

Contextual Picking of Volumetric Structures

Peter Kohlmann*
Vienna University of Technology

Stefan Bruckner†
Vienna University of Technology

Armin Kanitsar‡
AGFA HealthCare

M. Eduard Gröller§
Vienna University of Technology

ABSTRACT

This paper presents a novel method for the interactive identification of contextual interest points within volumetric data by picking on a direct volume rendered image. In clinical diagnostics the points of interest are often located in the center of anatomical structures. In order to derive the volumetric position which allows a convenient examination of the intended structure, the system automatically extracts contextual meta information from the DICOM (Digital Imaging and Communications in Medicine) images and the setup of the medical workstation. Along a viewing ray for a volumetric picking, the ray profile is analyzed for structures which are similar to predefined templates from a knowledge base. We demonstrate with our results that the obtained position in 3D can be utilized to highlight a structure in 2D slice views, to interactively calculate centerlines of tubular objects, or to place labels at contextually-defined volumetric positions.

Keywords: Contextual Visualization, Smart Interaction, Linked Views, Medical Visualization, Feature Selection.

Index Terms: I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction Techniques; J.3 [Life and Medical Sciences]: Medical Information Systems—

1 INTRODUCTION

Current computer hardware allows to display volume data with different rendering techniques simultaneously and in real time. For a certain medical diagnostic task in the clinical routine, a digital hanging protocol defines how the data is reformatted and arranged on the screen. For some examinations, e.g., in mammography, the hanging protocol is highly standardized, whereas, e.g., in vascular examinations, more often customized hanging protocols are preferred. Frequently, Multi-Planar Reformatting (MPR) is the technique of choice to provide sectional renderings. With a Curved Planar Reformation (CPR) the whole extent of a tubular structure is displayed within a single image. These 2D renderings are often accompanied by a Direct Volume Rendering (DVR) technique like raycasting. The examination of a structure in its three-dimensional setting often provides better insights into contextual information.

A typical hanging protocol arranges a side-by-side presentation of different views of the volumetric data. The physician performs various interactions during the examination of the data. Examples of frequently recurring interactions are scrolling through slices (2D images), zooming, panning, labeling, windowing (2D/3D images), or viewpoint selection and clipping (3D images). The synchronization of the different views is quite challenging because it is not trivial to determine if an interaction in one view leads to changes within another view. In earlier work [7, 8] we presented the LiveSync interaction metaphor, a solution for the live synchronization of a

2D slice view and a 3D volumetric view. The viewing parameters for the 3D view are derived automatically from a picking of the anatomical structure of interest on the 2D slice. In that work the viewing parameters are viewpoint, zoom factor, clipping planes, and transfer function setup.

This paper presents a new approach to handle the picking of a structure which is directly performed on the 3D volumetric view in a context-sensitive way. Therefore, the first step is the identification of the 3D position of interest within the volumetric data. For each point on a 2D slice its exact 3D position can be calculated easily, whereas the picking of a point in the 3D volumetric view is not well defined. This is due to the fact that for each pixel on the screen ray casting is performed, where opacity and color is accumulated along a ray from the eye point through the volume. Potentially each position along this ray might be the desired volumetric position. A simple solution to this problem is the definition of a *first-hit position* as the location along the ray where a certain opacity threshold is exceeded. For some cases this might be sufficient, but often a different, contextually-defined position is of interest. We propose ray-profile templates which are designed to represent anatomical structures like, e.g., a vessel, the aorta, the airway, or a vertebra. The ray profile which is calculated for each picking on the 3D volumetric view is then scanned and analyzed to find similarities to the defined ray-profile templates in a knowledge base.

To narrow down the number of anatomical structures which are of relevance for a certain examination, the presented method extracts contextual information automatically from the DICOM data. The DICOM standard [16] provides a detailed specification of a format for medical images which includes meta information like parameters of the scanner and patient information. Further context is provided by the setup of the medical workstation and the selection of clinical tools. For the examination of anatomical structures, frequently center points are of special interest. The *first-hit* solution cannot provide these positions because of self-occlusion or occlusion by other structures. Our proposed method can either return the first hit of the determined structure of interest or its center along the analyzed ray.

When the 3D position of interest is located, the next important step is the presentation of the result to the physician. This paper depicts the highlighting of the results by synchronized 2D slice views as a straight-forward solution. Also the placement of labels at the appropriate 3D positions is demonstrated. For instance, the whole spine can be labeled by picking each vertebra directly in the 3D volumetric view. Further it is shown that approximate centerlines can be calculated interactively by tracing along tubular structures in the 3D view. These centerlines could be utilized to display CPR renderings of the structure or to guide a segmentation process.

2 RELATED WORK

In medical visualization some techniques have been developed to ease the interaction with multiple views of a certain data set. Kohlmann et al. [7, 8] presented the first attempt to combine optimal viewpoint estimation and synchronized views for the visualization of medical volume data. Götzelmann et al. [6] presented an approach where 3D visualizations are linked with textual descriptions, e.g., from medical textbooks. Their approach focuses on an educational purpose and supports students to learn the terminology and to understand textual descriptions of complex objects.

