



Martin Eisemann,  
Bernd De Decker,  
Anita Sellent,  
Edilson De Aguiar,  
Naveed Ahmed,  
Hans-Peter Seidel,  
Philippe Bekaert,  
Marcus Magnor

Floating Textures

Technical Report / Computer Graphics Lab, TU  
Braunschweig ; 2007-10-4

<http://www.digibib.tu-bs.de/?docid=00021424>



MARTIN EISEMANN

Institut für Computergraphik, TU Braunschweig

BERND DE DECKER

Expertise Centre for Digital Media, Universiteit Hasselt

ANITA SELLENT

Institut für Computergraphik, TU Braunschweig

EDILSON DE AGUIAR

Max-Planck Institut für Computergraphik

NAVEED AHMED

Max-Planck Institut für Computergraphik

HANS-PETER SEIDEL

Max-Planck Institut für Computergraphik

PHILIPPE BEKAERT

Expertise Centre for Digital Media, Universiteit Hasselt

MARCUS MAGNOR

Institut für Computergraphik, TU Braunschweig

# Floating Textures

**Technical Report 2007-10-4**

October 1, 2007

Computer Graphics Lab, TU Braunschweig



# Abstract

We present a novel multi-view, projective texture mapping technique. While previous multi-view texturing approaches lead to blurring and ghosting artefacts if 3D geometry and/or camera calibration are imprecise, we propose a texturing algorithm that warps (“floats”) projected textures during run-time to preserve crisp, detailed texture appearance. Our GPU implementation achieves interactive to real-time frame rates. The method is very generally applicable and can be used in combination with many image-based rendering methods or projective texturing applications. By using Floating Textures in conjunction with, e.g., visual hull rendering, light field rendering, or free-viewpoint video, improved rendering results are obtained from fewer input images, less accurately calibrated cameras, and coarser 3D geometry proxies.



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Previous Work</b>	<b>5</b>
<b>3</b>	<b>Problem Description</b>	<b>7</b>
<b>4</b>	<b>Floating Textures</b>	<b>11</b>
4.1	Acceleration for Static Scenes . . . . .	13
4.2	Soft Visibility . . . . .	15
<b>5</b>	<b>GPU Implementation</b>	<b>19</b>
<b>6</b>	<b>Results, Discussion</b>	<b>21</b>
<b>7</b>	<b>Discussion &amp; Conclusions</b>	<b>25</b>
<b>8</b>	<b>Acknowledgements</b>	<b>29</b>

**CONTENTS**

# Chapter 1

## Introduction

To take advantage of the continuing progress in graphics hardware capabilities for realistic rendering, ever more detailed model descriptions are needed. Because creating complex models with conventional modeling and animation tools is a time-consuming and expensive process, direct modeling techniques from real-world examples are an attractive alternative. By scanning, or reconstructing, the 3D geometry of an object or scene and capturing its visual appearance using photos or video, the goal of direct modeling techniques is to achieve photo-realistic 3D rendering results at interactive frame rates.

Texture mapping was introduced in computer graphics as early as 1974 as a very effective means to increase visual rendering complexity without the need to increase geometry details [Cat74]. Later on, projective texture mapping overcame the need for explicit texture-to-surface parameterization and enabled applying conventional photographs directly as texture [SKvW<sup>+</sup>92]. To texture 3D geometry from all around and to reproduce non-Lambertian reflectance effects View-dependent texture mapping [DTM96] and Unstructured Lumigraph Rendering [BBM<sup>+</sup>01] use multiple photographs taken from different viewpoints. Given exact 3D geometry and sufficiently many, well-calibrated and registered images, such image-based modeling techniques achieve impressive 3D rendering results [WAA<sup>+</sup>00, LKG<sup>+</sup>03].

Unfortunately, acquiring highly accurate 3D geometry and calibrated images turns out to be at least as tedious and time-consuming as model creation using software tools. In response, a number of different image-based rendering (IBR) techniques have been devised that make do with more approximate geometry. With a mere planar rectangle as geometry proxy, Light Field Rendering arguably constitutes the most “puristic” image-based rendering technique [LH96]. Here, many images are needed to avoid blurring or ghosting artefacts [CCST00]. If dense depth maps are additionally available, Lumigraph rendering can compensate for parallax between images to yield convincing results from less images [GGSC96, BBM<sup>+</sup>01]. Other image-based rendering approaches implicitly or explicitly recover approxi-



Figure 1.1: Comparison of standard linear interpolation using the Unstructured Lumigraph weighting scheme (left) and the Floating Textures approach (right) for the same input images and visual hull geometry proxy. Ghosting along the collar and blurring of the shirt's front, noticeable in linear interpolation, are eliminated on-the-fly by Floating Textures.

mate 3D geometry from the input images to which the images are applied as texture [MBR<sup>+</sup>00, VBK05, CTMS03]. In general, however, the price for contending with approximate 3D geometry is that (many) more input images must be available, else rendering quality degrades and artefacts quickly prevail, Fig. 1.1. Another unsolved challenge is that, while image-based rendering can compensate for coarse 3D geometry if sufficiently many input images are available, all techniques still require well calibrated cameras and input images. Thus, time-consuming camera calibration procedures must precede image-based object acquisition every time. But even with the utmost care taken during acquisition, minute camera calibration inaccuracies, tiny 3D scanning holes, and small registration errors can occur and visibly degrade rendering quality of the model. The only option to rescue rendering quality then is to "fix" the model, registration, or calibration by hand, or to repeat the acquisition process all over again.

What is needed is a multi-image texturing algorithm that achieves best-possible results from imprecisely calibrated images and approximate 3D geometry. In the following, we propose a GPU-based multi-view texturing algorithm towards this goal. Because the algorithm runs independently on the graphics card, it can be used in conjunction with many image-based modeling and rendering (IBMR) techniques to improve rendering outcome.

