



# Volumetric Clipping Surface: Un-occluded Visualization of Structures Preserving Depth Cues into Surrounding Organs

Bhavya Ajani, Aditya Bharadwaj, and Karthik Krishnan<sup>(✉)</sup>

Samsung Research India, Bangalore, India  
bhavya.ajani@samsung.com

**Abstract.** Anatomies of interest are often hidden within data. In this paper, we address the limitations of visualizing them with a novel dynamic non-planar clipping of volumetric data, while preserving depth cues at adjacent structures to provide a visually consistent anatomical context, with no-user interaction. An un-occluded and un-modified display of the anatomies of interest is made possible. Given a semantic segmentation of the data, our technique computes a continuous clipping surface through the depth buffer of the structures of interest and extrapolates this depth onto surrounding contextual regions in real-time. We illustrate the benefit of this technique using Monte Carlo Ray Tracing (MCRT), in the visualization of deep seated anatomies with complex geometry across two modalities: (a) Knee Cartilage from MRI and (b) bones of the feet in CT. Our novel technique furthers the state of the art by enabling turnkey immediate appreciation of the pathologies in these structures with an unmodified rendering, while still providing a *consistent* anatomical context. We envisage our technique changing the way clinical applications present 3D data, by incorporating *organ viewing presets*, similar to *transfer function presets* for volume visualization.

**Keywords:** Ray tracing · Focus+Context · Occlusion

## 1 Introduction

3D datasets present a challenge for Volume Rendering, where regions of interest (ROI) for diagnosis are often occluded. These ROIs usually cannot be discriminated from occluding anatomies by setting up a suitable transfer function. Clipping planes, Cropping and scalpel tools have been widely used to remove occluding tissue and are indispensable features of every Medical Visualization workstation.

However, none of the existing techniques render the ROIs un-occluded while maintaining depth continuity with the surrounding. In this paper we address this limitation by introducing a novel dynamic non-planar clipping of volumetric data. Matching depth between the ROIs and surrounding for improved depth perception, while still supporting an un-occluded, un-modified visualization of

the ROIs with no user interactions in real-time is the key contribution of this work. Additionally, we describe our technique in the context of Cinematic rendering.

### 1.1 Focus+Context

Focus+Context (F+C) is well studied in visualization [1]. It uses a segmentation of the data to highlight the ROIs (*Focus*) while still displaying the surrounding anatomies (*Context*). The principle of F+C is that for the user to correctly interpret data, interact with it or orient oneself, the user simultaneously needs a detailed depiction (Focus) along with a general overview (Context). Existing F+C techniques resort to a distortion of the *visualization space*, by allocating more *space* (importance sampling, various optical properties, viewing area etc.) for the Focus [1, 2]. Methods include *cut-aways* (where fragments occluding the view are removed) [3, 4], rendering the context with different optical properties [3], *ghosting* of the context (where contextual fragments are made more transparent) [4] and importance sampling with several forms of sparsity [4, 5] and exploded views and deformations to change the position of context fragments. Wimmer et. al [6] extended ghosting techniques to create a *virtual hole cut-away* visualization using various clipping functions such as box and sphere so as to create a vision channel for deep seated anatomies.


Humans determine spatial relationships between objects based on several depth cues [2]. In surgical planning, correct depth perception is necessary to understand the relation between vessels and tumors. State of the art ray tracing methods use various techniques including shadows to highlight foreground structures for improved depth perception. Ultimately, depth perception often necessitates interactions such as rotation.

### 1.2 Clinical Application

We demonstrate our technique across two different modalities: Knee cartilage in T1w MRI and complex bones of the ankle and foot in CT.

Knee MR scans are the third most common type of MRI examination [7] and Knee Osteo-arthritis is the leading cause of global disability. Lesions shows up as pot holes; varying from full-thickness going all the way through the cartilage, to a partial-thickness lesion. Subtle cartilage lesions are notoriously difficult to detect. A considerable number of chondral lesions (55%) remain undetected until arthroscopy [8]. The mean thickness of healthy cartilages in the knee varies from 1.3 to 2.7 mm [9]. Visualization of the cartilage and its texture enables better diagnosis. However, un-occluded visualization along with context to appreciate the injury and the degradation of a structure that is so thin, curved and enclosed by several muscles and bones in the knee is challenging.

