

Multiscale Moment-Based Painterly Rendering

Diego Nehab¹ Luiz Velho²

¹PUC-Rio – Pontifícia Universidade Católica, Rio de Janeiro

²IMPA – Instituto de Matemática Pura e Aplicada

diego@tecgraf.puc-rio.br lvelho@impa.br

Abstract

In this paper we present a new method for painterly rendering of images. Our method extends the image-moment stroke placement algorithm in two ways: we employ a multiscale scheme for computing strokes and we provide a parametrized mechanism for controlling stroke distribution. In addition, we introduce a flexible implementation framework based on the extension language Lua.

Keywords: Non-photorealistic rendering, Painting, Dithering, Lua

1 Introduction

The evolution of computer graphics led naturally to the development of different types of visualization techniques. Initially, the focus was on photorealistic rendering, where the goal is to generate synthetic images that are indistinguishable from real photographs [17]. More recently, there has been a growing interest in non-photorealistic rendering techniques that emphasize the expressive depiction of visual information [3].

Non-photorealistic rendering is, by definition, a very broad area of research that encompasses many rendering styles in various application contexts. Two useful criteria for classification of non-photorealistic rendering techniques are: the type of source data; and the nature of the simulated process.

Techniques are classified according to source data into *object-space* methods that use the 3D model of a scene to create the rendered image [11, 14], and *image-space* methods that work directly on the 2D image [4, 5]. Hybrid methods take advantage of both 3D and 2D data to produce the final result [9].

Most non-photorealistic techniques are inspired in traditional visual art forms, such as painting (oil [5], watercolor [1, 18]), drawing (pen-and-ink [7, 2, 22], pencil [19, 21], charcoal [23]), and printing (etching, engraving [12, 13]). Here, both the physical process and the medium provide a paradigm for computation and interaction.

In this paper we present a new method for painterly rendering of images that improves upon previous work in the area.

1.1 Related Work

Painterly rendering simulates the appearance of painted images. The basic primitive in this technique is a brush stroke. Images are generated by applying a sequence of brush strokes to a 2D canvas. A brush stroke has various attributes, such as position, shape and color [20].

In object-space methods, brush strokes are first associated with the 3D geometry of objects in a scene, and then projected to the image plane defined by a virtual camera [11]. In image-space methods, brush strokes are placed on the output image, based on 2D information derived from input images [4].

Interactive methods allow the user to guide the rendering process in a manual or semi-automatic manner, by indicating where strokes should be placed [4]. Non-interactive methods render the image automatically based on input parameters and data analysis [5].

Some methods process a sequence of images exploiting temporal coherence [10, 6].

The technique described in this paper is a non-interactive, image-space method. It is based on the image moment-based stroke placement algorithm [15, 16]. The main original contributions in our work are: a multiscale scheme for computing the strokes, a parametrized mechanism for controlling stroke distribution, an image abstraction specially optimized for the algorithm, and a flexible implementation framework based on the extension language Lua [8].

1.2 Overview

This paper is organized as follows. Section 2 reviews the moment-based painterly rendering method. Section 3 describes the additions proposed in the new method. Section 4 shows some results of using the method. Section 5 concludes with final remarks and a discussion of ongoing work.

2 Review of the image-moment based painterly rendering

Given a source image and a stroke template image, the painterly rendering algorithm outputs a painted version of the source image. The method proceeds as an artist who progressively strokes a canvas trying to reproduce the source image on it.

In this paper, we use gray-scale images, but the method presented extends naturally to color images.

```
-----
-- Painterly renders an image
-- Input
-- Source: source image
-- Stroke: stroke template image
-- S: neighborhood size
-- Output
-- painted version of source image
-----
function PainterlyRender(Source, Stroke, S)
    local List = ComputeStrokeList(Source, S)
    local w, h = GetWidth(Source), GetHeight(Source)
    return PaintStrokeList(Stroke, List, Canvas(w, h))
end
```

Program 1: Painterly rendering algorithm.

