a pine tree? How many more to depict a pine forest? When NPR is applied to image compression, it challenges the criteria for deciding what needs to be conveyed for a visual message to be well received by its audience. For instance, 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 effectively. When painterly abstraction takes the place of conventional artifacts in image rendering, such as aliasing and noise, it appears to diminish the visual impact of conventional constraints on im-

age representation, such as the Nyquist sampling rate. Where photorealistic reconstruction tends to draw attention to its flaws, stylized rendering appeals to the viewer's imagination to make an incomplete description seem believable. In this way, image stylization enables the complexity of image representation to be reduced.

Effective graphic design demands that an image convey the intended impression on its viewer, appearing true to its purpose rather than merely faithful to its subject. However, 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. 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. In graphic design, an important distinction is made between content and style, text and typeface. Likewise, in computer graphics, as in object-oriented programming, model and view often demand separate and independent specifications to allow the same data to serve different purposes in different contexts. We have applied this fundamental principle to color images. Our image representation explicitly separates the description of image content from the specification of image style, allowing content and style to be saved, changed and reused independently. For electronic imaging, the separation of representation from presentation has important consequences. By relying on procedural rendering primitives, our system can render stylized images at any desired output resolution. For instance, we can synthesize the image grain when printing or displaying a low resolution image on a high resolution device. Also, by keeping the continuous coordinates of the color samples independent from the discrete pixels of the display device, we can ensure that commonly applied spatial image transformations, such as scaling, rotation, and projection, affect only the mapping between the two coordinate systems and leave sampled colors unchanged. In this way, we are able to prevent spatial image transformations from degrading the saved image data, a common problem for images stored as pixel arrays.

Every image has a resolution at which stylized depiction becomes inevitable. This situation becomes easily noticeable when an image is magnified or compressed. The common artifacts of photorealistic image reconstruction, such as blocking, blurring, ringing, and anisotropic distortion, are a reflection of computational expediency and not necessarily human preference. Our work is based on the idea that, when 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 graphic designer 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. Intentional stylization endows an image with a visual heritage. We give graphic designers the tools required to express their creative intentions by developing personalized rendering styles that are especially well suited for displaying compressed imagery. Through exercising control over the compression and interpolation artifacts, the graphic designer may be able to improve the perceived visual quality of an image by shifting the viewer's expectations from a photographic reproduction to an artistic expression.

Directing the viewer's attention through stylized presentation allows for selective emphasis. To convey a scene at a glance, a comprehensible depiction need not be comprehensive. For instance, the discrepancy between realism and photorealism can be observed in the way photographers use lenses to softly blur the background to lure the viewer's eye to focus on the foreground. In our system, such visual effects can be produced using an importance map. Sometimes, the deliberate omission of detail can stimulate the viewer's interest. As discussed by Strothotte and Strothotte,⁶⁸ psychological studies have found that ambiguity invites scrutiny, thereby improving memory performance. In another experiment, a rough line sketch was found to be more effective at promoting discussion among its viewers than a precisely shaded rendition of the same scene. Abstraction can serve to engage the imagination. It encourages the viewer to fill in the empty spaces between the brush strokes with projections of his or her own expectations, and thus 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. When Schmidhuber³⁷ explored interactive drawing techniques designed to support a concise description, he speculated that a picture's ability to capture the essence of its subject may be related to how closely its visual complexity reflects the minimal description length required given the viewer's prior knowledge. Though a picture may be worth a thousand words or a thousand kilobytes, it only takes less than a thousand bits of pixel data to render a face instantly recognizable.⁶⁹ Actually, the limited span of human attention can only process and recognize around 30 to 60 bits of visual information at a time.⁷⁰ This disparity of information capacity between image representation and visual perception offers great scope for future research in applying stylized rendering to image compression.

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 develop a novel adaptive sampling algorithm and a novel point-based rendering framework for image stylization.

Clearly, these methods are not aimed at image compression applications for which objective visual fidelity is all that matters. In practice, they are most appropriate in contexts where images communicate ideas or illustrate narra-

tives. NPR techniques abandon the conventional goal of exact reproduction in pursuit of the evocative capacity for visual communication. For an efficient image representation, where some visual information needs to be implied rather than encoded, stylized rendering has the advantage of making abstraction and simplification appear legitimate. In visual communication, clarity can be more valuable than completeness. As this fundamental observation challenges the assumptions of conventional image compression, it points the way for future research into the design of expressive, effective, and efficient image representations. There is an opportunity for image compression techniques to embrace the venerable aesthetic principle that less is more.

