

Smart Super Views — A Knowledge-Assisted Interface for Medical Visualization

Gabriel Mistelbauer*
Vienna University of
Technology, Austria

Arnold Köchl¶
Kaiser-Franz-Josef Hospital
Vienna, Austria

Hamed Bouzari†
Austrian Academy
of Sciences

Stefan Bruckner||
Vienna University of
Technology, Austria

Rüdiger Schernthaner‡
Medical University
of Vienna, Austria

Milos Sramek**
Austrian Academy
of Sciences

Ivan Baclija§
Kaiser-Franz-Josef Hospital
Vienna, Austria

Meister Eduard Gröller††
Vienna University of
Technology, Austria




Figure 1: Smart Super Views: The relevance of the views is mapped to a super-elliptical shape and decreases from left to right. The subsequent views are shown, starting from the left: The oblique slice view following the selected vessel in green, the coronal, sagittal and axial slice views in white, the bone view in blue, the tissue view in yellow and the vessel view in red. The border colors additionally distinguish the different views.

ABSTRACT

Due to the ever growing volume of acquired data and information, users have to be constantly aware of the methods for their exploration and for interaction. Of these, not each might be applicable to the data at hand or might reveal the desired result. Owing to this, innovations may be used inappropriately and users may become skeptical. In this paper we propose a knowledge-assisted interface for medical visualization, which reduces the necessary effort to use new visualization methods, by providing only the most relevant ones in a smart way. Consequently, we are able to expand such a system with innovations without the users to worry about when, where, and especially how they may or should use them. We present an application of our system in the medical domain and give qualitative feedback from domain experts.

Keywords: Visualization, Fuzzy Logic, Interaction.

Index Terms: H.5.2 [Information Interfaces and Presentation]: User Interfaces—Interaction styles; I.2.3 [Artificial Intelligence]: Deduction and Theorem Proving—Uncertainty, “fuzzy”, and probabilistic reasoning; I.3.6 [Computer Graphics]: Methodology and Techniques—Interaction techniques

1 INTRODUCTION

With progress in data acquisition modalities, the amount of information is increasing and more versatile data is available. In the medical domain, these data have to be processed by clinicians,

leading to a high workload and long reporting times. Various approaches try to reduce the number of images to inspect, but with advancing technology new modalities are likely to come and lead to even more images. Hence, the information must be structured and brought to clinicians in a meaningful way.

Nowadays, there are several different kinds of visualization areas, each with its own set of techniques suited for specific purposes. In the field of radiology, established methods such as Maximum Intensity Projection (MIP) are widely applied and can be, for example, extended by Maximum Intensity Difference Accumulation (MIDA) [4]. Specific methods, such as Curved Planar Reformation (CPR) [8] or Multipath Curved Planar Reformation (mpCPR) [9], are tailored to the investigation of calcifications on vessel walls. Clinicians must be aware of which technique fits to which region in the data and to which purpose. For example, it may not be appropriate to use MIP for the assessment of a stent treatment, because a vessel can appear completely blocked even if it is just calcified at the vessel wall. However, blood might still be able to flow, which could be revealed when inspecting a slice or a CPR. This possible inadequate use should be avoided. It becomes more likely with an increasing number of available visualization methods, because not every physician might be aware of their intended application areas. Clinicians may not know all possible pitfalls, potentially leading to negative consequences.

Common user interfaces exhibit several problems. For example, menus can become very long and cluttered and therefore it is often difficult to find the desired option. Toolbars can contain too many buttons and although one is able to adjust them manually, this is quite cumbersome. Popup menus try to focus on the most important option, but suffer from the same drawbacks if overloaded. With increasing information and possibilities it has become more difficult to choose the most relevant options. Furthermore, interfaces may consist of a lot of technical terms and expressions not everyone is aware of.

We propose a concept, where an image acts as a menu itself. The visualization becomes the main user interface, by augmenting it with dynamically generated integrated views. Once a user picks a Region-of-Interest (ROI), the most suitable visualization techniques are determined. We call this approach *knowledge assisted*

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

†e-mail: hamed.bouzari@oeaw.ac.at

‡e-mail: ruediger.schernthaner@meduniwien.ac.at

§e-mail: ivan.baclija@wienkav.at

¶e-mail: a.koechl@wienkav.at

||e-mail: bruckner@cg.tuwien.ac.at

**e-mail: milos.sramek@oeaw.ac.at

††e-mail: groeller@cg.tuwien.ac.at




Figure 2: Various aspects of our contribution.

sparse interaction. The subsequent section gives an overview of related work concerning visualization, interaction and knowledge representation. Our concept is outlined in Section 3 and detailed in Section 4. Visual mapping and interaction are discussed in Section 5 and results are presented in Section 6. Implementation details are provided in Section 7. Qualitative domain expert feedback is given in Section 8 and discussed in Section 9. The paper is concluded in Section 10.

2 RELATED WORK

Our approach aims to simplify the interaction with medical visualization systems by including additional knowledge. We incorporate aspects of several different research directions in our approach (Fig. 2).

Medicine/Radiology. Ota et al. [20] give a comprehensive evaluation concerning the reliability of multi-detector Computed Tomography Angiography (CTA) in comparison to Digital Subtraction Angiography (DSA) as gold standard, by assessing the importance of observing axial images. They conclude that multi-detector CTA is a reasonable alternative to common diagnostic techniques for patients with lower extremity arterial occlusive diseases. Portugaller et al. [21] analyze different visualization techniques for quantitative lesion assessment in lower extremity arteries with an area reduction of $\geq 70\%$, such as stenoses or occlusions. The result of their study shows that MIP together with axial images is most accurate in detecting cross-sectional lesions. Northam et al. [19] assess in their study the reduction of the field-of-view from the whole data set to a specific region, in order to increase the spatial resolution. However, possible pathologies outside this limited view will be missed. Schertler et al. [30] demonstrate in their study the significant impact of non-vascular diagnosis even when using vessel CPRs. They show that axial images must always be taken into consideration. It is not enough to use CPR alone, as 27% of pathologies would remain unrecognized.

