Real-time Surface Analysis and Tagged Material Cleansing for Virtual Colonoscopy

Christoph Russ[1,2], Christoph Kubisch[2], Feng Qiu[1], Wei Hong[1], Patric Ljung[1]

[1]Siemens Corporate Research, Princeton, NJ, USA
[2]Department of Simulation und Graphics, Otto-von-Guericke University Magdeburg, Germany

Abstract

Virtual Colonoscopy is an important procedure for screening and detecting colorectal cancer. It increases patient comfort and reduces risks compared to optical colonoscopy. Oral contrast is used to emphasize the soft tissue border and avoid the need for physical colon cleansing. In order to ensure a reliable diagnosis, it is currently necessary to remove the fecal tagging in a time consuming pre-processing step. As the result can include artifacts and may effect polyp size, this paper proposes a novel technique that allows realistic visualization of the surface boundary based on unmodified CT images. A combined iso-surface reconstruction and direct volume rendering approach is developed to handle partial volume artifacts efficiently and allow on-the-fly surface reconstruction. The algorithm supports real-time analysis of detected surfaces and can differentiate material transitions between air, soft tissue and fluid. The surface-based rendering furthermore allows photo-realistic visualization through screen space shading to support procedure planning and interactive training.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing and texture

1. Introduction

Volume rendering techniques are of high interest to medical imaging. Approximately 70 million computed tomography (CT) scans are performed annually [dGMK*09], which underlines the use of volumetric imaging in medical applications and creates a need for advanced rendering methodologies. Due to the amount of data and complexity of visible structures, research in this area is generally focusing on new visualization techniques to allow fast detection of diseases and enhancement of areas of interest. State-of-art medical procedures use volumetric datasets to analyze and visualize patients inner organs. An important application in this field is virtual endoscopy as described by Hong *et al.* [HMK*97]. Instead of performing an actual intervention, a CT scan allows non-invasive diagnosis. Even though we apply our research to the lower abdomen, the proposed rendering methods in this paper can also be used for other applications involving virtual endoscopy and organ visualization.

Patient preparation for optical colonoscopy involves liquid diets and physical colon cleansing to ensure a clear view during the procedure. Insufflation of the organ is required to examine the extended surface and find polyps, which are abnormal and potentially cancerous tissue growths. A computer-generated, 3D flythrough to find such deformations can prevent the traditional screening procedure and is considered more comfortable for patients and less difficult for physicians. However, virtual endoscopy also involves disadvantages. Besides the increased amount of radiation, only a traditional colonoscopy can unveil all surface details and allow treatment of diseases. Nevertheless, for the detection of polyps larger than 5 mm, clinical studies indicate that sensitivity and specificity of virtual and real colonoscopy are similar [FNS*99].

Within virtual colonoscopy patient preparation and colon insufflation are also necessary. To ensure minimal personal discomfort caused by physical colon cleansing and enhance the visible tissue boundary, oral contrast agents are given to patients. The resulting images show a more significant border between air and soft tissue. Fluid and remains of stool inside the colon are now represented by high intensity values,
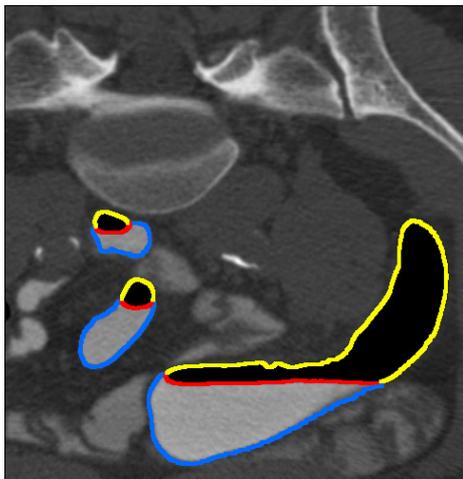
Figure 1: *CT image taken for virtual colonoscopy. Red: air-fluid transition, which is misclassified as a soft tissue boundary due to the partial volume effect. Yellow: border between air and soft tissue. Blue: fluid to soft tissue boundary.*

leading to glaring artifacts and the displacement of the actual boundary. Figure 1 highlights the areas of interest to a physician. Polyps could be located anywhere along the yellow and blue transition areas. The aim of digital bowel cleansing is to identify and remove the tagged material.

As the definition of an accurate surface boundary is important for the detection of small geometric deformations along this border, we invent a novel method to reconstruct and visualize a complete inner colon surface at real-time. Our method allows the detection of potential artifact areas and highlights them implicitly. Initial segmentation and fluid removal is not required as we can visualize the surface without modifications to the data. Advanced shading methods allow high quality rendering results and visual enhancement of polyps, which supports the diagnostic process and can be used for training or procedure planning purposes.