Acknowledgments

We wish to thank Malcolm Sabin, Carsten Moening, Alan Blackwell, Peter Robinson, and Victor Ostromoukhov for their time, insight, and advice. In conducting this research, Mark Grundland gratefully acknowledges the financial support of the Natural Sciences and Engineering Research Council of Canada, le Fonds Québécois de la Recherche sur la Nature et les Technologies, the Celanese Canada Internationalist Fellowship, the British Council, the Overseas Research Student Award Scheme, the Cambridge Commonwealth Trust, the Royal Academy of Engineering, the University of Cambridge, and Peterhouse. The images were provided by FreeFoto.com and the Waterloo Brag Zone.

References

- O. Egger, P. Fleury, T. Ebrahimi, and M. Kunt, "High-performance compression of visual information—a tutorial review, part I: still pictures," *Proc. IEEE* **87**(6), 974–1013 (1999).
- M. Grundland, C. Gibbs, and N. A. Dodgson, "Stylized rendering for multiresolution image representation," *Proc. SPIE* **5666**, 280–292 (2005).
- Synthetic Software, *Studio Artist*, **3.5** (2006), see <http://www.synthetic.com/>.
- Corel Corporation, *Corel Painter*, **10.1** (2007), see <http://www.corel.com/painter>.
- Informatix Software International, *Piranesi*, **5** (2007), see <http://www.informatix.co.uk/piranesi/>.
- P. Haeberli, "Paint by numbers: abstract image representations," *Proc. SIGGRAPH*, pp. 207–214 (1990).
- A. Hertzmann, "A survey of stroke-based rendering," *IEEE Comput. Graphics Appl.* **23**(4), 70–81 (2003).
- F. Durand, V. Ostromoukhov, M. Miller, F. Duranleau, and J. Dorsey, "Decoupling strokes and high-level attributes for interactive traditional drawing," *Proc. Eurographics Workshop on Rendering*, pp. 71–82 (2001).
- A. Hertzmann, "Painterly rendering with curved brush strokes of multiple sizes," *Proc. SIGGRAPH*, pp. 453–460 (1998).
- J. Hays and I. Essa, "Image and video based painterly animation," *Proc. Intl. Symp. Non-photorealistic Animation Rendering*, pp. 113–120 (2004).
- L. Streit and J. Buchanan, "Importance driven halftoning," *Proc. Eurographics*, pp. 207–217 (1998).
- M. Ashikhmin, "Fast texture transfer," *IEEE Comput. Graphics Appl.* **23**(4), 38–43 (2003).
- A. Hertzmann, C. E. Jacobs, N. Oliver, B. Curless, and D. H. Salesin, "Image analogies," *Proc. SIGGRAPH*, pp. 327–340 (2001).
- V. Ostromoukhov and R. D. Hersch, "Artistic screening," *Proc. SIGGRAPH*, pp. 219–228 (1995).
- H. Takagi, "Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation," *Proc. IEEE* **89**(9), 1275–1296 (2001).
- P. Bentley, *Evolutionary Design by Computers*, Morgan Kaufmann, San Francisco (1999).
- G. R. Greenfield, "Evolving expressions and art by choice," *Leonardo* **33**(2), 93–99 (2000).
- E. Baker and M. Seltzer, "Evolving line drawings," *Proc. Graphics Interface*, pp. 91–100 (1994).
- T. He, L. Hong, A. Kaufman, and H. Pfister, "Generation of transfer functions with stochastic search techniques," *Proc. IEEE Visual.*, pp. 227–234 (1996).
- R. Poli and S. Cagnoni, "Genetic programming with user-driven selection: experiments on the evolution of algorithms for image enhancement," *Proc. Genetic Program.*, pp. 269–277 (1997).
- J. Dalton, "Adaptive learning of aesthetic imaging transformation," *Proc. DICTA* **2**, 659–665 (1993).
- J. Dalton, "Image similarity models and the perception of artistic representations of natural images," *Proc. SPIE* **3016**, 517–525 (1997).
- J. C. Dalton, "Perceptual image analysis for graphical rendering and digital libraries," *Proc. SPIE* **4662**, 226–234 (2002).
- K. Sims, "Interactive evolution of equations for procedural models," *Visual Comput.* **9**(8), 466–476 (1993).
- D. S. Ebert, F. K. Musgrave, D. Peachey, K. Perlin, and S. Worley, *Texturing and Modeling*, 2 ed., AP Professional, San Diego (1998).
