



FORENSIC CASE ANALYSIS: FROM 3D IMAGING TO INTERACTIVE VISUALIZATION

Forensic case analysis: from 3D imaging to interactive
visualization

Martin Urschler, Alexander Bornik, Eva Scheurer, Kathrin Yen, Horst
Bischof, Dieter Schmalstieg

*Inst. for Computer Graphics and Vision
Graz University of Technology, Austria*

Technical Report
ICG-TR-xxx
Graz, March 1, 2011

Abstract

Recently forensic investigators have shown an increasing interest in the benefits of 3D medical imaging modalities to improve forensic case analysis by enabling observation of body interior. While Computed Tomography (CT) allows excellent depiction of bony structures, Magnetic Resonance Imaging (MRI) provides good contrast in soft tissue structures and involves no exposure to harmful radiation. We present an integrated, interactive framework to aid forensic investigators in case analysis and courtroom presentation tasks. Our framework makes extensive use of advanced computer graphics and computer vision techniques and is implemented entirely on the GPU. It enables a forensic analysis work-flow from the raw data, over segmentation techniques to enhance structures of interest, to the presentation via focus and context enabled multi volume rendering techniques, resulting in interactive scenes, videos, or images suitable for the courtroom. Two selected case studies demonstrate the practical applicability of our framework on real forensic cases.

Keywords: *visualization of forensic images, interactive 3D segmentation*

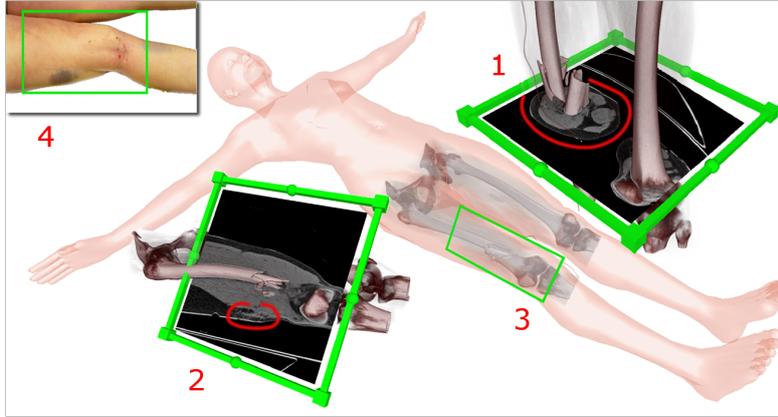


Figure 1: Example of a forensic case analysis and courtroom presentation. After a motor vehicle accident, a CT scan was performed. We use a volume rendering to depict bone structures and also present 2D slices showing the original grayvalues (1,2). The red circle regions highlight the fractured bone (1) and a hematoma (2). Further, we place the volumetric scan into a reference body model to provide a context for the data (3). Finally a photograph from an external examination is placed into the presentation (4).

1 Introduction

The analysis of forensic cases heavily relies on digital information for documentation purposes, especially to reconstruct accident and crime scenes and to present forensic findings in court [1]. While forensic investigations have made use of photography for many decades now, recent years have brought an increasing interest in 3D medical imaging modalities like computed tomography (CT) and magnetic resonance imaging (MRI) [2, 3]. The obvious benefit of CT and MRI devices is the possibility of imaging the whole 3D anatomy from the inside. This enables an investigator to retrieve additional forensic information, which is often invisible from the outside.

While the forensic analysis and presentation of 2D photographs is an established method despite potential difficulties in interpretation, working with 3D data requires more sophisticated tools compared to photos. Direct presentation of the 3D data in the form of stacks of 2D images (i. e., slices) is not feasible in court, since the amount of slices is usually very large, and even more importantly, it requires a trained radiologist to interpret 2D slices from volumetric data. To be useful for presentation in the courtroom, the complexity of 3D volumetric data has to be reduced, and the forensic findings have to be highlighted after preparation by a forensic expert.

3D volume rendering [4] addresses the need for appealing visualization

of the CT/MRI scans. However, in clinical forensics, the scans are usually restricted to relatively small regions of the body, to limit the victim’s exposure to harmful radiation in CT or to limit the scanning time in MRI. Moreover, the need for high spatial resolution opposes full-body scans, as injuries of forensic interest are often locally restricted. Nonetheless, injuries can be distributed over several body regions, and the investigator ends up with a number of locally restricted scans. These partial scans must be related to a full-body reference model to be useful for in-court presentation.

In this work, we present a software framework dedicated to the analysis and presentation of forensic cases involving 3D volumetric data from medical imaging modalities. This software framework is a work in progress for the Ludwig-Boltzmann-Institute for Clinical-Forensic Imaging at Graz University of Medicine. The institute pursues an inter-disciplinary research agenda that brings together experts from forensic medicine, radiology, jurisdiction, physics, computer graphics, and computer vision. The institute’s purpose is to establish protocols and tools to implement the use of volumetric medical imaging modalities in forensic medicine for assisting forensic experts and judges in decision making. While also being able to provide forensic expertise on post-mortal cases, the institute’s focus lies on clinical forensic imaging of the living, where the problem of avoiding harmful radiation arises. Thus, especially soft-tissue diagnosis builds upon MRI as its preferred imaging modality.

