Non-Photorealistic Rendering: From a general view to Suggestive contours

Giorgio Conte and Massimo De Marchi

EVL Creative Coding
CS 488 - Computer Graphics
University of Illinois at Chicago

November 2014

Abstract

Non-Photorealistic Rendering (NPR) is a very important branch of computer graphics. It was born in 1990s and from that time it has gained more and more popularity among the academic community. NPR contains applications in many different contexts which makes this field as interesting as stimulating. In this report, we have focused first our attention on understanding the main steps done in the history of this discipline, then we have investigated some of the most well-known techniques and, in the end, we focused our attention on suggestive contours. After having analyzed the theoretical aspects, we have implemented one image-based algorithm to render this kind of effect.

1 Introduction

Computer Graphics is concerned with the development of tools for generating various forms of imagery. People and artists are interested in these tools for the production of images to communicate visual information and the important point is that appropriate form of imagery for a given task depends upon the nature of the communication. Computer graphics has an amazing success story considering realistic rendering. Conversely, there is a branch of computer graphics that seeks for something different, namely principles and techniques known to artists for many centuries, that allows to get rid of unnecessary details and that bias the user experience and perception. Non-Photorealistic Rendering (NPR) is the name often used to refer to all the techniques that follows this principles and in the last twenty years has became such a popular topic that lot of efforts have been spent to find new algorithms.

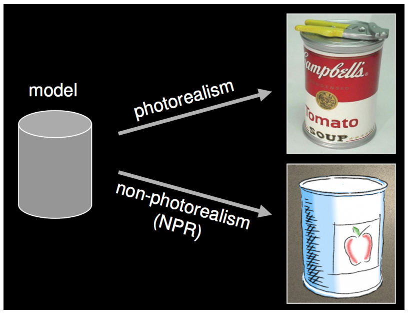

Figure 1: NPR vs Photorealism

Long before computer graphics became famous and popular, artists knew the distinction between photorealism and non-photorealism. Looking at figure 2 you could imagine it as a photo while a painting like the one in figure 3 exhibits a level of abstraction that remarkably departs from realism.

Figure 2: Bouguereau, The Little Shepherdess, 1891

Photorealism is certainly good for the job of documentation, where the aim is to record scenes or objects in all their details. Crime scene photography is a perfect example. Should new evidence or theory prompt a reinvestigation, crime scene photos may be probed for new clues. NPR, such as a sketch of the scene, could record the relative positions of the things present in the site, but any details not drawn by the artist at that time is going to be permanently lost.

In addition, photorealistic computer graphics is uniquely suited to the situation where one would like to create live-action scenes (e.g. movies) but where its difficult to stage the real scene. A possible task could be to simulate a special effect such as huge explosions, or a flood over Chicago lakeshore requiring photorealistic computer graphics imagery for convincing effect.

Even if we know that these effects are not real, cartoonish NPR could not be exploited to render those scenes because it would destroy the illusion necessary for the audience to break of disbelief and enjoy the story.

Figure 3: Van Gogh, The Starry Night, 1889

NPR may be the appropriate

means for communication for some tasks, but basically it depends on the situation. Since computer graphics has traditionally focused his attention mostly on photorealistic rendering, there are still lots of unexplored opportunities for new forms of content for authors. However, there are also various applications where NPR is recommended, for example when image is used for an *explanation* such as in figure 4. Given that, very likely the user has the real tool in front of him, a photographic image contains lots of unnecessary details; instead, the author can better explain the shape and function of the item by disambiguating different parts using distinct colors and remove physically accurate shading that reduce the clarity of the image. Silhouette and feature lines are often used to increase the objects description.

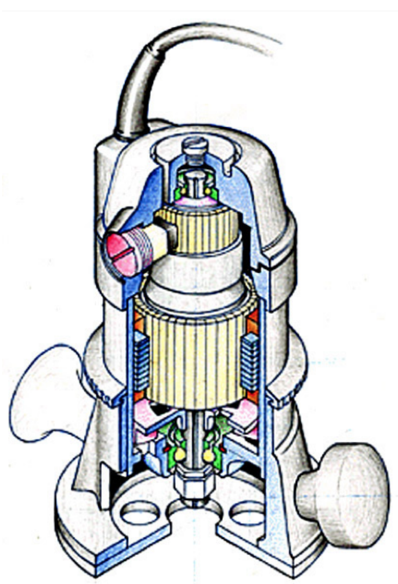

Figure 4: Technical illustration of a mechanical tool

Stylization is frequently used to *illustrate* abstract concepts using very concise image. In figure 5 the artist communicates an abstract idea of the role of higher education in our society. If we imagine to replace this with a similar photoreal image, the abstract concept would not have been as effective as it is.