Multi-View Visualization. Visualization applications frequently make use of multiple linked views to simultaneously depict different representations of the data. The concept of linking and brushing connects these views by interactively highlighting selections in all views [27]. We also draw inspiration from superviews, introduced by Motro [17], which provide a homogeneous view while accessing multiple databases. Any interaction with the superview is decomposed into queries for each individual database. The outcomes are then assembled to provide the result for the initial query. Bier et al. [2] introduced toolglasses and magic lenses, which enable the depiction of alternative data representations in interactive see-through widgets. In the context of volume visualization, different types of magic lenses have been explored by LaMar et al. [14] and Wang et al. [36]. Inspired by traditional illustrations, Bruck-

ner and Gröller [3] presented integrated contextual views in their VolumeShop system. Based on this concept, Taerum et al. [32] used contextual close-ups to present high-resolution sub-volumes of medical volume data. Ropinski et al. [29] proposed interactive close-ups for the visualization of multimodal data. Their work also presented a layout algorithm for the placement of these views. Balabanian et al. [1] describe integrated views in a graph layout in order to navigate through a volume hierarchy. They combine various types of interactions and visualizations in order to explore different regions of interest. Smart views, as presented by Radloff et al. [23], provide a refined view management in a multi-display environment. Recent work by Steinberger et al. [31] visually links selections in multiple applications to preserve context. In our work, we not only focus on how alternative views are presented, but also which views are best suited for different parts of the data.

Smart Navigation Techniques. As orientation can be difficult when investigating three-dimensional data, several approaches have been presented to simplify navigation. McGuffin et al. [16] used deformations and a set of interactive widgets allowing users to virtually browse through different tissues of a medical volume data set. Tietjen et al. [33] presented LiftCharts as a simple tool for easing navigation when interacting with slice views of segmented medical data. Viola et al. [35] discussed a method for retaining context while refocusing on different structures of interest in a volume data set. Kohlmann et al. [11, 12] proposed LiveSync, a method for interactively synchronizing a 3D view with 2D slice views. They automatically determine a good 3D view for a position selected on a 2D slice using local data properties. In further work, they also describe how to determine a 2D slice position from a single position on a 3D rendered image in different scenarios by matching ray profiles against a database [13]. Diepenbrock et al. [6] used a fish-eye view in combination with preview images to ease navigation in virtual fly-throughs.

Knowledge-Assisted Visualization. The aim of knowledge-assisted visualization approaches is to improve visual exploration and analysis by utilizing information about the visualization process itself (e.g., a user’s chosen visualization parameters and abstractions), and information about the scientific data to be visualized (e.g., high level abstract characterization and findings) [5, 37]. Tzeng et al. [34] use a painting metaphor to classify and visualize volumetric data. They feed the extracted information to an intelligent system, using a neural network or a support vector machine, in order to apply the classification to a new data set in a similar region. Rezk-Salama et al. [26] incorporated expert knowledge to develop a high-level interface for transfer function design. Their approach used a principal component analysis to extract high-level parameters from a set of user-defined transfer functions. The work of Nam et al. [18] proposed a method for extracting and indexing features to enable the automatic categorization of volume data sets. A knowledge-assisted system for the analysis of geological data was presented by Kadlec et al. [7]. Their approach captures the properties of seismic features to improve segmentation. One way to formalize domain knowledge is fuzzy logic [38, 40], which enables the specification of rules and relations using linguistic variables. Rautek et al. [24], for example, used a fuzzy rule base to map volumetric attributes to visual properties. They also extended their approach to incorporate interaction [25]. We also employ a fuzzy logic rule base, which enables us to represent the suitability of certain types of visualizations for different spatial regions based on domain expert knowledge.

3 OVERVIEW

We propose a smart user interface that intuitively presents only relevant options in contrast to, for example, a list view showing all possibilities. We are motivated by the fact that clinicians often base their reporting on checking a large number of images. We enhance