We see our main technical contribution in providing an integrated and interactive framework that supports all necessary steps to progress from raw CT and MRI scans to presentations suitable for the courtroom. An important requirement is that all components work in such a way that the elevated storage requirements and computational complexity are hidden from the user. Forensic experts can use the tool freely to explore and manipulate the data with immediate feedback and without incurring frustrating waiting times, as if they were just processing 2D images from their digital camera. This high performance requirement can be successfully addressed by leveraging recent improvements in parallel processing on the GPU through languages such as CUDA, and with the help of relatively inexpensive multi-GPU workstations.

Through the use of a flexible scene abstraction, the framework also lends itself to the preparation of appealing illustrative visualizations suitable for the courtroom, in the form of images, videos or interactive demonstrations. In Fig. 1 we illustrate an example of a case prepared with our framework, where a person was overrun by a car resulting in a fractured femur and hematoma. Our work-flow enables us to visualize all of the involved data sources. In Section 4, we show further case studies that demonstrate the practical applicability of our proposed work-flow.

2 Requirements and system overview

Supporting complex medical imaging applications requires convergence of computer vision and computer graphics, so that the solutions from these two disciplines can be put at the expert’s command without imposing unnatural and unproductive restrictions. Thus, we have designed a framework by analyzing the goals of forensic experts and deriving requirements from their work-flow. The framework is intended to assist the forensic expert in *visualization* of different sources of information, *segmentation* and highlighting of forensically relevant information and preparation of intuitive *presentations* for the courtroom. In the following, we summarize these requirements and approaches.

2.1 Visualization

The core visualization component must be flexible in its ability to work with different sources of data, e.g. 3D volumes, 3D geometry and 2D photographs. Furthermore, it must support working with multiple data sets concurrently. These needs have led to the development of a rendering engine that integrates volume rendering with surface geometry rendering in a way that is transparent to the user. It is based on GPU accelerated ray-casting [5]. This ray-casting engine allows the concurrent visualization of multiple volume or geometry data sets at interactive frame-rates, thus it very suitable for combining different sources of forensic information. See Appendix A for more information on the volume rendering approach.

2.2 Segmentation

To enhance presentations, findings of forensic relevance have to be extracted and clearly separated from the rest of the volumetric data. In computer vision, this process is referred to as segmentation. The analysis of forensic cases requires a generic, interactive 3D segmentation framework which is not designed for specific application cases, as is commonly the case in medical image segmentation, but allows flexibility in the structures to be extracted. Depending on a specific case, structures of forensic interest may include bones, muscle tissue, hematoma, vascular structures, and organs with injuries. We use an extensible segmentation mechanism in order to adapt to different degrees of difficulty of the segmentation task. Complexity is added by including increasing amounts of prior knowledge into the segmentation models. Like the visualization, the segmentation component relies on a highly parallel GPU implementation. See Appendix B for an in-depth treatment of this

module.

2.3 Presentation

In-court presentation should be performed in an intuitive manner to present forensic findings as well as their anatomical context in a way that is easy to understand for people without a radiological background by means of a reference body model. The presentation of different sources of information requires tools that align the different coordinate systems of 3D volumes, 2D photographs and reference model geometry into a common coordinate frame by means of registration algorithms. Besides registration, presentation also benefits from a focus & context visualization paradigm [6], where specific details of a forensic case can be highlighted while showing them in the context of the remaining data set and the reference model. This is easily possible due to our flexible, multi-volume visualization component. Animations to highlight relevant structures create an improved cue for forensic case interpretation, and can be presented in the form of videos or interactive demonstrations in the courtroom.

3 Work-flow

Our proposed system is inspired by the analysis and in-court presentation requirements of forensic experts and maps these requirements into a work-flow consisting of four stages (see Fig. 2). First, the *acquisition* stage is concerned with capturing data from a medical imaging modality; additionally photographs may be taken. Second, a *preprocessing* stage deals with simplifying the data representation and improving image quality if necessary. Third, the *3D forensic analysis* stage extracts additional information from the given data. Fourth, the final *3D forensic presentation* stage embeds all sources of data into a single visualization, using reference models if required. From this presentation tool, videos, still images or interactive demonstrations for the courtroom are created. Implementation details on the software components may be found in Appendix C. Note that both the analysis tool and presentation tool are based on the same core visualization component, described in Appendix A.