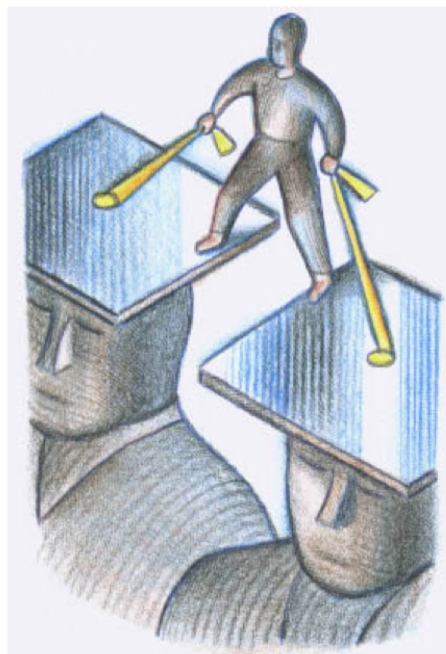

Figure 5: Sutter

Finally, *storytelling* is the last example where stylized and abstract image can help to communicate scenes which exist only in the author's imagination. The effect would be totally different with comparable photoreal drawing. Moreover, children, the common audience in this case, generally respond best to brightly colored illustrations, and somehow find them more engaging than photos.

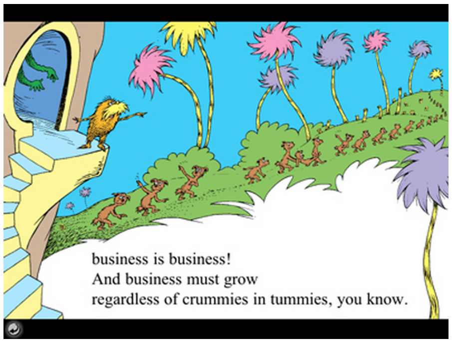

Figure 6: Dr. Seuss

In this document we would like first to give a brief overview of some of the best known techniques and algorithms present in the literature then describe more in the detail an image-space algorithm to render suggestive contours and show the result obtained with the implementation. The report is structured as follows. First in section 2 we briefly outline the history of NPR and in section 3 we present some non-photorealistic algorithms and techniques found in the literature. In section 4 we describe in details how to render suggestive contours and in section 5 we presents the results obtained. Finally in section 6 we outline the current research issues and we draw some conclusions in section 7.

2 History of NPR

Non-photorealistic rendering is a relatively young field that emerges in the 1990s as an alternative to photorealism. The term *non-photorealistic rendering* was probably coined by George Winkenbach and David Salesin in their paper dated 1994 [19] and it has been criticized for many reasons among which the fact that the term *photorealism* is ambiguous and graphics researchers and artists refer to it with different meanings. Moreover many of the techniques that are used in non-photorealistic rendering are not actually *rendering* techniques; in fact many of the algorithms in the literature are based on image post-processing. Research efforts were directed to simulate hand drawn illustrations, namely attempt to simulate what an artist might do in order to enhance comprehensibility of visual information. However, even before the work done by Winkenbach, in 1990, Takafumi Saito [17] proposed a new rendering technique that tries to achieve this by implementing an algorithms for discontinuities, edges, contour lines, and curved hatching.

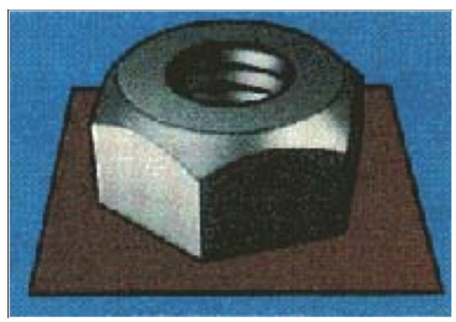

Figure 7: Saito 1990

Subsequently many other techniques have been proposed among which Pen & Ink [19], watercolor [6] and curved brush strokes [10].

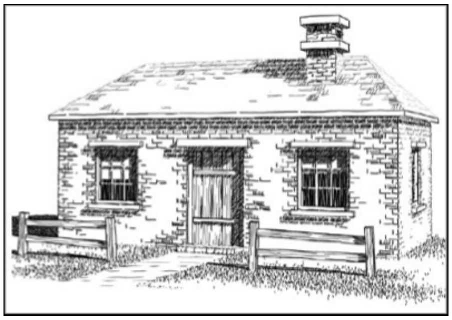

Figure 8: Winkenbach 1994

Figure 9: Curtis 1997

Figure 10: Hertzmann 1998

All these techniques were focusing on static imagery, but in the latter part of the decade researchers started to be interested in dynamic NPR, that is allowing the production of imagery that would be otherwise tedious to do by hand. One of the main issues was to pursue *frame to frame coherence* as described by Meier [15] and Litwinowicz [13] works.

