Computers & Graphics 50 (2015) 47–60

Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Technical Section

Multimodal volume illumination

Erik Sundén [a,1], Sathish Kottravel [a,2], Timo Ropinski [a,b,3]

[a] Linköping University, Department of Science and Technology, Interactive Visualization Group, SE-601 74 Norrköping, Sweden
[b] Ulm University, Institute for Media Informatics, Visual Computing Research Group, James-Franck-Ring, 89081 Ulm, Germany

ARTICLE INFO

Article history:
Received 22 October 2014
Received in revised form
7 May 2015
Accepted 10 May 2015
Available online 19 May 2015

Keywords:
Volume rendering
Volumetric illumination
Multimodal visualization

ABSTRACT

Despite the increasing importance of multimodal volumetric data acquisition and the recent progress in advanced volume illumination, interactive multimodal volume illumination remains an open challenge. As a consequence, the perceptual benefits of advanced volume illumination algorithms cannot be exploited when visualizing multimodal data – a scenario where increased data complexity urges for improved spatial comprehension. The two main factors hindering the application of advanced volumetric illumination models to multimodal data sets are rendering complexity and memory consumption. Solving the volume rendering integral by considering multimodal illumination increases the sampling complexity. At the same time, the increased storage requirements of multimodal data sets forbid to exploit precomputation results, which are often facilitated by advanced volume illumination algorithms to reduce the amount of per-frame computations. In this paper, we propose an interactive volume rendering approach that supports advanced illumination when visualizing multimodal volumetric data sets. The presented approach has been developed with the goal to simplify and minimize per-sample operations, while at the same time reducing the memory requirements. We will show how to exploit illumination-importance metrics, to compress and transform multimodal data sets into an illumination-aware representation, which is accessed during rendering through a novel light-space-based volume rendering algorithm. Both, data transformation and rendering algorithm, are closely intervened by taking compression errors into account during rendering. We describe and analyze the presented approach in detail, and apply it to real-world multimodal data sets from biology, medicine, meteorology and engineering.

1. Introduction

With the advancements in sensing technology the analysis of multimodal data sets has become more common in many scientific domains. As a consequence, visualization became a key technology that enabled exploration and understanding of the relationships between different modalities. While progress in the area of advanced volume illumination has led to improved spatial comprehension by depicting the relationships between spatial structures more clearly [1,2], interactive multimodal volume illumination is still an open challenge, and will continue to be as data set sizes increase. This is especially a limitation when dealing with multimodal acquisitions, where the relationship between the structures across modalities is one of the main benefits, and therefore should be preserved for better comprehension during exploration. When for instance performing cranial surgery, the overall structure of the brain, often best visible in MRI data, should be seen within the context of the supplying vessels, often imaged through contrast-enhanced CT acquisitions. Only when spatial relations can be understood across modalities, can the full potential of multimodal data be exploited during interactive exploration.

Several interactive volumetric illumination algorithms have been proposed [3], which increase the degree of realism in single modality visualizations, and many of these algorithms are known to improve the spatial comprehension [1,2]. While the separation of structures through separate transfer functions can be made for different modalities, a clear spatial comprehension between structures across modalities as well as for each separate modality is not easily achievable without introducing global additions such as occlusion or illumination, designed to enhance depth cues and separate the different structures in the data. Furthermore, the incorporation of light sources, which enable the illumination, does not affect the semantics encoded in a transfer function, and the general user interface for light sources can be considered superior to the user interface of transfer functions, in aspects such as intuitiveness and usability. It should be clearly noted that lighting does not replace the usage or importance of transfer functions, but complements it for a better separation of the structures in the data. However, the introduction of global illumination cannot be

E-mail addresses: erik.sunden@liu.se (E. Sundén), sathish.kottravel@liu.se (S. Kottravel), timo.ropinski@uni-ulm.de (T. Ropinski).
[1] Tel.: +46 11 36 36 26; fax: +46 11 36 32 70.
[2] Tel.: +46 11 36 31 17; fax: +46 11 36 32 70.
[3] Tel.: +49 731 50 24200; fax: +49 731 50 31302.

http://dx.doi.org/10.1016/j.cag.2015.05.004
0097-8493/

made without the global knowledge of what all modalities at a specific grid point (voxel) as a context of all the occlusion from a point towards a respective light, or around the neighborhood of each voxel is at some point required.

When designing such algorithms several challenges need to be overcome, in order to provide interactivity when changing rendering parameters, such as the camera or the transfer function. Due to the global nature of advanced volumetric illumination, each sampling operation is dependent on several other sampling operations, which are used to determine the influence of surrounding structures. As these additional sampling operations result in a severe performance bottleneck, a common approach has been to trade memory complexity for computing time. By using different forms of illumination caches, with varying precomputation or update times (see, for example, [4–6]) it became possible to store the relevant information in the graphics processing unit's (GPU) memory to enable advanced illumination of single volumetric data sets at interactive frame rates. Due to the increased sampling complexity of multimodal volumetric data, this scenario would benefit from precomputation in order to reduce rendering times. However, as these data sets comprise several volumes, in the future these enlarged memory footprints forbid exploitation of additional illumination caches during rendering. Thus, higher sampling complexity and memory demands go hand in hand for multimodal volume illumination, and must be tackled in an integrated manner.