3.1 Image acquisition

Our work focuses on forensic investigations of the living as opposed to the concept of virtual autopsies, which is dealt with extensively in [3]. CT has

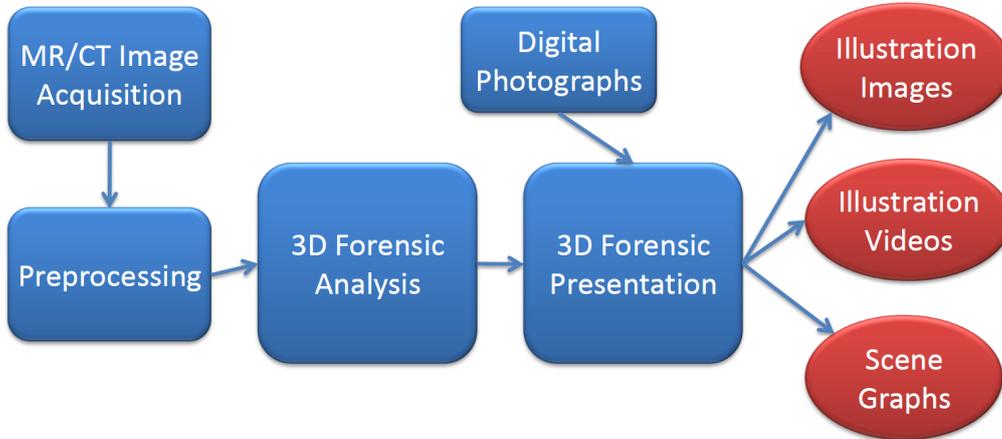


Figure 2: Overview of our forensic analysis and presentation work-flow. After image acquisition we perform a preprocessing step to enhance the 3D volume(s) if necessary. Next, our analysis step extracts forensically relevant structures from the 3D volume(s). Finally, the presentation step combines all available data, places it into a reference coordinate system and allows the output of images, videos or a scenegraph describing the combined scene. The output can be used for courtroom demonstrations.

the advantage of a very high resolution and is often used in examinations involving bone injuries. The main drawback of CT is the harmful ionizing x-ray radiation. MRI does not involve harmful radiation and has a wide range of applications in soft-tissue imaging, due to the large number of possible scanning protocols. However, this flexibility is also a disadvantage, as an experienced physicist is needed to help in designing protocols for a given question, e. g., protocols to depict hematoma. Volumetric datasets are frequently stored in the DICOM format. However, for our purposes, we convert DICOM to the simpler Analyze format, and represent it at runtime in a volume scene graph that is suitable for interactive manipulation.

3.2 Preprocessing

Depending on the quality of the acquired volumetric data sets, some preprocessing may become necessary. We perform a denoising step with the Rudin-Osher-Fatemi model [7] that implements a robust, edge-preserving reconstruction based on a total variation energy minimization. While this model is well-suited for a fast GPU implementation, denoising is executed as

an offline step for convenience. After denoising, an optional resampling using tricubic interpolation reduces the size of input images for efficiency reasons and to compensate for the anisotropic resolution of the acquisition (typically the in/plane resolution is higher than the slice thickness).

3.3 3D forensic analysis

Visualization of volumetric data requires a transfer functions, which maps intensity ranges from the medical imaging source (typically 12 bits) to color and opacity values. This mapping of intensities and intensity gradients (for lighting calculations) enables high-quality 3D volume renderings that are much better suited to be presented in court than the original stack of 2D grayscale images. Unfortunately, applying a transfer function alone makes it impossible to discriminate anatomical structures that lie in the same intensity range, or to emphasize structures of interest (e. g., hematoma, bones or injured organs) while retaining sufficient contextual structures of the visualized data.

For the extraction of such interesting structures, we require segmentation followed by multi-object visualization. The visualization is responsible for displaying original volume data, segmentation results, and additional user-provided information, e.g., arrows indicating the direction of an impact. Another benefit of segmentation is that it allows deriving quantitative indices like mass or volume of a structure. The segmentation uses either a standard 3D region growing approach or an energy minimization algorithm based on geodesic active contours (see Appendix B).

Unlike most established medical analysis software, our framework lets the user perform both segmentation and visualization concurrently and directly in the three-dimensional view. This integrated approach (Fig. 3) avoids manual switching between separate segmentation and visualization tools and thereby accelerates the workflow by providing immediate feedback on the segmentation. Thus, interactive change of parameter choices for the segmentation algorithm becomes feasible.

The GUI provides the user with a 3D view and optional 2D views (axial, coronal or sagittal) on the data. Interaction, e.g., specification of the seed regions, selection of regions of interest, and segmentation refinement is possible using painting tools. To define the structure that has to be segmented, the segmentation tool lets the expert paint foreground and background seed regions directly on the volumetric structure or an embedded cutting plane. For quick specification of regions of interest, we provide a space carving tool that allows pruning space by defining a screen-space region from which a geometric extrusion is constructed. The extruded volume is voxelized and used