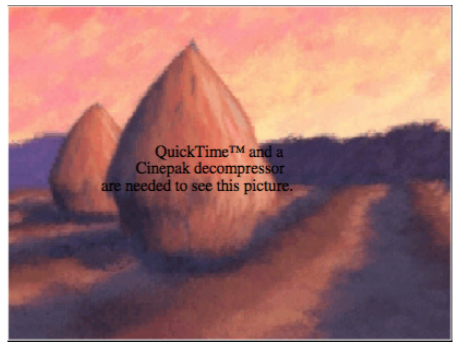

Figure 11: Meier 1996

Figure 12: Litwinowicz 1997

These result were achieved offline, but various works among which Kolwasky [11], Gooch [8] and Praun [16], thanks also to the evolution of graphics hardware, opened the door for realtime animations. This was important because for the first time it was possible to create interactive illustrations, otherwise impossible to do by hand.

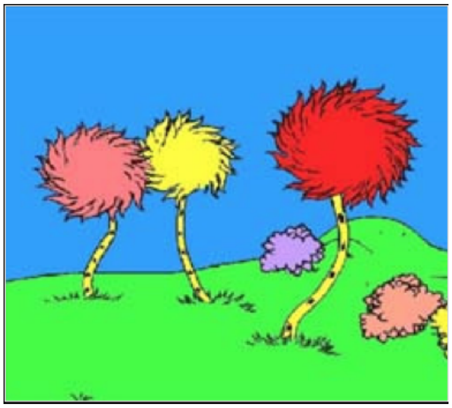

Figure 13: Kowalski 1999

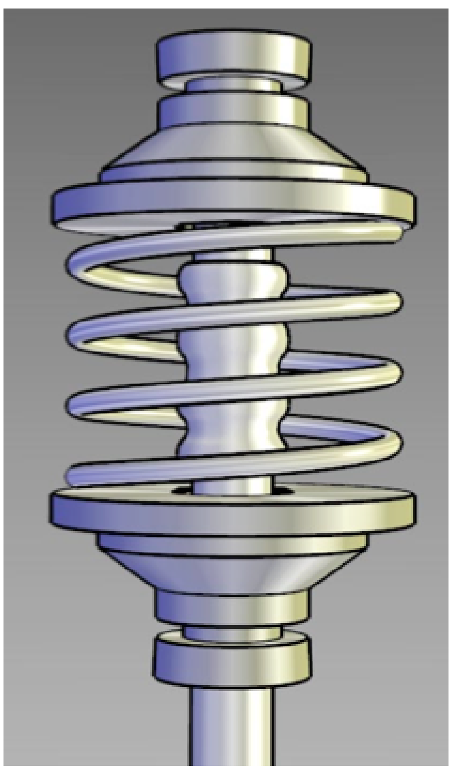

Figure 14: Gooch 1998

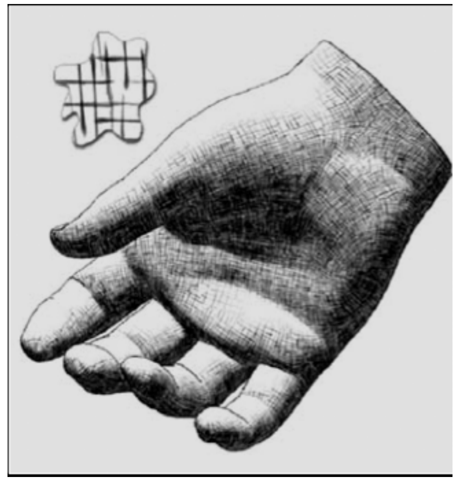

Figure 15: Praun 2001

Nowadays, more recent researches are trying to exploit NPR for visualization, especially for the production of medical illustrations useful for many applications.

3 Techniques

In this section we would like to present and briefly describe some techniques which made the history of Non-Photorealistic rendering.

3.1 Pen and Ink

A very interesting and effective technique is called "pen and ink" and it was presented first by Winkenbach and Salesin in [19], 1994. This was one of the first NPR technique and it has became quite famous in the academic literature. The basic idea is to render object and images as they were drawn by an artist just using pen and ink. The pen and ink effect has a very long story, since it first appeared in manuscripts of the Middle Ages, but

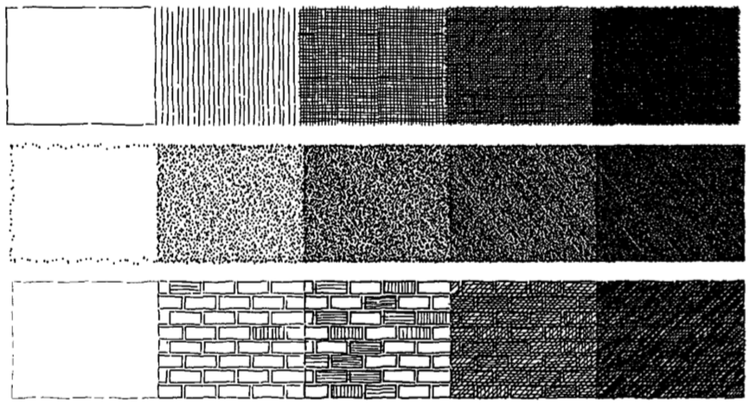

Figure 16: A collection of stroke textures presented by [19]