*Cinematic Rendering* which uses MCRT has advanced state of the art in medical visualization [10]. It has been used in the clinic to generate high quality realistic images primarily with CT, but also using MR. Advances in Deep



**Fig. 1.** Visualization of knee on a T1w 1.5T MRI ( $300 \times 344 \times 120$  voxels, sagittal acquisition, resolution  $0.4 \times 0.4 \times 1$  mm). (a) Rendering the whole volume completely occludes the cartilage (b) Clipping plane requires manual adjustment, yet has a poor cartilage coverage due to its topology (c) Cutaway view of the cartilage, where occluding fragments are removed. Note the depth mismatch between the cartilage and surrounding anatomy causing perceptual distortion. Also note poor lighting of the focus. (d) Our proposed VCS method, with the same viewing parameters shows the cartilage with maximal coverage and smoothly extrapolates depth onto surrounding structures, allowing for improved appreciation of contextual anatomy in relation to the focus. (e) Focus (cartilage) is outlined in yellow. Boundary points of this are sampled as indicated by the control points in green to compute a clipping surface spline. It is worthwhile mentioning that an accurate cartilage segmentation is not typically necessary since the bone is hypo-intense compared to the cartilage. (f) CT foot (64 slice CT VIX, OsiriX data) using the proposed VCS method. The bones were segmented by a simple threshold at 200HU. Note that the method captures the non-planar structure of bones of the feet successfully.

Learning have made possible computation of accurate segmentations. There is a need for visualization techniques to generate high quality Focus specific F+C renderings to enable faster diagnosis.

## 2 Existing Techniques

### 2.1 Clipping Planes

Clipping planes can be used to generate an un-modified dissection view. Figure 1b shows this visualization of the knee cartilage using MCRT depicting lesions in the femoral cartilage. The placement of the clipping plane requires significant interaction. Since the cartilage is a thin structure covering the entire curved joint, the clipped view results in low coverage of the cartilage.

### 2.2 Cut-Aways

*Cut-aways* were first proposed in [4]. The idea is to cut away fragments occluding the Focus. Occluding fragments are rendered fully transparent, therefore unlike ghosting it provides an unmodified view of the cartilage. This is possible using two render passes. The Focus depths at the current view are extracted in the first pass, by checking if the ray intersects with its segmentation. A second pass renders all data. The starting locations of the rays that intersect the Focus are set to the depth extracted from the first pass so that Focus is rendered un-occluded.

A cut-away of the cartilage is shown in Fig. 1c. Note the boundaries of the cartilage where there is a clear depth mismatch resulting in *cliffs* in the visualization causing a perceptual distortion. Also note the poor lighting of the cartilage, with shadows of the context cast onto the focus, making a contralateral assessment difficult.

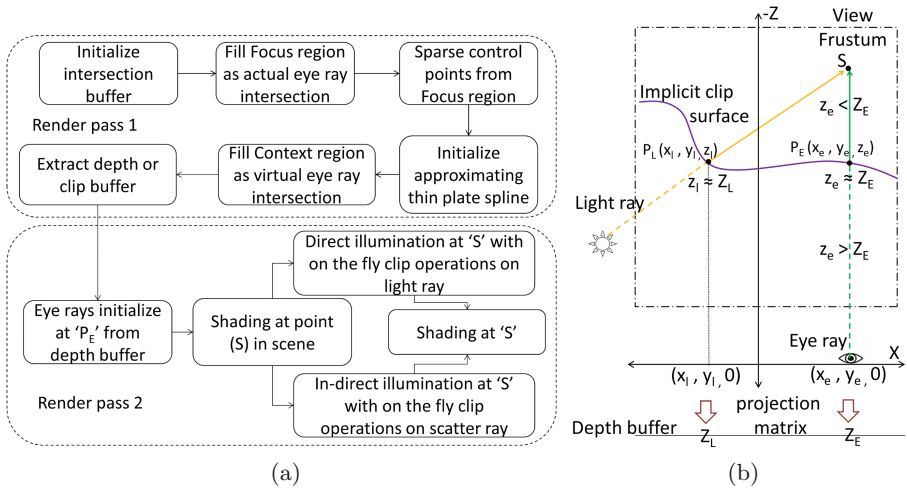
## 3 A Real-Time Depth Contiguous Clipping Surface

Similar to other F+C techniques our method requires prior semantic segmentation to define Focus and Context. Automatic segmentation of Focus (Cartilage) in T1w MRI is derived using deep learning as explained in our previous work [11].

### 3.1 Methodology

We extrapolate depth for Context from Focus. We use approximating Thin Plate Splines (TPS) [12] to provide a smooth, differentiable depth through the Focus onto the rest of the view frustum. Figure 1d shows the MCRT from the same viewpoint using the proposed method. The Volumetric Clipping Surface (VCS) is implicitly defined through a depth buffer in an orthographic view space or frustum.