The process can be divided into two phases: analysis and synthesis. In the analysis phase, a list of strokes is calculated from the source image. In the synthesis phase, the strokes are painted over a blank canvas. Both phases can be seen in Program 1.

The algorithm outlined above, and detailed in the following sections, generates images similar to that shown in Figure 1, and is the result of previous work [16]. We proceed with the description of the




Figure 1: Painterly rendering process

synthesis process, which will make clear to the reader the requirements to be fulfilled by analysis process, explained subsequently.

2.1 The synthesis phase

The algorithm works with strokes that are described by the following set of attributes: color, location, orientation and size. According to these parameters, painting a stroke on a canvas corresponds to the process of scaling, rotating, and using the stroke template image as a transparency map to be blended in the output image, in the correct position and with the appropriate color.

```

-----
-- Paints a stroke on a canvas
-- Input
-- Stroke: stroke template image
-- params: stroke parameter set
-- Canvas: canvas to receive strokes
-- Output
-- painted version of source image
-----
function PaintStroke(Stroke, params, Canvas)
    local Scaled = Scale(Stroke, params.w, params.h, New())
    local Rotated = Rotate(Scaled, params.theta, New())
    Blend(Rotated, params.xc, params.yc, params.color, Canvas)
end

-----
-- Paints a stroke list on a canvas
-- Input
-- Stroke: stroke template image
-- List: stroke parameter list
-- Canvas: canvas to receive stroke
-- Output
-- painted version of source image
-----
function PaintStrokeList(Stroke, List, Canvas)
    for i = 1, Length(List) do
        PaintStroke(Stroke, List[i], Canvas)
    end
    return Canvas
end

```

Program 2: The synthesis phase.

Program 2 shows the complete implementation of the synthesis phase of the algorithm. The partial result for a stroke list can be seen in Figure 2.




Figure 2: Example of stroke painting, for different stroke templates.

When there are enough strokes and each of them has the appropriate attributes, a result such as that of Figure 2 can be achieved. The computation of such a list of strokes is the task of the analysis phase of the algorithm.

2.2 The analysis phase

The first step in the creation of a stroke list is the definition of the stroke distribution. In a second step, each stroke in the distribution has its parameters computed.

The definition of a stroke distribution is based on the observation that high-frequency details in the source image should be represented by many small strokes, whereas low-frequency regions should be represented by fewer large strokes.

The computation of the stroke parameters results from the analysis, by the image moments theory, of the neighborhood in the source image where the stroke is to be placed. Since the same tools provide the information needed for the computation of the stroke distribution, their use is introduced first.

2.3 Computing stroke parameters

The goal of each stroke is to approximate a neighborhood of the source image. From each neighborhood, the image-moment based approach determines the corresponding stroke parameters in two steps. The first step computes a color difference image between the region and the color at the center of that region. The second step determines the remaining stroke parameters based on the image moments of the color difference image created in the first step.

2.3.1 The color difference image

The color difference image attempts to measure the distance between the color of the stroke and the color of each point in the source image neighborhood being considered. Ideally, the resulting image shows a picture of the shape that a stroke of that color should have if it was to approximate the region. In other words, the operation isolates the segments of the region that can be better represented in the stroke color. An algorithm that generates such result is shown in Program 3.

```
-----
-- Computes the difference between Region and color
-- Input
-- Region: source image region
-- color: stroke parameter list
-- Diff: color difference image buffer
-- Output
-- color difference image
-----
function ColorDifferenceImage(Region, color, Diff)
    local d0 = 0.2
    local f = function(d)
        if d < d0 then return (1 - (d/d0)^2)^2
        else return 0 end
    end
    Clear(Diff, 0)
    for x = 0, GetWidth(Region)-1 do
        for y = 0, GetHeight(Region)-1 do
            local d = abs(color - GetColor(Region, x, y))
            SetColor(Diff, x, y, f(d)) end
        end
    end
    return Diff
end
```

Program 3: The color difference image. Function f increases the contrast of the result.