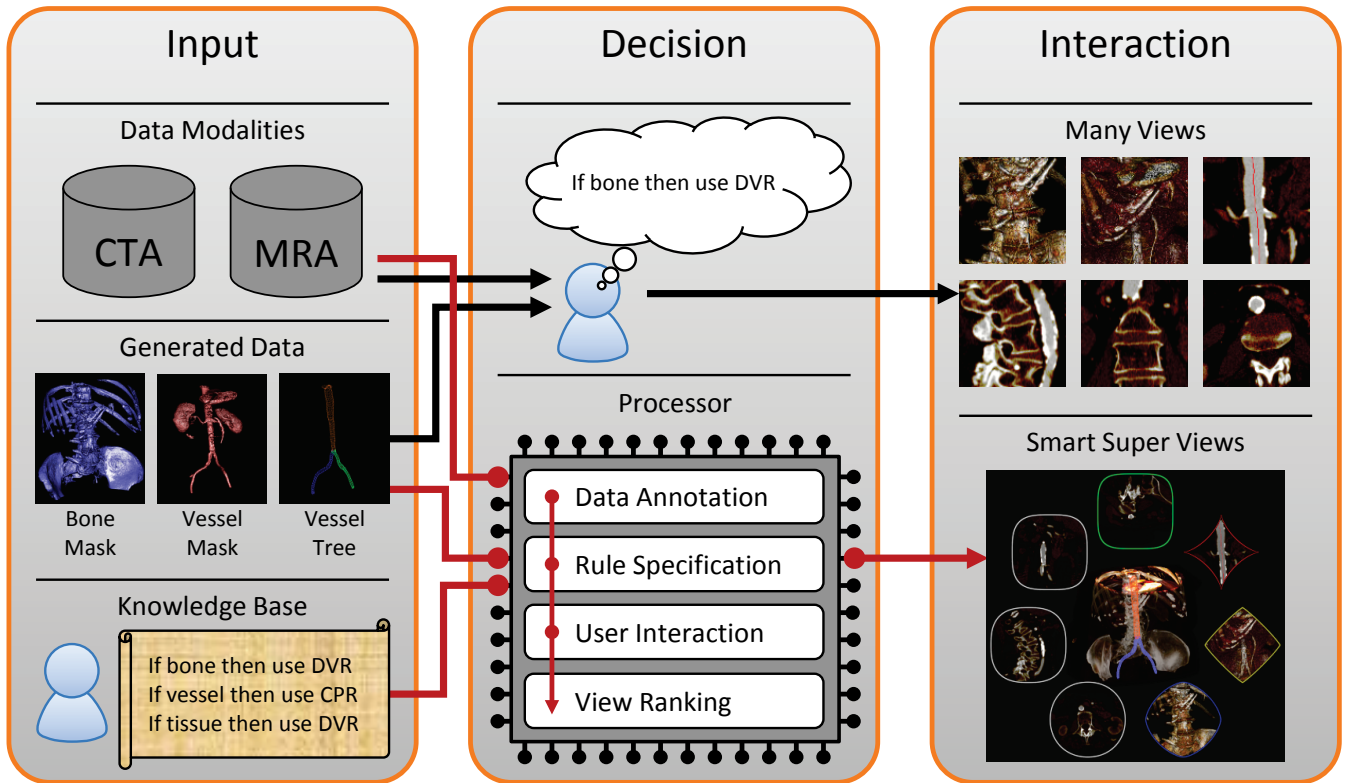


Figure 3: Usual workflow of medical reporting (black arrows). The clinician has to know which method is applicable to which region of the data and for which purpose it has been designed. Our system (red arrows) supports clinicians by formalizing this expert knowledge as rules in the input step. A smart processor decides for a user-specified ROI which visualizations are suited and should be presented. The clinician has then the possibility to interact with those suggested views in order to inspect the ROI in more detail.

our views with knowledge in order to provide the user only relevant visualizations for interaction. Fig. 3 presents the common workflow of medical reporting (black arrows) and shows how our proposed system integrates therein (red arrows). We distinguish three steps and each will be subsequently outlined briefly. The most important and most notable difference between the conventional and our approach is that the user no longer needs to be aware which method has to be used in a specific region, because this knowledge is provided in a rule base defined by domain experts.

Information is acquired and gathered from different sources in the **first step**. In addition to a volume data set, our application offers as input a bone mask, a vessel mask and a vessel tree. The bone and vessel segmentation is usually done by radiological assistants in order to generate the respective masks. First, they create a vessel tree by manually tracking vessels and specifying centerlines and radii accordingly. These vessels consist of segments spanned between branchings or endings, as shown with different colors in the vessel tree image in the generated data in Fig. 3. This vessel information is used to differentiate bones from vessels and separate them. Since the automatic process might not produce a perfect separation, manual fine tuning is done afterwards. This step is rather time consuming, but is done for every patient and therefore we can use this additional information. The most significant difference compared to the conventional approach is that a domain expert has to define a set of rules, which maps the input to specific output. This offers a flexible solution of extending the output by incorporating new rules and visualizations.

Based on all the input information, a user decides, in the **second step**, where a ROI is located and which specific visualization technique is applicable there. This requires significant expertise

from the user, because he or she must be aware of the algorithmic internals and the intent. Our proposed approach assists the user by employing an inference engine that processes the input together with a knowledge base. If the user selects a ROI, the most suitable visualization techniques are suggested automatically.

In the **third step**, the user interacts with the provided visualizations and inspects the ROI in more detail. In contrast to the common scenario, where the user must choose a visualization technique for interaction, our system provides support by suggesting suitable choices. Additionally, interaction is reduced, because only a small set of relevant visualizations is suggested in each specific case. Among these relevant techniques for medical diagnosis are CPR, MIDA, MIP and Direct Volume Rendering (DVR) as well as axial, coronal, sagittal slices and an oblique slice following the currently selected vessel tree segment.

4 SMART VIEW INFERENCE

When the user is browsing through the three-dimensional data, our system determines the most relevant visualization techniques in the corresponding region. The processor shown in Fig. 3 outlines how this is achieved by using four modules. Fig. 4 gives a more detailed view of these modules, where the general workflow is top-down following the black arrows. The thinner black lines illustrate the dependency graph between the input and the output.

The first module is the data annotation, where semantic layers are extracted from various sources of input. Rules are specified or fetched from an external source, like a database, in the second module. These rules define the connection between the input and output for a specific region the user is currently exploring. The third module bundles the information of the semantic layers according to the