As particular contributions, our paper presents:

- a novel texturing algorithm that constitutes a symbiosis between classical linear interpolation and optical flow-based warping refinement;
- a novel weighting and visibility scheme which significantly reduces artefacts at occlusion boundaries;
- a general algorithm that can be applied in conjunction with many IBMR techniques to improve rendering quality;
- an efficient GPU-based implementation of the proposed algorithm which achieves interactive to real-time rendering frame rates;
- a simple extension to our Floating Textures approach for static scenes which reduces the actual rendering part to a simple texture look-up.

The remainder of the paper is organized as follows. After reviewing relevant previous work in Section 2 we examine the underlying problem of ghosting and occlusion artefacts in multi-view projective texture mapping in Section 3. In Section 4 we describe our Floating Textures as a way to eliminate ghosting, occlusion and calibration artefacts. Implementation details are given in Section 5, and experimental evaluation results for a variety of different test scenes and IBMR techniques are presented in Section 6 before we discuss limitations and conclude with Section 7.

## 1 Introduction

4

---

## Chapter 2

# Previous Work

**Sampling Problem:** In light field rendering, a novel view is generated by appropriately re-sampling from a large set of images [LH96, MP04]. Given sufficiently many input images, the synthesized novel view can be of astonishingly high quality. Else, ghosting artefacts or, at best, blurring degrade light field rendering quality. Several researchers have investigated how many input images are minimally needed to create artefact-free light field rendering results [CCST00, LS04]. Alternatively, rendering quality can be enhanced by considering not only input camera positions but also the current viewpoint [ESM07] or by adding back in high frequency components [SYGM03]. Nevertheless, these approaches still introduce at least some image blur if the scene is undersampled.

**Geometry aided IBR:** Instead of relying on sampling density, another approach to increase rendering quality from sparsely sampled image data is to make use of a geometry proxy representing the scene. Prominent examples are Lumigraph Rendering by Gortler *et al.* [GGSC96], Unstructured Lumigraph Rendering by Buehler *et al.* [BBM<sup>+</sup>01] and Debevec *et al.*'s View-dependent Texture Mapping [DTM96]. Isaaksen *et al.* [IMG00] showed that if scene depth estimation is precise enough and no occlusion occurs, any single scene element can, in theory, be reconstructed without ghosting artefacts. Several other researchers built upon this insight to improve rendering results [ZKU<sup>+</sup>04, VBK05]. An interesting approach which not only blends colours on a geometry proxy but attempts to reconstruct a consistent, view-dependent geometry of the scene with the aid of bilateral filtering was only recently presented by Waschbüsch *et al.* [WWG07]. In image-based modeling approaches, high-quality 3D geometry scans of real-world objects are used and textured with a collection of photographs [RCMS99, WAA<sup>+</sup>00, Bau02, ZWT<sup>+</sup>05]. Acquiring these detailed models is time-consuming and, of course, is possible only for scenes that hold still during acquisition. In any case, image calibration inaccuracies, subcritical sampling, and geometry acquisition errors remain potential sources of

rendering artefacts.

**Occlusion handling:** In settings with very sparse camera setups, occlusion and registration artefacts become annoyingly obvious. Carranza *et al.* [CTMS03] therefore proposed to use a visibility map, computing visibility for every vertex from several camera views that are slightly displaced in the image plane. Lensch *et al.* [LKG<sup>+</sup>03] search for depth discontinuities to discard samples close to them, as they are prone to errors. In Virtual Viewpoint Video [ZKU<sup>+</sup>04] these discontinuities are rendered separately using Bayesian image matting to compute fore- and background colours along with opacities. All of these approaches either assume very precise underlying geometry, or they need extensive pre-processing to achieve interactive rendering frame rates.

**Warping Techniques:** If complete correspondences between image pixels can be established, accurate image warping becomes possible [CW93, MB95]. Often, however, current methods for automatic camera calibration and depth acquisition are too imprecise. For very similar images, optical flow techniques have proven useful [HS81, LK81]. They can be employed to create smooth morphs between images, as was used, e.g., by Vedula *et al.* for spatio-temporal view interpolation [VBK05].

The work with the closest resemblance to ours was proposed by Aliaga *et al.* in the context of creating virtual walkthroughs [AFYC02, AYFC03]. Aliaga *et al.* capture a very dense set of omnidirectional images, with an average distance between capturing positions of 1.5 inches. For image synthesis, they identify corresponding features in the different images, triangulate them in one reference view, and warp this mesh according to the current viewpoint and the interpolated feature positions. Our approach differs markedly in a number of important points. First, Floating Textures do not rely on previously tracked features which can be difficult to establish, especially for natural scenes. Second, we explicitly take occlusion effects into account, instead of relying on dense sampling. Third, instead of using only a sparse set of features to establish the necessary information for warping, we search for a globally optimal solution, which makes our approach less dependent on scene properties.

## Chapter 3

# Problem Description

In a slightly simplified version, the plenoptic function  $\mathbf{P}(x, y, z, \theta, \phi)$  describes radiance as a function of 3D position in space  $(x, y, z)$  and direction  $(\theta, \phi)$  [AB91]. The notion of image-based rendering now is to approximate the plenoptic function with a finite number of discrete samples of  $\mathbf{P}$  for various  $(x, y, z, \theta, \phi)$  and to efficiently re-create novel views from this representation by making use of some sort of object geometry proxy.

Any object surface that we choose to display can be described as a function  $\mathbf{G} : (x, y, z, \theta, \phi) \rightarrow (x_o, y_o, z_o)$ , i.e., by a mapping of viewing rays  $(x, y, z, \theta, \phi)$  to 3D coordinates  $(x_o, y_o, z_o)$  on the object's surface. Of course, the function  $\mathbf{G}$  is only defined for rays hitting the object, but this is not crucial since one can simply discard the computation for all other viewing rays. With  $\mathbf{G}_O$  we denote the function of the true surface of the object, and with  $\mathbf{G}_A$  we denote a function that only approximates this surface, Fig. 3.1.