*e-mail: kohlmann@cg.tuwien.ac.at

†e-mail: bruckner@cg.tuwien.ac.at

‡e-mail: armin.kanitsar@agfa.com

§e-mail: groeller@cg.tuwien.ac.at

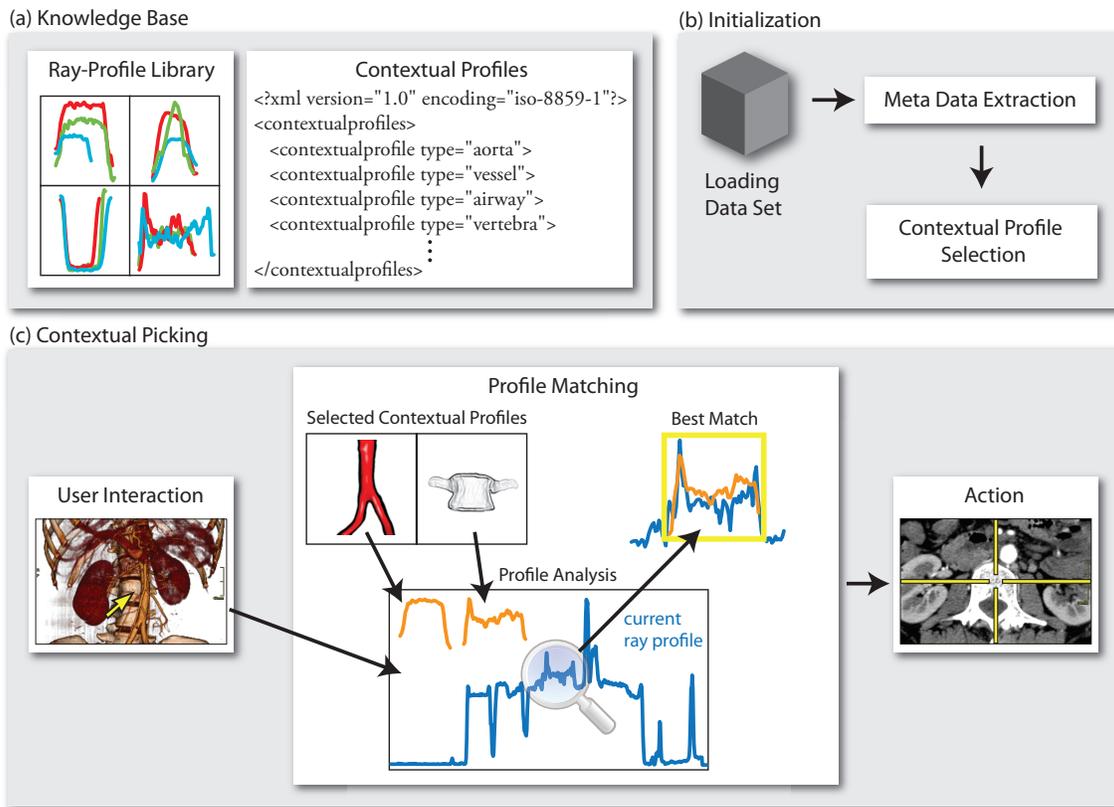


Figure 1: Contextual picking overview: (a) A knowledge base provides a library of ray profiles of anatomical structures together with an XML description of available contextual profiles. (b) An initialization step is performed during data set loading. Meta information is extracted to automatically select contextual profiles from the knowledge base. (c) For a picking on the volume rendered image, the current ray profile is analyzed to detect anatomical structures which are represented by the selected contextual profiles. If a good match is given, then this can be utilized for instance to highlight an interest point (e.g., the center of a vertebra) in the slice view.

Related to this work, IBM is currently developing the Anatomic and Symbolic Mapper Engine [3]. This technology uses a 3D model of the human body which is linked to medical records. Whenever the doctor clicks on a certain part of the body, a search of the medical records is triggered to extract the relevant information. The Medical Exploration Toolkit presented by Tietjen et al. [18] bundles various concepts for loading, visualizing, and exploring segmented medical data sets. Critical distances to pathological structures are computed, and displayed in a synchronized manner in 2D and 3D. The integrated LIFTCHART [17] displays the overall extents of structures within the volume in a narrow frame as color bars. A structure can be selected in the LIFTCHART and the corresponding slice is displayed in the slice viewer. Another clinical application they integrated is the NECKSURGERYPLANNER [19]. Segmented structures can be enabled and disabled by a textual representation. They are synchronously highlighted in the 3D and 2D views to support operation planning for neck dissections.

A lot of research has concentrated on the extraction of certain anatomical structures. Multi-scale filtering approaches are very popular for the detection of curvilinear structures. Vessel enhancement filters based on eigenvalue analysis of the Hessian matrix have been proposed, e.g., by Sato et al. [13] and Frangi et al. [5]. Tek et al. [14] presented an approach which focuses on the segmentation of vessel cross sections. A single click inside the vessel on a slice initiates mean shift-based ray propagation to detect the boundary of the vessel. Tschirren et al. [20] presented an airway segmentation algorithm based on fuzzy connectivity. Their method uses small adaptive regions of interest which follow the airway branches

during the segmentation process. Kovács et al. [9] developed a system for automatic segmentation of the entire aorta without any user interaction for treatment planning of aortic dissections. The segmentation is based on a Hough transformation to detect the approximate circular shape of the aorta. To fit this shape more closely to the actual contour of the aortic lumen an elastic mass-spring deformable model is utilized. An interesting concept for the detection of tubular objects in medical imaging is the *Gradient Vector Flow (GVF) snake* introduced by Xu and Prince [21]. This method first calculates a field of forces (GVF forces) over the image domain. These forces drive the snake to fit to the boundaries of an object. Bauer and Bischof [1, 2] utilize the properties of the GVF for the detection of tubular objects and the extraction of curve skeletons, e.g., for virtual endoscopy. They argue that conventional tube detection or line filters which use local derivatives at multiple scales have problems with undesired diffusion of nearby objects. Their GVF-based method allows an edge-preserving diffusion of gradient information. Malik et al. [10] presented a rendering algorithm called feature peeling. They analyze peaks and valleys of intensity ray profiles for a given viewpoint to detect features inside the volume data. By classifying a number of feature layers it is possible to scroll through the layers to inspect various structures.

Sketch-based techniques are employed to classify and segment volume data by painting directly on the volume rendering. Owada et al. [11] presented a sketching interface which allows an intuitive segmentation of volumetric regions. The user has to draw 2D strokes along the contour of the 3D target region to initiate a constraint segmentation process. Related to this approach, Chen et

al. [4] applied sketching to seeded region growing for volume segmentation. Initially, the user specifies a region of interest by placing a closed free-form sketch on the volume rendering. Then a region of interest is extruded to facilitate the definition of seed points for the region growing. Ropinski et al. [12] proposed an interface for the design of 1D transfer functions which is based on the drawing of strokes on the volume rendered image. Features of interest can be identified by strokes which are close to their silhouettes. Based on a histogram analysis, component transfer functions are automatically generated for the identified features. The user can then decide which of the selected features should be integrated into the final transfer function.