The quality of the computed parameters depends on the quality of the segmentation produced by the color difference image. In particular, low contrast images may consistently give rise to similar stroke parameters. To increase the contrast, a function is used that maps color difference values into the intensity values actually stored in the resulting image.

2.3.2 Using image moments

Image moments are summations over all pixels of the image. The notions of area, position and orientation are captured by the quantities M_{00} , M_{01} , M_{10} , M_{11} , M_{20} , M_{02} , as defined by the following equation:

$$M_{im} = \sum_x \sum_y x^i y^m I(x, y) \quad (1)$$

The stroke parameters corresponding to a rectangle that, when rendered, creates an image that has the same image moments than the color difference image of the region being approximated can be computed from equation (1) — See [16] for details and formulas. Program 4 shows the corresponding implementation.

```
-----
-- Computes stroke parameters for a neighborhood
-- Input
-- Source: stroke template image
-- x, y: neighborhood center
-- S: neighborhood size
-- Diff: color difference image buffer
-- Output
-- a set with the parameters:
-- xc, yc: center of stroke
-- color: stroke color
-- w, h: stroke size
-- theta: stroke orientation
-----
function StrokeParameters(Source, x, y, S, Diff)
    local Region = Share(Source, x, y, S, S, New())
    local color = GetColor(Source, x, y)
    ColorDifferenceImage(Region, color, Diff)
    local m00, m01, m10, m11, m02, m20 = Moments(Diff)
    if m00 < S*S/100 then return nil end
    local dxc = m10/m00;
    local dyc = m01/m00;
    local a = m20/m00 - dxc*dxc;
    local b = 2*(m11/m00 - dxc*dyc);
    local c = m02/m00 - dyc*dyc;
    local t1 = a - c
    local theta = atan2(b, t1) / 2;
    local t2 = sqrt(b*b + t1*t1);
    local t3 = a+c
    if t3 + t2 < 0 then return nil end
    local w = sqrt(6 * (t3 + t2));
    if t3 - t2 < 0 then return nil end
    local h = sqrt(6 * (t3 - t2));
    local xc = x + floor(dxc - S/2 + 0.5);
    local yc = y + floor(dyc - S/2 + 0.5);
    if w < 1 or h < 1 then return nil end
    return { xc = xc, yc = yc, w = w, h = h, theta = theta, color = color }
end
```

Program 4: The stroke parameters.

2.4 Determining the stroke distribution

The frequency information needed for the definition of a stroke distribution is obtained with the computation of a stroke area image — an image in which the value of each pixel corresponds to the area of a stroke approximating its neighborhood. The stroke distribution is given by a stroke positions image, in which each position is marked by a dot. This image is generated from the stroke area image by a special dithering algorithm.

2.4.1 The stroke area image

The area of a stroke associated with any position in the source image can be estimated by the M_{00} of the color difference image between the color and the neighborhood of the position. Program 5 shows the implementation of this idea and Figure 3 shows a sample of the result it produces.

```

-----
-- Computes the stroke area image
-- Input
-- Source: source image
-- S: neighborhood size
-- Output
-- stroke are image
-----
function StrokeAreaImage(Source, S)
    local w, h = GetWidth(Source), GetHeight(Source)
    local Output = BlankCanvas(w, h)
    local Diff = BlankCanvas(S, S)
    local Region = New()
    for y = 0, h-1 do
        for x = 0, w-1 do
            Share(Source, x, y, S, S, Region)
            local wr, hr = GetWidth(Region), GetHeight(Region)
            ColorDifferenceImage(Region, GetColor(Source, x, y), Diff)
            SetColor(Output, x, y, Moment00(Diff, wr, hr)/(wr*hr))
        end
    end
    return Output
end

```

Program 5: The stroke area image.




Figure 3: Example of stroke area image.

2.4.2 The stroke positions image

The stroke positions image is a monochrome dithered version of the stroke area image. The algorithm used to create the stroke distribution must be designed to concentrate strokes around the dark regions of the stroke area image, and to avoid large regions without strokes. To this end, the original study used a modified version of a space-filling curve dithering algorithm, in which the accumulated intensity values were inversely proportional to the area of the stroke.




