

Innovations Report

Painterly Motion Blur for Moving images

Marie Borgesson (c1173284)

BACVA3

NCCA 2005/06

Abstract

In this report a series of steps for applying a painterly motion blur to an object in motion is presented. The techniques used can in general be applied to any sequence of 2D images where the subject can be separated from the background. These steps were developed from a series of experiments generated from research within the area of non-photorealistic rendering and motion blur for animation and games. It uses a combination of motion vectors and a number of displacements by filmed footage.

1 Introduction

Computer generated images gives an easy and automated way of rendering every part of an image with fine detail. More and more painterly rendering methods are produced trying to achieve the creative look of hand made art. Motion blur is one of the areas where a photorealistic approach is still often used. The aim of this project was therefore to create a motion blur that makes the images as expressive and full of motion as a still image can be, but being able to use it on 3D generated moving images.

1.1 What is Motion Blur?

What makes humans able to see is the specialised brain tissue in the back of our eyes, the retina, which is able to convert light into electrochemical signals that then get sent to the brain. Only a very small area of the retina has a high visual accuracy. This area, the fovea, covers about 2 degrees of a person's field of view, which means that only objects within these 2 degrees will appear clear in focus. To keep something in focus the brain therefore moves the eyes to counteract the movement of the object to keep it on the fovea. If the object moves faster than the eyes can keep up with, the object will slip onto the retina and therefore appear blurred. [9]

In film, the blurriness of a moving object in a scene is determined by the shutter speed of the camera. If the subject being filmed is moving a noticeable distance during the time the shutter of the camera is open, the subject will appear blurred along the movement relative to the camera; this is what we call motion blur. Since this blur occurs in much the same way it would in a normal eye, we see this as being natural. Increasing the film sensitivity however, using film with a higher ISO number needed when using a faster shutter speed, leads to less clarity due to coarse film grain (with digital sensors, from thermal noise). [2]

This means that in film, motion blur essentially is created due to a weakness of the medium; the trade-off in exposure settings, having to keep a balance between sensitivity and clarity. Therefore when we move over to different media that do not have the same "weakness", such as stop motion animation and 2D or 3D animation, motion blur is not automatically created. In the case of these media, if motion blur is not added afterwards, a low frame rate (such as 24 frames/second) could result in the loss of fluidity of motion.

1.2 Photo-realism not equal to believability

The way things look in real life and the way we perceive them are not necessarily the same. This is one reason why images do not have to be photo-real to still be believable.

For most people the logical half of our brain, normally the left hemisphere, is the most dominant side of the brain. This side deals with language and the recognition of symbols while the right side of the brain handles perception of shapes and lines. As children we all develop a system of symbols through drawing. These symbols are simplified shapes that help the logical side of the brain to recognise things we see around us and associate them with the right word. Since the symbols are already simplified this means that we do not have to see all the details of an object in real life to associate it with the right symbol and therefore the corresponding word. [3]

This simplification, where the lack of detail in an image automatically makes the viewer fill in the rest themselves, means that artists do not have to depict every detail in a drawing or painting, letting the viewer in on the interpretive process. Still images are often meant to tell a story. By abstracting the scene and only putting details in certain areas, the artist can direct the eyes of the viewer without deciding for them what to think. In this way the viewer is able to imagine what has happened before and after this snapshot in time. There might be several possible scenarios or the artist might lead the viewer to see only the one that he or she has intended.

An artist can also add to the drama of the image or make the image more expressive in other ways. The artist might, for example, choose to only use colours within a restricted colour palette or exaggerate the lighting. Another way is to communicate force or speed through the quality and orientation of the brush strokes or pencil lines.

The nature of CG generated images is that everything is shown with an equal amount of clarity. This is often more than we normally perceive in real life and it also does not direct the viewer to interpret the image/images further. This makes it important to work with directing the viewers' eyes when dealing with CG generated images.

This lead me to look into ways of creating a non-photorealistic motion blur.

2 Related Work

2.1 Painterly techniques applied to moving images

To achieve a hand drawn look for animation in computer graphics there are two main difficulties to overcome. One is the effect where the animation appears as if it was moving behind a patterned image (swimming textures). This is due to the brush strokes sticking to the view plane. The other is the effect of noise, where the change of each brush stroke from frame to frame is so large that it becomes distracting to the viewer.

