

The Use of Optic Flow in the Painterly Rendering of Animated Models

David Maclay
Collaborative Visual Computing Laboratory
Dept. of Computer Science
University of Cape Town
dmaclay@cs.uct.ac.za

Edwin Blake
Collaborative Visual Computing Laboratory
Dept. of Computer Science
University of Cape Town
edwin@cs.uct.ac.za



Fig.1. An example of San rock art of the sort that this project hopes to animate.

ABSTRACT

Image based painterly rendering has always had difficulty enforcing frame-to-frame continuity in animations. This paper lays out an approach that uses a form of optic flow to ensure continuity in the painterly rendering of 3D models characteristic of virtual environments. The proposed technique will be used to emulate the rock painting style of the San people of southern Africa.

KEYWORDS – Painterly Rendering, Optic Flow, Animation, Billboards, Image-based, Non-Photorealistic.

1. INTRODUCTION

"Computer graphics has long been defined as a quest to achieve photorealism. As it gets closer to this grail, the field realizes that there is more to images than realism alone." [1]

This research sets out to investigate the use of optic flow techniques in the painterly rendering of animated 3D models. Extensive research has already been done on the use of Optic flow in the photorealistic rendering of animated scens. Even more research has been done on the painterly rendering of 3D models, both animated and static. The use of both Optic flow

and painterly rendering in one integrated system has not been extensively explored.

2. THE BENEFITS OF PAINTERLY RENDERING

Most current methods of rendering 3D scenes involve the use of models comprised of polygons. While this technique is excellent for urban environments, it is difficult to portray scenes comprising predominantly organic forms. Methods that attempt to exactly depict such complex structures often produce an effect which appears wrong to human observers.

"The focus of most rendering research in the last two decades has been on the creation of photorealistic imagery. These methods are quite sophisticated, but tend to create imagery that is mechanical-looking because detail is represented very accurately. Recently there has been a movement toward more creative and expressive imagery in computer graphics but few techniques that provide ways to achieve different looks, especially for animation." [2]

Painterly rendering produces images which don't have a mechanical appearance. This is particularly important for natural scenes, which contain many irregularities in the real world.

The realistic modeling and rendering of natural scenes is computationally expensive due to their complexity, therefore a deliberately non-photorealistic approach with rendering shortcuts offers great benefits.

A further advantage of painterly rendering is that very simple models can produce rich images, a tree modeled as a sphere mounted on a cylinder could be complex and interesting if portrayed using the correct brush strokes, and no two trees would appear the same.

Finally, it must be taken into account that artistic painting is an established field, with a vast and varied cultural heritage.

By using a painterly style of rendering it is possible to exploit established analogies, a brush stroke can represent a leaf or even the entire foliage of a tree. The same analogies also exploit a certain amount of 'artistic license' - perfect accuracy is not expected.

Established styles of painting vary from one culture to the next. By selecting a style of painting familiar to members of a particular culture it may be possible to produce more engaging virtual environments.

3. FRAME TO FRAME COHERENCE

One of the biggest problems in painterly animation is to produce images that are coherent over time. The earliest methods of painterly rendering took individual still images as input with no real provision for animation [3]. The pseudo-random nature of most methods for generating brush strokes means that the distribution of brush strokes would change randomly from one frame to the next producing an image that appeared as if viewed through a shower door [2].

Meier addressed this problem by extending a particle rendering system developed by Reeves and Blau [4]. Objects are modeled as persistent 3D point clouds, which are rendered using 2D brush strokes. This technique achieves perfectly coherent animation, the one drawback is that the brush strokes are limited to rather uniform lozenges, one per particle.

To generate more varied brush strokes it is necessary to generate the shape of the brush strokes based on pre-rendered 2D images of the sort a human artist would observe. Such image-based rendering techniques permit long flowing brush strokes which are more visually pleasing [5]. However, to achieve coherent animation without using a particle system some means of correlating elements of the 2D image with the underlying 3D model is required.

4. THE USE OF OPTIC FLOW

Optic flow is used mainly in the field of computer vision. An optic flow field is generated based on changes in successive images of the observed scene. Information about the structure and motion of elements of the scene may be inferred from the optic flow field. Optic flow is a useful tool for creating continuity in image-based animation.

4.1. CONVENTIONAL OPTIC FLOW

The computer vision approach to optic flow has been used successfully to process video for painterly rendering. Because video is used as input, the 3D geometry is unknown and any transformations of reused frames must be based on information extracted from 2D frames. The processing need not be performed in real-time.

To achieve visual consistency from one frame to the next, both [6] and [7] use an optic flow field calculated from consecutive video frames to translate individual brush strokes. Where the brush strokes become too sparse or dense, brush strokes are added or removed respectively.