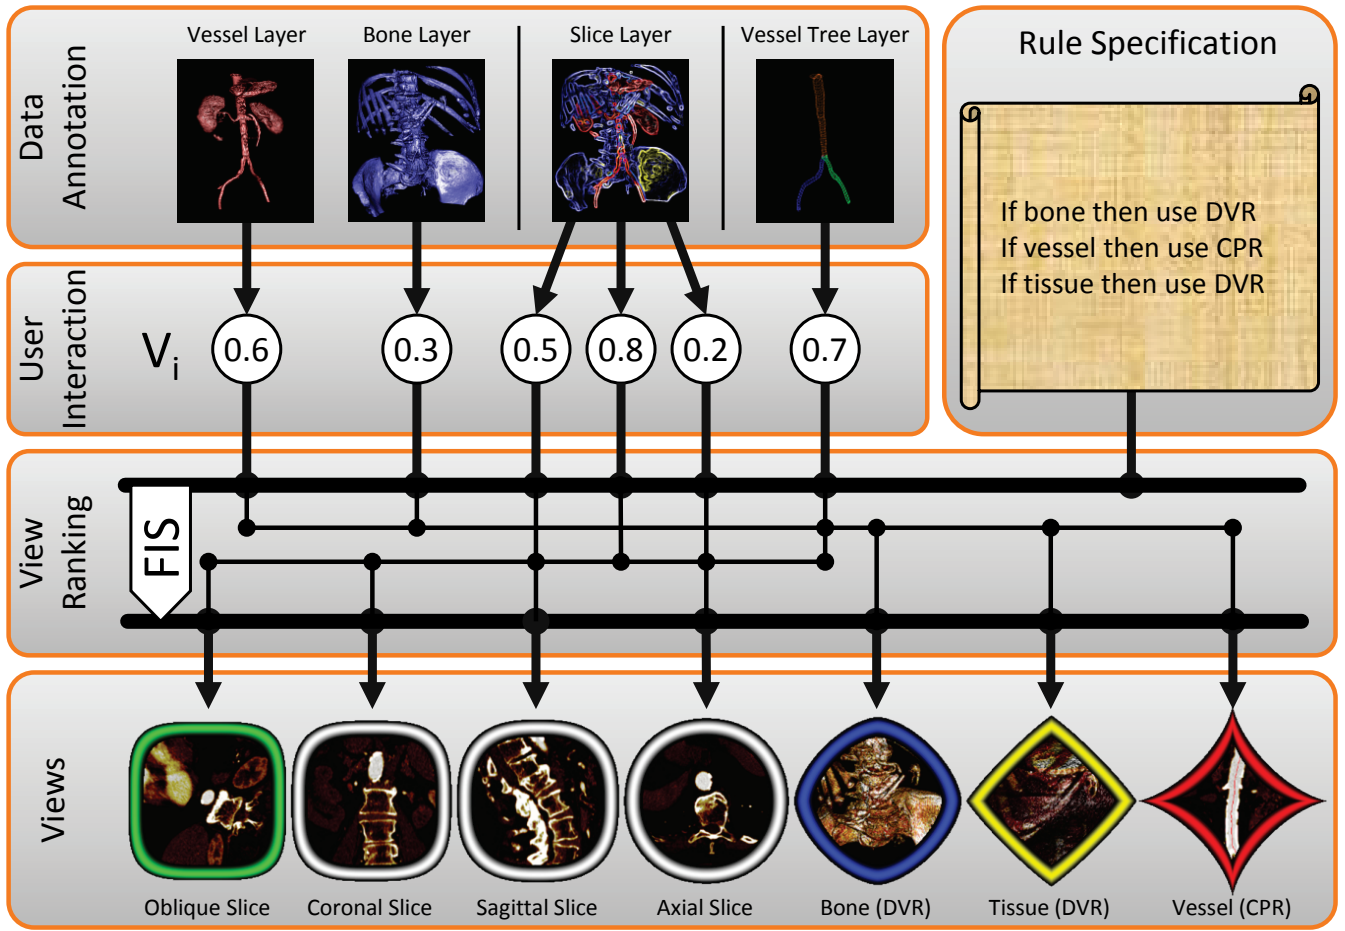


Figure 4: Workflow (black arrows) of the processor consisting of four modules. First, semantic layers are extracted by data annotation and, second, a rule base is defined or loaded. In the third module, the information of every semantic layer is condensed to a single value, according to user interaction. In the next module these values are fed to the Fuzzy Inference System (FIS) where they are fuzzified in order to evaluate the rules. The results of these rules are aggregated and used to determine a ranking of suitable visualizations. The two thick lines in the view ranking module indicate the input and output respectively. Once an output is derived and defuzzified, the resulting views are sorted. They exhibit a super-elliptical shape with their relevance mapped to the exponent of the super-quadratics definition.

user interaction and feeds it to the fourth and final module, the view ranking. In this module, a Fuzzy Inference System (FIS) combines the condensed information using fuzzy logic and provides a ranking of important views only.

4.1 Data Annotation

The goal of this module is to extract meaningful information in the form of semantic layers from the raw input data. Examples are given in the images in the data annotation of Fig. 4. The input can either originate from data acquisition modalities such as CTA or Magnetic Resonance Angiography (MRA), or can be manually specified, like the bone and vessel masks, by, for example, radiological assistants. Since these semantic layers (L_{Bone} , L_{Vessel} , L_{Slice} and $L_{VesselTree}$) are the basis of all subsequent operations, they should be determined carefully.

Bone & Vessel Layers. For bones and vessels, we need to know the depth value. Therefore we use ray-casting and determine the first hit with the segmentation masks of bones and vessels during rendering in real-time. These masks are provided by radiological assistants in a semi-automatic procedure. We store the depth values of the bone and vessel masks in separate semantic layers.

Slice Layer. We have two categories of slices. First, the axis aligned slices such as the *axial*, *sagittal* and *coronal* slices, and,

second, the *oblique* slices that move along selected vessel tree segments. In the case of the axis aligned slices, we use ray-casting and estimate the importance of a slice along a ray. This is done in real-time by determining the first hit of a ray with an object, either vessel or bone. Then, the dot product between the corresponding slice axis and the gradient vector of the hit position determines the importance. If the dot product equals zero, the importance is highest, because the gradient vector lies within the slice plane. Otherwise the importance decreases with increasing dot product until it reaches zero in case the gradient vector coincides with the slice axis. The semantic layer for the axis aligned slices stores a 3D vector for every pixel, where each coordinate depicts the maximum importance of the corresponding slice along the ray. The (x,y,z) coordinates represent the importance of the sagittal, coronal and axial slices respectively.