Figure 4: Example of stroke positions image.

Figure 4 shows an example of how the stroke positions image should look. This image was created from the stroke area image of Figure 3 by our own parametrized dithering algorithm. This new algorithm will be explained in due time, in Section 3.

2.4.3 Computing the stroke list

Completing the implementation of the painterly rendering algorithm, Program 6 shows the procedure used to compute the stroke list. The only addition to what was previously discussed is that, before being returned to the caller, the stroke list is sorted by stroke area. This is to prevent fine details, represented by small strokes, from being overwritten by larger strokes.

```

-----
-- Returns a list of strokes
-- Input
-- Source: stroke template image
-- S: neighborhood size
-- Output
-- a list with strokes, sorted by order
-----
function ComputeStrokeList(Source, S)
    local List = {}
    local w, h = GetWidth(Source), GetHeight(Source)
    local Area = StrokeAreaImage(Source, S)
    local Position = Dither(Area, S)
    local Diff = Alloc(S, S, New())
    for y = 0, h-1 do
        for x = 0, w-1 do
            if GetColor(Position, x, y) == 0 then
                Append(List, StrokeParameters(Source, x, y, S, Diff))
            end
        end
    end
    Sort(List, function(a,b) return a.w*a.h > b.w*b.h end)
    return List
end

```

Program 6: The stroke list.


3 Original contributions to the algorithm

The ideas presented so far describe a complex process to create images that resembles human hand painting, as seen in Figure 1. Some aspects of the process can be improved, other parts can be implemented in a way that deserves documentation. This session describes what was added by our research.


3.1 Multi-resolution analysis

As expected, the stroke parameters, computed with the aid of the color difference images and the image moments theory, correctly approximate local source image neighborhoods. Unfortunately, although small details can be captured within a neighborhood, features that occupy more than the size of a single neighborhood cannot, and must therefore be represented by a group of smaller strokes. Unless the used stroke template image has a low opacity overall, this composition becomes evident. Furthermore, a needlessly large amount of small strokes must be used to represent what could be approximated by fewer large strokes.


In order to capture strokes over a wider range of sizes, we use a multi-resolution approach during the analysis phase of the algorithm. Stroke lists are collected for each resolution in a pyramid built from the source image. The painted result at each level is obtained by blending its strokes over the result of the lower resolution level. Program 7 shows the implementation of these ideas.



(a) Level 1



(b) Level 2



(c) Level 3



(d) Level 4

Figure 5: Strokes at different resolution levels.

```

-----
-- Painterly renders an image, with multi-resolution
-- Input
-- Source: source image
-- Stroke: stroke template image
-- S: neighborhood size
-- L: number of levels
-- Output
-- painted version of source image
-----
function MultiResolutionPainterlyRender(Source, Stroke, S, L)
    local Pyramid = Source
    local l = 2
    while l <= L do
        local w, h = GetWidth(Pyramid[l-1])/2, GetHeight(Pyramid[l-1])/2
        Pyramid[l] = Copy(Scale(Pyramid[l-1], w, h, New()), New())
        l = l + 1
    end
    local Canvas = BlankCanvas(GetWidth(Pyramid[L])/2,
                                GetHeight(Pyramid[L])/2)
    l = L
    while l >= 1 do
        local w, h = GetWidth(Pyramid[l]), GetHeight(Pyramid[l])
        Canvas = Copy(Scale(Canvas, w, h, New()), New())
        local Sp, E = Spread(S, L, l), Enhance(S, L, l)
        local List = ExtendedStrokeList(Pyramid[l], S, Sp, E)
        Canvas = PaintStrokeList(Stroke, List, Canvas)
        l = l - 1
    end
    return Canvas
end

```

Program 7: The multi-resolution painter algorithm.



(a) Level 6

Figure 6: Composition of strokes at different resolution levels.

Figure 6 shows an example of painterly rendered image created by the multi-resolution method. Figures 5(a) to (d) depict the strokes at each resolution level that are composed together to create the final image in Figure 6. The strokes were generated from blurred images coming from the multiresolution pyramid. The images are shown scaled to the same resolution to simplify comparison, and to illustrate the steps followed by the algorithm.