3 CONTEXTUAL PICKING OVERVIEW

The overview of related works indicates that often highly specialized methods are used to detect various anatomical structures within the volumetric data. This paper presents a generalized system which is not limited to a certain type of structure. An example illustrates one potential application area for the presented system: A frequently needed task during an orthopedic examination of the spine is the labeling of individual vertebrae. The vertebral column consists of 7 cervical, 12 thoracic, 5 lumbar, 5 sacral, and 3 to 5 coccygeal vertebrae. If the labeling is performed in 2D slice views only, quite some scrolling through the slices is involved to navigate to a meaningful 3D position where a single label is placed. A good position for the label is the center of the vertebral body. By picking on the vertebrae in the 3D view, contextual picking allows a convenient labeling. If the label for the first picked vertebra and the labeling direction is given, then a single contextual picking on the following vertebrae is sufficient to add the appropriate labels at a central position within the vertebral body.

Figure 1 shows the building blocks to achieve contextual picking. A knowledge base consists of a ray-profile library and contextual profiles. The ray-profile library holds ray-profile samples (intensities and gradient magnitudes) of various anatomical structures. A contextual profile for a certain structure bundles the needed information to react to a contextual picking operation. In the XML format it describes the following components: The type of the structure, a list of keywords, minimal and maximal extent of a structure, a representative mean ray profile built from the samples in the ray-profile library, and the default reaction to a picking operation. An initialization step is performed whenever a new data set is loaded into the workstation. The DICOM header as well as the workstation environment is analyzed to extract the relevant meta data for selecting the applicable contextual profiles.

Contextual picking is initiated by positioning the mouse cursor on the 3D view and pressing a hot-key. Whenever the physician picks on a structure of interest in the 3D view, the following steps are performed: First, information from the current picking, which includes the intensity and gradient magnitude values along the viewing ray, as well as accumulated opacities and positions of clipping planes, are collected. Second, the representative mean ray profiles of the selected contextual profiles are compared to the current ray profile. Finally if a good match is detected, this result is utilized to highlight the anatomical structure of interest in an appropriate way, e.g., in MPR views.

4 KNOWLEDGE BASE

Most approaches to automatically detect features in medical volume data need a considerable amount of user interaction to set up the needed parameters. Besides, they are often very specialized on a certain type of anatomical structure as well as on a specific extent of the feature. Two important preconditions for the presented system are that it has to be as generic as possible and easy to extend. For these reasons a ray-profile library was set up which is supported by an easy-to-use interface for the generation of new ray-profile

samples. Together with the contextual profiles this knowledge base provides all information to react to a contextual picking.

4.1 Ray-Profile Library

The ray-profile library is implemented as an XML file which stores ray-profile samples for various anatomical structures. A ray-profile sample for a certain structure consists of a sample id, a textual description, the spacing along the ray in mm, the extent of the structure in mm, and a list of intensity and gradient magnitude values. The system provides a convenient user interface to add new ray-profile samples to the ray-profile library. Figure 2 illustrates the generation of a ray-profile sample of a contrast-enhanced blood vessel. After picking on a vessel in the 3D view, a ray profile is generated and displayed. A plot of the ray profile shows the intensity values (blue) and the values of the gradient magnitude (green) along the ray. To ease the selection of a ray-profile sample, the color transfer function for the corresponding scalar values is displayed in a small horizontal bar below the profile plot.

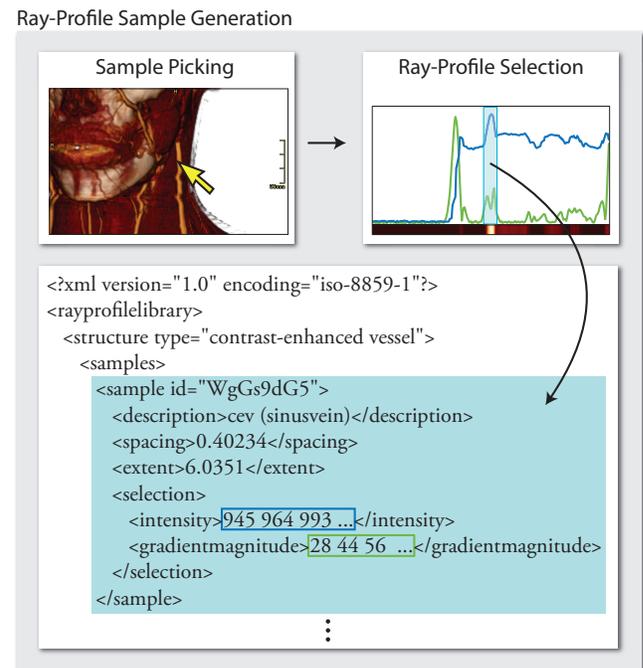


Figure 2: The generation of a ray-profile sample for a contrast-enhanced vessel is initiated by picking on the structure in the 3D view. A ray profile is displayed which shows the intensity and the gradient magnitude values along the viewing ray. From a selected subset of this ray profile, which represents the picked vessel, a sample is generated and stored in the ray-profile library.

By dragging the mouse with the left button pressed a semi-transparent window is painted over the ray profile which represents a selection along the ray. By confirming this selection, a sample for a new or an already existing structure is written to an XML file. For a certain anatomical structure it is recommended to generate several representative samples. The intensity values of an anatomical structure can vary slightly, e.g., because of the patient's age or sex. Further the size of the structures varies because of the mentioned factors. A set of multiple samples in a contextual profile for example enables to detect vessels of a certain diameter range. The generation of ray-profile samples has to be done only once for each anatomical structure. They are added to the library by a domain expert. This step is not visible to the physician who just uses the contextual picking.

4.2 Contextual Profiles

The contextual profiles are stored in an XML file which contains all needed instructions to react to a contextual picking. Listing 1 shows the XML skeleton of the contextual profile for an anatomical structure. First of all it has a *type* entry which has to match with the corresponding *structure types* in the ray-profile library. A list of keywords describes the type of data sets and/or the setup of the medical workstation in which the defined structure is of interest. The *extent* defines a range of the typical extent of the anatomical structure. In the *meanrayprofile* entry, a mean ray profile is stored which is generated from the available ray-profile samples in the ray-profile library. Finally, *return* defines which position will be returned (e.g., the center of the structure) and which default action shall be performed (e.g., highlighting of the obtained position in MPR views).