Vessel Tree Layer. In order to determine the user-selected vessel segment (the orange segment of the vessel tree of the data annotation module in Fig. 4) color picking is performed. The vessel tree is rendered by a mesh renderer where color encodes the segments. By reading back the color at the pixel under the cursor in the current frame buffer, the segment can be identified. The semantic layer for the vessel tree is a binary mask that indicates if a segment was hit or not. It is not important where the vessel is exactly located spatially,

but only if the user hovers over a vessel or not. This is the reason why the spatial information is not used.

4.2 Rule Specification

The rule specification module concerns the definition of the connections between the input data and the output visualizations. In order to model these connections we use a collection of *if-then* clauses and store these rules in a human readable form in an external file. The rules are defined by domain experts in order to provide specific mappings of input to output. These mappings should reflect, of course, the clinical protocol of their daily routine.

The antecedents of the rules can contain several logical combinations of the input semantic layers by using logic operations such as *and*, *or* and *not*. The consequent of every rule consists of a set of applicable visualizations. Linguistic variables, such as *bone*, *vessel* or *tissue*, are used to define the rules in a human readable form. This makes it easier to extend and maintain the system by non computer scientists. Due to the fact that our linguistic variables do not exhibit a binary state, either true or false, we use fuzzy logic [38] to process them accordingly.

This externally stored domain knowledge offers the possibility of incorporating changes easily, by simply adding or removing rules. Another advantage is the exchange of this knowledge in order to provide different modes for several application areas, for example, interdisciplinary medical discussion rounds between radiologists, clinicians and surgeons.

4.3 User Interaction

If the user interacts and requests suggestions for a specific region in the data, the proposed visualization techniques should correspond to the data contained therein. So far, we have the semantic layers and a set of rules that define the input to output mapping. What we still need to do is to determine the input for the rules themselves. This input is a value that is subsequently fuzzified in order to be used in the evaluation of the antecedents of the rules. To distinguish between non-fuzzy and fuzzy values, the former ones are referred as crisp values.

Region-of-Interest (ROI). The user defines a 2D Region-of-Interest by the cursor position on the screen and a specified radius. Taking a small neighborhood into account avoids high variational changes of the suggested smart views due to data coherence.

The information of the semantic layers is used inside the ROI to derive a single crisp input value in order to evaluate the rules. However, not every value of the semantic layers inside this region should have the same influence. The main focus should be at the cursor position. Thus, the crisp scalar input values, V_i , for every linguistic input variable are computed over the ROI by

$$V_i = \frac{1}{N} \cdot \sum_{(x,y) \in ROI} [C_i(x,y) \cdot S(x,y)] \quad (1)$$

where (x,y) is a pixel inside the ROI, $C_i(x,y)$ is a conditional term for modeling correspondence of multiple semantic layers, $S(x,y)$ is a spatial influence function and N is the normalization factor, the area of the region in pixels. The following equations show, how the conditional term is computed at a pixel $p(x,y)$ using the information of the different semantic layers (L_{Bone} , L_{Vessel} , L_{Slice} and $L_{VesselTree}$).

$$C_{Bone}(x,y) = \begin{cases} 1, & \text{if } L_{Bone}(x,y) \leq L_{Vessel}(x,y) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

$$C_{Vessel}(x,y) = \begin{cases} 1, & \text{if } L_{Vessel}(x,y) \leq L_{Bone}(x,y) \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

$$C_{SagittalSlice}(x,y) = L_{Slice}(x,y) \cdot x \in [0, 1] \quad (4)$$

$$C_{CoronalSlice}(x,y) = L_{Slice}(x,y) \cdot y \in [0, 1] \quad (5)$$

$$C_{AxialSlice}(x,y) = L_{Slice}(x,y) \cdot z \in [0, 1] \quad (6)$$

$$C_{VesselTree}(x,y) = L_{VesselTree}(x,y), 0 \vee 1 \quad (7)$$

The value for bone, for example, needs the information of the vessel and bone semantic layers, because both store the depth of the respective masks. C_{Bone} will be one, only if the depth value in its semantic layer is less than the one in the vessel semantic layer at the same pixel position. In the case of the slices, only the value for the corresponding axis is taken. For the vessel tree, either one or zero is returned, because only a binary mask is stored. This mask indicates if a vessel is located under the current pixel.

Spatial Influence. Not every position inside the ROI should have the same influence on the result. Information close to the cursor position should contribute more. For example, when the user hovers over the aorta, the ROI might contain some bone, such as the spine. If the focus is on the aorta, then the bone should have less influence than the vessel and this should be reflected in the output visualizations. We model this by computing a weight for every pixel $p(x,y)$ of the ROI, with cursor position \mathbf{c} , according to the following function

$$S(x,y) = \frac{1}{(1 + \frac{\lambda}{r} \|\mathbf{p} - \mathbf{c}\|_2)^\sigma} \quad (8)$$

with r being the radius of the ROI in pixels, σ controlling the speed of the influence decay and λ providing the weight at the boundary pixels of the region. The influence decreases with distance from the center. Experiments have shown that using $\sigma = 2$, $\lambda = 0.3$ together with $r = 10$ provides a good response, where the suggested views reflect the underlying data.

4.4 View Ranking

Once the crisp input values V_i (see Equation 1) have been determined for all linguistic input variables, they are fed to the inference engine together with the rules. Since we use fuzzy logic, our inference engine is a Fuzzy Inference System (FIS) as outlined in Fig. 5.

Fuzzification. The crisp input values are fuzzified using fuzzy membership functions. Such functions could have a triangular, trapezoidal or Gaussian shape. Throughout this paper, triangular membership functions are used.

