Computerized Medical Imaging and Graphics 51 (2016) 40–49

Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

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

Efficient visibility-driven medical image visualisation via adaptive binned visibility histogram

Younhyun Jung[a,b,*], Jinman Kim[a,b], Ashnil Kumar[a,b], David Dagan Feng[a,b,c], Michael Fulham[d,e]

[a] The Institute of Biomedical Engineering and Technology, University of Sydney, Australia
[b] BMIT Research Group, School of Information Technologies, University of Sydney, Australia
[c] Med-X Research Institute, Shanghai Jiao Tong University, China
[d] Sydney Medical School, University of Sydney, Australia
[e] Department of Molecular Imaging, Royal Prince Alfred Hospital, Australia

ARTICLE INFO

Article history:
Received 7 October 2015
Received in revised form 14 March 2016
Accepted 12 April 2016

Keywords:
Adaptive binning
Visibility histogram
Cluster analysis
Transfer function
Medical imaging

ABSTRACT

'Visibility' is a fundamental optical property that represents the observable, by users, proportion of the voxels in a volume during interactive volume rendering. The manipulation of this 'visibility' improves the volume rendering processes; for instance by ensuring the visibility of regions of interest (ROIs) or by guiding the identification of an optimal rendering view-point. The construction of visibility histograms (VHs), which represent the distribution of all the visibility of all voxels in the rendered volume, enables users to explore the volume with real-time feedback about occlusion patterns among spatially related structures during volume rendering manipulations. Volume rendered medical images have been a primary beneficiary of VH given the need to ensure that specific ROIs are visible relative to the surrounding structures, e.g. the visualisation of tumours that may otherwise be occluded by neighbouring structures. VH construction and its subsequent manipulations, however, are computationally expensive due to the histogram binning of the visibilities. This limits the real-time application of VH to medical images that have large intensity ranges and volume dimensions and require a large number of histogram bins. In this study, we introduce an efficient adaptive binned visibility histogram (AB-VH) in which a smaller number of histogram bins are used to represent the visibility distribution of the full VH. We adaptively bin medical images by using a cluster analysis algorithm that groups the voxels according to their intensity similarities into a smaller subset of bins while preserving the distribution of the intensity range of the original images. We increase efficiency by exploiting the parallel computation and multiple render targets (MRT) extension of the modern graphical processing units (GPUs) and this enables efficient computation of the histogram. We show the application of our method to single-modality computed tomography (CT), magnetic resonance (MR) imaging and multi-modality positron emission tomography-CT (PET-CT). In our experiments, the AB-VH markedly improved the computational efficiency for the VH construction and thus improved the subsequent VH-driven volume manipulations. This efficiency was achieved without major degradation in the VH visually and numerical differences between the AB-VH and its full-bin counterpart. We applied several variants of the *K*-means clustering algorithm with varying *K*s (the number of clusters) and found that higher values of *K* resulted in better performance at a lower computational gain. The AB-VH also had an improved performance when compared to the conventional method of down-sampling of the histogram bins (equal binning) for volume rendering visualisation.

1. Introduction

The advent of efficient volume rendering algorithms and powerful graphical processing units (GPUs) has enabled the introduction of direct volume rendering (DVR) visualisation that provides three-dimensional (3D) views and interactive navigation. DVR has been

* Corresponding author at: The Institute of Biomedical Engineering and Technology, University of Sydney, Australia.
E-mail address: yjun6175@uni.sydney.edu.au (Y. Jung).

http://dx.doi.org/10.1016/j.compmedimag.2016.04.003
0895-6111/

broadly applied to medical images in clinical diagnosis (Kim et al., 2004), planning for surgery/radiotherapy (Kruger et al., 2008), and training (Georgii et al., 2007). In DVR, transfer functions (TF) are used to control the optical properties of data value (intensity) in a volume and play an important role in defining a satisfactory visualisation. In practice, however, it is difficult and time-consuming to define TFs due to the lack of objective feedback mechanisms to quantitatively measure the influence of the TF manipulations on the resulting visualisation. Thus unintuitive and iterative tweaking of the TF parameters is inevitable. Correa and Ma (2011) proposed a visibility histogram (VH), which represents how much of each structure (intensity) is visible to the user (see the green bars in Fig. 1), to provide additional feedback to the user during TF manipulations. The VH was effective and meaningful in providing visual feedback to the user about the impact of specific volume manipulations on the resulting visualisation.

Many investigators have leveraged the VH to enhance the TF design by automatically optimising (maximising) the visibility of regions of interest (ROIs) (Cai et al., 2013; Ruiz et al., 2011; Wang et al., 2011), e.g. in multi-modality medical images (Jung et al., 2013).