Listing 1: XML skeleton of a contextual profile

```
<contextualprofile type="">
  <keywords>...</keywords>
  <extent>...</extent>
  <meanrayprofile>
    <spacing>...</spacing>
    <intensity>...</intensity>
    <gradientmagnitude>...</gradientmagnitude>
  </meanrayprofile>
  <return>
    <position>...</position>
    <reaction>...</reaction>
  </return>
</contextualprofile>
```

Mean ray-profile generation: The generation of a mean ray profile is motivated by the variation of intensities and extents of anatomical structures. To obtain a good representation of a structure, all samples from the ray-profile library which correspond to the type of the contextual profile are collected. Figure 3 (left) shows three intensity ray-profile samples of the aorta which were all captured using different data sets. The similarities of the three samples are clearly visible. The intensity values are in a range of Hounsfield Units between about 880 and 1440 (shown on the y-axis of the plots) and the extents of the samples differ in a range between about 17.5 and 28.5 mm (shown on the x-axes of the plots). In general the samples start with a steep intensity ascent followed by a plateau and a steep intensity descent. The generated mean sample (right) shows the same characteristics with the advantage that outliers are filtered out. The algorithm for the generation of the mean ray profile first calculates the mean extent of the available samples. Then the ray-profile samples are scaled horizontally. The scaling factor is determined by dividing the mean extent by the extent of the current sample. Afterwards the mean of the corresponding intensity values is calculated at uniformly distributed positions along the x-axis. Analogously, a mean ray profile is also generated for the gradient magnitudes. Taking the mean extent of the ray-profile samples seems to be an appropriate approach because we assume that there is an approximate Gaussian distribution of the extents for multiple samples of a single structure.

The described algorithm for the generation of the mean ray profiles is well suited to preserve slopes which are characteristic for a certain structure. This is due to the fact that the steepness of slopes in shorter samples is decreased and the steepness of slopes in longer samples is increased by the horizontal scaling. Taking the mean of the intensity values results in a mean ray profile which represents the anatomical structure accordingly. Mean calculations have to be performed only when a new sample is added to the ray-profile library.

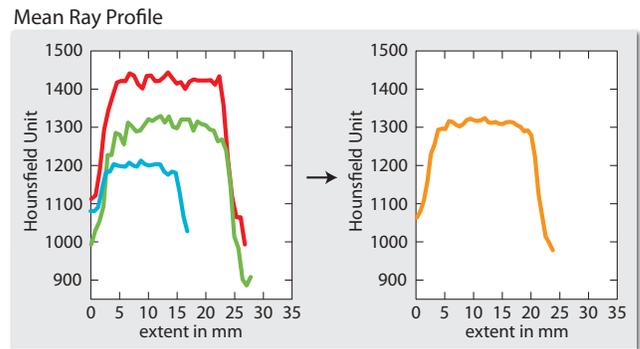


Figure 3: Three samples of the aorta (left) with varying extent and intensity range are collected from the ray-profile library. These samples are utilized to construct a representative mean ray profile (right).

5 INITIALIZATION

The major challenge for a system which allows contextual picking of anatomical structures in the 3D view is to gain as much information as possible about the structure of interest. Typically for a certain type of examination only a small number of structures is relevant for the diagnosis. For instance in vascular examinations veins and arteries are of special interest whereas in an orthopedic examination the spine and bones are more important structures. There are two reasons to narrow down the number of target structures in a certain examination. First, our approach is based on a matching between ray-profile samples from a knowledge base and the ray profile extracted for the current picking. In volume data, structures are occluding each other. Thus the analysis of the current ray profile with respect to all structures given in the contextual profiles, might lead to ambiguous results. For example a vessel in front of a bone is usually not relevant for an orthopedic examination. In such cases the contextual picking has to respond appropriately. Second, the analysis of the current ray profile through matching with a lot of structures of varying extents can lead to high computational costs.

As this work aims to minimize the overhead of the user interaction, the presented system extracts valuable meta information directly from the DICOM headers and the current setup of the medical workstation. The DICOM format is widely used in hospitals as a standard for handling, storing, printing and transmitting information in medical imaging. DICOM files contain the image pixel data of the scanned slices combined with header information. A *Patient* entity contains relevant data about the patient like name, age, and sex. Data concerning the examination like the name of the physician or the description of the examination is stored in a *Study* entity. A *Series* entity represents a logical unit of information about the imaging modality (*Equipment*) and information about the spatial relations of images within a series (*Frame of Reference*). Finally, the image itself and meta data about the image are stored in an *Image* entity. Grouping the image and the meta data prohibits that the image data gets separated from the meta information [16]. To narrow down the number of structures the physician might be interested in, we identified a suitable set of entries in the DICOM header:

- (0018, 0015) - Body Part Examined
- (0008, 1030) - Study Description
- (0008, 103E) - Series Description
- (0040, 0254) - Performed Procedure Step Description
- (0018, 1030) - Protocol Name

After a new data set is loaded, our system extracts the textual description which is stored for these entries. Entries which contain details about the patient like sex, age, and weight are further candidates which could be utilized to gather information about the size and the intensity range of a specific structure. The system also considers if the medical workstation is only used for certain examinations. For instance often a medical workstation with reduced functionality is available, e.g., as a vascular, orthopedic, cardiac, or mammography workstation.

The automatically extracted information is used to select the suitable contextual profiles from the knowledge base. Each contextual profile contains a list of classified keywords to decide if the represented structure is relevant for the currently loaded data set with the current setup of the workstation. Listing 2 shows exemplary keywords in a contextual vertebra profile.