.

only in the last century this kind of effect has been developed as an art form by itself. Pen-and-ink illustrations are very simple and effective, this simplicity makes them especially attractive. Moreover, each stroke could also gain more expressiveness by adding small irregularities in its path and pressure. It is worth to remind that in this kind of illustrations, the drawer could only vary the tone just by changing the density of the strokes.
To render this effect, the authors have introduced the notion of "stroke textures" which are a set strokes used to produce both texture and tone. In figure 16, some of them are depicted. Each texture has a given priority and the author mix those textures in a post processing phase, after having applied a simple Phong shader. The tone depends basically on the diffuse light. Then to achieve a better effect they have developed an algorithm that allows the artist to "indicate" which area should be rendered as white. This technique is called "indication". Stroke textures can also be used to render shadows as it is depicted in figure 17.

3.2 Cel shader

Cel shading or toon shading is a well know technique to render objects with the style of cartoon or comic books. This technique is somewhat recent and it started to appear from around the beginning of twenty-first century. Basically, there are two main techniques that are used to create a toon effect: a tone-based technique and texture-based one. The first shading approach is based on the diffuse light taken from Phong shader. The simple idea is to discretize the result computed by diffuse light with just three or four values. Thus, it is possible to reduce the number of tones rendered and make the scene look like a cartoon. The second method is based on 1D or 2D textures. The principal idea behind this procedure is to create a texture with just a bunch of colors. Then, the color of each pixel is taken from this tex-

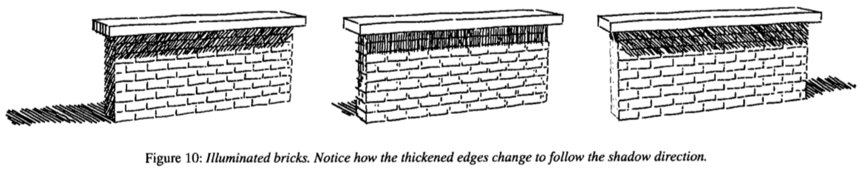

Figure 17: An image rendered using strokes textures and shadows. Image taken from [19]
.

ture using the diffusion coefficient. It is also possible to extend this method using a 2D texture. Then, to choose the right matching color we could use a new parameter which, in the vast majority of the case, is depth. Figures 18 and 19 explain better the two techniques described so far. To get a more "realistic" result, it is possible also to render a black contour around the shapes. Many algorithms both object-space and image-space have been described in the literature, but we would like to point out the one which we describe in section 4.

Figure 18: This picture shows an example of 2D texture that could be used to render a cel shaded object.

3.3 Stylized highlights

A more recent technique used to render cartooned scene is called "stylized highlights" and it was presented by Anjyo et al. in [1]. This technique focuses its attention in how to render the specular light in a toon scene. In particular, the authors explain how to slightly modify the result of the Blinn's specular model [4] in order to obtain a better effect. The algorithm proposed interactively modifies the initial shape through geometric, stylistic and Boolean transformations. In fact in cel shading, the highlight must be a semantic notation rather than a part of the physics. The effect achieved by author is represented in figure 20, while a scene rendered using this shading technique is depicted in figure 21.

Figure 19: This picture represents how to create toon shading using a one dimensional textures.

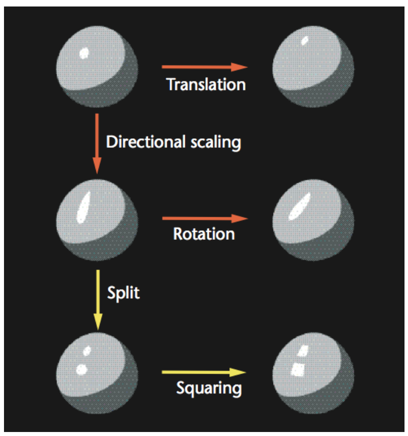

Figure 20: This picture shows all the transformation that can be applied using [1]. On the lower right it is possible to see the final result.

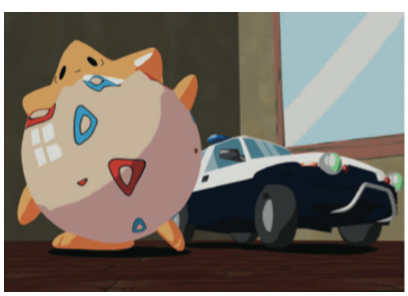

Figure 21: A character and a scene rendered using "stylized highlights" shading.