2. Previous and Related Work

The concept of virtual colonoscopy and the digital exploration of CT scans for polyp detection has been studied intensively. Early investigations by Hong *et al.* [HMK*97] and Rubin *et al.* [RBA*96] show that volume visualization techniques can be used to identify the colon surface and allow the detection of geometrical deformations. Electronic colon cleansing has been studied intensively [WLL*06], [CLW*00] and commonly involves image segmentation procedures to detect or enhance the tissue boundary as shown by Näppi *et al.* [NY08]. Lakare *et al.* [LWSK00] introduces a technique that allows ray-based analysis of a volume data set to detect material boundaries. Their technique removes contrast-enhanced voxels as a pre-processing step based on segmentation and ray-profile analysis.

Partial volume effects arising from the discretization of the captured images can lead to falsely classified voxels and visualization artifacts. This effect has been analyzed by Serlie *et al.* [STF*03] to create a multi-dimensional transition model to improve polyp detection rates within virtual colonoscopy. Their approach concentrates on three-material transition areas and uses similar characteristics for the analysis as presented in Lum *et al.* [LM04], who uses 2D transfer functions (TFs) with data sample pairs along the gradient direction for advanced volume lighting.

An adaptive empty space leaping procedure for translucency rendering in colonography was implemented by Lee *et al.* [LLL*09]. It operates on voxel level and improves the visual quality and efficiency of volume ray-casting by alternating between empty space skipping and a refined rendering step. This approach is useful to display translucent areas, but does not handle partial volume effects or tagged materials in general. Scharsach *et al.* [SHN*06] presents a rendering procedure, which combines direct volume rendering with surface shading. However, their approach simply renders any material hidden underneath an occluder. They do not allow the definition or distinction between materials, which limits the control over the transparency of the various surface layers, which is required for the removal of tagged materials. Hadwiger *et al.* [HSS*05] proposes a technique to generate high quality iso-surfaces by using volume ray-casting and post-processing effects in screen space. Their algorithm is extended with additional screen-space effects in this paper to add more realism to the visualization. A similar technique was applied to endoscopic rendering by Krüger *et al.* [KKSP08]. However, they do not handle partial volume effects and can not detect contrast-enhanced fluid.

Our proposed rendering framework does not only reconstruct partial volume layers and detects contrast-enhanced areas correctly, but also executes in real-time without the need of initial segmentation. Thus, we eliminate tedious pre-processing steps and visualize the captured data without modifications.

3. Soft-tissue boundary Detection and Rendering

Due to its flexibility and speed, ray-casting is used for rendering. The core algorithm that allows the real-time identification and reconstruction of contrast-enhanced areas within the volumetric dataset, consists of four main processing steps:

1. An iso-surface is detected initially to represent the potential soft tissue border.
2. These position coordinates are then analyzed to differentiate between different material transitions and effectively find fluid boundaries.
3. Dependent on the estimated fluid probability, direct volume rendering with a simple 2D TF setup is used to overcome partial volume artifacts and reconstruct misclassified surface areas.

4. Finally, a second iso-surface defines the actual boundary between the fluid and soft tissue.

Assuming the camera position is located within the volume and more specifically within air or fluid, the computation of start and end positions for the ray-casting initialization is not necessary. The origin of every ray is set to the current view location, which can be inside fluid or air. The individual ray direction is determined by camera parameters. In this example a fisheye lens distortion setup is used for an authentic, wide angle view as known from real endoscopes. In the following section we will explain the surface generation in detail and then focus on the shading effects.

3.1. Potential Soft Tissue Detection

To allow the use of advanced surface rendering techniques (e.g. deferred shading, screen space filtering) the definition of surface coordinates is required. Every ray is traversed with a fixed sample distance and terminates as its intensity is larger than the soft tissue threshold. To avoid staircase artifacts, hitpoint refinement [Sch05] enables the generation of a smooth surface by using a bisectional search algorithm. The final ray position then defines the surface coordinate for every pixel. Experiments have shown that ten search steps create visual results with sufficient continuity.

However, this initial surface can not represent the soft tissue boundary accurately as a single threshold would classify soft tissue and fluid samples the same way. Even though another higher threshold could separate them, the partial volume border at the transition from air to fluid remains misclassified. Figure 2 shows an example histogram for CT scans of the lower abdomen with a 1D trapezoid TF that differentiates air, soft tissue and fluid. Values within the low intensity range refer to air, while the medium intensity values are mostly related to soft tissue. The small peak within the high intensity range is related to contrast agents and bones. A transition from air to contrast-enhanced fluid will therefore contain soft tissue samples from the intermediate intensity range, which leads to rendering artifacts. Visualizing these areas appropriately requires the detection of a probability with which every position belongs to the air-soft tissue boundary or the air-fluid boundary.