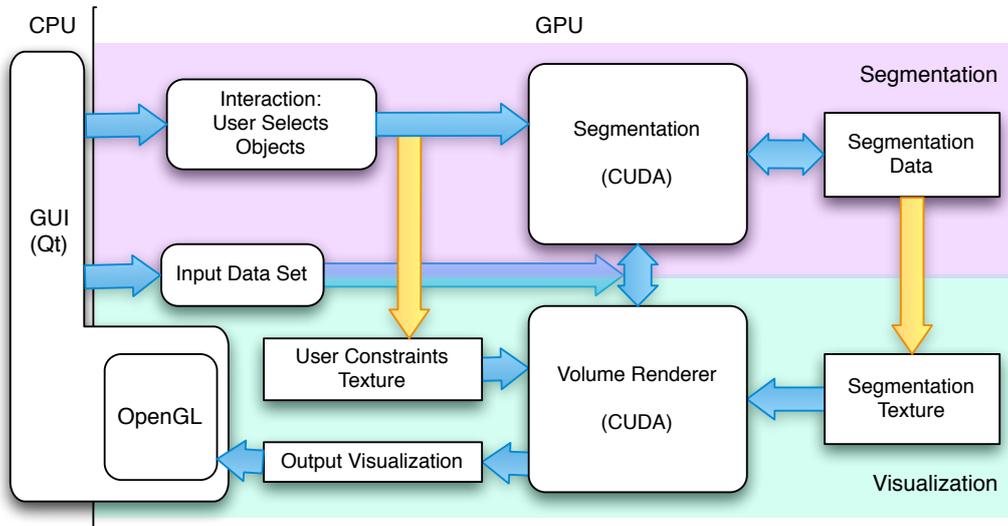


Figure 3: Unlike conventional approaches that separate segmentation and visualization into sequential tools, our integrated segmentation and visualization system provides immediate feedback of the visualization on intermediate or final segmentation results and enables more efficient interaction.

to restrict segmentation operations to a volumetric region of interest. Results from several space carving steps may be combined with Boolean operations to support more complex region of interest geometries.

3.4 3D forensic presentation

The 3D forensic presentation component uses the same core visualization module as the analysis stage. However, its objective is not to identify and highlight relevant forensic findings, but to to combine and arrange different sources of information in a still, animated or interactive illustration. All elements – volume data, segmentations, supplementary geometry like a reference manikin and photos – are arranged in a scene-graph which can be manipulated as needed with 3D direct manipulation widgets.

Elements embedded in the scene graph can be instanced multiple times with different parameters such as geometric transformation or transfer function. Moreover, the volume rendering can support a full set of Boolean operations on volumes. This makes it easy to create in-place focus & context techniques such as a magic lens that makes the skin transparent above a region of interest to reveal the interior.

If a volumetric data set encompasses a very limited portion of the body,

we can further register this data set to a generic reference manikin¹. This also serves as a means to focus on forensic data in the context of the whole body, an issue that is crucial for intuitive presentation.

The registration makes use of deliberately placed or anatomically derived markers, which are available on both the reference model and the data set. This provides a rough rigid registration, which is refined using the surface-based iterative closest point algorithm. After registration, the surfaces from the reference model and the outer surface of the 3D volumetric scan are aligned, and the structures of forensic interest are placed into the reference coordinate system.

4 Case studies illustrating the work-flow

4.1 Broken clavicle

In the first case, a thorax CT of the victim was available, containing the fractured right clavicular bone. The 20-year old male subject was involved in a motor vehicle accident (driver, frontal car crash) and had a fastened seatbelt. The heavy impact broke the right clavicular bone. To demonstrate the injuries, we decided to visualize both clavicular bones, the fractured and the unharmed one, in the context provided by the volumetric CT data. Note that in this case it is not necessary to place the CT scan into a reference body model for courtroom presentation, since the body context is clearly visible.

We denoised the CT input volume and cropped a portion of the data set, so that the body surface of the CT scan still gives a good indication of the overall location of the injury (see Fig. 4a). We had to downsample the data set from its original resolution to $256 \times 256 \times 256$ voxels for further processing.

First, we loaded the CT data into the 3D forensic analysis tool (see Fig. 4a) to interactively segment the clavicular bones. To segment the right clavicular bone, we quickly specified a region of interest around it (see Fig. 4b) with the help of space carving. This sets up the rough location for the following detail segmentation, which is very important to distinguish the bone from a close-by tubular structure with similar density.

We painted foreground (green spheres) and background (red spheres) seed regions to specify the structure that we wanted to segment (Fig. 4c). The result of the segmentation process (Fig. 4d), which is available with a short delay of 0.5-1sec, is further refined by removing unwanted structures. Finally, we store the refined segmentation (Fig. 4e) for later use. We repeat the same procedure for the unharmed left clavicular bone. After producing the

¹Provided by the Makehuman project, <http://www.makehuman.org>, Sep. 2010.

segmentations, we combine all of our information sources (volumetric data set and segmentation data sets) into a single scene.