3.4 Technical drawings

As we were saying before, NPR has became very important also for representing technical illustration. In fact, it has been showed that non photo-realistic images are more useful for explaining how a mechanical piece is done. Using a NPR, it is possible to eliminate all the unnecessary details that bias the user. The lighting model proposed by [8] allows shading to occur only in mid-tones so that edge line are more visible. This technique gives a clearer picture of the shape, structure and material than the traditional computer graphics techniques do. Basically, the technique presented aims to render shape with a shading algorithm based on cool-to-warm tones. From a mathematical point view, the color of each pixel is given by the following formula:

$$I = (\frac{(1+\vec{(}I)\cdot\vec{(}N)}{2})k_{cool} + (1 - \frac{1+\vec{(}I)\cdot\vec{(}N)}{2})k_{warm} \tag{1}$$

where $\vec{I}$ is the light vector, $\vec{N}$ is the normal of the surface in that point, k_{cool} is a cold color and k_{warm} is a warm color. In the example provided by the authors they use a palette that goes from blue to yellow so k_{cool} and k_{warm} are computed as follows:

$$k_{cool} = k_{blue} + \alpha k_d \tag{2}$$
$$k_{warm} = k_{yellow} + \beta k_d \tag{3}$$

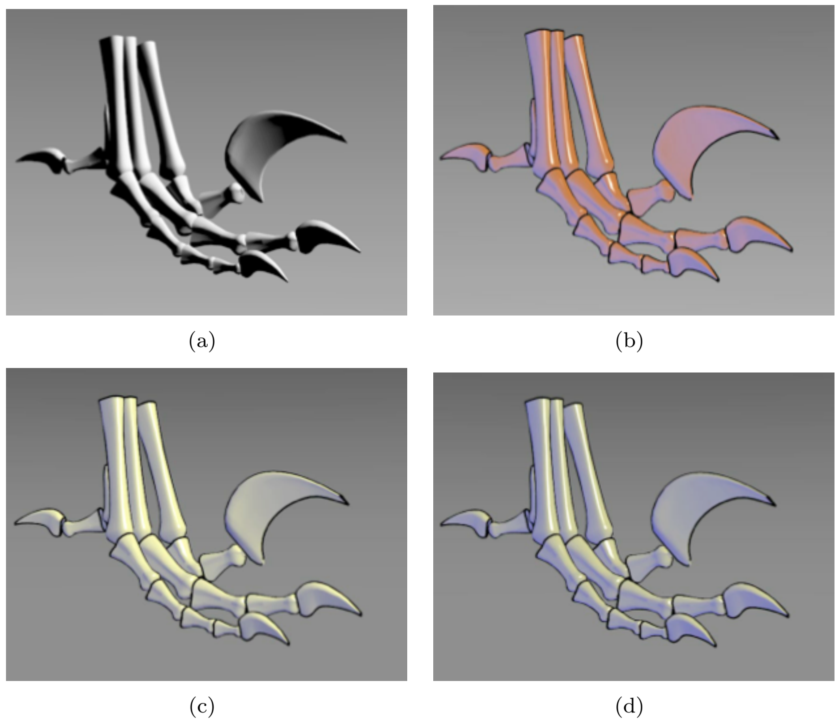

Figure 22: In figure (a) a phong shaded object is rendered while in (b), (c) and (d) the warm cool effect is applied with different configurations.

Where k_{blue} and k_{yellow} are the tones of the blue and yellow respectively, while k_d is the color of the object according to the diffuse shading light. α and β will determine the prominence of the object color and the strength of the luminance shift and they are free parameters. In figures 22(a), 22(b), 22(c) and 22(d) it is possible see the same object rendered with this technique, but using different configurations.

The same technique has also been used to render metallic material in a non photorealistic way. Usually drawers represent a metallic surface by alternating dark and light bands. Lines are drawn in the direction of the axis of minimum curvature. To simulate this kind of effect, the authors map a collection of twenty stripes of varying intensity along the axis of maximum curvature. The stripe intensities vary between 0.0 and 0.5 and the stripes closest to the source light are overwritten with white. Within the stripes the colors are blended together in a linear way. Then the cool-warm effect described before is added, while the edge lines are generated interactively using the technique presented in [14]. Figure 23 shows the results obtained with the technique described.

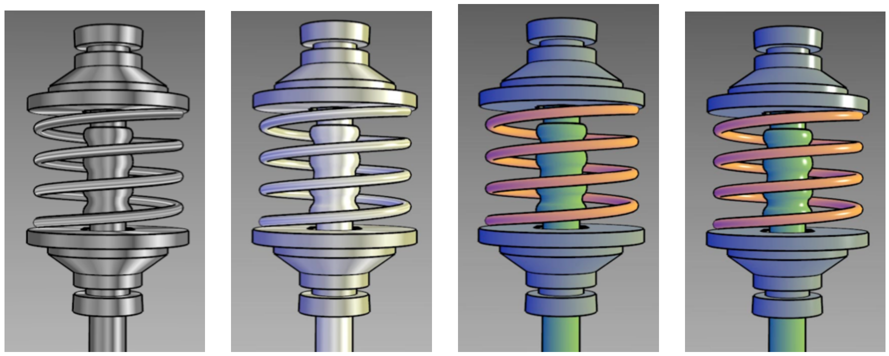

(a) Achromatic metallic shading (b) Metal effect with cool-to-warm hue (c) Colored metal object (d) Colored metal object with specular light

Figure 23: This set of pictures shows how different effects can be created with the technique described presented in [8]. In subfigure (a) only the metal effect is rendered while in subfigure (b) also the tone shift is added. In subfigures (c) and (d) the object color is also rendered and in subfigure (d) a specular light is added.