Listing 2: Keywords of the contextual vertebra profile

```
<keywords>
  <strong>
    Workstation=Orthopedic
    BodyPartExamined=*SPINE ...
  </strong>
  <medium>
    BodyPartExamined=ABDOMEN ...
  </medium>
  <weak>
    BodyPartExamined=HIP ...
  </weak>
  <kickout>
    Workstation=Cardiac
    Workstation=Vascular ...
  </kickout>
</keywords>
```

A list of keywords is categorized into the classes *strong*, *medium*, *weak*, and *kickout*. The information which is extracted during loading the data set is compared with these entries to decide if a certain anatomical structure is relevant for the current examination. In the presented example a vertebra is strongly relevant when using an orthopedic workstation and/or when the examined body part is the spine. Within a cardiac or a vascular workstation a vertebra is typically not a structure of interest. With this approach it is possible to select from a ranked list the contextual profiles which are suitable in the given environment. The ranking is given by a comparison of the extracted information with the classified keywords.

6 CONTEXTUAL PICKING

After one or more contextual profiles are automatically selected, the physician can perform the picking directly on the 3D view and the system provides an immediate feedback. For each contextual picking, the current ray profile is analyzed to find close similarities in the selected contextual profiles. This analysis is done by a profile-matching algorithm which evaluates a cost function to measure the degree of similarity. As our main focus was to provide a framework for contextual visualization we did not deeply investigate the applicability of established pattern classification methods from fields like signal processing, computer vision, or neurocomputing. Our empirical approach gives satisfactory results and the comparison to other methods has to be done in future work. Based on the outcome of the profile matching the respective action is taken. Three actions have been implemented so far. The default action is the highlighting of the center of the picked anatomical structure in MPR views. Further, contextual picking is integrated into a spine labeling system to demonstrate its potential to place labels at meaningful 3D positions. Finally the system allows the calculation of approximate centerlines of picked tubular structures.

6.1 Profile Matching

The profile matching to detect the anatomical structure of interest is repeated for each of the selected contextual profiles. A contextual profile provides a mean ray profile which represents a structure together with the information about a minimal and a maximal extent of the structure. To allow the search for structures within this range along the current ray profile, a non-uniform scaling of the mean ray profiles is performed. In Figure 4 the minimal (left) and the maximal scaling (right) is shown for the mean ray profile (middle) of the aorta to cover a structure extent from about 16 to 32 mm. The scaling algorithm only performs a scaling on positions where the gradient magnitude is low. This avoids an undesirable change of the steepness of the major slopes.



Figure 4: The mean ray profile of the aorta (middle) is non-uniformly scaled between a minimal (left) and a maximal range (right) to match the corresponding structure in the current ray profile.

The implementation of the profile matching is described in Algorithm 1. In this algorithm the mean ray profile and the current ray profile are represented by intensities as well as gradient magnitudes. Thus, the summed Euclidean distances are calculated for the corresponding intensities and gradient magnitudes. The matching

Algorithm 1 Structure detection by profile matching.

```
1: set  $pos = 0$ ,  $width = 0$ ,  $minCost = MAX\_VALUE$ 
2: for each scaling of the mean ray profile  $m$  do
3:   for each step  $s$  along the current ray profile  $c$  do
4:     if  $s + m.length < c.length$  then
5:       sum up squared Euclidean distances (SEDs) at the corresponding positions along  $m$  and  $c$  ( $s$  to  $s + m.length$ )
6:     end if
7:   end for
8:   set profile matching  $cost$  to  $sum\_of\_SEDs/m.length$ 
9:   if  $cost < minCost$  then
10:    set  $pos = s$ ,  $width = m.length$ ,  $minCost = cost$ 
11:   end if
12: end for
```

algorithm detects the section along the current ray profile which is most similar to the mean ray profile at different scales. A length normalization of the fitting costs allows the comparison between different scales, as well as the comparison of responses between different contextual profiles. This is especially important if multiple contextual profiles are selected and thus, the ray profile is scanned for different structures. For instance, if an abdominal data set is loaded the contextual aorta profile and the contextual vertebra profile might be selected. To decide which structure is detected, a trade off between the cost function and the suitability of the contextual profile for the current environment has to be calculated. In the following, low costs of the profile matching are equivalent to a good or high response of a contextual profile.

Optimizations: To decrease ambiguities of the contextual profile response and to increase the performance of the system several

optimizations were applied to the algorithm. First of all, the positions of the clipping planes are considered before the profile matching is performed. Only the part of the volume which is not clipped away is analyzed for matching regions. On the one hand this lowers the computational cost and on the other hand it reduces the chances of ambiguities. For example in the case of the contextual vessel profile, a vessel which is located in front of another vessel along the ray but clipped away could lead to unintended results. A second optimization which is applied for similar reasons takes the opacity transfer function into account. The analysis of the ray profile starts only at a sample position where a small opacity threshold is exceeded. From this position on the remaining part of the ray (if not clipped away) is considered. Third, the cost function for the profile matching is implemented so that the costs slightly increase along the ray. With this adjustment, structures which are closer to the viewer, e.g., vessels which are in front of other vessels, return lower costs. To apply this modified cost function line 8 in Algorithm 1 is replaced by the equation

$$cost = (sum_of_SEDs/m.length) \times (1 + 0.5 \times s/c.length), \quad (1)$$

where sum_of_SEDs are the summed-up squared Euclidean distances, $m.length$ is the sample size of the current matching profile, s is the current sample position along the ray profile, and $c.length$ is the total sample size of the ray profile. Our experiments have shown that the multiplication of the costs with a penalty factor of up to 1.5 for distant structures leads to good results for features which appear multiple times along the viewing ray. Despite of the improvements achieved by Equation 1 two overlapping target structures which are very close to each other might still be problematic. An example is if the aorta due to the current viewpoint of the volumetric view is right in front of some vertebrae and contextual profiles for both structures are active.

If the user is continuously tracing along the aorta it would be quite disturbing if the contextual vertebra profile and the contextual aorta profile alternate in generating the better response. The proposed solution to this problem is to deactivate all but one contextual profile as long as a continuous tracing goes on. As soon as the hot-key is released the other contextual profiles are re-activated. Finally, a default contextual profile is implemented which returns the *first-hit position*. The opacity is accumulated along the ray until a certain opacity threshold is reached. This contextual profile becomes active if the cost function of the profile matching returns too high values, which means that no contextual profiles are detected along the current ray profile.