To generate high quality data values for every 3D sampling point, linear interpolation of the underlying data can lead to noise-related artifacts. We can generate smoothened data samples, by using a convolution-based implementation of tri-cubical filtering on the graphics processing unit (GPU) presented in [SH05]. This method furthermore allows on-the-fly calculation of high quality first and second order derivatives, which can significantly improve the surface visualization and allow enhancement of polyps through curvature. Surface normals have also been generated via the central difference of surrounding samples. The results are consistent as the gradient is only calculated along robust material boundaries.
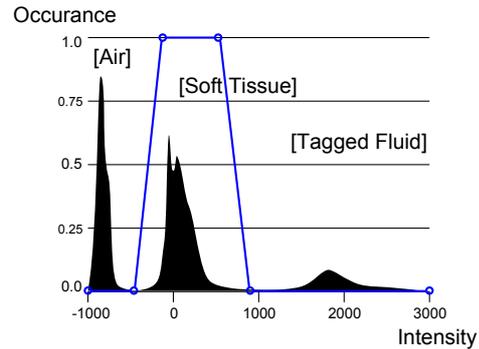

Figure 2: *The histogram of a lower abdomen CT scan shows how a one dimensional TF can separate air (low intensity) and contrast-enhanced fluid (high intensity). However, partial volume samples at the intersection of air to fluid are located within the soft tissue range (medium intensity).*

3.2. Air-Fluid Border Detection

The previously generated surface coordinates are located within the local area of a material transition. It is therefore possible to calculate a robust gradient for every location. Larger gradient magnitudes indicate higher contrast and therefore hint towards the air-fluid boundary. However, this measure alone is not considered reliable as it is influenced by noise and might not be differentiable to the actual soft tissue boundary.

To detect high intensity values underneath the surface, we use a maximum intensity projection of *n* samples along the normal direction. Even though this method can suffer from inaccuracies in the gradient computation, it provides an initial estimate of the air-fluid surface probability. For larger amounts of fluid, it can furthermore be assumed that the gradient of this transition is aligned with the gravity direction, which increases the likeliness of a partial volume layer in our algorithm. Even though noise effects this estimation, experiments have shown that the large intensity difference between soft-tissue and contrast-enhanced fluid allows a robust detection within four to six voxel layers.

Figure 3 shows an example of intensity and gradient magnitude profiles along a ray hitting an air to fluid boundary directly. To reconstruct thin surface layers, we combine the iso-surface rendering approach with an intermediate direct volume rendering step. Potentially misclassified surface areas and partial volume artifacts are handled through a 2D TF as described in the following section.

3.3. Partial Volume Effect Removal

To successfully remove partial volume artifacts, which lead to undesirable surfaces at air to fluid transitions, we use a direct volume rendering approach. A simple 2D TF (Figure 5), using the intensity and gradient magnitude as input, is applied for an adaptive number of samples into the volume.
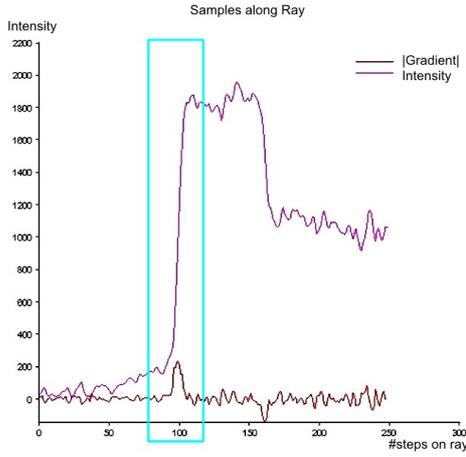
Figure 3: *This intensity profile relates to sample values along a gradient at the air-fluid boundary. The transition from low to high intensity is characterized through a large gradient magnitude that increases and decreases consistently.*
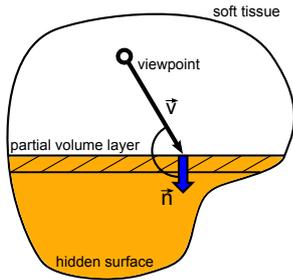

Figure 4: *After the iso-surface detection, an air-fluid transition probability p is found by searching for high intensity values along the local normal N. The partial volume reconstruction and additional iso-surface definition is then executed for areas with $p > 0$.*

The actual ray direction is not changed to avoid distortion, while the number of samples s is weighted proportional to the angle between normal $\mathbf{N}$ and view direction $\mathbf{D}$:

$$s = \lfloor k \times (\mathbf{N} \cdot \mathbf{D}) \rfloor ; \{k \in \mathbb{Z}, k = const\} \qquad (1)$$