Next, a variety of camera calibration techniques exist to establish the projection mapping  $\mathbf{P}_i : (x, y, z) \rightarrow (s, t)$  which describes how any three-dimensional point  $(x, y, z)$  is mapped to its corresponding 2D-position  $(s, t)$  in the  $i$ -th image. From its projected position  $(s, t)$  in image  $\mathbf{I}_i$ , the 3D point's colour value  $(r, g, b)$  in that image can be read out,  $\mathbf{I}_i : (s, t) \rightarrow (r, g, b)$ . Then, any novel view  $\mathbf{I}_{\text{Linear}}^V$  from a virtual viewpoint  $V$  synthesized by a weighted linear interpolation scheme, as employed by most IBR systems, can be formulated as

$$\mathbf{I}_{\text{Linear}}^V(x, y, z, \theta, \phi) = \sum_i \mathbf{I}_i^V(x, y, z, \theta, \phi) \omega_i \quad (3.1)$$

with

$$\mathbf{I}_i^V(x, y, z, \theta, \phi) = \mathbf{I}_i(\mathbf{P}_i(\mathbf{G}_A(x, y, z, \theta, \phi))) \quad (3.2)$$

$$\omega_i = \delta_i(\mathbf{G}_A(x, y, z, \theta, \phi)) w_i(x, y, z, \theta, \phi) \quad (3.3)$$

and  $\sum_i \omega_i = 1$ . Note that  $\delta_i$  is a visibility factor which is one if a point on the approximate surface  $\mathbf{G}_A$  is visible by camera  $i$ , and zero otherwise.




Figure 3.1: Top left: Geometry inaccuracies cause ghosting artefacts. Point  $P$  on the original surface  $\mathbf{G}_O$  is erroneously projected to 3D-position  $P_1$  from camera  $C_1$  and to 3D-position  $P_2$  from camera  $C_2$  if only the approximate geometry  $\mathbf{G}_A$  is available. Top right: Small imprecisions in camera calibration can lead to false pixel projections (red lines, compared to correct projections displayed as blue lines). This leads to a texture shift on the object surface and subsequently to ghosting artefacts. Bottom: Visibility errors. Given only approximate geometry  $\mathbf{G}_A$ , point  $P$  is classified as being visible from  $C_2$  and not visible from camera  $C_1$ . Given correct geometry  $\mathbf{G}_O$ , it is actually the reverse, resulting in false projections.

### 3 Problem Description

9

The notation  $\mathbf{I}_i^V$  is used to denote the image rendered for a viewpoint  $V$  by projecting the input image  $\mathbf{I}_i$  as texture onto  $\mathbf{G}_A$ .  $w_i$  is the weighting function which determines the influence of camera  $i$  for the every viewing ray.

Note that (3.1) is the attempt to represent the plenoptic function as a linear combination of re-projected images. Unfortunately,  $\mathbf{P}$  is not guaranteed to lie in the solution space of  $\mathbf{I}_{\text{Linear}}^V$  for all viewpoints and directions. For several reasons, weighted linear interpolation cannot be relied on to reconstruct the correct values of the plenoptic function:

1. Typically,  $\mathbf{G}_O \neq \mathbf{G}_A$  almost everywhere, so the input to (3.2) is already incorrect in most places, Fig. 3.1 top left.
2. Due to calibration errors,  $\mathbf{P}_i$  is not exact, leading to projection deviations and, subsequently, erroneous colour values, Fig. 3.1 top right.
3. In any case, only visibility calculations based on the original geometry  $\mathbf{G}_O$  can provide correct results. If only approximate geometry is available, visibility errors are bound to occur, Fig. 3.1 bottom.

In summary, in the presence of even small geometry inaccuracies or camera calibration imprecisions, a linear approach is not able to correctly interpolate from discrete image samples.

### **3 Problem Description**

**10**

---

## Chapter 4

# Floating Textures

In the following, we describe our approach to remove blurring and ghosting artefacts caused by geometry and calibration inaccuracies using an adaptive, non-linear approach. In a nutshell, the notion of Floating Textures is to correct for local texture misalignments by determining the optical flow between projected textures and warping the textures accordingly in the rendered image domain. Both steps, optical flow estimation and multi-texture warping, can be efficiently implemented on graphics hardware to achieve interactive to real-time performance.

As input, the algorithm requires nothing more but a set of images, the corresponding, possibly imprecise, calibration data, and a geometry proxy. For simplicity, we will first assume an occlusion-free scene and describe how occlusion handling can be added in Sect. 4.2. Without occlusion, any novel viewpoint can, in theory, be rendered from the input images by warping. To determine the warp fields, we are safe to assume that corresponding pixels in different images have a set of similar properties, like colour or gradient constancy, so that the following property holds:

$$\mathbf{I}_j = \mathbf{W}_{\mathbf{I}_i \rightarrow \mathbf{I}_j} \circ \mathbf{I}_i \quad , \quad (4.1)$$

where  $\mathbf{W}_{\mathbf{I}_i \rightarrow \mathbf{I}_j} \circ \mathbf{I}_i$  warps an image  $\mathbf{I}_i$  towards  $\mathbf{I}_j$  according to the warp field  $\mathbf{W}_{\mathbf{I}_i \rightarrow \mathbf{I}_j}$ . The problem of determining the warp field  $\mathbf{W}_{\mathbf{I}_i \rightarrow \mathbf{I}_j}$  between two images  $\mathbf{I}_i, \mathbf{I}_j$  is known as optical flow estimation [HS81, LK81]. If pixel distances between corresponding image features are not too large, algorithms to robustly estimate per-pixel optical flow are available [BBPW04].