Using the presentation tool, the scene can be arranged further. A spherical geometry object is placed into the main volume and configured with a different transfer function. While the main volume shows the skin as context, the spherical focus uses a low-opacity transfer function and reveals the fractured clavicular bone (Fig. 4f). By animating the sphere, the cue of motion further improves the presentation result (see accompanying video). For further examinations, we freely arrange the segmented clavicular bones (Fig. 4g). This can aid further investigations, such as the estimation of likely direction of the force that broke the bone, and in turn the chain of events that led to the accident.

4.2 Hematoma

The second case consists of two MRI volumes, showing a rough localization and a detailed depiction of a hematoma in the left gluteus after a sports accident. The first volume was acquired using a T1 weighting in a spin-echo sequence. It provides an overview of the left hip and upper thigh region (see Fig. 5a); it however does not show considerable contrast in the blood pool, and the hematoma is hard to be diagnosed in this image. The second volume is a proton density scan, again with a spin-echo sequence and fat saturation enabled to suppress the MR signal in the fatty tissue. This gives a good contrast to localize the fuzzy blood pool structure, since the blood pool remains unchanged by the fat saturation. This second scan (Fig. 5b) shows solely the left gluteus in higher resolution, so we refer to it as the detail scan.

The forensic interest in this case is to visualize the blood pool indicating a hematoma which may be invisible from the outside. Unfortunately the blood pool is a very local structure with a fuzzy appearance, which makes it hard to delineate. Radiologists only have few experiences with imaging of subcutaneous tissue lesions, as these usually have no clinical relevance. However, for forensic experts such lesions can give important clues for the analysis and forensic reconstruction of a case. Note that if only slices from the MRI detail scan are presented, the anatomy is very hard to interpret for non-radiologists. A further motivation for volumetric analysis and presentation of hematoma is the need for privacy, which may restrict showing photographs of injuries.

For the preparation of this case, we first performed a segmentation of the fuzzy blood pool structure (Fig. 5b, (1)) from the detail MRI scan. Since the blood pool shows a good contrast to the surrounding tissue, we use the

region growing segmentation algorithm for this task. The green structure in Fig. 5c shows the segmentation result.

For the presentation, we created a scene containing the segmentation result, the T1 weighted MR scan of the hip and upper thigh region, and our reference body model. The reference body model and the MR data sets were registered by a surface-based iterative closest point registration. Since there were no markers present in the data sets, we required a manual initialization of the registration. This visualization can be used as the basis for determination of force directions that led to a hematoma, or to investigate dependencies between internal and external injuries. Furthermore, it was possible to derive quantitative indices from this representation. For example, in this case the blood pool volume is 4.13 ml.

5 Conclusion and outlook

In this work, we have presented a prototype of a novel framework for forensic case analysis and presentation using volumetric 3D data from MR/CT imaging modalities. CT and MR imaging are increasingly used in forensic case analysis and reconstruction of the sequence of events. However, forensic expertise requires specific analysis different from the clinical needs, and the imaging data have to be translated into an easily understandable manner when being presented to non-medical experts in court. The possibilities of visualization techniques can profoundly support such analyses and in-court presentations and thereby improve the quality of forensic imaging expertise, which again is an important factor for legal certainty.

The analysis of forensic cases requires flexible tools to allow extraction of relevant structures and visualization of different data sources. We build upon an integrated system for segmentation and multi-volume visualization that performs these demanding tasks on consumer GPUs, thus profiting from their power and competitive pricing. We focus on the aspects of interactive, expert-driven case analysis with immediate feedback on analysis actions, and the notion of focus & context for the presentation of forensic findings. Our case studies successfully demonstrated our system.

We are currently in the process of extending our framework to a more generic analysis tool and investigate segmentation models that involve stronger prior knowledge. In this way, we intend to deal with forensic analysis tasks where our framework currently has limitations, such as the investigation of tubular structures and structures with strong shape constraints.