Implication Evaluation. Fuzzy logic offers the possibility to specify knowledge in the form of linguistic rules and terms rather than with exact numbers [39]. For the specification of the rules,



Figure 5: Workflow of our FIS. First, crisp input values are fuzzified using fuzzy membership functions. Second, the antecedents of the *if-then* rules are evaluated using fuzzy set operators. The implications are evaluated with the minimum operator [15]. Third, the fuzzy implication results are aggregated and, fourth and finally, defuzzified using the centroid method to provide a crisp output.

we use the concept of the *if-then* rule with fuzzy antecedents and consequents. For every rule, its antecedent is evaluated using fuzzy set operators and the consequent is computed using the minimum operator [15].

Aggregation. Once all fuzzy consequents are obtained, they are aggregated by combining their fuzzy membership functions. Again, fuzzy set operators are used to aggregate the fuzzy consequents. The result of the aggregation step is one fuzzy output.

Defuzzification. In order to obtain a crisp output value from the aggregated fuzzy one, defuzzification is needed. Common defuzzification approaches are the centroid method, maxima decomposition, center of maxima calculation or height defuzzification. In all subsequent examples of our work, the centroid technique is used.

The final crisp output value is used to determine the membership value of every linguistic output variable. This membership value is the relevance of the respective suggested visualization. Subsequently, the proposed views are sorted according to their relevance and presented to the user.

5 VISUAL MAPPING AND INTERACTION

The output of the smart view inference is a ranked list of suggested visualizations that need to be presented to the user in a suitable way. Every visualization will be shown in a separate view, embedded into the current user interface, analogue to a Head-up Display (HUD).

A HUD is a transparent display presenting data where the viewer does not need to look away from the usual focus area. It was originally designed for pilots in order to alleviate them from looking down on their instruments. It allows them to look forward with the head positioned up. Hence the name Head-up Display. We use this analogy to present the relevant views embedded into the visualization the user is interacting with. This avoids interaction with other windows or menus and the focus remains in the same window. While the user is still able to see the ROI, the original view is enriched by additional so called *smart views*. The term *smart view* has been chosen to reflect that the views have been determined in a knowledge-assisted, i.e., smart, way.

5.1 Visual Mapping

In order to visually convey relevance, we incorporate it into the shape of a view. Several properties of glyphs, such as shape, size, color and orientation are described by Ropinski and Preim [28] with regard to perception. They state that shape and size of glyphs are perceived *pre-attentively*, within the first quarter of a second. Hence, we propose to apply properties of glyphs to our smart views as well. Concerning different shapes, Kindlmann [10] states, that super-shapes convey information more clearly and precisely. Therefore, we use a super-elliptical shape for our smart views and additionally color the boundary of the shapes differently. Simultaneously, the relevance is encoded in the size of the area covered by the shapes on the screen. This is intended to allow an unambiguous identification of the most relevant views. Combining all these aspects, our views are suggested according to some underlying knowledge and exhibit a super-elliptical form, hence the name *smart super views*.

A super-ellipse is defined as the set of points (x, y) with

$$\left| \frac{x}{a} \right|^n + \left| \frac{y}{b} \right|^n = 1 \quad (9)$$

where $a, b > 0$. Since differences in shape and size are perceived pre-attentively [28], we map the relevance of a visualization to the exponent n of the super-ellipse. This leads to the following equation for our views

$$\left| \frac{x - x_{pos}}{size} \right|^{m(rel)} + \left| \frac{y - y_{pos}}{size} \right|^{m(rel)} \leq 1 \quad (10)$$

where *size* is the chosen size of the view, $m(rel)$ is the relevance shape-mapping function with *rel* being the relevance of every suggested view and (x_{pos}, y_{pos}) is the position of the view. We empirically determined mapping function $m(rel)$ as

$$m(rel) = (0.5 + 2 \cdot rel)^\alpha \quad (11)$$

where α controls how fast the shapes of the suggested views become rectangular. Using $\alpha = 1.5$ covers most of the dominant shapes, i.e., rectangles, circles and stars. Additionally, we omit very thin star shapes by shifting with 0.5 and scaling the relevance by two. Due to the α exponent, highly relevant views will have a rectangular shape, whereas the least important ones will have a star-like form. With this mapping, shape and area vary related to the relevance obtained from the view ranking.

Spatial Arrangement. Once the shape of the views is determined to reflect their relevance, they need to be spatially arranged. We have implemented two possible layout strategies:

- **Linear Layout.** All views are arranged along a horizontal line, starting with the most relevant view at the left and proceeding to the right with decreasing relevance (Fig. 6). In order to avoid overlapping with the ROI, they are positioned at the bottom of the view they are embedded in.
- **Radial Layout.** The views are radially positioned around the ROI. They are placed in a counter clockwise fashion. The most relevant view is placed above the ROI. This avoids visual clutter in the case of many views. Fig. 7 shows an example of views with varying relevance and their corresponding arrangement.

5.2 Interaction

The user can hover over a specific smart super view and interact with it in the same way as with the overview visualization they are integrated in. No manual switching between the different views is required. This makes the interaction sparse and allows the user to immediately get more insight into a specific region. We distinguish between two states of interaction. In the overview state, the user navigates in the overview visualization and identifies a ROI. In order to explore this region in detail, the user can switch to the inspection state. For every interaction in this state, the user gets immediate feedback in the form of smart super views. If a view is selected by hovering over it, the user is able to inspect the visualization of the view by, for example, zooming, panning and rotating. Additionally, the size and border of a selected view are enlarged in order to give the user a visual feedback. When finished with inspection, pressing a key will return to the overview state and all smart super views will disappear.

Legend. In order to provide an overview of the entire range of the relevance, a legend of possible shapes is presented. This conveys the relevance visually and provides a context. If the user selects a smart super view, the corresponding glyph is indicated in the legend at the position of the respective relevance. Additionally, a connection band between the glyph in the legend, the view and the ROI is shown. Following this band, one will quickly perceive the overall layout and arrangement. Fig. 6 shows smart super views arranged in a linear layout, together with the legend on the left. The least relevant view (tissue) with the yellow border is selected and following its connection band, one can clearly see the corresponding relevance. The legend shows high relevance at the top and the least relevance at the bottom. In order to see the position of the ROI better, horizontal and vertical rulers are drawn (Fig. 6).