3.5 Painterly rendering with brush strokes

Another technique often used to produce images for graphic design and illustrations employs brush strokes to simulate an hand-painted effect. This approach is very useful because it allows to reproduce various artistic styles and give emphasis to what is important in a given photograph. All the algorithms present in the literature try to address the various challenges of this approach; the first of these is *physical simulation*, and impressive works have been done by [6] and [5]. Another challenge is to provide artists with user-friendly tools that take care of the mechanical details of image creation while providing the user parametric styles to produce various effects. This is useful especially for animation, where the animator can select the style to be applied and the algorithms take care of all the tedious work. The first researchers that addressed this challenge are Litwinowicz [13] and Hertzmann [10]. Basically this technique works in way similar to how artists paint, namely starting with rough strokes and going back over the painting with smaller strokes to add details. Each layer is painted with different stroke size and this helps to emphasized certain details that the artist wants to be highlighted. In [10] Hertzmann describes an algorithm that follows this approach; more in detail the algorithm produces the final result adding thinner strokes layer by layer in a pyramid manner. Each layer starts upon a blurred version of the original photograph and the brushes are applied to areas that are as large as the brush size; according to this method importance is given to the area of the image that contains most detail since it is where the smaller brush strokes are applied. In order to add expressiveness to the algorithm, Hertzmann modeled the brush strokes as anti-aliased cubic B-splines, arguing that curved strokes better express silhouette and light effects.

4 Suggestive Contours

In this section we are going to describe first what suggestive contour are, then the algorithm that we use to create them.

4.1 Overview

Suggestive contours [7] is the topic on which we mostly focused our attention. Suggestive contours are lines different from creases and traditional contours. In fact, this new kind of line presented by the authors are lines drawn in visible parts of the shape that could be traditional contours from points of view close to the real camera position. As it was explained in the paper, suggestive contours could live consistently with true contours, because they anticipate and extend the true ones. Suggestive contours are described by the authors as "features where a surface bends sharply away from the viewer, yet remains visible" so that they are almost contours and actually they become contours in nearby views. Figure 25 better represents the concepts described so far.

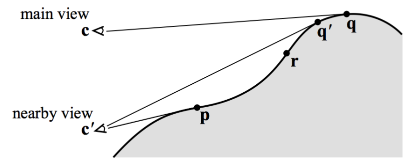

Figure 25: A frame showing both contours (q) and suggestive contours (p), as seen by the main camera c. When the camera moves from c to c', p becomes a contour, while the contour in q slides along the surface to q'. Image taken from [7].

Mathematically speaking, the authors gave three equivalent definitions of the suggestive contours. Here we report the one that, in our opinion, is simpler and clearer to be understood. The definition says: "the suggestive contour generator is the set of minima of $\mathbf{n} \cdot \mathbf{v}$ in the direction of $\mathbf{w}$". Where $\mathbf{n}$ is the normal defined in that point, $\mathbf{v}$ is the view vector and $\mathbf{w}$ is the projection of the view vector on the tangent plane in $\mathbf{p}$.

4.2 Algorithm

Two algorithms are presented by the authors to render suggestive contours consistently. The first is object-based, that means that the algorithm works directly using the shape of the objects and it does not need a post processing phase, while the second one is image-based, that means that it is a filter applied to an image of the rendered scene. For reason of simplicity and effectiveness, we concentrate our effort in the second one. In fact, it is possible to detect suggestive contours directly from a rendered image. The solution implemented reflects the definition we have reported in the subsection 4.1 and the

(a) Original photograph (b) Impressionist style applied

Figure 24: An example of what the algorithm by Hertzmann can produce

algorithms works as follows. First, we render the entire scene using a diffuse light located in the camera origin and then, in the rendered image, the algorithm detects suggestive contours as steep valleys in intensity. Since it is not possible to detect the radial direction from the rendered image, the algorithm just find all the valleys. Despite many algorithms have been proposed to solve this problem, the authors suggest a very simple and effective filter and the idea is as follows. For each pixel, the shader retrieves the brightness of neighbour pixels, we actually do not consider a circle window as it has been suggested by the authors, but we use a square. Across the squared window the greatest intensity is p_{max}. Then the pixel is considered a valley if the following two conditions are met:

- $p_{max} - p_i$ exceeds a fixed threshold d
- no more than a certain percentage s of the pixels in the window are strictly darker than p_i

Finally we just add the true contours and the suggestive contours to render the final scene. In figure 26 the results obtained by the authors are reported.