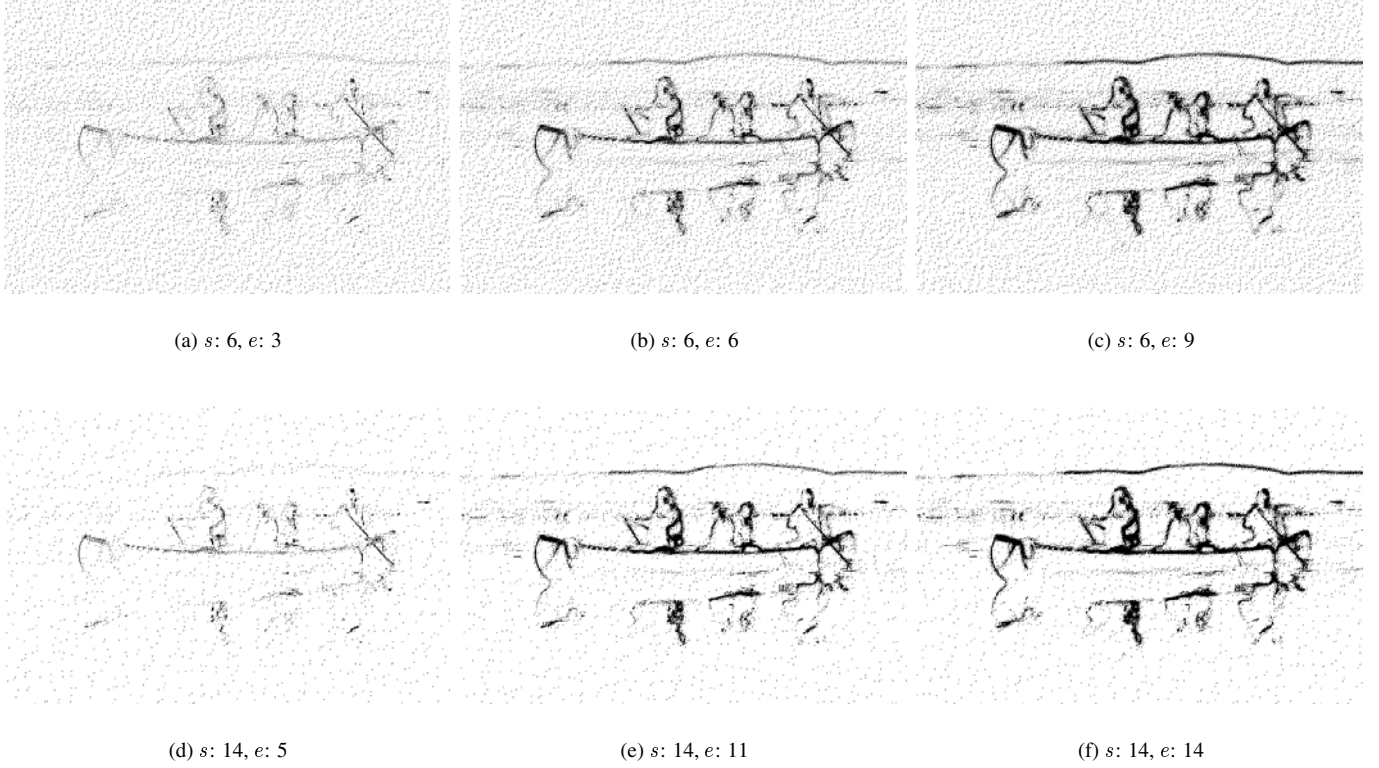


Figure 7: Parametrized stroke positions images.

3.2 Parametrized stroke positions image

The stroke list for higher resolution levels should not only concentrate strokes on high frequency areas of the source image, but also avoid placing strokes over lower frequency areas. Otherwise, strokes coming from lower resolution levels would be consistently obscured by the strokes coming from higher resolution levels, producing a result no better than the single-resolution approach.