We render the scene in two render passes. In the first render pass we compute the depth buffer that maintains depth continuity. In a second render pass we



**Fig. 2.** (a) Overview of MCRT rendering in two render passes with a Focus specific clipping surface. Render pass 1 is done once. Render pass 2 is carried out multiple times to refine the MCRT render estimate. (b) On the fly clip operations against an implicit clip surface (marked in purple). Eye ray is marked in green and Light ray is marked in yellow. The scatter ray is not shown to avoid clutter. The clipped portion of rays are shown as a dashed line.

render the scene by clipping the eye and light rays based on this computed depth buffer (see Fig. 2b); thereby implicitly clipping them with the VCS in a warp view space. MCRT is an iterative rendering process, where several iterations are used to arrive at the estimate of the scene for a given set of viewing parameters. The first render pass is carried out once, while the second render pass is a part of each MCRT iteration.

In the first render pass, we compute an intersection buffer which stores the points of intersection (in view space) for all eye rays. The intersection buffer is conceptually divided into three distinct regions. These are regions where the eye rays (a) intersect the Focus ROI (b) intersect the Context ROI (c) do not intersect the model bounding box. The first render pass consists of two phases. In the first phase, the Focus region of the intersection buffer is filled. This contains the *actual intersection points* for all eye rays intersecting with Focus (i.e. those that intersect the segmentation texture). In the second phase, we estimate *virtual intersection points* for the Context region, from the computed *actual intersection points* in the Focus region. This provides us with a C0, C1, C2 continuous intersection buffer in view space, which we call an intersection surface.

This is done as shown in Fig. 2a. We select a sparse set of control points falling on boundary of the Focus region of the intersection buffer, by uniformly sampling the boundary contour points. In this work, we use  $N = 50$  control points. Using these points, we initialize an approximating TPS taking all  $(x, y)$  co-ordinates of control points (i.e. origin of corresponding eye ray) as data sites and its  $z$  co-ordinate (i.e. intersection depth in view space) as data value. We choose a

spline approximating parameter  $p$  as 0.5. The computed surface spline is used to extrapolate *virtual intersection depths* (i.e.  $z$  co-ordinate of intersection of an eye ray with origin  $(x, y)$ ) maintaining continuity with intersection points on Focus region boundary. This fills the *Context region intersection buffer*. We discard the rays that do not intersect the model bounding box. The depth buffer comprises the  $z$  co-ordinates (or depths) of all intersection points.

The second render pass is carried out multiple times as part of each render estimate of the MCRT rendering pipeline. In an orthographic projection, (as is commonly used for medical visualization), as eye rays are parallel to the Z axis, an eye ray with origin  $(x_e, y_e, 0)$  is clipped by moving it's origin to a point  $P_L(x_e, y_e, z_e)$  (see Fig. 2b) such that  $z_e = Z_E$ , where  $Z_E$  is the corresponding depth value for that eye ray. In MCRT, shading at any point  $S$  involves computation of both the direct and the indirect (scatter) illumination [13]. Hence, light rays also need to be clipped appropriately for correct illumination. As shown in Fig. 2b, a light ray is clipped at the intersection point  $P_L(x_l, y_l, z_l)$  with the implicit clipping surface. This intersection point is computed on the fly using ray marching as a point along the light ray whose  $z_l$  co-ordinate (in view space) is closest to the corresponding clipping depth value,  $Z_L$ , in the depth buffer. To estimate the corresponding depth value for any point  $P(x, y, z)$  on the light path, the point is first mapped into screen space (using the projection matrix) to get continuous indices within the depth buffer. The depth value is then extracted using linear interpolation from the depth buffer with these continuous indices.


### 3.2 Computational Complexity

With the application of this technique to MCRT, the addition of the first clipping depth computation pass amounts to roughly one additional iteration, out of typically 100 iterations used to obtain a good image quality. Therefore, it comes at a low computational complexity. On a system (Win7, Intel i7 3.6 GHz dual core, 8 GB RAM, NVidia Quadro K2200) an extrapolated depth buffer for a viewing window of size  $512 \times 512$  is computed in 0.1 s for the dataset in Fig. 1d.

## 4 Simulation

To appreciate our proposed method and to visually validate it, we render a simulated model. The model is a volume of size  $512 \times 512 \times 512$  voxels. Its scalar values are the  $z$  indices, in the range  $[-255, 255]$ . The scalar values are chosen to spatially vary smoothly across the data (for simplicity along the  $z$  axis) to enable an understanding of the continuity of the clipping surface both spatially and in depth by examining the shape of the surface and the scalar values across it. The *focus* (segmentation mask) is a centered cuboid ROI of size  $255 \times 255 \times 391$  voxels. The voxel spacing is such that the model scales to a unit cube. The color transfer function maps  $-255$  to blue and  $255$  to red. We render this scene using a cut plane, cut-away and our method.