It is necessary to adjust the number of steps through the partial volume layer to ensure the complete coverage of the this area as shown in Figure 4. The generation of a well defined surface with direct volume rendering requires high frequencies in the TF. To increase the number of TF samples, but prevent expensive oversampling of the 3D data, we introduce a novel approach based on linear interpolation between two sucessive tri-cubically filtered points along the ray (Figure 5). The intermediate values are calculated without the need of additional tricubical interpolation, which reduces the total number of all samples through the partial vol-
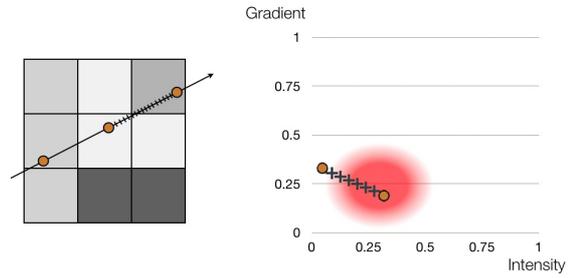

Figure 5: *Linear interpolation between two successive data samples can remove undersampling artifacts in the TF. Left: Large points show tricubical sampling positions. A fixed number of subsamples shown as black lines are calculated through linear interpolation between those. Right: Corresponding lookup in 2D TF to account for high frequencies.*
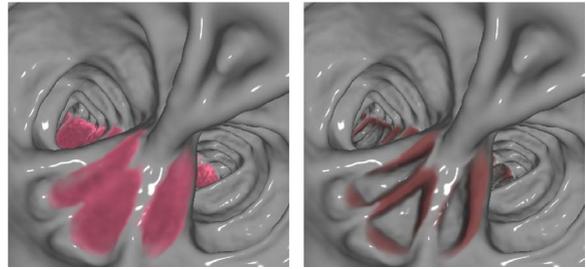

Figure 6: *Left: Areas colored in bright red are likely to be fluid transitions. Right: Brown areas show the occurred surface reconstruction, while all gray areas are defined by surface coordinates.*

ume area. This technique is comparable with pre-integrated TFs [EKE01], but does not require additional memory and works well for 1D and 2D cases. It is furthermore less complex than local peak detection [KHW*09]. By using a parametric TF implicitly within the shader, we can even prevent expensive texture lookups. The 2D TF itself suppresses high gradient samples and renders medium intensity values. It specifically does not include the contrast-enhanced high densities and large gradient values. This procedure, to overcome partial volume layers, is only applied to the previously identified, potential fluid surface areas and allows the implicit reconstruction of soft tissue values that have been falsely classified as shown in Figure 6.

The ray either terminates after an adaptively defined number of steps s (Equation 1) or after full opacity is reached. In case the sampling position actually traversed s times along the ray and no surface has been rendered, a secondary iso-surface search step is initiated, which starts at the last sampling position. Figure 7 describes how the partial volume layer is passed and the different rendering techniques have been combined into a hybrid approach.
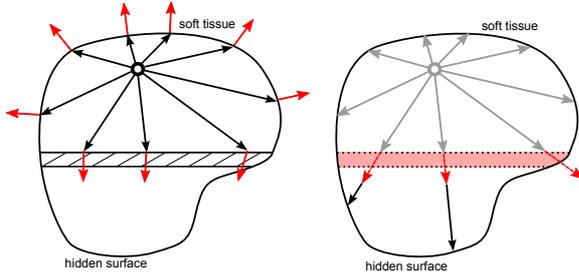
Figure 7: *As a first step all visible iso-surface intersections are defined and adjusted with hitpoint refinement. The second stage (red arrows) show the surface analysis along the gradients. The light red area has been reconstructed and the final iso-surface is defined as an additional surface search.*

Figure 8: *a: Detected high intensity values create a mask for further computation. b: The reconstructed boundary is masked here. c: Initial iso surface rendering results. d: Depth map of the completed fluid removal algorithm. e: Surface normals. f: Position coordinates for every fragment.*

3.4. Final Surface Composition

To define a continuous colon wall the initial iso-surface algorithm is applied to previously identified fluid areas. Every individual ray is traversed from its current sampling position after passing the partial volume layer. A second threshold is required to find the boundary between high intensity fluid values and the medium intensity soft tissue range. Threshold comparison and the calculated gradient have to be inverted as we are now looking at the fluid to soft tissue transition area. The presented surface detection algorithm requires final composition of the generated locations and gradients. The results are interpolated to ensure smooth transitions based on the following measures.

1. The secondary iso-surface between fluid and soft tissue is combined with the direct volume rendering results via interpolation with the final opacity values for every ray after *s* number of steps (Equation 1).
2. The result is then combined with the rendered color and position values of the initial iso-surface detection by interpolating with the estimated fluid probability.
3. As the surface positions generated through the reconstruction process might include visual artifacts, their location is stored in a 2D mask (Figure 8), which allows the individual coloration of these areas as part of the surface shading.