To make the process more efficient, [7] calculates regions of change based on the difference between the current source frame and the previous source frame after optic flow has been used to warp it. This means that new brush strokes are used only when optic flow methods are not sufficient.

4.2. INVERSE OPTIC FLOW

An alternative form of optic flow is used when the source is a 3D model, as is the case in most virtual environments. In this case the problem is inverted – the geometry is known, but the 2D images must be generated. Information about changes to the 3D scene geometry may be used to transform existing images to approximate successive frames.

The inverse form of optic flow has been used mainly as a rendering shortcut in more photorealistic rendering. This technique was used to reduce the computation of rendering successive frames in the Talisman architecture [8]. Talisman adopts a layer based approach, where different objects are rendered into separate layers. These layers are composited to produce the final image. Optic flow transforms are applied to layers permitting their reuse from frame to frame. Webb [9] extended this work using a Taylor series decomposition of the optic flow field. The relative significance of the more computationally expensive higher order terms is evaluated, permitting their use only when necessary to avoid visible errors.

While advances in computer hardware have made most image-based rendering shortcuts

redundant, there is still great potential for the use of optic flow to solve problems in non-photorealistic animation. To date, some form of optic flow has proven to be the only means of enforcing continuity between consecutive frames of image-based artistic rendering systems.

The inverse approach has yet to be applied to automated painterly rendering, but similar techniques have been used to animate the creations of human artists. Inkwell [10] allows animators to draw onto a set of textures, which are then deformed and translated to provide motion. Harold [11] creates a 3D environment where users create a world by drawing onto a system of billboards located in 3D space. Both of the above systems require that images be created by a human artist.

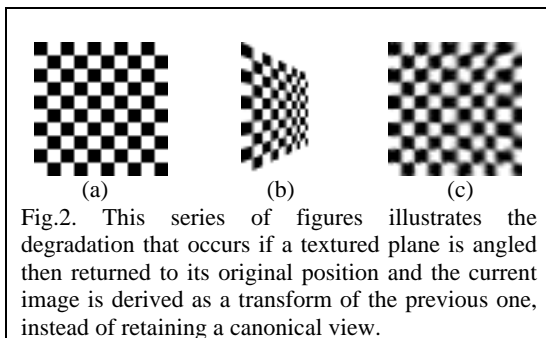
5. OUR TECHNIQUE

Our approach is to animate 3D polygonal models using an image-based technique to generate brush strokes on 2D textures with optic flow transforms to ensure continuity from frame to frame.

5.1. MAPPING PAINTED TEXTURES TO MODEL ELEMENTS

Each independently articulated element of a model has a (RGBA) textured quadrilateral assigned to it. Brush strokes are generated on the otherwise transparent texture. To create the brush strokes, a simple polygonal rendering of each element of the model is produced every frame, and brush strokes are created using this source image. Over successive frames the painted texture will be refined as more brush strokes are added. The quadrilateral carrying the texture will be manipulated in 3D to match the motion of the element it represents.

In order to prevent the steady degradation of the image, a canonical view of the texture will be retained at all times. As in [11] brush strokes will be projected back onto the quadrilateral anchored in 3D space using the inverse of the transform that would be required to produce a perspective view of the texture.



5.2. LIMITATIONS OF THE PROPOSED APPROACH

In approximating 3D objects with 2D images some errors will occur. The main problem comes when an object rotates, and surface area that was not previously visible comes into view, or existing surfaces are occluded.

Further more, it is desirable that the brush strokes appear to lie in the image plane, perpendicular to the line of sight, as real brush strokes would. This will be the case at the time of creation. Unfortunately as the model element moves the resulting transformation of the texture will deform the existing brush strokes so that they no longer have the correct proportions. The simple solution to both these problems is to fade existing brush strokes in response to model motion by reducing the alpha values on the texture. At the same time the rate of new brush stroke generation must be increased.

5.3. PARTICULAR ADVANTAGES OF COMBINING OPTIC FLOW AND PAINTERLY RENDERING.

The main shortcomings of Optic flow are; its inability to handle occlusions properly, a slight flicker and some distortion. All of these flaws can be disguised using the deliberately non-photorealistic nature of painterly rendering.

In the case of painterly rendering of a 3D scene, one is working with a collage of 2D images (brush strokes) which mimic a 3D environment. This is uniquely suited to an Optic flow approach, which generates new frames by transforming 2D image fragments from previous frames.

A further potential benefit of the combination arises from the fact that the Optic flow component effectively requires that information about 3D movement in the scene be extracted. This information could be used by the painterly rendering algorithm to create more evocative brushstrokes.

6. PROJECT OVERVIEW

The project sets out to ultimately produce painterly animations that emulate the rock paintings of the San people. The objective being to demonstrate a more natural representation of an environment that is unsuited to the mechanical rendition of mainstream polygonal rendering.