- K. Perlin and L. Velho, "Live paint: painting with procedural multi-scale textures," *Proc. SIGGRAPH*, pp. 153–160 (1995).
- T. H. Morrin, "A black-white representation of a gray-scale picture," *IEEE Trans. Comput.* **23**(2), 184–186 (1974).
- S. Carlsson, "Sketch based coding of grey level images," *Signal Process.* **15**(1), 57–83 (1988).
- D. E. Pearson and J. A. Robinson, "Visual communication at very low data rates," *Proc. IEEE* **73**(4), 795–812 (1985).
- L. Roberts, "Picture coding using pseudo-random noise," *IRE Trans. Inf. Theory* **8**(2), 145–154 (1962).
- M. Koehler and M. Kunt, "Image data compression by contour texture modelling," *Proc. SPIE* **397**, 132–139 (1983).
- M. F. Barnsley, A. Jacquin, F. Malassenet, L. Reuter, and A. D. Sloan, "Harnessing chaos for image synthesis," *Proc. SIGGRAPH*, pp. 131–140 (1988).
- P. Salembier, P. Brigger, J. R. Casas, and M. Pardas, "Morphological operators for image and video compression," *IEEE Trans. Image Process.* **5**(6), 881–898 (1996).
- J. H. Elder, "Are edges incomplete?" *Int. J. Comput. Vis.* **34**(2–3), 97–122 (1999).
- J. A. S. Viggiano and N. M. Moroney, "Color reproduction algorithms and intent," *Proc. IS&T/SID Color Imag. Conf.*, pp. 152–154 (1995).
- I. Herman and D. Duke, "Minimal graphics," *IEEE Comput. Graphics Appl.* **21**(6), 18–21 (2001).
- J. Schmidhuber, "Low-complexity art," *Leonardo* **30**(2), 97–103 (1997).
- L. Markosian, B. J. Meier, M. A. Kowalski, L. S. Holden, J. D. Northrup, and J. F. Hughes, "Art-based rendering with continuous levels of detail," *Proc. Intl. Symp. Non-photorealistic Animation Rendering*, pp. 59–66 (2000).
- M. Salisbury, C. Anderson, D. Lischinski, and D. H. Salesin, "Scale-dependent reproduction of pen-and-ink illustrations," *Proc. SIGGRAPH*, pp. 461–468 (1996).
- L. Kovacs and T. Sziranyi, "Efficient coding of stroke-rendered paintings," *Proc. Intl. Conf. Patt. Recog.* **2**, 835–838 (2004).
- T. Sziranyi and Z. Toth, "Random paintbrush transformation," *Proc. Intl. Conf. Patt. Recog.* **3**, 151–154 (2000).
- F. Aurenhammer, "Voronoi diagrams: a survey of a fundamental geometric data structure," *ACM Comput. Surv.* **23**(3), 345–405 (1991).
- N. Ahuja, B. An, and B. Schachter, "Image representation using voronoi tessellation," *Comput. Vis. Graph. Image Process.* **29**(3), 286–295 (1985).
- I. Amidror, "Scattered data interpolation methods for electronic imaging systems: a survey," *J. Electron. Imaging* **11**(2), 157–176 (2002).
- L. Darsa and B. Costa, "Multiresolution representation and reconstruction of adaptively sampled images," *Proc. SIBGRAPI*, pp. 321–328 (1996).
- M. Kashimura, Y. Sato, and S. Ozawa, "Image description for coding using triangular patch structure," *Proc. ICCS/ISITA*, pp. 330–334 (1992).
- L. Rila, "Image coding using irregular subsampling and delaunay triangulation," *Proc. SIBGRAPI*, pp. 167–173 (1998).
- F. Anton, D. Mioc, and A. Fournier, "Reconstructing 2-D images with natural neighbour interpolation," *Visual Comput.* **17**(3), 134–146 (2001).
- J. A. Robinson, "Image coding with ridge and valley primitives," *IEEE Trans. Commun.* **43**(6), 2095–2102 (1995).
- C. S. Kaplan, "Voronoi diagrams and ornamental design," *Proc. Symp. Intl. Soc. Arts, Math., Arch.*, pp. 277–283 (1999).
- A. Klein, P. P. Sloan, A. Colburn, A. Finkelstein, and M. F. Cohen, "Video cubism," *Microsoft Research Tech. Report*, MSR-TR-2001-45 (2001).
- J. P. Collomosse and P. M. Hall, "Cubist style rendering from photographs," *IEEE Trans. Vis. Comput. Graph.* **9**(4), 443–453 (2003).
- O. Deussen, S. Hiller, C. van Overveld, and T. Strothotte, "Floating points: a method for computing stipple drawings," *Proc. Eurographics*, pp. 41–50 (2000).
- S. Hiller, H. Hellwig, and O. Deussen, "Beyond stippling: methods for distributing objects on the plane," *Proc. Eurographics*, pp. 515–522 (2003).