4. Photo-realistic Visualization Components

Besides the proposed rendering algorithm for colon cleansing, we also improve the surface rendering by creating a photo-realistic on-the-fly visualization. The generated image buffers (Figure 8) are used for deferred shading as an efficient visualization method for polyp detection. As part of our concept, we add screen space ambient occlusion to improve the visibility of polyps (Figure 9), while texturing and lighting effects add more realistic rendering features. This can be especially helpful for training of physicians and pre-procedural planning.
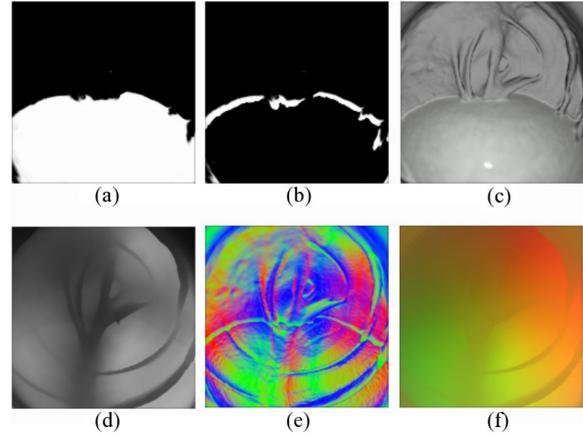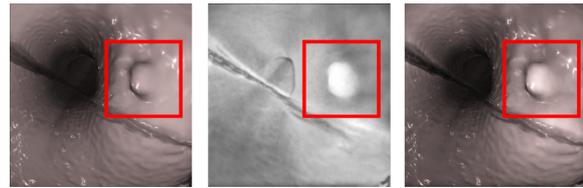

Figure 9: *Left: Basic Phong shading of the colon surface. Center: Ambient occlusion generated from depth map in screen space. Right: Screen space ambient occlusion applied to final compositing.*

4.1. Shading and Texturing

The diffuse light reflection is calculated after the Blinn-Phong shading model. Specular reflections in real colonoscopies however, appear more complex. A thin mucosa layer covers the surface and introduces distorted and slightly arbitrary light reflections. A similar effect can be created by distorting the generated normals with Perlin noise.

To allow a photo-realistic visualization of the colon without an existing texture parameterization, two different techniques have been combined to include surface details. A procedural formula to generate organ like structures from Perlin noise was implemented for a smooth and continuous color contribution throughout the colon. As the surface location is known for every pixel, a texture can be projected onto each axis and blended together based on the tri-planar mappings depended on the local normal [KKSP08]. This technique allows us to add blood vessels to the visualization, which are difficult to generate procedurally.
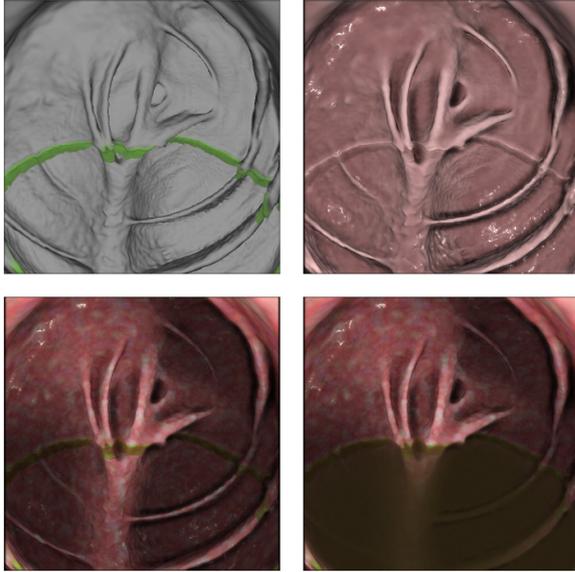
Figure 10: *Top left: A simple diffuse shading of the extracted iso surfaces. Top right: We apply screen space ambient occlusion. Bottom left: Adding detail textures and light attenuation. Bottom right: Fluid can be visualized transparently.*

Figure 11: *Real-time, high quality curvature calculation allows on-the-fly enhancement of possible polyp structures as an example extension of our surface extraction algorithm.*

4.2. Shadows, Occlusion and Translucency

To avoid additional render passes for shadow calculation we can employ the fact that the only light source within colonoscopies is attached to the actual camera. The hereby created shadow appears as a hallow around the geometric highlights of the organ. Screen space occlusion testing [Mit07] is a fast method to estimate the ambient occlusion of arbitrary complex scenes in 2D. It requires a depth map as input and allows the fast calculation of a global illumination estimation. By using this procedure within virtual colonoscopy, we gain multiple benefits. As polyps are naturally less occluded objects, their appearance is highlighted and their visibility enhanced. The lighting of the scene appears much more complex and gains significantly in depth. Finally, an alternative parameterization of the algorithm allows to enhance ambient occlusion, caused by depth differences, and generate halo shadow areas around depth discontinuities, which imitate real shadows very efficiently. Figure 10 shows the intermediate rendering results after application of this technique.