Barbara J. Meier presents, in her paper *Painterly Rendering for Animation* [5], a way to achieve a painterly look in moving images whilst maintaining frame to frame coherence. She does this by combining the areas of painterly rendering of still images and particle rendering. Meier's rendering algorithm first creates particle sets that represent the geometry in the scene. A number of reference images are then rendered to obtain colour, orientation and size; all attributes that are needed to draw the brush strokes. The colour reference image is rendered with a smooth shader. This image takes into account the lighting of the scene and broad colour changes. To get the orientation of each brush stroke the scene is then rendered with a special shader that encodes surface normals into the image. The final reference image encodes x and y scaling information into the picture to determine the size of each brush stroke. After this information has been stored in each particle, they are sorted in order of their distance to the camera and rendered, starting with the particles furthest away. Figure 1 shows an overview of this process.

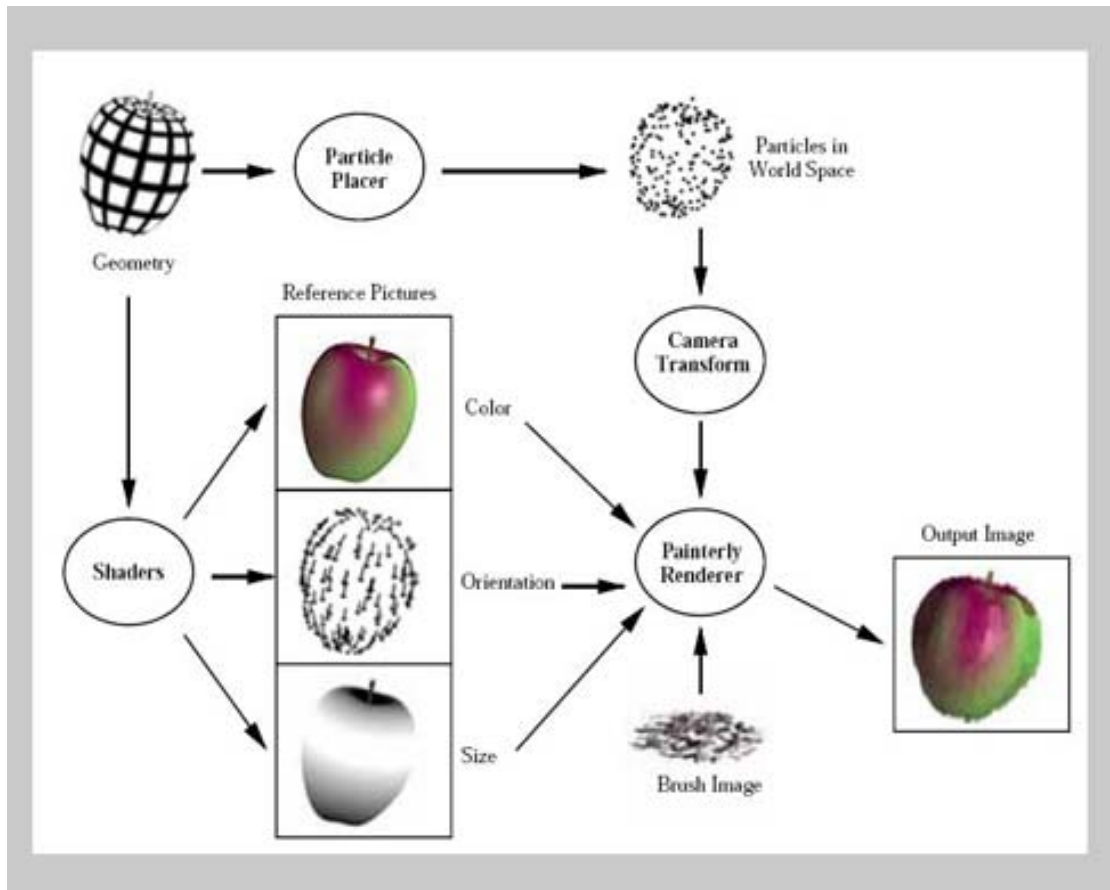


Figure 1: Barbara J. Meier's Painterly Render (Image from Barbara J. Meier [5]).

The idea of this rendering algorithm is that because particle sets are being used in the place of the geometry the brush strokes “stick” to the surfaces instead of floating around in no relation to the geometry. Using this method creates another problem; when the brush strokes move with the objects, the planes/images of the brush strokes move around under each other breaking the illusion of them flowing into each other dynamically and showing that the brush strokes are just 2D planes. The problem arises because the particles do not move in the direction of the brush stroke orientation. It also becomes apparent that there is too little randomisation among the strokes and they appear to all be the same image.