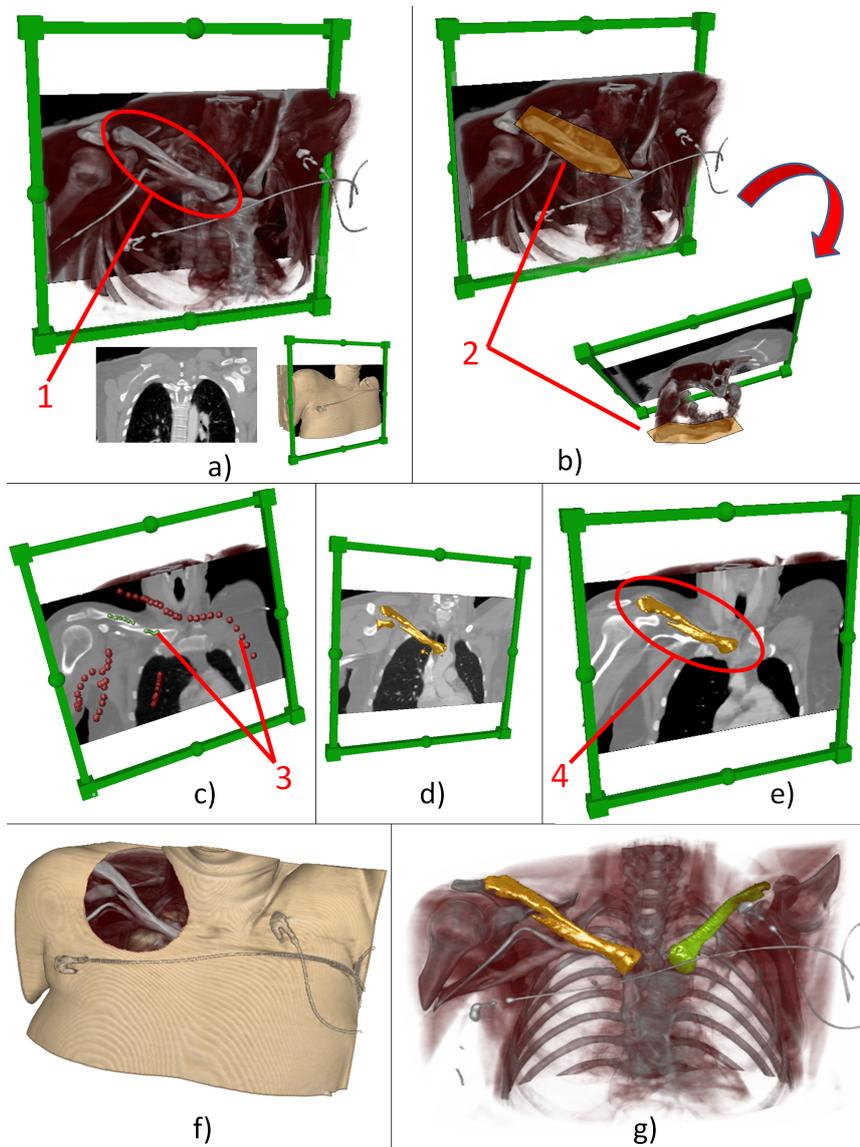


Figure 4: The work flow of the clavicle case. The volume rendering in a) already shows the fractured clavicular bone (1), the context of the body surface and the slice through the volume (green rectangle) is depicted in the smaller visualizations. b) illustrates the space carving approach to specify a region of interest (2) restricting the segmentation. In c) seed regions (3) are placed to mark background (red) and foreground objects (green). d) shows the cluttered segmentation result, which is refined to produce the fractured clavicular bone (4) in e). f) presents the data in a focus & context fashion, while in g) the two segmented bone structures are visualized together with the volume rendering, where they can be manipulated and investigated.