The depth difference between the initial iso-surface extraction and the final surface definition is equivalent to the local depth of detected fluid layers. It can therefore be used to visualize the present fluid areas as shown in Figure 10. As the depth of fluid is known through the detection of primary and secondary iso-surfaces, a diffuse blending can simulate its translucent appearance.
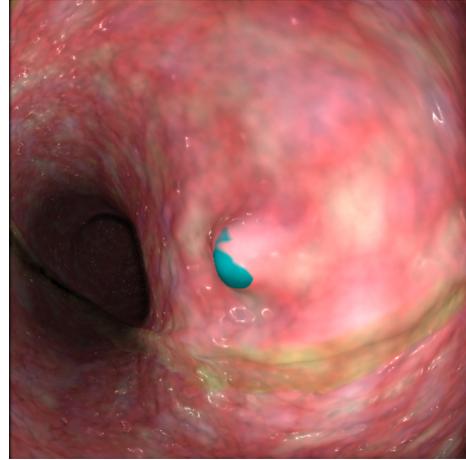
5. Additional Features for Virtual Colonoscopy

The presented algorithm for real-time tagged material removal was developed within the focus of a medical application. Our prototype completes the proposed concept for virtual colonoscopy and shows the useful implementation of this work. As a central aspect we focus on polyp localization and enhancement.

5.1. Curvature Generation and Polyp Detection

Using a convolution-based filtering method [SH05], we can generate the curvature for the extracted surface coordinates at runtime on the GPU. As shown by Zhao *et al.* [ZBB*06] the curvature is a useful measure for polyp identification. As we are focusing on real-time solutions, curvature information can be used to enhance the polyp appearance. A 2D texture mapping can then identify polyp caps using the maximum and minimum principle curvatures as input. Furthermore, our fluid detection algorithm can support polyp identification as they cover a slightly different intensity range compared to soft tissue as shown in Hong *et al.* [HQK06] and Näppi *et al.* [NY08]. The enhancement of identified polyps can be seen in Figure 11.

5.2. Intensity Rendering

For manual surface analysis purposes, an interactive plane can be rendered within the three dimensional endoscopic view, which displays the original intensity values of the CT scan data. These values can be compared directly to the visualized surface. The plane can be rotated interactively and moves with the camera. A transparent border allows depth impressions, while the center area cuts through the surface and displays the actual intensity values as the result of an applied 1D TF. Figure 12 shows an example.
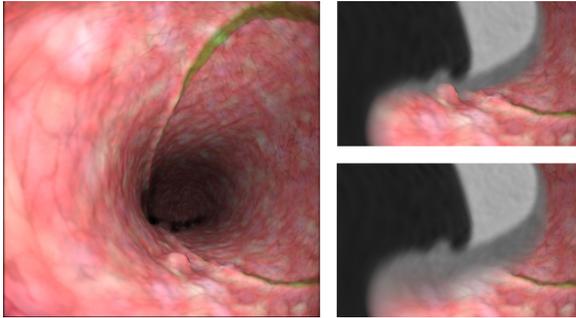
Figure 12: *A flexible, depth sensitive MPR overlay can help to analyze the original surface structure and identify polyps. It also shows the original fluid boundary.*

6. Results

We use OpenGL and multiple GLSL shaders for the implementation of the proposed surface extraction algorithm and the additional visualization techniques. The different effects can be toggled comfortably through a graphical user interface. The eXtensible Imaging Platform (XIP) [XIP] provided the underlying visualization framework and handling of the scenegraph. The test setup for our implementation consists of a 2.8 Ghz Intel Core2Duo Processor, 2GB RAM and a single nVidia GeForce 8800 GTX. The 17 test datasets provide prone and supine scans at image resolutions of 512×512 pixel with 400 to 600 individual slices. They are encoded with 12 bit accuracy and were provided by the National Biomedical Imaging Archive at the National Cancer Institute. For evaluation, all datasets include optical and virtual colonoscopy reports with polyp size and location, which are used for the confirmation of identified polyps. Every scan contains tagged fluid as well as stool and at least one polyp. Even though the overall performance depends mainly on the initial iso-surface detection, an average framerate of more than 30 fps could be maintained at a screen resolution of 512×512 pixel within the experiments.

As clearly visible on the left in Figure 13, the partial volume layer between fluid and air has been removed successfully. However, due to the remaining high intensity values inside the dataset, the normals are distorted along the three-material transition area. This effect was expected and could partially be reverted by recalculating the gradient based on the extracted surface values in screen space. However the effects are not severe and are already corrected partially through the surface reconstruction step of our algorithm. It is therefore not necessary to adjust the gradient at these locations.