5 Results

The results achieved using the algorithm described before are very good. To better compare the suggestive contours with the true contour we rendered the same object using the two different shaders in the same scene. As it is possible to see from the image 27, the effect created allows users to better understand the shape of the object rendered. In fact, much more details are drawn with respect to the basic contour shader.

6 Current research issues

At the state of the art there are many challenges in the field of Non-Photorealistic Rendering. Gooch in his paper [9] revised some of the major challenges and in this section we are going to look at them.
The first challenge is studying how to encode what makes something beautiful, and Gooch refer to that as "Algorithm aesthetics". Part of this challenge involves the stimulation of visual area of the brain.
The second important challenge is how to capture the essence of an im-

(a) Image rendered with true contours

(b) Image rendered with suggestive contours

Figure 26: A comparison between the same object rendered without (a) and with suggestive contours (b) .

age through abstraction, like Antoine Saint-Exupéry says:

> *"You know you've achieved perfection in design, Not when you have nothing more to add, But when you have nothing more to take away."*

This is a very important goal because it allows to easily recognize the most salient information. A closely related challenge is represented by the new area of Visual analytics which faces the problem to map salient feature in the data with suitable visual representations. There are many fields of study which can take advantage from visual analytics such as medical visualization.

An interesting topic that combines *medical aspects* and *non photorealistic rendering* is the one proposed by Lee et al. in [12]. In this paper the authors proposed a real time rendering tool to help surgeons in minimally invasive interventions. In particular, some techniques to track tissue deformation, 3D reconstruction and augmented reality for robotic assisted surgery are discussed and it is also highlighted the importance of integrate pre-operative data and intra-operative data in order to obtain a good and accurate surgical navigation framework.

The fourth challenge is referred by Gooch as "interaction that enables right-brain thinking", namely enable user with poor artistic skills to create compelling visuals and let artist express themselves without having system parameters get into the way. The fifth challenge is about CG imagery evaluation, and this can be done through a sort of "Turing test".

Finally one of the toughest challenges is to explore new artistic media, and many techniques are moving toward this direction such as Analysis/Synthesis [18] and Evolutionary Art.

Evolutionary art is a recent research topic which try to apply artificial in-

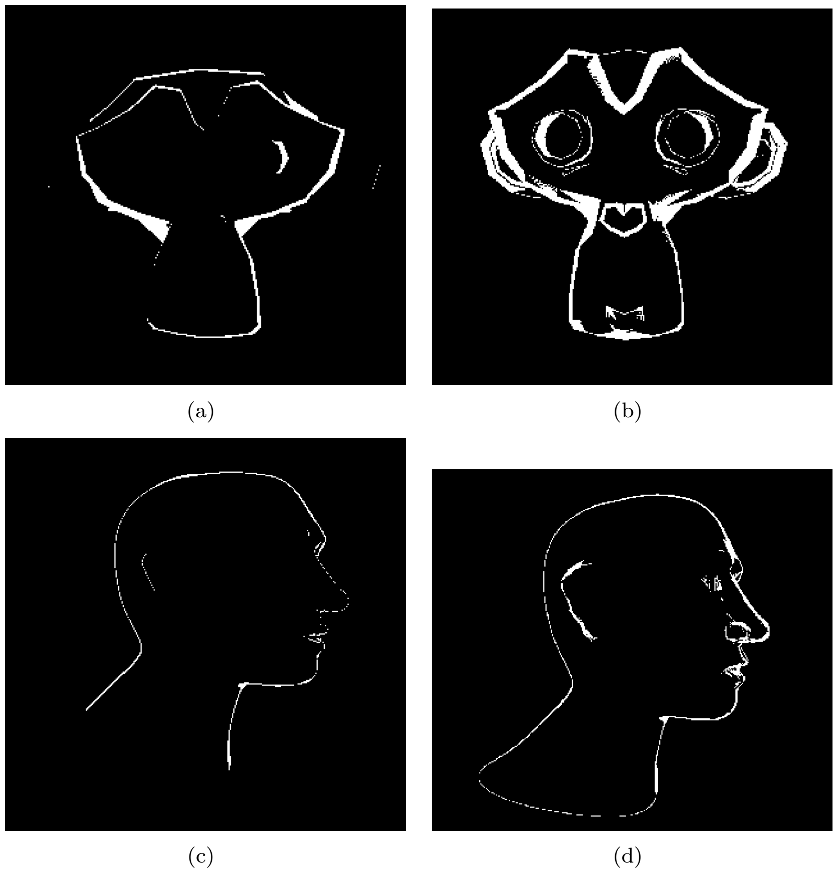

Figure 27: In figure (a) a monkey model is rendered using simple contours while in figure (b) the same model is rendered using also suggestive contours. Another comparison is shown with a face model in figure (c) and figure (d).

telligence algorithms and techniques in the NPR field. The topic was introduced first in [3] and within this context Baniasadi and Ross have presented in [2] a new approach using genetic algorithms. With this system, a source image is taken as input, and a genetic program is evolved that will re-render the image with non-photorealistic effects. The effect rendered is depicted in figures 28(a) and 28(b).