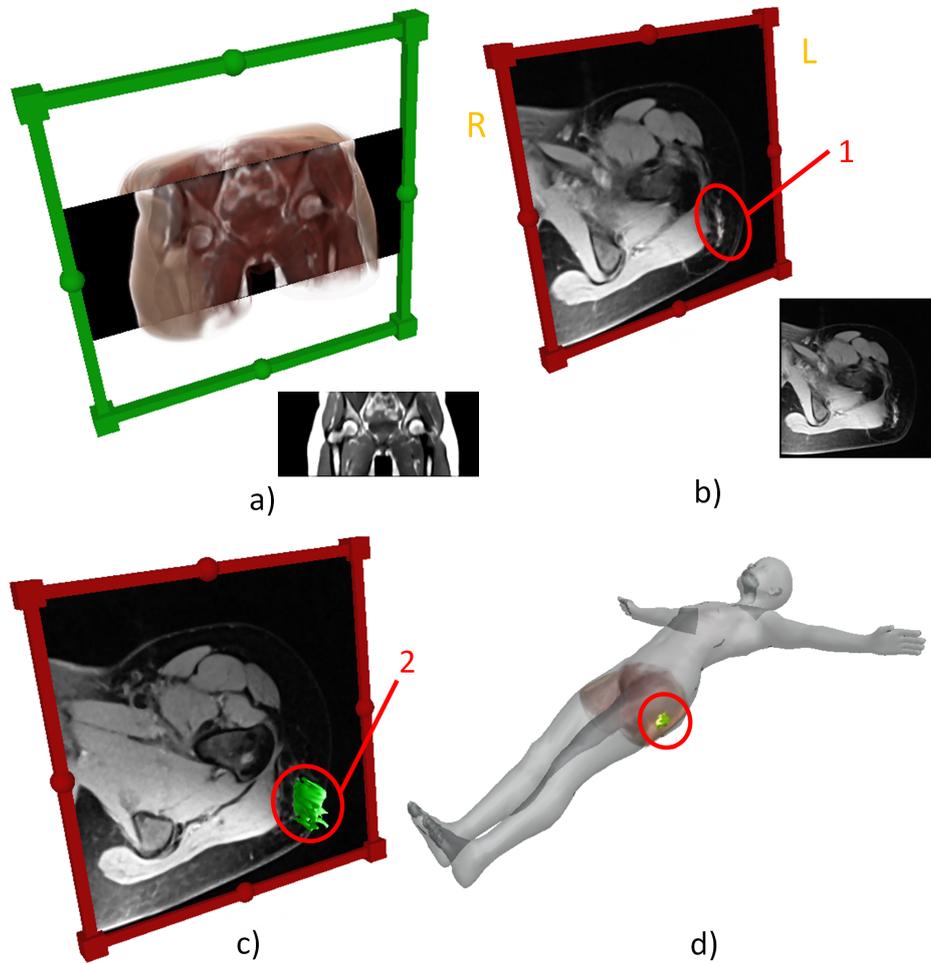


Figure 5: The work flow of the hematoma case. The volume rendering in a) shows the T1 weighted scan and a coronal slice through the volume. In b) the proton density scan shows a good contrast in the hematoma. The depicted slice is hard to interpret for a non-radiologist. In c) the segmentation of the fuzzy blood pool is shown (2), which is visualized together with the volume rendered T1 scan and a virtual reference body model in d).

A Multi-volume rendering

Direct volume rendering (DVR), mostly in the form of ray-casting, has recently achieved interactive performance using GPU programming. In our work, we build on the work in [5], which allows real-time rendering of multiple volumes with arbitrary polyhedral boundaries. This approach, further on referred to as *polyhedral DVR*, supports a large number of simultaneous volumes, complex translucent and concave polyhedral objects as well as Boolean operations of volumes and geometry in any combination.

Polyhedral DVR scales only with the memory footprint of the volumetric dataset, and its performance does not strongly depend on the number of volumetric intersections in the scene. It is based on a software rendering pipeline written entirely in CUDA, which allows to circumvent the limitations of the conventional fixed-function graphics pipeline on the GPU. Low-latency local memory on the GPU is used to accelerate the two computationally intensive stages of the ray-casting procedure. First, all polyhedral boundaries are rasterized using a hierarchical tiling approach with coverage masks. Second, all fragments produced in the rasterization that cover a single pixel are sorted by depth. The result is forwarded to the sampler, which steps along the ray and can adjust its sampling strategy at the boundaries of each depth segment.

Unlike depth peeling [8] strategies, which are conventionally used for direct rendering of multiple volumes, *polyhedral DVR* traverses the scene only once and it is also fully flexible concerning the interpretation of the current ray. Moreover, with *polyhedral DVR*, the whole scene can be defined as a tree of Boolean operations with arbitrary depth.

This behavior is crucial for rendering of multi-variate or multi-volume datasets. It also makes it easy to mix rendering of volumes and polygonal surface geometry. For example, a segmentation obtained by our framework can be used in both polygonal and voxelized form to constrain or alter the rendering of a larger enclosing volume.

B Two-label volume segmentation

Extraction of forensically relevant details requires a flexible segmentation technique. As our core segmentation tool, we use an interactive foreground-background segmentation formulated as an energy minimization. The underlying mathematical formulation describes a geodesic active contour (GAC) model that separates fore- and background, i. e., the hypersurface at the border of the two labels [9]. The GAC model may be extended to include prior knowledge on the gray-value distribution, the texture characteristics,

and shape constraint information of the foreground and background region. We use a continuous formulation of GAC energy minimization in a variational framework [10]. It was shown in [11] that the minimization of the GAC energy is equivalent to solving the weighted total variation (TV) model, with the benefit that the energy formulation is convex and converges to a global optimum representing the desired segmentation result. Traditional segmentation approaches like level-set methods are prone to get stuck in local minima. The weighting g is related to edges in the input image $I(\mathbf{x})$ with g close to 0 for edge locations and close to 1 for homogeneous regions. The GAC energy formulated as a weighted TV minimization is given as:

$$\min_u \left\{ \int_{\Omega} g(I(\mathbf{x})) |\nabla u(\mathbf{x})| d\mathbf{x} + \lambda \int_{\Omega} u(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} \right\} \quad (1)$$

Given an input image $I(\mathbf{x})$ in the domain $\Omega \subset \mathbb{R}^3$, we seek u , a binary labeling of the image into foreground ($u = 1$) and background ($u = 0$). The first term in (1) is the weighted TV regularization term penalizing discontinuities in the segmentation via the gradient of the labeling u . The second term of (1) is a pointwise data-term, where a positive $f(\mathbf{x})$ forces $u(\mathbf{x})$ to be background, and a negative $f(\mathbf{x})$ forces $u(\mathbf{x})$ to be foreground. We refer to f as our seed image.