Unfortunately, the high computational cost and large memory requirements limit the application to medical image volumes with large intensity ranges and volume dimensions. Fig. 1 shows the computation of the VHs for a medical image volume; it is necessary to compute the visibilities of every voxel and structure, including the lungs, muscle, and bone, as well as a histogram that represents their distributions along the intensity range (histogram bins). During histogram construction, the number of bins is the parameter that has the major impact on the computational cost. Prior investigators (Correa and Ma, 2011; Cai et al., 2013; Ruiz et al., 2011; Wang et al., 2011; Jung et al., 2013) have improved the computational performance by reducing the bin size and partitioning the intensity ranges to a smaller number of fixed bin sizes. Such an 'equal' binning (EB) technique, however, is not able to accurately represent the underlying distribution of its full-bin (FB) counterpart as shown by the missing peaks and valley in the 64 bin EB-VH in Fig. 1. The dissimilarity of the EB and FB distributions means that users gain inaccurate feedback to guide their subsequent TF manipulation, thus resulting in unwanted visualisations. For instance, using 64 EB-VH, users would need to control the red point (marked with (2) in Fig. 1) to decrease the corresponding visibility peak and increase the visibilities of the lungs and bone. In contrast, the FB-VH guides them to control green (1) or blue points (3), which are the dominant visibility peaks in the histogram distribution.

In this study, our plan was to compute VHs by adaptively binning (AB) them with a cluster analysis algorithm. In our algorithm, the voxels in a volume are 'clustered' into a smaller number of related intensity bins. The clustered bins are then used to build an AB-VH that can adaptively allocate bins according to the distribution of the intensity range. As outlined in Fig. 1, our 64 bin AB-VH allocated more bins (black circle in the VHs) to intensity ranges that have large variations, e.g. the muscle, when compared to the EB counterpart. Our AB thus preserved the dominant peaks and valleys that are missing with EB. As such, the visual feedback provided by the AB during TF manipulation is consistent with the FB approach. Due to the reduced number of bins, we are able to leverage the Multiple Render Targets (MRT) extension from the GPU architecture. This allows a greater number of bins to be stored in memory simultaneously, thereby enabling histograms with 32 bins to be constructed in a single GPU execution (rendering pass) instead of eight rendering passes. We will show the application of our AB algorithm in a variety of common medical image modalities including high-resolution computed tomography (CT), magnetic resonance (MR) imaging and multi-modality positron emission tomography-CT (PET-CT). We adopted well-known and robust generic K-means clustering algorithms. K-means clustering is capable of partitioning medical images into groups of related voxels that are homogeneous with respect to the voxel's attributes, e.g. grey-scale intensity (Pal and Pal, 1993; Pham et al., 2000) and it has also been used to segment CT (Bae et al., 1993; Lee et al., 2008), PET (Lee, 2010) and MR studies (Singh et al., 1996; Yan and Karp, 1994). We note that for VH computation, the primary emphasis is not on the accuracy of the resulting segmentation, but in the ability to accurately represent the distribution of the intensity ranges in the histogram.

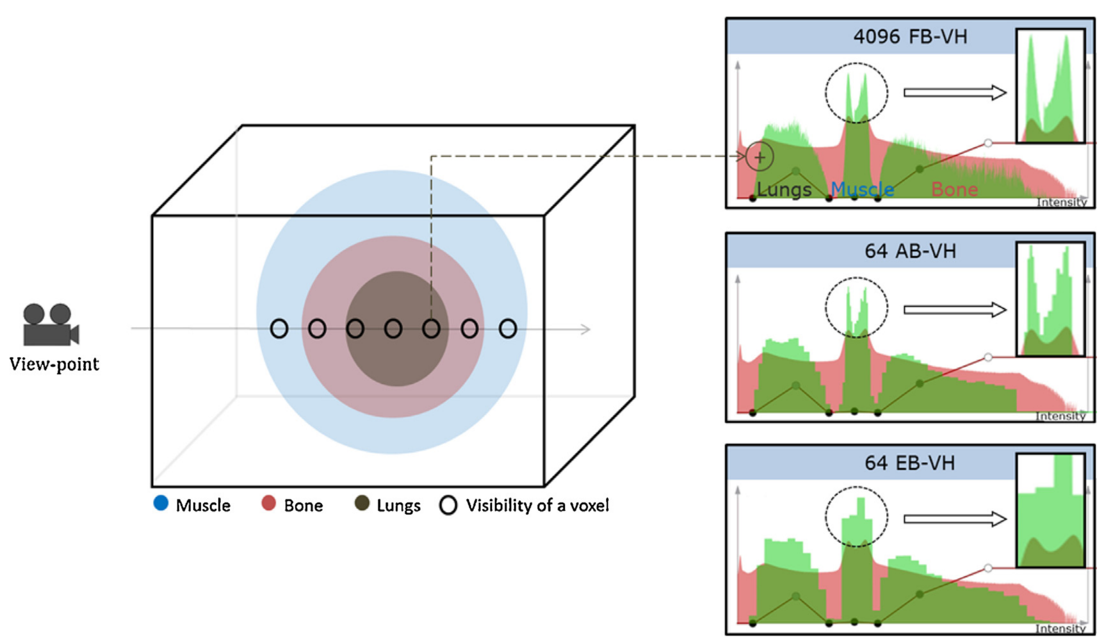

Fig. 1. VH construction with 3 different binning schemes. Note that AB with 64 bins (64 AB-VH) is able to preserve the major distribution of the intensity ranges when compared to its FB counterpart with 4096 bins (4096 FB-VH). EB with 64 bins (64 EB-VH) is unable to represent the local peaks and valley as indicated by the black circle. (For interpretation of the references to colour in this text, the reader is referred to the web version of this article.)