In this paper, we introduce a novel approach for multimodal volume illumination which reduces the number of sampling operations as well as their complexity, while at the same time reducing the GPU memory footprint. The key idea is to limit illumination computations only to those samples which are visible to the camera as well as the light source. While this can reduce the number of sampling operations drastically, determining these samples is not easy as it is an order-dependent process. While many volume illumination algorithms work in a two pass process (first a light pass and second a camera pass) to deal with this order-dependence, our approach exploits a visibility function which can be queried in constant time in order to defer illumination computations and perform them only if necessary. This order-independent processing enables direct compositing into the camera view without requiring an intermediate illumination cache. Hence we can benefit from the advantages of deferred shading and exploit it for semi-transparent direct volume rendering of multimodal data. As a consequence our approach can incorporate multiple light sources (see Fig. 1) at interactive frame rates without increasing GPU memory requirements. While shadows in single light source setups are known to improve the spatial comprehension of volume rendered scenes [1], multiple light source setups come with the benefit that shadowed regions neither appear to be pitch black, nor suffer from a low contrasts introduced by intense ambient lighting. As it can be seen in Fig. 2, this leads to vivid shadows which do not cover critical structures, as for instance a tumor in an MRI scan. This observation is also in line with the findings made by Halle and Meng [7], which indicate that multiple light source setups are beneficial in visualization. Furthermore, Lee et al. [8] designed a system for calculating the optimal placement of multiple light sources for rendering of scientific data sets. Thus, to fulfill this goal we need efficient sampling and further reduced memory requirements. Therefore, we transform the multiple modalities into a single illumination-aware data structure from which we can draw samples in most cases in constant time. The error of this compression process is taken into account during rendering, which interlinks data compression and rendering algorithm in a tight manner. To solve the challenge of multimodal volume illumination at interactive frame rates we make the following contributions:

- A novel light-space volume rendering algorithm which supports multimodal volume illumination with multiple light sources by deferring illumination computations.

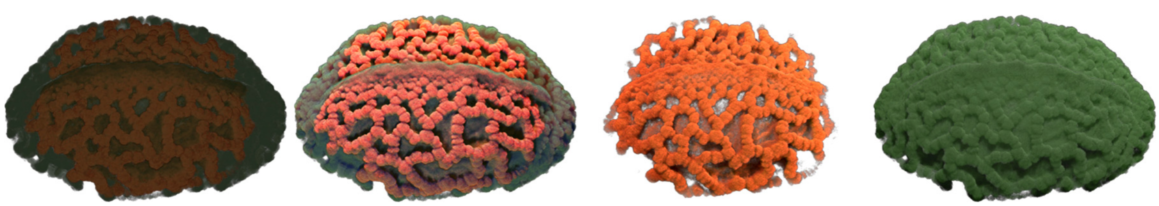

Fig. 1. Multimodal volume illumination applied to a two modality raster microscopy scan of a Lily pollen (367 × 331 × 135 voxels). While standard multimodal volume rendering does not show relations across modalities (first), multimodal volume illumination with two light sources depicts these relations (*second*). Showing the individual modalities next to each other with advanced volumetric illumination does also not convey their spatial relationship (third and fourth).

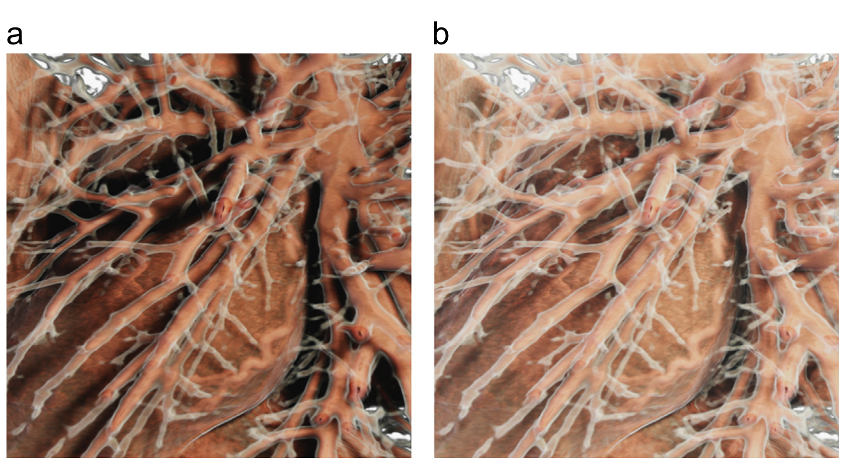

Fig. 2. Comparison between single light source setup (a) and multiple light source setup (3 point lights) (b).