For our Floating Textures approach, we propose a symbiosis of linear interpolation and optical flow-based warping. We first project the photographs from cameras  $C_i$  onto the approximate geometry surface  $\mathbf{G}_A$  and render the scene from the desired viewpoint  $V$ , creating the intermediate images  $\mathbf{I}_i^V$ , Fig. 4.1. Note that while corresponding image features do not yet exactly line up in  $\mathbf{I}_i^V$  they are much closer together than in the original photos (disparity compensation). We can apply optical flow estimation to




Figure 4.1: Rendering with Floating Textures. The input photos are projected from camera positions  $C_i$  onto the approximate geometry  $G_A$  and on to the desired image plane from viewpoint  $V$ . The resulting intermediate images  $I_i^V$  exhibit mismatch which is compensated by warping all  $I_i^V$  based on the optical flow to obtain the final image  $I_{\text{Float}}^V$ .

the intermediate images  $\mathbf{I}_i^V$  to robustly determine the pairwise flow fields  $\mathbf{W}_{\mathbf{I}_i^V \rightarrow \mathbf{I}_j^V}$ .

To compensate for more than two input images, we linearly combine the flow fields according to Eq. (4.3), apply these to all intermediate images  $\mathbf{I}_i^V$  and blend them to obtain the final rendering result  $\mathbf{I}_{\text{Float}}^V$ . To reduce computational cost, instead of establishing for  $n$  input photos  $(n-1)n$  flow fields, it often suffices to consider only the 3 closest input images to the current viewpoint.

We use an angular weighting scheme as proposed in [BBM<sup>+</sup>01, DTM96] because it has been found to yield better results for coarse geometry than weighting schemes based on normal vectors [CTMS03]. Our Floating Textures are, in fact, independent of the weighting scheme used as long as the different weights sum up to 1 for every pixel, which is necessary to ensure that corresponding features coincide in the output image.

The processing steps are summarized in the following functions and visualized in Fig. 4.1:

$$\mathbf{I}_{\text{Float}}^V = \sum_{i=0}^k (\mathbf{W}_{\mathbf{I}_i^V} \circ \mathbf{I}_i^V) \omega_i \quad (4.2)$$

$$\mathbf{W}_{\mathbf{I}_i^V} = \sum_{j=0}^k \omega_j \mathbf{W}_{\mathbf{I}_i^V \rightarrow \mathbf{I}_j^V} \quad (4.3)$$

$\mathbf{W}_{\mathbf{I}_i^V}$  is the combined flow field which is used for warping image  $\mathbf{I}_i^V$ . Eq.(4.2) is therefore an extension of (3.1) by additionally solving for the non-linear part in  $\mathbf{P}$ .

Note that our Floating Textures deliberately do not satisfy the epipolar constraint anymore. To make use of epipolar geometry constraints one has to presume perfectly calibrated cameras, which is seldom the case. Instead, by not relying on epipolar geometry, Floating Textures can handle imprecise camera calibration as well as approximate geometry, while the minute shifting of the texture on the surface is visually absolutely unnoticeable.

## 4.1 Acceleration for Static Scenes

For static scenes it might seem unnecessary to re-compute the flow fields for every frame. But for a coarse geometry proxy, one cannot simply assign constant texture coordinates to every vertex and every input image.

Instead, we propose a slight variation of our Floating Texture generation which can be computed during preprocessing.

Instead of computing flow fields between input images after they have been projected into the desired viewpoint, we render the scene from each camera position  $C_i$  and project all other input images into its viewpoint, i.e.,




Figure 4.2: Rendering with Floating Textures for static scenes: in a preprocessing step, the images  $I_i^{C_i}$  of every camera are projected onto the approximate surface  $G_A$  and into every other input camera  $C_j$ , resulting in the intermediate images  $I_j^{C_i}$ . Then the warp field between these images in every camera view is calculated. For rendering the image  $I_{\text{Float}}^V$  from viewpoint  $V$ , the warp field of each camera is queried for the texture coordinate offset (black arrows) of every rendered fragment, and the corrected texture value (blue dots in image plane) is projected back onto the object and into the novel view.

we render  $\mathbf{I}_j^{C_i}$ ,  $j \in \{1, \dots, n\}$ . The flow fields  $\mathbf{W}_{I_i^{C_i} \rightarrow I_j^{C_i}}$  are then established between the rendered image  $\mathbf{I}_i^{C_i}$  and every  $\mathbf{I}_j^{C_i}$ , with  $j \neq i$ . As the views from the cameras do not change over time for static scenes, image synthesis for a new viewpoint reduces to simple projective texturing using warped texture coordinates, Fig. 4.2:

$$I_{out}^V = \sum_{i=0}^k \left( \left( \sum_{j=0}^k (\omega_j \mathbf{W}_{I_i^{C_i} \rightarrow I_j^{C_i}}) \right) \circ \mathbf{I}_i^{C_i} \right)^V \omega_i \quad (4.4)$$

Note that in comparison to viewpoint-centered warping (4.2) rendering quality may be slightly reduced. On the other hand, the online rendering computations are reduced to two simple texture lookups per fragment and camera.

## 4.2 Soft Visibility

Up to now, we have assumed only occlusion-free situations, which is seldomly the case in real-world scenarios. Simple projection of imprecisely calibrated photos onto an approximate 3D geometry model typically cause unsatisfactory results in the vicinity of occlusion boundaries, Fig. 4.3 top left: texture information from occluding parts of the mesh project incorrectly onto other geometry parts. With respect to Floating Textures, this not only affects rendering quality but also the reliability of flow field estimation.