In the short film *Dreamkeeper*, *Glassworks* [6] use another kind of non-photorealistic rendering to create a blurred motion trailing effect. They shot the footage with stereo lenses, producing two images for every 4 perforations on the film, the space one frame normally takes. The two images were compared to create a depth map and forces acting like increased air pressure was simulated by detecting changes in velocity of the optical flow. Half the images were discarded and from this a warped space effect that stayed locked to the features in shot was produced. This base layer was done using Revision FX optical flow algorithms driven by the optical flow data of the sequence. By combining the depth map and the air pressure simulation, a pseudo-3D vector flow was generated. This was used to control a very large amount of particles, the colours of which were determined by the corresponding pixel colour of the colour corrected base layer. This particle pass was added to the base layer to produce the final image.

The particles in this sequence look like they move from the motion of the subject being filmed. This interaction works well to convey motion and make the scene believable, even though it is non-photorealistic.

Despite *Glassworks* results, using particles, fitting in with the desired effects for the motion blur, the complexity of the techniques are beyond the scope of this project

2.2 Methods for “faked” motion blur in animation and games

In the paper Image-Based Motion Blur for Stop Motion Animation, Gabriel J. Brostow and Irfan Essa [1] present a technique to automatically simulate realistic motion blur for footage where motion blur is not produced due to the medium. This approach is a post process that can be applied to images from stop motion or video. These kinds of media do not have any 3D transformation data available like other mesh-based animation renders do. To detect motion paths, which motion blur should later be applied along, the pixels are tracked in simplified groups from frame to frame. As a result of transformations other than translation and rotation, lighting changes and the effects of perspective projection, a map of corrective motion vectors is needed for more realistic pixel transformations. The pixels are blurred according to the corrected motion vectors of an amount controlled by the user.

The benefit of this technique is that the amount of blur can be controlled interactively by the user by changing the setting of time, corresponding to the shutter speed in filming. Another way of achieving this is by using Reel Smart Motion Blur [7]. This is a plugin that can be used for a number of different 2D and 3D packages and post production software. By rendering out a sequence of motion vector images from an animation in a 2D or 3D package, the motion in x is stored in the red channel and the motion in y in the green. The range of pixel values is then mapped to the motion vector values using a RSMB Vectors 3 node in post production software. The amount of blur can be controlled by changing the MotionBlurAmt and the MaxDisplace of this node.

Up until quite recently the lack of motion blur in games has been compensated for by a high frame rate, ideally 60 frames/second or over. However this does not look natural to us, but it helps the motion to not feel broken. More recent computer games try to simulate motion blur by blending the current frame with faded previous frames, making the previous frames appear behind the current. This does not simulate natural motion blur properly, but creates a motion trailing effect that helps keep the motion fluid.

2.3 Ways to get motion in a still image

To artists using traditional materials to create still images conveying motion, a wide variety of methods are available. A few examples are smearing, brush stroke directions, motion trails and elongating.

In photography motion can be captured in a still photo by using long exposure times. This creates a trail of motion, where the beginning and end are equally visible (figure 2).



Figure 2: Star trails around the north celestial pole. (Image and caption (c) Harald Edens, reproduced with permission [2])

In a Photoshop tutorial, “Photoshop Displacement Effects” [4], Harrison describes how to bring motion into a still image through distorting displacement filters in Photoshop. He starts off with a photo or an image where the subject is in focus. A displace filter is then applied to the original image. The displacement uses another image, greyscaled and high contrasted, featuring an interesting pattern to distort the original image (figure 3). This distorted image is layered on top of the original and where the distortion is not desired, it should be erased to reveal the underlying original (figures 4).




Figure 3: Original image displaced with a black and white image. (Black and white image from Peter Harrison [4])



Figure 4: Original image layered with the displaced image where parts have been erased.

3 Painterly Motion Blur

In the following subsections, the progress of creating the painterly motion blur is described, starting with early tests, showing the progress to the most recent version. This latest test is then described in the steps needed to create it and how the result can be varied.

3.1 Progress through experimentation

As with many creative processes, the latest version of this image sequence was created through a trial and error process. Research helped with the generation of ideas and what to test, but the testing phase was ultimately what decided if a technique was going to be used for the final sequence or not.