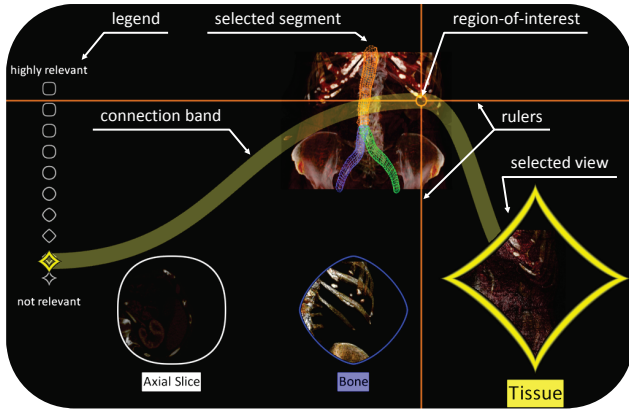


Figure 6: Smart super views arranged in a linear layout at the bottom of the window with the overview visualization being MIP. The shape of the views changes with decreasing relevance from left to right. Tissue (yellow) is the least relevant view. It is selected and therefore enlarged. The yellow band connects the selected view with the ROI and the corresponding glyph in the legend.

6 RESULTS

Fig. 1 presents an overview of possible views in our system. The relevance decreases from left to right. In this example, the most relevant view is the oblique slice view that is followed by the coronal, sagittal and axial slice views. Then, bone using DVR precedes tissue using DVR and the last view is the vessel view showing a CPR. High relevance leads to an almost rectangular shape, whereas the least important view will look like a star. The border colors additionally distinguish the various views, whereas the coronal, sagittal and axial slice views share the same color.

Fig. 7 shows a data set rendered with MIP together with the vessel tree rendered in wire-frame mode. When hovering over the aorta, it will be displayed in the suggested views, as shown in the vessel view on the bottom of the image with the red border. This smart super view offers the user the possibility to rotate the selected part of the aorta using CPR. In this scenario, tissue is ranked the least relevant view. This can be seen from the star-like shape of the view. The other two views with the white border are slices with almost identical relevance, reflected in their circular shape.

Fig. 8 gives an example where the user is moving over a stent. Because MIP is the underlying overview visualization, it initially seems that the stent completely blocks the vessel. However, when inspecting the vessel in the corresponding view with the red border using CPR, it is found out that this is actually not the case and the blood is still able to flow.

Fig. 9 shows several smart super views for different regions of a human lower extremity data set using MIP as overview visualization. If the user moves over the aorta, slice, vessel and tissue views are suggested, as shown in Fig. 9(a). These views are placed around the ROI in a radial layout and the sagittal slice is selected. The user can interact with this view by, for example, scrolling through the slices or changing the windowing function. Another suggestion is given in Fig. 9(b), where the focus is on the pelvic bone. The bone view (blue border) is selected and has the highest relevance, followed by an axial and sagittal slice view. Only the bones defined by the bone mask are shown with DVR. Next, moving to the knee region, the vessel view and the selected oblique slice view (green border) are shown in Fig. 9(c), because the user moves over a vessel. If moving a bit away from the vessel, other types of tissue will be taken into account. This leads to the example in Fig. 9(d), where tissue and slice views are suggested. The tissue view (yellow) is selected and the knee is inspected more precisely using DVR.




Figure 7: Smart super views presented in a radial layout around the ROI, while moving over the aorta of a human lower extremities data set. The overview visualization is MIP.




Figure 8: Detailed investigation of a stent in the vessel view (red border) using CPR. The user is able to rotate the vessel and check whether the blood can still flow through this vessel.

7 IMPLEMENTATION

We implemented our smart super view system into an angiography framework written in C++ and taking Qt4 for the user interface. The system provides functionality for loading data sets, defining segmentation masks and semi-automatically tracking vessels. The framework is used in the daily routine of two hospitals and we have extended it by implementing a plugin for the proposed smart super views. These are not used in the daily routine yet.

All visualizations, the overview one as well as the ones of the smart super views, are rendered into separate frame buffer objects using CUDA. The super-elliptical shapes together with the views and the shape legend are rendered using GLSL shaders. The rulers, the outline of the ROI and the connection bands are drawn using the painting capabilities provided by Qt4.