A common approach to handle the occlusion problem is to establish a binary visibility map for each camera, multiply it with the weight map, and normalize the weights afterwards so they sum up to one. This efficiently discards occluded pixels in the input cameras for texture generation [CTMS03, LKG<sup>+</sup>03]. One drawback of this approach is that it must be assumed that the underlying geometry is precise, and cameras are precisely calibrated. In the presence of coarse geometry, the usage of such binary visibility maps can create occlusion boundary artefacts at pixels where the value of the visibility map suddenly changes, Fig. 4.3 left column.

To counter these effects, we create a “soft” visibility map  $\Omega$  for the current viewpoint and every input camera using a distance filter on the binary map:

$$\Omega(x, y) = \begin{cases} 0 & \text{if } \delta(x, y) = 0 \\ \frac{occDist(x, y)}{r} & \text{if } occDist(x, y) \leq r \\ 1 & \text{else} \end{cases} \quad (4.5)$$

Here  $r$  is a user-defined radius, and  $occDist(x, y)$  is the distance to the next occluded pixel. If  $\Omega$  is multiplied with the weight map, (4.5) makes sure that occluded regions stay occluded, while hard edges in the final weight map are removed. Using this “soft” visibility map effectively removes the above mentioned occlusion artefacts, Fig. 4.3 bottom right.




Figure 4.3: Top left: Projection errors if occlusion is ignored. Top right: Optical flow estimation goes astray if occluded image regions are not properly filled. Second row, left: Visualization of a binary visibility map for three input cameras. Second row, right: Visualization of a soft visibility map for three input cameras. Third row, left: Weight map multiplied with the binary visibility map. Third row, right: Weight map multiplied with the soft visibility map; note that no sudden jumps of camera weights occur anymore between adjacent pixels. Bottom left: Final result after texture projection using weight map with binary visibility. Bottom right: Final result after texture projection using weight map with soft visibility. Note that visible seams and false projections have been effectively removed.

## 4.2 Soft Visibility

17

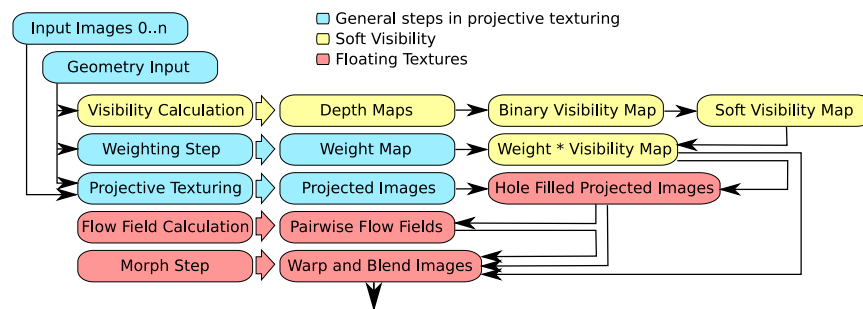


Figure 4.4: Complete overview of our Floating Textures algorithm on GPU. See text for details.

To improve optical flow estimation, we fill occluded areas in the projected input images  $\mathbf{I}_i^V$  with the corresponding colour values from that camera whose weight  $\omega$  for this pixel is highest. Otherwise, the erroneously projected part could seriously influence the result of the Floating Texture output. With hole filling, the quality of the flow calculation is strongly improved. One could also think about using already warped texture values from surrounding “secure” areas for iterative hole filling.



## Chapter 5

# GPU Implementation

The following is a description of our complete algorithm for dynamic scenes. A block diagram is given in Fig. 4.4. The extension to static scenes is straightforward. We assume that camera parameters, input images and a geometry proxy is given. The geometry representation can be of almost arbitrary type, e.g., a triangle mesh, a voxel representation, or a depth map (even though correct occlusion handling with a single depth map is not always possible due to the 2.5D scene representation).

First, given a novel viewpoint, we query the closest camera positions. For sparse camera arrangements, we typically choose the two or three closest input images. We render the geometry model from the cameras' viewpoints into different depth buffers. These depth maps are then used to establish for each camera a binary visibility map for the current viewpoint, similar in spirit to [CTMS03]. We use these visibility maps as input to the soft visibility shader. The calculation of  $\Omega$  can be efficiently implemented in a two-pass fragment shader. Next, a weight map is established by calculating the camera weights per output pixel. We use an angular weighting scheme similar to [BBM<sup>+</sup>01]. The final camera weights for each pixel in the output image are obtained by multiplying the weight map with the visibility map and normalizing it so that the weights sum up to 1.

To create the input images for the flow field calculation, we render the geometry proxy from the desired viewpoint several times into multiple render targets, in turn projecting each input photo onto the geometry. If the weight for a specific camera is 0 for a pixel, the colour from the input camera with the highest weight at this position is used instead.

To compute the optical flow between two images we rely on a GPU-optimized implementation of the optical flow technique by Brox *et al.* [BBPW04]. We found this algorithm not only very accurate but also quite robust to noise. Optical flow computation time depends on image resolution as well as on the amount of texture mismatch. Per rendered frame and three input images, we need to compute six flow fields. Even though this processing step

is computationally expensive and takes approximately 90% of the rendering time, we still achieve between 5 and 24fps at  $1024 \times 768$ -pixel rendering resolution on an NVidia GeForce 8800GTX. Compared to the results of a CPU implementation presented in [BBPW04], we observe a speedup factor of 70. While we also experimented with faster optical flow algorithms, like a multi-scale implementation of the well known technique by Horn and Schunck [HS81], we found that results were not as satisfactory. Tests have shown that the limiting speed factor is, in fact, not computational load, but the high number of state changes necessary to compute the flow fields.

Once all needed computations have been carried out, we can combine the results in a final render pass, which warps and blends the projected images according to the weight maps.