In the cut plane view (Fig. 3a), both Focus and Context get clipped. In the cut-away view (Fig. 3b), there is a clear depth mismatch between the Focus and



**Fig. 3.** Visualization of the unit cube model with a cuboid focus within it, using (a) Cut plane (b) Cut-away (c) our method. (d) Computed clipping surface (rotated by  $90^\circ$  for visual appreciation). Color bar indicates depth values in mm. The continuity of the surface with the object boundary and the consistency of scalar values across it indicates that the surface is smooth with continuity both spatially and in depth from the focus onto the background.

Context, which results in a perceptual distortion. In addition, the focus is poorly lit due to shadows being cast by the context. Contrast this with the visualization using our method (Fig. 3c) where the entire Focus region is made visible while keeping depth continuity with the surrounding context.

Figure 3d shows the clipping surface that was computed, (rotated by  $90^\circ$  as indicated by the axes legend) for purposes of visualization. Note that, in our actual rendering pipeline, we do not explicitly compute the clipping surface (this is implicitly computed via the depth buffer as described in the previous section).

## 5 Conclusions

Advances in visualization enable better appreciation of the extent of injury/disease and its juxtaposition with surrounding anatomy. We introduce the novel idea of an on the fly computed Focus specific clipping surface.

Although we use it in the context of MCRT, these techniques are applicable to Direct Volume Rendering. We differentiate our work from other F+C works in four ways: (a) With our method, the structures of interest are rendered unoccluded and un-modified, which is essential for diagnostic interpretation, (b) The Focus transitions to the Context seamlessly by maintaining continuity of depth between the Focus and Context regions, thereby aiding interpretation, (c) There is no user interaction required to view the Focus and (d) The Focus is rendered with the same optical properties as the Context. We do not distort the visualization space or use tagged rendering and do not propagate errors in the segmentation to the visualization.

We believe that this work will change the way clinical applications display volumetric views. With the increasing adoption of intelligence in clinical applications, that automatically compute semantic information, we envisage these applications incorporating *organ presets*, similar to *transfer function presets*. We envisage uses of this technique in fetal face visualization from obstetric ultrasound scans and in visualization for surgical planning and tumor resection.

## References

1. Card, S.K., Mackinlay, J.D., Shneiderman, B. (eds.): Readings in Information Visualization: Using Vision to Think. Morgan Kaufmann, Burlington (1999)
2. Ware, C.: Information Visualization: Perception for Design. Kaufmann, Pittsburgh (2004)
3. Li, W., Ritter, L., Agrawala, M., Curless, B., Salesin, D.: Interactive cutaway illustrations of complex 3D models. ACM Trans. Graph. **26**(3) (2007)
4. Viola, I.: Importance-driven feature enhancement in volume visualization. IEEE Trans. Vis. Comput. Graph. **11**(4), 408–418 (2005)
5. Weiskopf, D., et al.: Interactive clipping techniques for texture-based volume visualization and volume shading. IEEE Trans. Vis. Comput. Graph. **9**(3), 298–312 (2003)
6. Wimmer, F.: Focus and context visualization for medical augmented reality. Chair for Computer Aided Medical Procedures, Technical University Munich, thesis (2007)
7. GoKnee3D RSNA 2017. [www.siemens.com/press/en/pressrelease/2017/healthineers/pr2017110085hcen.htm](http://www.siemens.com/press/en/pressrelease/2017/healthineers/pr2017110085hcen.htm)
8. Figueroa, D.: Knee chondral lesions: incidence and correlation between arthroscopic and magnetic resonance findings. Arthroscopy **23**(3), 312–315 (2007)
9. Hudelmaier, M., et al.: Age-related changes in the morphology and deformational behavior of knee joint cartilage. Arthritis Rheum. **44**(11), 2556–2561 (2001)
10. Comaniciu, D., et al.: Shaping the future through innovations: from medical imaging to precision medicine. Med. Image Anal. **33**, 19–26 (2016)
11. Raj, A., Vishwanathan, S., Ajani, B., Krishnan, K., Agarwal, H.: Automatic knee cartilage segmentation using fully volumetric convolutional neural networks for evaluation of osteoarthritis. In: IEEE International Symposium on Biomedical Imaging (2018)
12. Bookstein, F.L.: Principal warps: thin-plate splines and the decomposition of deformations. IEEE Trans. Pattern Anal. Mach. Intell. **11**(6), 567–585 (1989)
13. Kroes, T., Post, F.H., Botha, C.P.: Exposure render: an interactive photo-realistic volume rendering framework. PLOS one **7**(7), 1–10 (2012)