6.2 Contextual Picking Action

The implemented default action to react to a contextual picking is the highlighting of the detected interest point in the MPR views. For each picking on the volumetric view, a three-dimensional position within the data set is computed. This position can be the center of the target structure along the current viewing ray or the first-hit position if no target structure is detected. The structure's center can be calculated easily as the start and the extent of the structure along the viewing ray is determined by the profile-matching algorithm. To show the obtained position in the MPR views, the axial, coronal, and sagittal cross sections for this volumetric position are displayed. The centering of this position on the slices, as well as the overlay of crosshairs is used to highlight the target structure.

Another proposed action following a contextual picking is the labeling of anatomical structures. Often the labeling is performed in the slice views alone although the volumetric view can be very well suited for this task. For instance by utilizing the contextual vertebra profile, the physician gets more flexibility in the spine labeling process. The whole spine can be easily labeled in the 3D view. A single contextual picking on each vertebra determines the exact

three-dimensional position of each label. Finally, the estimation of feature center points during the continuous tracing along a structure can be utilized to calculate approximate centerlines of tubular structures like vessels. If the obtained approximation is not sufficient it can be a helpful initial input for more accurate centerline-detection algorithms.

7 PERFORMANCE AND RESULTS

The contextual picking is implemented in Java and the contextual profiles as well as the ray-profile library are stored in the XML format. For parsing and manipulation of the XML files within the Java classes the JDOM API [15] is used. The contextual picking is integrated into a real-world medical workstation which is under development by our collaborating company partner. All contextual picking-related computations are performed interactively. For the generation of the ray-profile samples in our ray-profile library eight different CT data sets were used. Three samples were taken for each anatomical structure from suitable data sets. The data sets shown in the result images of this section are different from the data sets which were used to establish the ray-profile library.

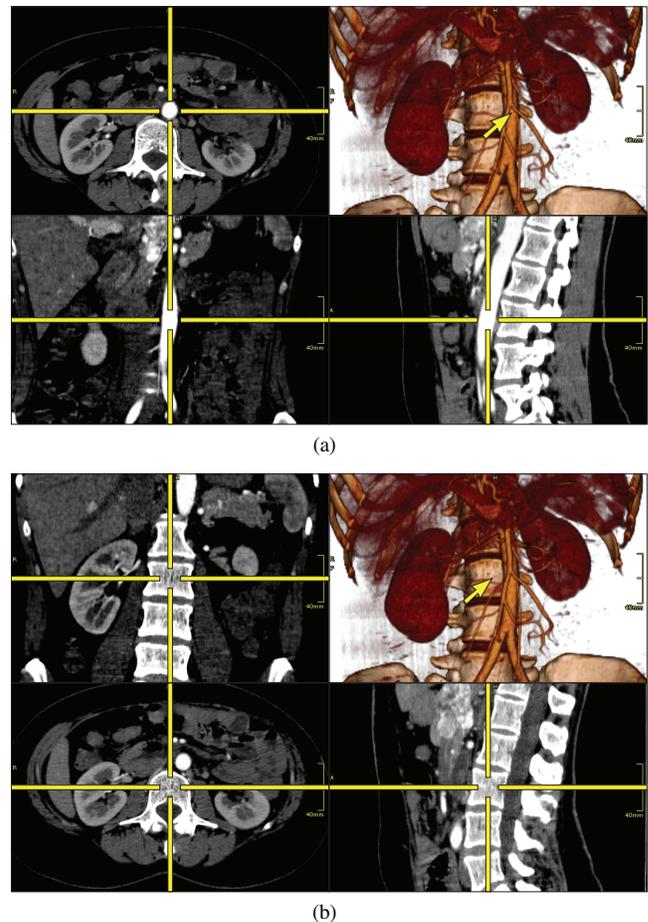


Figure 5: Contextual picking on a thoracic-abdominal CT data set. The 3D position which is returned by the contextual profile with the best response is used to provide meaningful MPR views of the picked aorta (a) and the picked vertebra (b).

In Figure 5 the contextual picking is illustrated for a thoracic-abdominal CT data set. After the data set is loaded into the workstation, meta information is extracted according to the description in Section 5. Based on this information, the contextual vertebra

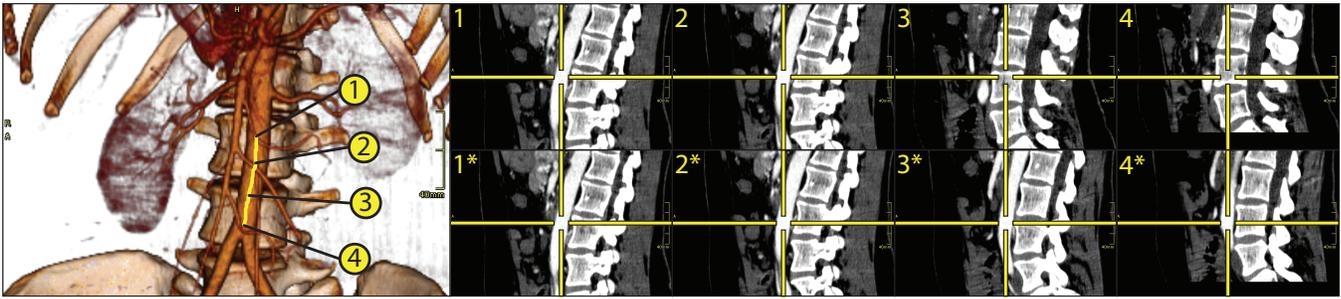


Figure 6: Continuous trace path along the aorta (left). This might lead to unintended responses for the vertebra if the contextual profiles of the aorta and the vertebra are selected (top-right [1-4]). The automatic temporary deactivation of the contextual vertebra profile during the tracing leads to a continuous capturing of the aorta (bottom-right [1*-4*]).

profile and the contextual aorta profile are selected automatically. When the aorta is picked as in Figure 5(a), the contextual aorta profile gives the best response. The detected center position of the aorta along the viewing ray is utilized to set up the axial, coronal, and sagittal MPR views. Four different views on the picked anatomical structure are provided to the physician. The picking of a close-by vertebra in Figure 5(b) leads to analogous results.