## Chapter 6

# Results, Discussion

To evaluate the proposed texturing approach, we have tested Floating Textures in conjunction with a number of different image-based rendering approaches. All tests were carried out using an OpenGL/GLSL implementation of our algorithm on a NVidia GeForce 8800GTX graphics card. Floating Texture frame rates vary between 5 and 24 fps, depending on the number of input images used and the amount of mismatch between textures which influences the number of iterations needed for the optical flow estimation algorithm to converge. The different IBR approaches with which we evaluated Floating Textures are (Figure 7.2, from top to bottom):

1. Synthetic Data Set: 49 input images synthesized from a textured 3D model, ground truth available;
2. Polyhedral Visual Hull Rendering: shape-from-silhouette reconstruction [FB03];
3. Free-Viewpoint Video: a parameter-fitted high-resolution 3D human model [CTMS03];
4. SurfCap: high-resolution geometry reconstructed using silhouette, feature, and stereo cues [SH07];
5. Light Field Rendering: a sub-sampled version of the Stanford Buddha light field data set [LH96].

Figure 7.1 depicts the corresponding geometry proxies for each IBR method. For each of these five different IBR techniques, we compared four different texturing approaches (Figure 7.2, from left to right):

1. Bandlimited Texturing [CCST00],
2. Filtered Blending [ESM07],
3. Unstructured Lumigraph Rendering [BBM<sup>+</sup>01], and

#### 4. Floating Textures.

**Synthetic data set** The ground truth model of the Stanford Bunny consists of 65k triangles. As input images, we rendered 49 views from random directions applying a coloured checkerboard pattern as texture. We then reduced the mesh to 948 triangles to use it as coarse geometry proxy. Bandlimited reconstruction as well as Filtered Blending introduce considerable blurring along texture discontinuities. Unstructured Lumigraph rendering, on the other hand, leads to ghosting artefacts (aliasing). Floating Textures, instead, is able to compensate for most texture mismatch and generates a crisp texture.

**Polyhedral Visual Hull Rendering** We tested different texturing approaches for the exact polyhedral visual hull reconstruction approach by Franco and Boyer [FB03], both in an off-line setup as well as in a real-time "live" system. Our setup of the "live" system consists of 8 cameras with a resolution of 1024x786 pixels that are connected in pairs of two to 4 PCs (the camera nodes), arranged in an approximate quarter-dome. One PC is used to calculate the approximate geometry of the scene from silhouettes obtained from the camera nodes. The camera nodes calculate these silhouettes by downsampling the input images and performing background subtraction. The walls of the scene are covered in green cloth to facilitate this process. The approximate geometry is sent to a PC which renders the final image. The rendering algorithm takes the images from 3 cameras to texture the approximate geometry so only these three images are sent over the network to conserve bandwidth. The system allows to capture and render scenes at approximately 10 fps and 640x480-pixel output resolution. Even though the reconstructed visual hulls are only very approximate geometry models, the Floating Textures method is able to remove the ghosting artefacts prevalent in Unstructured Lumigraph rendering, Fig. 1.1.

In the offline acquisition setup, we recorded a dancer using eight cameras arranged in a full circle. Excessive blurring is the result if bandlimited and filtered blending are applied, Fig. 7.2. With linear blending as applied in Unstructured Lumigraph rendering, ghosting and projection errors degrade rendering quality of the face and at the shoulders. These artefacts are efficiently removed by Floating Textures without introducing any additional blur.

#### **Free-Viewpoint Video**

For Free-Viewpoint Video acquisition, eight cameras are regularly spaced around a full circle [CTMS03]. Due to the cameras' far spacing, bandlimited and filtered blending eliminate all texture details, Fig.7.2. Since in Free-Viewpoint Video, a generic, 3D model is fit to the video streams by adapting only a good handful of animation parameters, the model surface corresponds only approximately to the person's actual 3D geometry, even though model geometry is very detailed. This causes noticeable ghosting

artefacts if Unstructured Lumigraph rendering is applied. The Floating Textures approach corrects for the projective texture mismatch and yields well-defined facial details, Fig.7.2.

**SurfCap** This data set was kindly provided to us from the SurfCap: Surface Motion Capture project [SH07]. Again, eight cameras are regularly spaced all around a full circle. Far camera spacing prevents bandlimited and filtered blending to preserve details. In computationally elaborate off-line processing, a highly tessellated, smooth geometry mesh is reconstructed, Fig. 7.1. Even though the mesh consists of 260+k triangles, ghosting and occlusion artefacts still degrade rendering quality if Unstructured Lumigraph rendering is applied. With Floating Textures, in contrast, virtually artefact-free rendering results are obtained.

**Light Field Rendering** We down-sampled the Stanford Buddha light field data set to  $8 \times 8$  images. While bandlimited rendering indiscriminately blurs away all details, more details are preserved in filtered blending, Fig. 7.2. Unstructured Lumigraph rendering (which corresponds to quadrilinear interpolation for light field rendering) introduces ghosting. With Floating Textures, in contrast, we achieve rendering results that are visually almost indistinguishable from the ground-truth. By using the Floating Textures approach in conjunction with light field rendering, comparable rendering results are obtainable from considerably fewer input images.



## Chapter 7

# Discussion & Conclusions

We have presented a new, general method to improve projective texturing using multi-view imagery in conjunction with some 3D geometry proxy. Our Floating Textures approach eliminates ghosting and occlusion artefacts and achieves improved rendering quality from coarser 3D geometry, fewer input images, and less accurate calibration. This saves memory, bandwidth, acquisition time, and money.