Compared to other colon cleansing methods, our algorithm does not only remove high intensity data values, but also highlights surface areas that might contain artifacts (Figure 14). This process can be seen as an important improvement to previous techniques and simplifies the image interpretation. Furthermore, the polyp visualization could
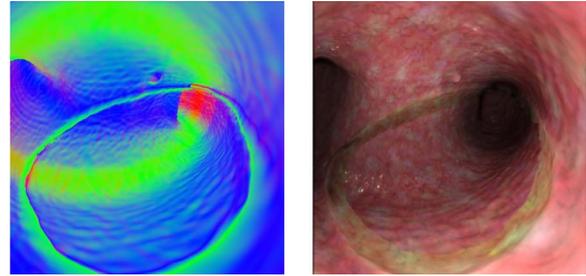

Figure 13: *Left: This normal visualization shows the distortion of gradients due to mixture of low, medium and high intensity values in volumetric dataset. Right: Highlighting of fluid border allows to hint potential artifact areas.*
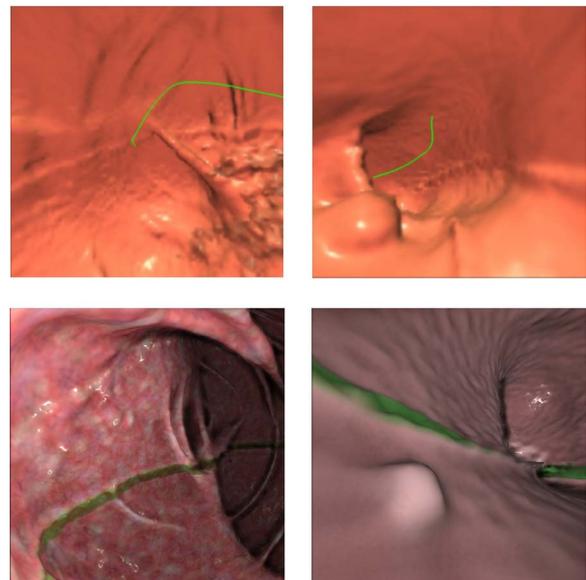

Figure 14: *Top: Existing solutions based on [LWSK00] show artifacts along the fluid boundary (Images courtesy of Viatronix, Inc. and Walter Reed Army Medical Center). Bottom: Our approach can highlight areas containing artifacts and uses advanced shading techniques for photo-realistic rendering and polyp enhancement through ambient occlusion.*

be enhanced through occlusion tests, which supports their identification. Advanced texture and lighting effects create a photo-realistic visualization, which can be used for training lessons or patient specific procedure planning. The surface reconstruction allows partial recovery of three-material transition areas, but could potentially be improved as artifacts remain visible.

7. Conclusion and Future Work

The presented rendering pipeline allows the direct visualization of fecal tagged colonoscopy data in real-time without the need of pre-processing and colon cleansing. We in-

troduced a virtual colonoscopy application that allows the photo-realistic visualization of the colonic lumen and provides various interactive tools for the viewer. The rendering procedure is applicable to different medical applications related to virtual endoscopy that require the removal of tagged material areas.

A novel algorithm for surface analysis and automated material interpretation generates high quality iso-surface coordinates, gradients and curvature values at runtime. Additional post-processing steps are used for advanced rendering techniques, adding screen space ambient occlusion and shadow silhouettes at low computational costs.

Our implementation does not require to modify the tagged CT input data. Thus, we can directly display the correlation between the presence of high intensity data values and possible distortions within the rendering. Critical areas around potential partial volume layers and areas with three-material transitions can be identified and highlighted to visualize increased artifact probability. We would like to perform a clinical study to analyze the use of the implemented procedures. Additional feedback from a broader range of users would help to measure the success of our surface generation and the acceptance of advanced visualization effects within virtual colonoscopy. It furthermore remains a challenging task to actually identify all polyps in real-time. The current enhancement through curvature values can be improved by taking additional geometric parameters into account. So far, the implemented application provides an appropriate platform for the interactive exploration of fecal tagged CT data and improves the current techniques in speed and visual quality.

References

[CLW*00] CHEN D., LIANG Z., WAX M. R., LI L. L. B., KAUFMAN A. E.: A novel approach to extract colon lumen from ct images for virtual colonoscopy. *IEEE Transactions on Medical Imaging 19*, 12 (2000), 1220–1226. 2

[dGMK*09] DE GONZÁLEZ A. B., MAHESH M., KIM K.-P., BHARGAVAN M., LEWIS R., METTLER F., LAND C.: Projected cancer risks from computed tomographic scans performed in the united states in 2007. *Archives of internal medicine 169*, 22 (2009), 2071–2077. 1

[EKE01] ENGEL K., KRAUS M., ERTL T.: High-quality pre-integrated volume rendering using hardware-accelerated pixel shading. In *HWWS '01: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware* (New York, NY, USA, 2001), ACM, pp. 9–16. 4