Figure 6 (left) shows the path (indicated by the yellow line) of a continuous tracing along part of the aorta. If the contextual aorta profile and the contextual vertebra profile are active, the contextual vertebra profile has the better response at some positions along the trace path although the user is interested in the examination of the aorta. Figure 6 (top-right [1-4]) shows the resulting sagittal slice views when only the best response is taken into account. The vertebra is captured as the prominent structure at the positions 3 and 4. Whenever a continuous tracing is performed, the assumption can be made that the user currently examines a single anatomical structure. Thus, just a single contextual profile is active during the continuous tracing and all the others are deactivated temporarily. This leads to the results shown in Figure 6 (bottom-right [1*-4*]). Along the trace path, the aorta is always captured as the prominent structure and jerky leaps in the MPR views between the aorta and a vertebra are avoided. The tracing along a tubular structure allows the computation of its approximate centerline.

Figure 7 depicts the result for the contextual picking of the airway in the 3D view (left) using a head CT data set. The contextual airway profile gives a better response than the contextual vessel profile which is also active. A highlighting of the corresponding position is performed in a 2D slice view (right). Occluding structures do not impede the detection of a central point within the airway.

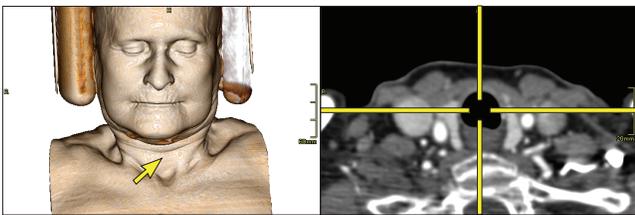


Figure 7: Contextual picking of the airway (left). The identified 3D position is used to provide a meaningful MPR view (right).

Figure 8 shows some results when the contextual picking is integrated into a spine labeling tool. With this tool the user has to specify the label for the first picked vertebra and a labeling direction (head-to-feet or feet-to-head). Then a single picking on each vertebra leads to the results shown in Figure 8(a). Figure 8(b) shows the labeling from another viewpoint if just the first-hit position is taken

for the 3D placement of the labels. If the placement is done by taking the positions determined by the contextual picking, the labels are in the center of the vertebral body as shown for the same viewpoint in Figure 8(c). The exact positions of the labels are depicted on the axial, coronal, and sagittal slices for the first-hit approach in Figure 8(d) and for the contextual picking approach in Figure 8(e).

8 CONCLUSION AND FUTURE WORK

In this paper we presented a novel method for the interactive identification of contextual interest points within volumetric data by picking on a direct volume rendered image. We built a knowledge base which holds characteristic ray-profile samples for different anatomical structures. New ray-profile samples can be added with an easy-to-use interface by domain experts. A contextual profile bundles information like for which kind of data sets it should be selected, the extent of the target structure, or the action which has to be performed if there is a high response to the contextual profile. Based on this knowledge base, the contextual profiles which are applicable in the current environment are selected automatically. Our profile-matching algorithm analyzes the viewing ray for each contextual picking on the 3D volumetric view. It returns a position within the volume with the best response to one of the selected contextual profiles. Based on this result, certain actions are performed interactively. In the simplest case, the obtained position is the center of the selected structure and is highlighted by crosshairs in MPR views. Because of the interactivity of the underlying computations, contextual picking is well suited to continuously trace the mouse pointer along volumetric structures. This allows to simultaneously examine the selected structure in multiple views. We have also demonstrated that the contextual picking can be easily integrated into a conventional spine labeling framework to increase its flexibility.

Classifying structures only along ray profiles is not a limitation of the contextual picking framework. Other local classification schemes could be added as well. More research is necessary to investigate if this could help to further improve the detection of a structure's interest point. Until now, our method is to a certain degree dependent on the chosen viewpoint of the 3D view. While this could be a problem in some cases, there are often default viewpoints for 3D diagnostic examination procedures. For this reason the viewpoint dependency was not a big issue in the application scenarios which are presented in this paper. Shape-based methods could be integrated into our framework to be more flexible in the selection of appropriate viewpoints. The challenge thereby will be to ensure interactivity. Alternatively, multiple contextual profiles could be provided for different viewpoints on a structure. The integration of techniques to display uncertainty information about the currently detected structure is another interesting direction for further research.