The minimization of our convex energy formulation (1) involves a relaxation of the non-convex binary labeling $u \in \{0, 1\}$ to the convex set $u \in [0, 1]$. We can solve the convex minimization by deriving and solving the associated Euler-Lagrange equations. The solution is globally optimal with respect to the user-specified foreground and background seeds given in f . We distinguish two types of seeds, weak and hard ones, respectively. Hard segmentation seeds are specified with $f = -\infty$ (foreground) and $f = \infty$ (background). These seeds can be used for interactive segmentation refinement by removing or adding structures to the foreground. Weak foreground seeds use $f < 0$ to model a tendency to develop the foreground label in the corresponding regions and in regions similar to the gray-values of the seed region. At these regions, the data term tries to make $u = 1$. However, depending on λ , the regularization can still work towards $u = 0$. A weak background seed $f > 0$ works equivalently for the background region.

The numerical solver of our partial differential equations uses a primal-dual algorithm, where the weighted TV energy is rewritten using a dual representation. A gradient descent on the primal unknown u combined with a gradient ascent on the dual variable result in a solution of the associated saddle point problem [12]. This algorithm can very efficiently be parallelized in CUDA.

C Implementation details

All components have to be integrated into a common software framework. The Qt library² is used to create the graphical user interface (GUI). Both the volume raycasting and the segmentation subsystems are implemented in CUDA, which gives significantly improved performance for these tasks compared to conventional GPU shading languages. The system therefore greatly benefits from the processing power of recent GPU technology. The rendering components are wrapped in nodes of the scene graph Coin3D³, which makes it easy to create, manipulate and store complex visual scenes.

The software intelligently distributes processing tasks to different GPUs if more than one is available. The test configuration used in the case studies consists of a workstation with a quad-core Intel i7 (6GB RAM) and three NVidia GTX 285 graphics cards (2GB RAM each). After initial uploading of the data set to the GPU memory, all user interactions, like specification of foreground and background seed regions or segmentation refinement, are performed directly on the GPU for rapid feedback. In our implementation of the forensic analysis and presentation tools, we use one available GPU solely for analysis tasks running its own thread, while we distribute the remaining GPUs to the core volume rendering in a parallelized fashion, where each GPU is responsible for rendering a part of the final frame. Scheduling prioritizes rendering over segmentation processing for the sake of interactivity.

We also provide a way to perform very demanding computations transparently over the network on a dedicated GPU server (NVidia Tesla S1070 in our environment), using the Ice software library⁴ for remote object communication.

References

- [1] Jack March, Damian Schofield, Martin Evison, and Noel Woodford. Three-Dimensional Computer Visualization of Forensic Pathology Data. *The American Journal of Forensic Medicine and Pathology*, 25(1):60–70, March 2004. [1](#)
- [2] Patric Ljung, Calle Winskog, Anders Persson, Claes Lundstroem, and Anders Ynnerman. Full body virtual autopsies using a state-of-the-art rendering pipeline. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):869–876, 2006. [1](#)

²<http://qt.nokia.com/>, Sep. 2010

³<http://www.coin3d.org>, Sep. 2010

⁴<http://www.zeroc.com>, Sep. 2010

- [3] Michael Thali, Richard Dirnhofer, and Peter Vock. *The Virtopsy Approach: 3D Optical and Radiological Scanning and Reconstruction in Forensic Medicine*. CRC Press Inc, 2008. [1](#), [4](#)
- [4] Marc Levoy. Display of surfaces from volume data. *IEEE Computer Graphics & Applications*, 8(3):29–37, 1988. [1](#)
- [5] Bernhard Kainz, Markus Grabner, Alexander Bornik, Stefan Hauswiesner, Judith Muehl, and Dieter Schmalstieg. Ray casting of multiple volumetric datasets with polyhedral boundaries on manycore gpus. *ACM Trans. Graph.*, 28(5):Article No. 152, 2009. [3](#), [13](#)
- [6] Lujin Wang, Ye Zhao, Klaus Mueller, and Arie Kaufman. The magic volume lens: An interactive focus+context technique for volume rendering. In *In Proc. of IEEE Visualization 05 (2005)*, pages 367–374, 2005. [4](#)
- [7] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Phys. D*, 60(1-4):259–268, 1992. [5](#)
- [8] Cass Everitt. Interactive order-independent transparency. Technical report, NVIDIA, 2001. [13](#)
- [9] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic Active Contours. *International Journal of Computer Vision*, 22(1):61–79, February 1997. [13](#)
- [10] Markus Unger, Thomas Pock, Werner Trobin, Daniel Cremers, and Horst Bischof. TVSeg - interactive total variation based image segmentation. In *Proceedings British Machine Vision Conference 2008*, Leeds, UK, September 2008. [14](#)
- [11] Xavier Bresson, Selim Esedoglu, Pierre Vandergheynst, Jean-Philippe Thiran, and Stanley Osher. Fast global minimization of the active contour/snake model. *Journal of Mathematical Imaging and Vision*, 28(2):151–167, 2007. [14](#)
- [12] Thomas Pock, Thomas Schoenemann, Gottfried Graber, Horst Bischof, and Daniel Cremers. A convex formulation of continuous multi-label problems. In *Proc European Conference Computer Vision (ECCV 2008)*, volume 5304 of *LNCS*, pages 792–805. Springer, 2008. [14](#)