[FNS*99] FENLON H. M., NUNES D. P., SCHROY P. C., BARISH M. A., CLARKE P. D., FERRUCCI J. T.: A comparison of virtual and conventional colonoscopy for the detection of colorectal polyps. *N Engl J Med 341*, 20 (1999), 1496–503. 1

[HMK*97] HONG L., MURAKI S., KAUFMAN A. E., BARTZ D., HE T.: Virtual voyage: interactive navigation in the human colon. In *SIGGRAPH* (1997), pp. 27–34. 1, 2

[HQK06] HONG W., QIU F., KAUFMAN A.: A pipeline for computer aided polyp detection. *IEEE Transactions on Visualization and Computer Graphics 12* (2006), 861–868. 6

[HSS*05] HADWIGER M., SIGG C., SCHARSACH H., BHLER K., GROSS M.: Real-time ray-casting and advanced shading of discrete isosurfaces. *Computer Graphics Forum 24*, 3 (Sep. 2005), 303–312. 2

[KHW*09] KNOLL A., HIJAZI Y., WESTERTEIGER R., SCHOTT M., HANSEN C., HAGEN H.: Volume ray casting with peak finding and differential sampling. *IEEE Transactions on Visualization and Computer Graphics 15* (2009), 1571–1578. 4

[KKSP08] KRÜGER A., KUBISCH C., STRAUSS G., PREIM B.: Sinus Endoscopy - Application of Advanced GPU Volume Rendering for Virtual Endoscopy. *IEEE Transactions on Visualization and Computer Graphics 14*, 6 (November/Dezember 2008), 1491–1498. 2, 5

[LLL*09] LEE T.-H., LEE J., LEE H., KYE H., SHIN Y. G., KIM S. H.: Fast perspective volume ray casting method using GPU-based acceleration techniques for translucency rendering in 3D endoluminal CT colonography. *Computers in biology and medicine 39*, 8 (Aug. 2009), 657–666. 2

[LM04] LUM E. B., MA K.-L.: Lighting transfer functions using gradient aligned sampling. In *VIS '04: Proceedings of the conference on Visualization '04* (Washington, DC, USA, 2004), IEEE Computer Society, pp. 289–296. 2

[LWSK00] LAKARE S., WAN M., SATO M., KAUFMAN A.: 3D digital cleansing using segmentation rays. In *VIS '00: Proceedings of the conference on Visualization '00* (Los Alamitos, CA, USA, 2000), IEEE Computer Society Press, pp. 37–44. 2, 7

[Mit07] MITTRING M.: Finding next gen: Cryengine 2. In *SIGGRAPH '07: ACM SIGGRAPH 2007 courses* (New York, NY, USA, 2007), ACM, pp. 97–121. 6

[NY08] NÄPPI J., YOSHIDA H.: Adaptive correction of the pseudo-enhancement of CT attenuation for fecal-tagging CT colonography. *Medical Image Analysis 12*, 4 (2008), 413 – 426. 2, 6

[RBA*96] RUBIN G., BEULLIEU C., ANNGIRO V., RINGLE H., NORBASH A., FELLER J., DAKE M., JEFFEREY R., NAPLE S.: Perspective volume rendering of CT and MR images: Applications for endoscopic imaging. *Radiology 199* (1996), 321–330. 2

[Sch05] SCHARSACH H.: Advanced gpu raycasting. In *Proceedings of the 9th Central European Seminar on Computer Graphics* (May 2005). 3

[SH05] SIGG C., HADWIGER M.: Fast third-order texture filtering. Addison-Wesley Professional, 2005. 3, 6

[SHN*06] SCHARSACH H., HADWIGER M., NEUBAUER A., WOLFSBERGER S., BÜHLER K.: Perspective isosurface and direct volume rendering for virtual endoscopy applications. In *EuroVis* (2006), pp. 315–322. 2

[STF*03] SERLIE I., TRUYEN R., FLORIE J., POST F. H., VAN VLIET L. J., VOS F.: Computed cleansing for virtual colonoscopy using a three-material transition model. In *MICCAI 2003, 6th Intern. Conf., Montréal, Proceedings, Part II* (2003), Springer, pp. 175–183. 2

[WLL*06] WANG Z., LIANG Z., LI X., LI L., LI B., EREMINA D., LU H.: An improved electronic colon cleansing method for detection of colonic polyps by virtual colonoscopy. *IEEE Transactions on Biomedical Engineering 53*, 8 (2006), 1635–1646. 2

[XIP] eXtensible Imaging Platform. http://openxip.org/. 7

[ZBB*06] ZHAO L., BOTHA C., BESCOS J., TRUYEN R., VOS F., POST F.: Lines of curvature for polyp detection in virtual colonoscopy. *IEEE Transactions on Visualization and Computer Graphics 12* (2006), 885–892. 6