7 Conclusion

Non-photorealistic rendering is a very interesting topic and it can be declined and applied to very different contexts. In fact, as we have described applications ranges from art to scientific visualization. This wide range of applications make the NPR field a very important branch of computer graphics. NPR is able to collect and to embed lots of different ideas from many contexts and this makes it a very appeal-

(a) Input image (b) Brush effect

Figure 28: An example of *evolutionary art*

ing field of research.
The results obtained from the implementation of the algorithm for suggestive contours are very satisfactory, because we are able to correctly reproduce the proposed technique. The effect obtained is stunning, especially if compared to the "true" contour technique. Since their birth, NPR techniques have always attempt to imitate great artworks of the past, but the ultimate challenge is to proof that NPR can be be an artistic medium on its own and it can create the art of the future.

References

[1] K-i Anjyo and Katsuaki Hiramitsu. Stylized highlights for cartoon rendering and animation. *Computer Graphics and Applications, IEEE*, 23(4):54–61, 2003.

[2] Maryam Baniasadi and Brian J Ross. Exploring non-photorealistic rendering with genetic programming. *Genetic Programming and Evolvable Machines*, pages 1–29, 2013.

[3] Perry Barile, Vic Ciesielski, and Karen Trist. Non-photorealistic rendering using genetic programming. In *Simulated Evolution and Learning*, pages 299–308. Springer, 2008.

[4] James F Blinn. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics*, volume 11, pages 192–198. ACM, 1977.

[5] Fractal Design Corporation. Fractal design painter.

[6] Cassidy J Curtis, Sean E Anderson, Joshua E Seims, Kurt W Fleischer, and David H Salesin. Computer-generated watercolor. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 421–430. ACM Press/Addison-Wesley Publishing Co., 1997.

[7] Doug DeCarlo, Adam Finkelstein, Szymon Rusinkiewicz, and Anthony Santella. Suggestive contours for conveying shape. *ACM Transactions on Graphics (TOG)*, 22(3):848–855, 2003.

[8] Amy Gooch, Bruce Gooch, Peter Shirley, and Elaine Cohen. A non-photorealistic lighting model for automatic technical illustration. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 447–452. ACM, 1998.

[9] Amy A Gooch, Jeremy Long, Li Ji, Anthony Estey, and Bruce S Gooch. Viewing progress in non-photorealistic rendering through heinlein's lens. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*, pages 165–171. ACM, 2010.

[10] Aaron Hertzmann. Painterly rendering with curved brush strokes of multiple sizes. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 453–460. ACM, 1998.

[11] Michael A Kowalski, Lee Markosian, JD Northrup, Lubomir Bourdev, Ronen Barzel, Loring S Holden, and John F Hughes. Art-based rendering of fur, grass, and trees. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 433–438. ACM Press/Addison-Wesley Publishing Co., 1999.

[12] Su-Lin Lee, Mirna Lerotic, Valentina Vitiello, Stamatia Giannarou, Ka-Wai Kwok, Marco Visentini-Scarzanella, and Guang-Zhong Yang. From medical images to minimally invasive intervention: Computer assistance for robotic surgery. *Computerized Medical Imaging and Graphics*, 34(1):33–45, 2010.

[13] Peter Litwinowicz. Processing images and video for an impressionist effect. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 407–414. ACM Press/Addison-Wesley Publishing Co., 1997.

[14] Lee Markosian, Michael A Kowalski, Daniel Goldstein, Samuel J Trychin, John F Hughes, and Lubomir D Bourdev. Real-time nonphotorealistic rendering. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 415–420. ACM Press/Addison-Wesley Publishing Co., 1997.

[15] Barbara J Meier. Painterly rendering for animation. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 477–484. ACM, 1996.

[16] Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. Real-time hatching. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, page 581. ACM, 2001.

[17] Takafumi Saito and Tokiichiro Takahashi. Comprehensible rendering of 3-d shapes. In *ACM SIGGRAPH Computer Graphics*, volume 24, pages 197–206. ACM, 1990.

[18] Javier Villegas and Angus Graeme Forbes. Analysis/synthesis approaches for

creatively processing video signals. In *Proceedings of the ACM International Conference on Multimedia*, pages 37–46. ACM, 2014.

[19] Georges Winkenbach and David H Salesin. Computer-generated pen-and-ink illustration. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 91–100. ACM, 1994.