Light-Field Supported Fast Volume Rendering

Clemens Birklbauer and Oliver Bimber*
Institute of Computer Graphics, Johannes Kepler University Linz, Austria

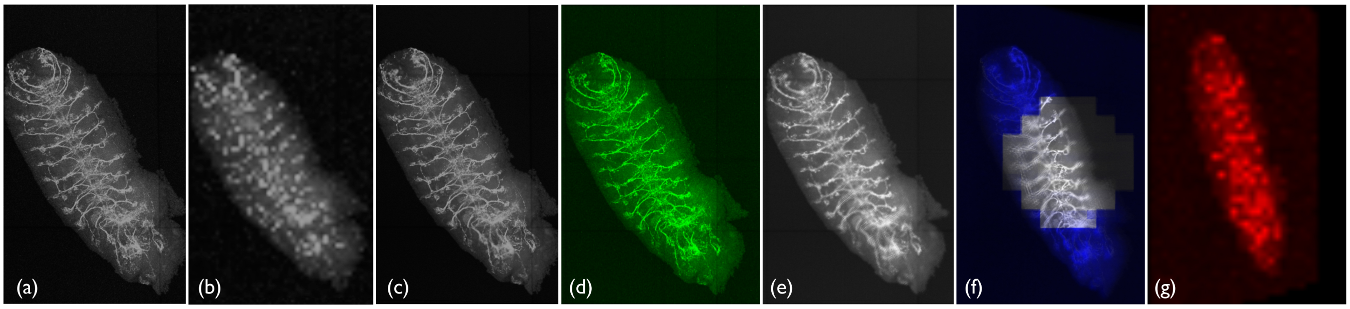

Figure 1: *Visualization of a drosophila: full-resolution volume rendering at 0.4 fps (a), volume rendering preview at 25 fps (b), our method at 25 fps (c). Color-coded contributions of different rendering sources during rotation (d-g): full-resolution volume rendering (green), volume rendering preview (red), full-resolution light-field rendering (gray), and low-resolution light-field rendering (blue). The visible seams are the result of the microscope's scanning process – not of visualization.*

Abstract

We present a combination of light-field rendering and volume rendering to enable the interactive exploration of large volumetric data sets. We recycle previously rendered images and use the idle times of the volume renderer for filling a cached-managed light field. The final images are then composed from both: light-field rendering and volume rendering – depending on the state of the light-field cache. Our method is never slower than the stand-alone volume render – but it accelerates significantly over time.

CR Categories: I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation—Display algorithms;

Keywords: light fields, volume rendering, image stacks, real-time

1 Introduction

Advances in imaging technology leads to a continues increase of image data sets. Modern scanning microscopes, for instance, produce image stacks with a megapixel lateral resolution and many hundreds to thousands slices in axial direction. This trend will continue – resulting in very large volumetric data sets that are difficult to explore interactively, since the complexity of volume rendering is proportional to spatial and lateral resolution of the data.

Light-field rendering is a fast and simple image-based rendering method that requires pre-computed image data. For volume rendering, each costly computed image is discarded after changing the viewing parameters, while it becomes idle if the viewing parameters are not changed and the visualization does not need to be updated.

We combine light-field rendering and volume rendering with two goals: We recycle previously rendered images and use the idle times for filling a cached-managed light field. The final images are then composed from light-field rendering and from volume rendering – depending on the state of the light-field cache. This leads to a significant increase in rendering performance and to the ability of exploring large volumetric datasets interactively.

2 Our Approach

For light-field rendering, we support a spherical light-field parameterization to enable surround navigation, as well as changing the field-of-view, aperture and focus. When not interacting with the volume, the volume renderer first computes, displays and caches a full-resolution image for the current viewing parameters, while it will then compute portions of the cached light-field data structure in the background. We apply dead reckoning to the users' interaction pattern to determine the priority-order of these portions. The volume renderer always computes the image portions at the necessary level of detail (LOD) – depending on the adjusted viewing parameters. They are cached and remain valid for light-field rendering until a higher LOD is required. In this case, the cache has to be updated by the volume renderer. For rendering the final image, the light-field renderer is used as much as possible. If the required portions of the light-field are cached, this results in a high frame-rate. Only those portions that cannot be provided by the light-field renderer have to be produced by the slow volume renderer. They will be cached as soon as they are available. In addition, we integrate the volume renderer's fast preview mode. Every time, the LOD produced by this mode is higher than the LOD in the light-field cache, we use the preview to deliver the corresponding parts of the final image. Furthermore, we manage a second, lower-resolution light-field cache that can be filled quicker than the full-resolution cache, and allow for a dynamic, on-demand increase or decrease of the light-fields' angular resolution. Thus, the final image is a piece-wise composition of image portions coming from four different sources: the full-resolution light-field, the low-resolution light-field, the volume renderer's preview and the volume rendered at full resolution.

*e-mail: {clemens.birklbauer,oliver.bimber}@jku.at

SIGGRAPH 2012, Los Angeles, California, August 5 – 9, 2012.
ISBN 978-1-4503-1435-0/12/0008