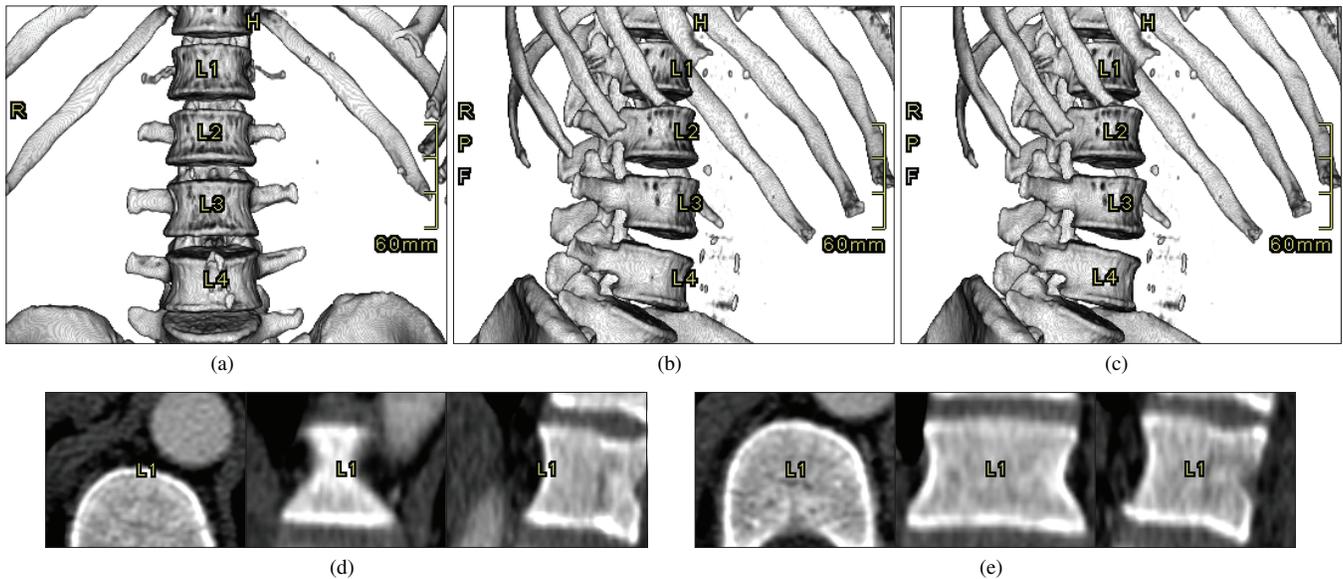


Figure 8: Labeling of lumbar vertebrae. (a) A single picking on each of the four lumbar vertebrae is performed for their labeling. (b) The result from another viewpoint when the first-hit position is taken for label placement. (c) The result for the same viewpoint when the contextual picking result is taken for label placement. (d) The exact labeling positions of L1 with just the first-hit approach. (e) The exact labeling positions of L1 with the contextual picking approach.

ACKNOWLEDGEMENTS

The work presented in this paper has been funded by AGFA HealthCare in the scope of the *DiagVis* project. We thank Rainer Wegenkittl, Lukas Mroz and Matej Mlejnek (AGFA HealthCare) for their collaboration and for providing various CT data sets. Additional data sets are courtesy of the OsiriX Foundation.

REFERENCES

- [1] C. Bauer and H. Bischof. Extracting curve skeletons from gray value images for virtual endoscopy. In *Proceedings of the 4th International Workshop on Medical Imaging and Augmented Reality 2008*, pages 393–402, 2008.
- [2] C. Bauer and H. Bischof. A novel approach for detection of tubular objects and its application to medical image analysis. In *Proceedings of the 30th DAGM Symposium on Pattern Recognition 2008*, pages 163–172, 2008.
- [3] R. N. Charette. Visualizing electronic health records with “Google-Earth for the body”. *IEEE Spectrum Online*, Jan. 2008. Available online at <http://www.spectrum.ieee.org/jan08/5854/>, September 2008.
- [4] H.-L. J. Chen, F. F. Samavati, M. C. Sousa, and J. R. Mitchell. Sketch-based volumetric seeded region growing. In *Proceedings of the Eurographics Workshop on Sketch-Based Interfaces and Modeling 2006*, pages 123–129, 2006.
- [5] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention – MICCAI 1998*, pages 130–137, 1998.
- [6] T. Götzelmann, P.-P. Vázquez, K. Hartmann, T. Germer, A. Nürnberger, and T. Strothotte. Mutual text-image queries. In *Proceedings of Spring Conference on Computer Graphics 2007*, pages 181–188, 2007.
- [7] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. LiveSync: Deformed viewing spheres for knowledge-based navigation. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1544–1551, 2007.
- [8] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. LiveSync++: Enhancements of an interaction metaphor. In *Proceedings of Graphics Interface 2008*, pages 81–88, 2008.
- [9] T. Kovács, P. C. Cattin, H. Alkadhi, S. Wildermuth, and G. Székely. Automatic segmentation of the vessel lumen from 3D CTA images of aortic dissection. In *Proceedings of Bildverarbeitung für die Medizin 2006*, pages 161–165, 2006.
- [10] M. M. Malik, T. Möller, and M. E. Gröller. Feature peeling. In *Proceedings of Graphics Interface 2007*, pages 273–280, 2007.
- [11] S. Owada, F. Nielsen, and T. Igarashi. Volume catcher. In *Proceedings of the ACM Symposium on Interactive 3D Graphics and Games 2005*, pages 111–116, 2005.
- [12] T. Ropinski, J. Pražni, F. Steinicke, and K. Hinrichs. Stroke-based transfer function design. In *Proceedings of the IEEE/EG Symposium on Volume and Point-Based Graphics 2008*, pages 41–48, 2008.
- [13] Y. Sato, C.-F. Westin, A. Bhalerao, S. Nakajima, N. Shiraga, S. Tamura, and R. Kikinis. Tissue classification based on 3D local intensity structures for volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 6(2):160–180, 2000.
- [14] H. Tek, D. Comaniciu, and J. P. Williams. Vessel detection by mean-shift based ray propagation. In *Proceedings of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis 2001*, pages 228–235, 2001.
- [15] The JDOM API Project Website. Available online at <http://www.jdom.org>, September 2008.
- [16] The National Electrical Manufacturers Association (NEMA). The DICOM Standard. Available online at <http://medical.nema.org/>, September 2008.
- [17] C. Tietjen, B. Meyer, S. Schlechtweg, B. Preim, I. Hertel, and G. Strauß. Enhancing slice-based visualizations of medical volume data. In *Proceedings of IEEE/Eurographics Symposium on Visualization 2006*, pages 123–130, 2006.
- [18] C. Tietjen, K. Mühler, F. Ritter, O. Konrad, M. Hindennach, and B. Preim. METK - The medical exploration toolkit. In *Proceedings of Bildverarbeitung für die Medizin 2008*, pages 407–411, 2008.
- [19] C. Tietjen, B. Preim, I. Hertel, and G. Strauß. A software-assistant for pre-operative planning and visualization of neck dissections. In *CURAC 2006*, pages 176–177, 2006.
- [20] J. Tschirren, E. A. Hoffman, G. McLennan, and M. Sonka. Intrathoracic airway trees: Segmentation and airway morphology analysis from low-dose CT scans. *IEEE Transactions on Medical Imaging*, 24(12):1529–1539, 2005.
- [21] J. C. Xu and J. L. Prince. Gradient vector flow: A new external force for snakes. In *Proceedings of the Conference on Computer Vision and Pattern Recognition 1997*, pages 66–71, 1997.