55. A. Hausner, "Simulating decorative mosaics," *Proc. SIGGRAPH*, pp. 573–580 (2001).
56. I. Ragnemalm, "Neighborhoods for distance transformations using ordered propagation," *CVGIP: Image Understand.* **56**(3), 399–409 (1992).
57. P. E. Danielsson, "Euclidean distance mapping," *Comput. Graph. Image Process.* **14**(3), 227–248 (1980).
58. K. E. Hoff, II, T. Culver, J. Keyser, L. Ming, and D. Manocha, "Fast computation of generalized voronoi diagrams using graphics hardware," *Proc. SIGGRAPH*, pp. 277–286 (1999).
59. A. S. Glassner, *Principles of Digital Image Synthesis*, Vol. 1, Morgan Kaufmann, San Francisco (1995).
60. C. Goodman-Strauss, "Aperiodic hierarchical tilings," *Proc. NATO-ASI E-354*, 481–496 (1999).
61. J. Patera, "Non-crystallographic root systems and quasicrystals," in *Proc. NATO-ASI C-489*, 443–465 (1997).
62. V. Ostromoukhov, C. Donohue, and P. M. Jodoin, "Fast hierarchical importance sampling with blue noise properties," *Proc. SIGGRAPH*, pp. 488–495 (2004).
63. J. Kopf, D. Cohen-Or, O. Deussen, and D. Lischinski, "Recursive Wang tiles for real-time blue noise," *Proc. SIGGRAPH*, pp. 509–518 (2006).
64. V. Ostromoukhov, "Sampling with polyominoes," *Proc. SIGGRAPH*, pp. 78:1–78:6 (2006).
65. Y. Eldar, M. Lindenbaum, M. Porat, and Y. Y. Zeevi, "The farthest point strategy for progressive image sampling," *IEEE Trans. Image Process.* **6**(9), 1305–1315 (1997).
66. H. Zhang, "Pattern generation with color map gouraud shading," *Comput. Graph.* **20**(1), 157–162 (1996).
67. M. Cadik, "Human perception and computer graphics," *Czech Technical University Postgraduate Study Report*, DC-PSR-2004-06 (2004).
68. C. Strothotte and T. Strothotte, *Seeing between the Pixels: Pictures in Interactive Systems*, Springer, Berlin (1997).
69. L. D. Harmon, "The recognition of faces," *Sci. Am.* **229**(5), 71–82 (1973).
70. P. Verghese and D. G. Pelli, "The information capacity of visual attention," *Vision Res.* **32**(5), 983–995 (1992).



Mark Grundland received his PhD in image processing from the University of Cambridge in 2007 and his Honors BA in computer science from McGill University in 2000. At the crossroads of computer graphics, computer vision, and visual art, his research seeks image processing algorithms for facilitating artistic visual expression. During his graduate studies at the University of Cambridge, he has pursued research in contrast enhancement, color correction, image compositing, and nonphotorealistic rendering.

Previously, at the Centre de Recherches Mathématiques de l'Université de Montréal, he was engaged in the scientific visualization of quasicrystals. He is a founder of Roleplay Technologies, where he is responsible for developing tools for authoring and simulating realistic interactive conversation for use in interactive entertainment and computer-based training. Currently, as a freelance consultant at Functional Elegance, his commercial work combines software design with market analysis and business strategy. His latest project involves information visualization for Internet search.



Chris Gibbs received his BA degree in computer science in 2003 from Cambridge University, after which he was a director for 2 years for a company providing SMS text messaging and mobile solutions to businesses. He is now self-employed and designing automated trading software to operate on the main European and American stock exchanges.



Neil A. Dodgson is a Reader in Graphics and Imaging in the Computer Laboratory at the University of Cambridge. He undertakes research in 3-D displays, 3-D modeling, and 2-D imaging. He received a BSc degree in physics and computer science from Massey University, New Zealand, in 1988, a PhD in image processing from the University of Cambridge, England, in 1992, and ScD, also from Cambridge, in 2007. He is a Fellow of Emmanuel College, Cambridge, a Fellow of the Institution of Engineering and Technology, and a chartered engineer.