To avoid this problem, the new procedure used to create stroke positions images accepts two parameters, referred to as the spreading and the enhancing factors. The spreading factor places an upper bound on the maximum distance between strokes, effectively controlling the overall stroke density. The enhancing factor controls the degree to which the density of strokes increases when close to the edges found in the stroke area image.

Before being considered, the value of each pixel is passed through the function defined by Equation (2), along with spreading and enhancing parameters.

$$se(v, s, e) = \frac{1}{(s^2 - 1)v^e + 1} \quad (2)$$

This function was designed to map the value 0 into 1 (small area values generate more strokes) and the value 1 into $\frac{1}{s^2}$ (even large areas contribute to the generations of strokes). The value $\frac{1}{s^2}$ was chosen so that a stroke position is issued after at most s^2 pixels. Furthermore, the effect of the enhancing factor is to accentuate small input values, which are exactly those close to the edges of the stroke area image.

In our implementation, the dithering proceeds by traditional error diffusion. For each pixel, the error accumulated due to truncation is spread to three of its adjacent pixels. However, in order

to avoid undesirable periodic artifacts, we shuffle the coefficients used to diffuse the error. Although the randomization of the dithering process may not be useful to produce high quality images, we are not interested in photo-realism. Therefore, this simple idea is enough to produce results with the desired properties, as seen in Figure 7.

3.3 Optimized Image abstraction

During the two phases of the algorithm, the basic computations performed over images are: stroke area image, stroke positions image, color image difference, image moments, scaling, rotation and blending. The choice of an appropriate image abstraction can simplify the task of efficiently implementing these operations. In particular, the abstraction should simplify computations involving image neighborhoods and avoid unnecessary memory allocations.

```
/* image data type */
typedef struct Tmono {
    float *buffer;
    int width, height, row;
    int reference;
} Tmono;
typedef Tmono *Pmono;
```

Program 8: C structure representing an image.

The data structure described by Program 8 can be used to store information about a newly allocated image (such as the `source` image of Figure 8) and can also store information representing part of an image (such as the `shared` image on the same figure). The `row` field always relate to the image that *owns* the buffer, and allows




Figure 8: Meaning of image structure fields.

routines to correctly determine the position of each pixel in shared references¹.

Our experience shows that adapting an image processing algorithm to deal with the above image representation transparently, either as input or output, is an effortless task. Furthermore, the adapted version usually suffers no measurable performance degradation. Therefore, every image processing function in the toolkit makes no assumptions whether the buffers are shared or owned by the images structures that point to them. A simple example in C of such function is presented by Program 9.

```

/*-----
 * Clears an image.
 * Input:
 * in: image to be cleared
 * c: clear value
 * Returns:
 * 0 in case of error, 1 if successful
 *-----*/
int mono_clear(Pmono in, float c)
{
    int skip = in->row - in->width;
    float *p = in->buffer;
    for (int y = 0; y < in->height; y++) {
        /* process row */
        for (int x = 0; x < in->width; x++) *p++ = c;
        /* skip part of row not belonging to reference */
        /* if image is not a reference, skip is 0 */
        p += skip;
    }
    return 1;
}

```

Program 9: Clearing an image represented by the abstraction structure.

This agreement allows the extensive use of references throughout all parts of the algorithm. Instead of supplying neighborhood limits

¹The reference field of the data structure distinguishes between owners and references. This information is only used by the memory management system and is otherwise irrelevant.

to each function in the API, a reference to the neighborhood can be created and the functions operate directly over them. The code is greatly simplified by concentrating all clipping logic in a single function and, since image parts are shared, there is no performance loss due to memory allocation.

3.4 The use of an embedded language

As with most artistic processes, the creation of a non-photorealistic filter involves a great deal of trial and error, specially because measuring the quality of the resulting method is a subjective matter. Each new idea brings new challenges and possibilities and to achieve good results it is important that as little time as possible is spent experimenting with them. Therefore, an expressive programming language should be used during the research process.