While we did not observe any problems during our evaluation experiments, it is obvious that strongly specular surfaces or badly colour-calibrated cameras will cause problems for the optical flow estimation. Also, if texture mismatch (ghosting) is too large, the optical flow algorithm might not be able to find correct correspondences. In general, however, the results with Floating Textures will not be worse than linear interpolation schemes. In our Floating Textures method, we deliberately disregard the epipolar constraint and allow textures to locally “float” in all directions. This is not a bug but a feature: this way, Floating Textures can compensate for imprecise camera calibration, while the small shifting of the texture on the surface is visually completely imperceptible.

One source for artefacts remaining in rendering dynamic scenes is that of temporally incoherent geometry. In the future, we intend to investigate how Floating Textures might be extended to compensate for temporal inconsistencies of the geometry proxy. Another line of research will be to improve object silhouette appearance by image-based means. Finally, for highly reflective or transparent objects, motion layer decomposition promises to be another interesting research direction.




Figure 7.1: Geometry proxies corresponding to the different IBR methods evaluated in our experiments, Fig. 7.2.

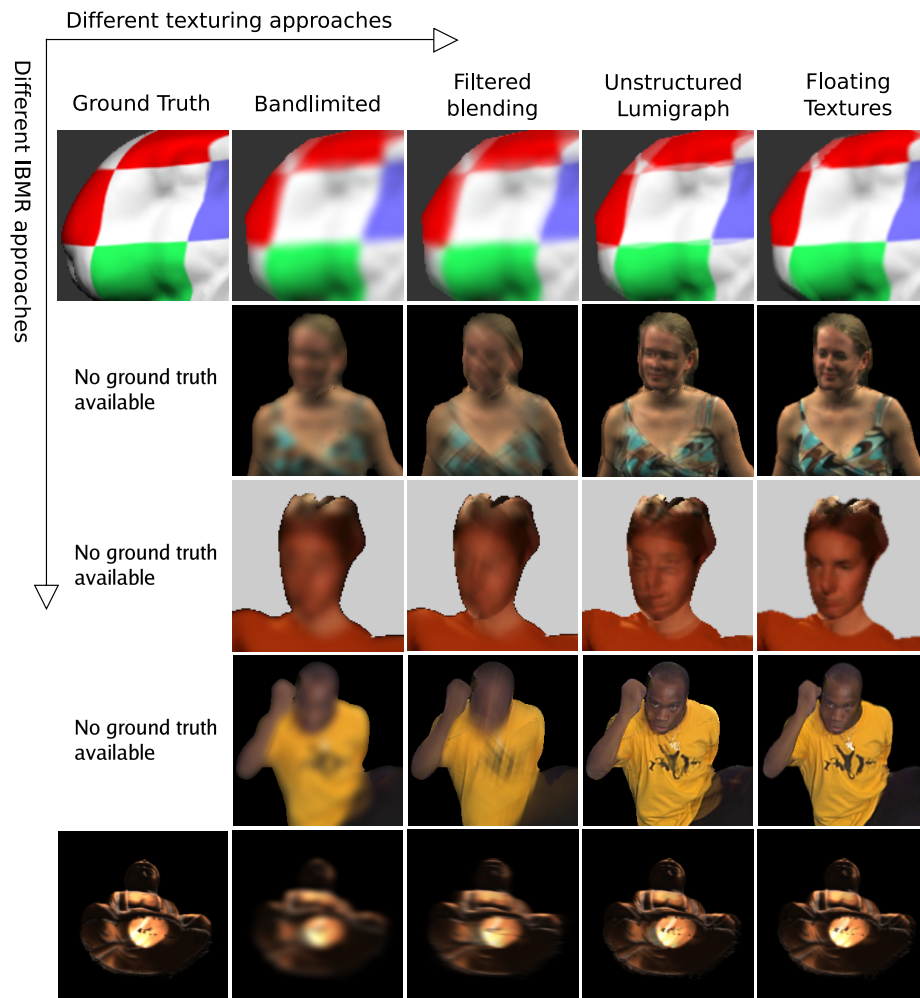


Figure 7.2: Comparison of different texturing schemes in conjunction with a number of IBMR approaches. From left to right: Ground truth image (where available), bandlimited reconstruction[CCST00], Filtered Blending[ESM07], Unstructured Lumigraph Rendering[BBM<sup>+</sup>01], and Floating Textures. The different IBMR methods are (from top to bottom): Synthetic data set, Polyhedral Visual Hull Rendering[FB03], Free-Viewpoint Video[CTMS03], SurfCap[SH07], and Light Field Rendering[LH96].

## 7 Discussion & Conclusions

28

---

## Chapter 8

# Acknowledgements

We would like to thank Jonathan Starck for providing us with the SurfCap test data (<http://www.ee.surrey.ac.uk/CVSSP/VMRG/surfcap.html>). The Max-Planck-Institute für Informatik for letting us use their studio for the dance scene, and the EDM at the university of Hasselt for allowing us to use their real-time acquisition studio.

## 8 Acknowledgements

30

---

# Bibliography

- [AB91] E. H. Adelson and J. R. Bergen. The Plenoptic Function and the Elements of Early Vision. *M. Landy and J. A. Movshon, (eds) Computational Models of Visual Processing*, pages 3–20, 1991.
- [AFYC02] Daniel G. Aliaga, Thomas Funkhouser, Dimah Yanovsky, and Ingrid Carlbom. Sea of images. In *VIS '02: Proceedings of the conference on Visualization '02*, Washington, DC, USA, 2002. IEEE Computer Society.
- [AYFC03] Daniel G. Aliaga, Dimah Yanovsky, Thomas Funkhouser, and Ingrid Carlbom. Interactive Image-Based Rendering Using Feature Globalization. In *I3D '03: Proceedings of the 2003 symposium on Interactive 3D graphics*, pages 163–170, New York, NY, USA, 2003. ACM Press.
- [Bau02] Adam Baumberg. Blending images for texturing 3d models. In Paul L. Rosin and A. David Marshall, editors, *Proceedings of the British Machine Vision Conference*, 2002.
- [BBM<sup>+</sup>01] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured Lumigraph Rendering. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 425–432, 2001.
- [BBPW04] Thomas Brox, Andrés Bruhn, Nils Papenberg, and Joachim Weickert. High accuracy optical flow estimation based on a theory for warping. In *Computer Vision - ECCV 2004, 8th European Conference on Computer Vision, Prague, Czech Republic, May 11-14, 2004. Proceedings, Part IV*, pages 25–36, 2004.
- [Cat74] E.E. Catmull. *A Subdivision Algorithm for Computer Display of Curved Surfaces*. PhD thesis, Departement of Computer Sciences, University of Utah, 12 1974.