One of the first tests combined interestingly patterned masks with motion blur. The idea was to mask out the areas the motion blur would not be applied to by using motion vectors. This could make an interesting pattern for where detail would appear on the image. Even though the outcome of this test was too random and not easily controlled by the user, it showed that realistic motion blur could be used to create non photorealistic effects by manipulating motion vectors.

Later tests were done using different kinds of footage as a source of displacement, with displacement nodes. Using different colour channels from the same image as displacement for x and y , proved to give a slightly asymmetrical displacement, while using the same channel would distort the image an equal amount in both directions.

A number of different pieces of footage were tested as displacement, both 3D and video. The conclusion was that high contrasted footage, such as coloured liquid on dark backgrounds, gave more of a trailing effect than smoother images, such as smoke.

Mixing and layering different tests lead to the latest image sequence.



Figure 5: Two pieces of footage of different kinds of liquid, used to displace as displacement in two layers.

3.2 How to recreate the motion blur

In the following 5 steps, Maya and Shake have been used, though other 3D and post production software with similar functions can be used. See appendix 1 for a graph of the node network.

Step 1 – Animated sequence

To create this kind of motion blur, an animated sequence with a subject moving relative to the camera is needed. The images generated from this sequence should not have any kind of blur due to motion (other kinds of blur, such as depth of field, are acceptable). This is the base pass.

Step 2 – Exaggerated “natural” motion blur

This step can be done in one of two ways. The second pass can be rendered out with motion vectors (using Reel Smart Motion Blur, for example, see section 2.2, second paragraph). Motion vectors allow you to interactively change the amount of motion blur in post production software (assuming you have the Reel Smart Motion Blur plugin). This makes it easier to see how much “natural” motion blur is needed to fit with the later passes, without having to render out more than the base pass (figure 6). Alternatively the animated sequence can be rendered out with “natural” motion blur

already added. Exaggerating it slightly will give a more dramatic effect and will enable this pass to fit better with the subsequent stages of this process. However, bear in mind that if this pass does not fit with later passes it will have to be rendered out again with different motion blur settings.



Figure 6: Base pass combined with motion vector pass giving, with the right settings, the exaggerated “natural” motion blur pass.

Step 3 – Displacement with irregular background

The next pass is generated using an IDisplace node. The IDisplace node has two inputs. The first is the image to be displaced and the second is the source of the displacement. The IDisplace node distorts the images of the first input in x and y according to the colour channels specified of the source of the displacement, the second input. Using different colour channels for the distortion in x and in y can create a slightly asymmetric distortion.

The first input used for this pass is the base pass. The second is a filmed grainy irregular background, where the red channel is specified to distort in x and the green channel in y (figure 7). This creates a slightly irregular distortion of the base pass that works well as a merger of the passes produced in previous and following step.



Figure 7: Base pass displaced with a grainy image as source.

Step 4 – Displacement with filmed liquid

This step also uses the IDisplace node. First create three instances of the base pass fileIn node, A, B and C. Leave A unchanged, but reposition nodes B and C by -1 and -2 frames respectively. Nodes B and C now appear behind A as a trail when the sequence is played back. The fileIn nodes are then plugged in to the first input of three separate IDisplace nodes.

The source of displacement is footage of a liquid running down a pane of glass. Depending on the quality of the footage, this fileIn node may need to have its contrast increased.

The filmed footage should be rotated and skewed to fit the direction and the perspective of the object in motion (figure 8). The output is then plugged into all three secondary inputs of the three IDisplace nodes and again the red channel is set to distort in x and the green in y .




Figure 8: Base pass displaced with filmed footage of liquid. Here the three displaced images are layered, with different opacities, into one.

Step 5 – Mixing and layering the different passes together

Now the different passes can be added together. First mix the exaggerated “natural” motion blur pass of step 2, with the displacement pass of step 3 and leave the opacity at 50% to show the two passes equally. To achieve a blurry effect on top of this, layer the output of the mix layer with the “natural” motion blur again in a max layer.

Combine B and C with a second mix layer, B as input one and C as input two. Use this output in a third mix layer together with A. Adjust the opacity of these two mix layers so that A is more visible than B and B more than C, to give an effect of the displacement fading away when played back.

To complete the last step of creating the painterly motion blur, combine the output of the max layer and the output of the third mix layer in a new max layer.