One possibility would be to use one of the many existing image processing toolkits, such as that provided by Matlab. However, the core operations needed by the painting algorithm, when not available directly from the toolkit, would have to be implemented in a low performance scripting language on which the toolkit is based. Furthermore, the use of most of these tools would present the authors with technical and, possibly, copyright problems when distributing a stand-alone version of the resulting filter. These considerations led us to use the Lua language [8].

Lua is a free language engine that can be easily embedded into any C application. It provides an API that allows the application to exchange data with Lua programs and also to extend Lua with C functions. Thus, the engine can be used to build domain-specific languages, such as the image processing toolkit developed for the present research. These extended languages inherit Lua's features, such as dynamic typing, automatic memory management with garbage collection and associative arrays, all accessible through a comfortable Pascal-like syntax. The simplicity of the language can be seen by Programs 1-7 which are actually real Lua source code, and a major part of our implementation.

Following these ideas, the painting algorithm was developed in a hybrid architecture: the core image manipulation functions were implemented in the C programming language, whereas the main algorithm was implemented in the Lua programming language. This approach brought together the performance of a compiled language and the simplicity of a rapid prototyping language. New ideas could be implemented in the Lua language and put into practice without even recompiling the program. Once acceptable results were reached, time critical operations were re-implemented in C.

4 Examples

The final aspect of a painterly rendered image is controlled by a series of factors. The artist has the freedom to choose the stroke template image, the local area size, and the number of multi-resolution levels.

In this section, we illustrate some of the results that can be obtained by presenting two examples. The stroke template image for both examples is a gaussian blob.

Figure 9 is the portrait of an old man. Figure 9(a) shows the original image and Figure 9(b) shows the result of applying the algorithm with the following parameters: *stroke size* = 10 and *number of levels* = 4.

Figure 10 is a photograph of a tiger. Figure 10(a) shows the original image and Figure 10(b) shows the result of applying the algorithm with the following parameters: *stroke size* = 15 and *number of levels* = 4.

5 Conclusions

In this work, we introduced a multi-resolution approach to the painterly rendering by local source image approximation method. The development of the research gave rise to a parametrized stroke distribution algorithm, easily implemented. The use of a high-level programming language presented a satisfactory framework to speed up the development process. Finally, a specially designed image abstraction simplified and optimized the implementation.

5.1 Future work

The image moment theory provides a powerful tool in the creation of local source image approximations. The quality of these approximations, however, is strongly dependent on the quality of the local color difference images. In a future work, we intend to investigate alternatives to the color difference images, attempting to produce better controllable results.

The encoding and storage of stroke lists will also be studied. By compacting the information provided by the lists, it is possible to represent painterly rendered images in a space efficient way. If the new segmentation techniques lead to substantial improvements, it may even be possible to encode photo-realistic images, transforming the scheme into an image compression algorithm.

Acknowledgments

This work was developed at the VISGRAF Laboratory of IMPA and TECGRAF of PUC-Rio.

We would like to thank Danilo Tuler de Oliveira and Diogo Vieira Andrade who participate in the early stages of this research.

The authors are partially supported by research grants from the Brazilian Council for Scientific and Technological Development (CNPq) and Rio de Janeiro Research Foundation (FAPERJ).