- [CCST00] Jin-Xiang Chai, Shing-Chow Chan, Heung-Yeung Shum, and Xin Tong. Plenoptic Sampling. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 307–318, 2000.
- [CTMS03] J. Carranza, C. Theobalt, M. Magnor, and H. P. Seidel. Free-viewpoint video of human actors. In *SIGGRAPH '03: Proceedings of the 30th annual conference on Computer graphics and interactive techniques*, pages 569–577, 2003.
- [CW93] Shenchang Eric Chen and Lance Williams. View Interpolation for Image Synthesis. *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 279–288, 1993.
- [DTM96] Paul E. Debevec, Camillo J. Taylor, and Jitendra Malik. Modeling and Rendering Architecture from Photographs: A Hybrid Geometry- and Image-Based Approach. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 11–20, 1996.
- [ESM07] Martin Eisemann, Anita Sellent, and Marcus Magnor. Filtered Blending: A new, minimal Reconstruction Filter for Ghosting-Free Projective Texturing with Multiple Images. *Proc. Vision, Modeling, and Visualization (VMV'07)*, Saarbrücken, Germany, 11 2007.
- [FB03] Jean-Sébastien Franco and Edmond Boyer. Exact polyhedral visual hulls. In *Proceedings of the Fourteenth British Machine Vision Conference*, pages 329–338, September 2003. Norwich, UK.
- [GGSC96] Steven J. Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F. Cohen. The Lumigraph. In *SIGGRAPH '00: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54, 1996.
- [HS81] B.K.P. Horn and B.G. Schunck. Determining Optical Flow. *Artificial Intelligence*, 16(1–3):185–203, August 1981.
- [IMG00] Aaron Isaksen, Leonard McMillan, and Steven J. Gortler. Dynamically Reparameterized Light Fields. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 297–306, 2000.
- [LH96] Marc Levoy and Pat Hanrahan. Light Field Rendering. In *SIGGRAPH '96: Proceedings of the 23rd annual conference*

- on Computer graphics and interactive techniques*, pages 31–42, 1996.
- [LK81] B.D. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proc. Seventh International Joint Conference on Artificial Intelligence*, pages 674–679, 1981.
- [LKG<sup>+</sup>03] Hendrik P. A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, and Hans-Peter Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.*, 22(2):234–257, 2003.
- [LS04] Zhouchen Lin and Heung-Yeung Shum. A Geometric Analysis of Light Field Rendering. *International Journal of Computer Vision*, 58(2):121–138, 2004.
- [MB95] Leonard McMillan and Gary Bishop. Plenoptic Modeling: An Image-Based Rendering System. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 39–46, 1995.
- [MBR<sup>+</sup>00] W. Matusik, C. Buehler, R. Raskar, S. Gortler, and L. McMillan. Image-based visual hulls. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 369–374, 2000.
- [MP04] W. Matusik and H. Pfister. 3D TV: A scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. In *SIGGRAPH '04: Proceedings of the 31st annual conference on Computer graphics and interactive techniques*, pages 814–824, 2004.
- [RCMS99] Claudio Rocchini, Paolo Cignomi, Claudio Montani, and Roberto Scopigno. Multiple textures stitching and blending on 3D objects. In *Proceedings of the Eurographics Workshop on Rendering*, pages 119–130, 1999.
- [SH07] J. Starck and A. Hilton. Surface capture for performance based animation. *IEEE Computer Graphics and Applications*, 27(3):21–31, 2007.
- [SKvW<sup>+</sup>92] Mark Segal, Carl Korobkin, Rolf van Widenfelt, Jim Foran, and Paul Haeberli. Fast Shadows and Lighting Effects using Texture Mapping. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 249–252, 1992.

- [SYGM03] J. Stewart, J. Yu, S. J. Gortler, and L. McMillan. A New Reconstruction Filter for Undersampled Light Fields. In *Proceedings of the Eurographics Workshop on Rendering*, pages 150–156, 2003.
- [VBK05] S. Vedula, S. Baker, and T. Kanade. Image based spatio-temporal modeling and view interpolation of dynamic events. *ACM Transactions on Graphics*, 24(2):240–261, April 2005.
- [WAA<sup>+</sup>00] Daniel N. Wood, Daniel I. Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H. Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 287–296, 2000.
- [WWG07] M. Waschbüsch, S. Würmlin, and M. Gross. 3D Video Billboard Clouds. *Computer Graphics Forum (Proceedings of the Eurographics 2007)*, 26(3), 2007.
- [ZKU<sup>+</sup>04] C. Zitnick, S.B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High-quality video view interpolation using a layered representation. In *SIGGRAPH '04: Proceedings of the 31st annual conference on Computer graphics and interactive techniques*, pages 600–608, 2004.
- [ZWT<sup>+</sup>05] Kun Zhou, Xi Wang, Yiyang Tong, Mathieu Desbrun, Baining Guo, and Heung-Yeung Shum. Texturemontage: Seamless texturing of arbitrary surfaces from multiple images. In *SIGGRAPH '05: Proceedings of the 32nd annual conference on Computer graphics and interactive techniques*, pages 1148–1155, 2005.