The proposed method of rendering is an approximation. It is important to establish the importance of different aspects of the animation system so as to use the optimal trade-offs in the final application. The project is therefore a series of experiments followed by a demonstration model utilising the results of the experiments.

The Rendering Process




Fig.3(a) The underlying 3D model.

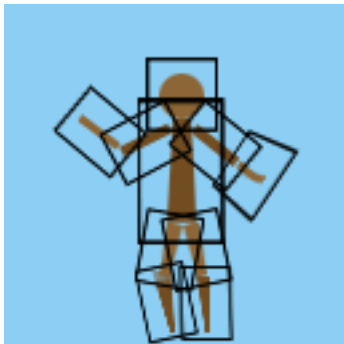


Fig.3(b) Textured 2D quadrilaterals are mapped to each model element.




Fig.3(c) New Brush strokes are projected onto the textured quadrilaterals each frame.



Fig.3(d) The Brush strokes are shown overlaid on a rock surface.

6.1. EXPERIMENTS

Because the effectiveness of the system is an aesthetic judgement, all experiments are by necessity user evaluations. The following questions will be assessed:

6.1.1. High vs. Low Order Optic Flow Significance

The relative impact of the higher and lower order optic flow terms on brush strokes will be evaluated by displaying animations utilising only some parts of the Taylor series decomposition. If the higher order terms are not significant, the rendering technique may be accelerated by their omission.

6.1.2. Accumulated Brush Stroke Distortion

As the textured quadrilaterals are transformed, the existing brush strokes will be distorted. What is an acceptable degree of distortion, and is it more problematic if distortions occur in only one of either the x or y dimensions?

6.1.3. Where to Add New Brush Strokes

Is it better to add new brush stroke in regions where there is a high rate of change in the source image, or in regions where brush stroke culling by alpha fading is greatest?

6.1.4. Special Cases of Rotation

In some cases a plane (the textured quadrilateral) is being mapped to a 3D object with an axis of symmetry (i.e. A tree modelled as a sphere mounted on a cylinder). When the object rotates about its axis of symmetry, is it better to have the quadrilateral ignore that component of the motion?

6.1.5. Model Coherence

A simple model of a human figure will be animated using textured quadrilaterals to represent each independently articulated element (i.e. torso, forearms and thighs). The object of the experiment is to verify that the composited fragments maintain a coherent complete image.

6.2. FINAL APPLICATION

The final application will be a real-time, animated, painterly rendering of 3D models of humans, eland and distant terrain features overlaid on a rock face background, in the style of indigenous San rock art. The human and eland are chosen because they are the two subjects most characteristic of San rock art.

7. CONCLUSION

We hope to demonstrate the optic flow approximations offer an effective means of ensuring the continuity of image-based painterly animations. This would allow for the use of very varied styles of painterly rendering, which would be particularly useful in the depiction of natural and non-industrialised virtual environments.

8. REFERENCES

[1] Durand Fredo, "An Invitation to Discuss Computer Depiction", Laboratory for Computer Science, MIT, *ACM/Eurographics Symp. NPAR'02*.

[2] Meier Barbara J., "Painterly Rendering for Animation", Walt Disney Feature Animation, © *ACM-0-89791-746-4/96/008*, *Proceedings SIGGRAPH 1996*.

[3] Haerberli Paul, "Paint by Numbers: Abstract Image Representations", Silicon Graphics Computer Systems, *Proceedings SIGGRAPH 1990*.

[4] Reeves William T., Blau Ricki, "Approximate and Probabilistic Algorithms for Shading and Rendering Structured Particle Systems", *In Computer Graphics (Proceedings SIGGRAPH 1985)*.

[5] Hertzmann Aaron, "Painterly Rendering with Curved Brush Strokes of Multiple Sizes", Media Research Laboratory Department of Computer Science Courant Institute of Mathematical Sciences New York University, *Proceedings SIGGRAPH 1998*.

[6] Litwinowicz Peter, "Processing Images and Video for An Impressionist Effect", Apple Computer, *Proceedings SIGGRAPH 1997*.

[7] Hertzmann Aaron, Perlin Ken, "Painterly Rendering for Video Animation", Media Research Laboratory Department of Computer Science Courant Institute of Mathematical Sciences New York University, *NPAR 2000*.

[8] Torborg J., Kajiya J.T., "Talisman: Commodity realtime 3D graphics for the PC", *Proceedings SIGGRAPH 1996*.

[9] Webb I., "An Extension to Optic Flow Analysis for the Generation of Computer Images", University of Cape Town, Masters Dissertation.

[10] Litwinowicz Peter, "Inkwell: A 2 1/2 - D Animation System", Apple Computer, *Proceedings SIGGRAPH 1991*

[11] Cohen Jonathan M., Hughes John F., Zeleznik Robert C., "Harold: A World Made of Drawings", Brown University, *NPAR 2000*.