References

- [1] Cassidy J. Curtis, Sean E. Anderson, Joshua E. Seims, Kurt W. Fleischer, and David H. Salesin. Computer-generated watercolor. *Proceedings of SIGGRAPH 97*, pages 421–430, August 1997.
- [2] Gershon Elber. Line Art Illustrations of Parametric and Implicit Forms. *IEEE Transactions on Visualization and Computer Graphics*, 4(1), January – March 1998. ISSN 1077-2626.
- [3] Stuart Green, David Salesin, Simon Schofield, Aaron Hertzmann, Peter Litwinowicz, Amy Gooch, Cassidy Curtis, and Bruce Gooch. Non-Photorealistic Rendering. *SIGGRAPH '99 Non-Photorealistic Rendering Course Notes*, 1999.
- [4] Paul E. Haeberli. Paint by numbers: Abstract image representations. *Proceedings of SIGGRAPH 90*, 24(4):207–214, August 1990.
- [5] Aaron Hertzmann. Painterly rendering with curved brush strokes of multiple sizes. *Proceedings of SIGGRAPH 98*, pages 453–460, July 1998. ISBN 0-89791-999-8. Held in Orlando, Florida.
- [6] Aaron Hertzmann and Ken Perlin. Painterly rendering for video and interaction. *NPAR 2000 : First International Symposium on Non Photorealistic Animation and Rendering*, pages 7–12, June 2000.
- [7] Aaron Hertzmann and Denis Zorin. Illustrating smooth surfaces. *Proceedings of SIGGRAPH 2000*, July 2000. Held in New Orleans, Louisiana.
- [8] Roberto Ierusalimsky, Luiz Henrique de Figueiredo, and Waldemar Celes. Lua: an extensible extension language. *Software: Practice & Experience*, 26(6):635–652, 1996.
- [9] Allison W. Klein, Wilmot W. Li, Michael M. Kazhdan, Wagner T. Correa, Adam Finkelstein, and Thomas A. Funkhouser. Non-photorealistic virtual environments. *Proceedings of SIGGRAPH 2000*, pages 527–534, July 2000. ISBN 1-58113-208-5.
- [10] Peter Litwinowicz. Processing images and video for an impressionist effect. *Proceedings of SIGGRAPH 97*, pages 407–414, August 1997. ISBN 0-89791-896-7. Held in Los Angeles, California.
- [11] Barbara J. Meier. Painterly rendering for animation. *Proceedings of SIGGRAPH 96*, pages 477–484, August 1996. ISBN 0-201-94800-1. Held in New Orleans, Louisiana.
- [12] Victor Ostromoukhov. Digital facial engraving. *Proceedings of SIGGRAPH 99*, pages 417–424, August 1999. ISBN 0-20148-560-5. Held in Los Angeles, California.
- [13] Yachin Pnueli and Alfred M. Bruckstein. Digidurer - a digital engraving system. In *The Visual Computer*, volume 10, pages 277–292, August 1994.
- [14] Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. Real-time hatching. *Proceedings of SIGGRAPH 2001*, pages 579–584, August 2001. ISBN 1-58113-292-1.
- [15] Michio Shiraishi and Yasushi Yamaguchi. Image moment-based stroke placement. Technical Report skapps3794, University of Tokyo, Tokyo Japan, May 1999.
- [16] Michio Shiraishi and Yasushi Yamaguchi. An algorithm for automatic painterly rendering based on local source image approximation. *NPAR 2000 : First International Symposium on Non Photorealistic Animation and Rendering*, pages 53–58, June 2000.
- [17] François X. Sillion. The state of the art in physically-based rendering and its impact on future applications. *Second Eurographics Workshop on Rendering (Photorealistic Rendering in Computer Graphics)*, pages 1–10, 1994. Held in New York.
- [18] David Small. Simulating watercolor by modeling diffusion, pigment, and paper fibers. *Proceedings of SPIE '91*, February 1991.
- [19] Mario Costa Sousa and John W. Buchanan. Observational models of graphite pencil materials. *Computer Graphics Forum*, 19(1):27–49, March 2000. ISSN 1067-7055.
- [20] Steve Strassmann. Hairy brushes. *Siggraph*, 20(4):225–232, August 1986.
- [21] Saeko Takagi, Masayuki Nakajima, and Issei Fujishiro. Volumetric modeling of colored pencil drawing. *Pacific Graphics '99*, October 1999. Held in Seoul, Korea.
- [22] Georges Winkenbach and David H. Salesin. Computer-generated pen-and-ink illustration. *Proceedings of SIGGRAPH 94*, pages 91–100, July 1994. ISBN 0-89791-667-0. Held in Orlando, Florida.
- [23] Eric Wong. Artistic rendering of portrait photographs. Master's thesis, Cornell University, 1999.




Figure 9: Oldman.



Figure 10: Tiger.