3.3 Variations of the technique

The steps in the previous subsection were applied to an image sequence generated from a 3D animation. However, the techniques described can be applied to images generated from different media. As long as the sequence of images is featuring a subject moving relatively to the camera without

any kind of blurriness due to motion and as long as the subject moving can be separated from the background, any sequence of 2D images can be used.

The result of the techniques can also be varied by using different sources for the displacement. In the previous example, milk running down a pane of glass with a black background was used. This gives a high contrast image which results in distinct trails in the displacement. If a more blurred displacement is desired, less contrasted footage, such as smoke, can be used.

As well as being used for non-photorealistic motion blur, these techniques can also be used to simulate something being emitted from the subject of the sequence, such as fire or a more philosophical meaning of something spiritual or dreamy. To get this effect, the “natural” motion blur could be toned down to separate the displaced passes from the subject in motion. It could also be done by using footage of static items.




Figure 9: Different variations of the technique, with the original to the far right.

4 Results

4.1 Future directions

This painterly motion blur does not handle the shadows of the moving object at the moment. This could be tackled by using a “natural” motion blurred shadow. However, it would be interesting to see what effects a more accurate shadow, to the shape of the object with this motion blur applied, would have.

The intention in the beginning of this project was to try to use particle trails for an accurate way of tracking motion. While it was decided that this motion blur would benefit from a simpler and faster technique, using filmed footage, I still believe that the motion blur can be developed further using particles.

In still images and paintings the artist sometimes applies different colours to the motion blurred section of the image. This can give a dramatic effect and is worth looking into further to improve the techniques used in this paper.

4.2 Conclusions

A series of techniques has been presented for creating a painterly motion blur. This approach is combining techniques for generating “natural” motion blur as a post process by using motion vectors and techniques for displacing images with other footage.

Even though the end result of this project did not turn out the way I had originally visualised it, this is a satisfying result. The direction I decided to take, lead to an interesting and different kind of motion blur. When applied to moving images an exaggerated distortion of the frame is achieved, which adds drama and focus on the motion. Although this painterly motion blur gives an interesting effect when used for moving images, even single frames convey motion. Despite having this as the only aim of the project in the start, I can now see how the motion blur could be developed further; by using the existing number of steps combined with techniques similar to Glassworks’ methods used for Dreamkeeper, but emitting particles only from the moving objects to generate the displacement data. I believe this would generate a displacement that looks more a part of the subject in motion.

5 Acknowledgement

I would like to thank the following people for the support they gave;

- My supervisor Phill Allen for providing guidance and helping me to get structured writing this report.
- My subject tutor and early supervisor, Stephen Bell, for providing lectures and notes on how to write a report and for giving feedback in the early stages of the project.
- Matt Birkett-Smith, for giving useful feedback and technical help throughout the project.
- Daniel Lim, for providing useful feedback and encouragement. His help with this report was also much appreciated.
- Bruce Steele and Peter Reilly at Glassworks, for help with research on the effects used to produce Dreamkeeper.

References

- [1] Brostow, Gabriel J. and Essa, Irfan - Image-Based Motion Blur for Stop Motion Animation, GVU Center / College of Computing Georgia Institute of Technology.
<http://www.cc.gatech.edu/cpl/projects/blur/> pages 561-566
- [2] *Edens, Harald* - Exposure times, aperture settings and film speed, Trade-off in exposure settings,
www.weather-photography.com
- [3] Edwards, Betty - Drawing on the Right Side of the Brain, London 2001, Harper Collins Publishers.
- [4] Harrison, Peter - 2006, Photoshop Displacement Effects, Computer Arts Magazine (issue 120)
- [5] Meier, Barbara J. - Painterly Rendering for Animation, Walt Disney Feature Animation, Pages 477-484.
- [6] Steele, Bruce (and Reilly, Peter), Glassworks - Techniques used to create Dreamkeeper, produced by Glassworks.
- [7] Reel Smart Motion Blur - Revision FX, www.revisionfx.com (/generalFAQsMotionVectors.htm)

[8] Wikipedia - Eye Movements, www.wikipedia.com

Further reading

Stephen Bell - Writing your innovations report.

Steph Greenberg - Why Non-Photorealistic Rendering? 1999 ACM Siggraph.

Academic services, Bournemouth University - How to: Cite references (The Harvard System).

Appendix 1 – Node network in Shake used to create the painterly motion blur.