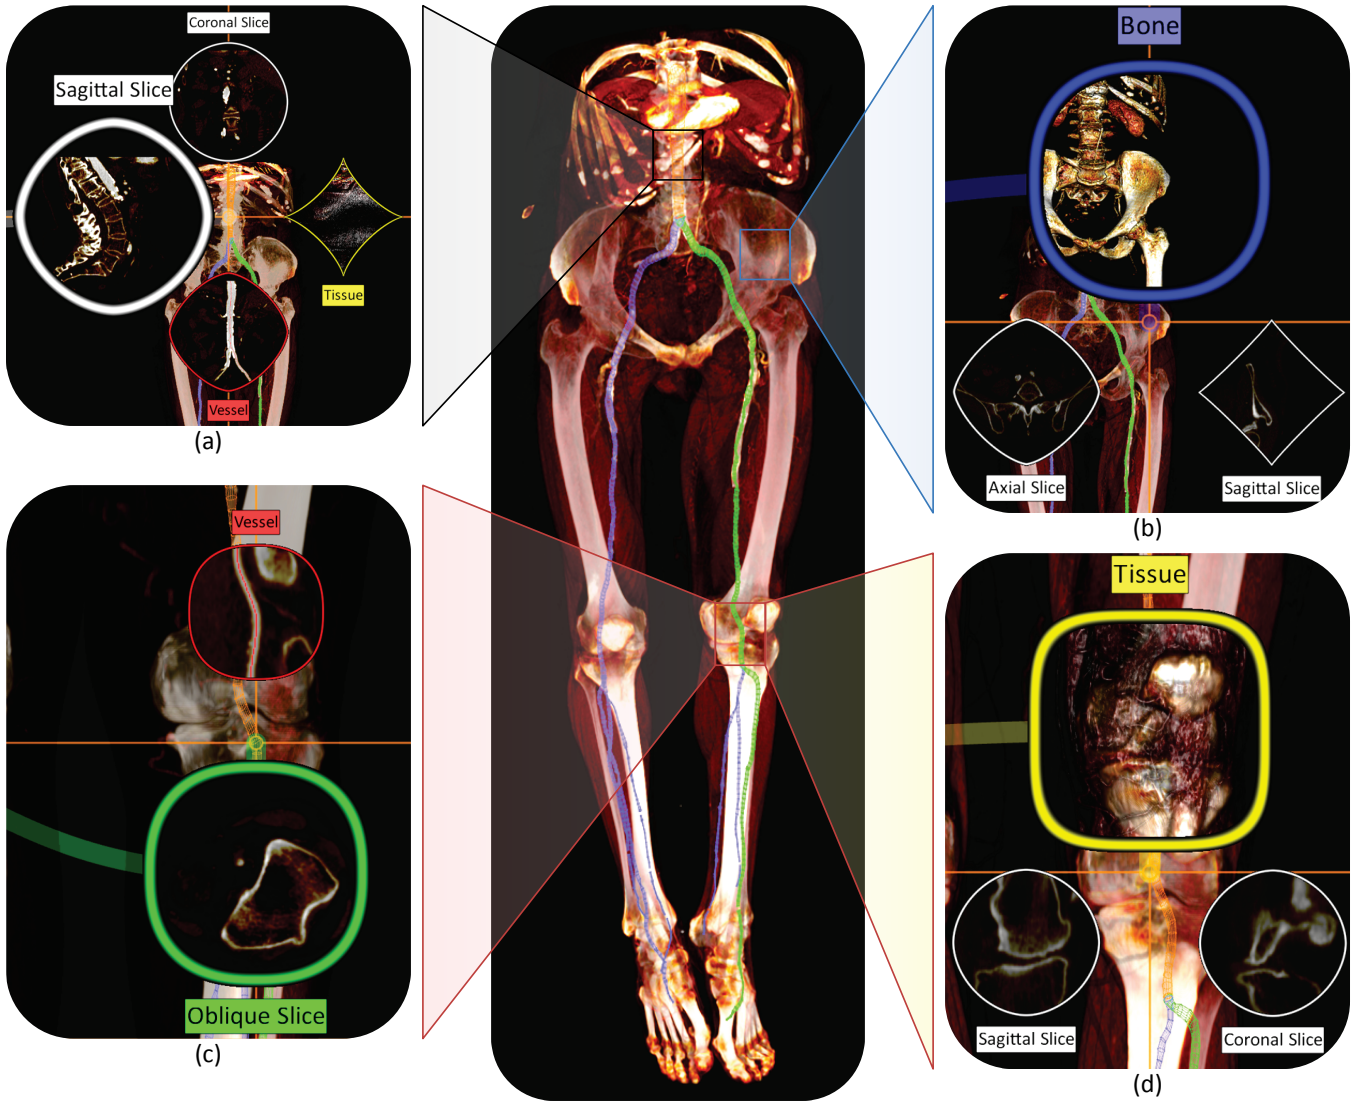


Figure 9: Several smart super views presented for various regions of a human lower extremity data set. The overview visualization is MIDA in all examples. Selected is (a) the sagittal slice view, (b) the bone view, (c) the oblique slice view and (d) the tissue view.

The user can interactively move the ROI to obtain smart super views and then inspect a specific one. The inspection can be done in real-time by rendering the active view enlarged in order to provide more space for investigation. Rotating, zooming, panning and changing the transfer and windowing functions are the default capabilities of every view.

We used the fuzzy-lite [22] library for integrating fuzzy logic into our application, because it provides an intuitive and fast C++ interface. The linguistic input and output variables as well as the membership functions are defined in the source code. The rules are stored in a text file and loaded when the smart super view plugin of our application starts.

8 EVALUATION

In discussions and semi-structured interviews with physicians, we acquired qualitative feedback about the applicability of our smart super views. We consulted two domain experts, i.e., radiologists. They told us, that the initial goal of a visual workflow is to reduce the number of slices to inspect. A common procedure is to investigate all slices of a patient data set in order to detect suspicious

regions. This works usually well, but consumes a considerable amount of time. For vessel investigation, CPR [8] and mpCPR [9] were proposed. They significantly reduce the number of images to inspect by a clinician.

With the increased resolution of current scanners, the data sets consist of more slices and the time that is required for an inspection rises. Additionally, the number of images increases as well with more techniques being available, even for specific purposes. For blood vessels such techniques are, for example, various stretched, straightened or untangled CPRs. For every technique images have to be generated according to a given set of viewing directions. The clinicians have to inspect all of these generated images, as some pathologies, such as vessel occlusions might remain unrecognized otherwise. This is a potential scenario of our proposed system, because it will automatically provide only relevant views. Additionally, the clinicians are able to view the ROI from various viewing directions, because smart super views offer more interaction than static images.

According to the study by Portugaller et al. [21] specific combinations of views lead to more accurate reporting results. They




Figure 10: Linguistic variables with their terms and corresponding triangular membership functions used in the FIS. The left and right most terms are shoulder terms. The colors correspond to those employed in the user interface. (a) Every linguistic input variable uses the terms *Low*, *Medium*, *High*. (b) The linguistic output variable with the terms for the vessel, bone and tissue views. (c) The linguistic output variable with the corresponding terms for the slice views.

conclude that using only axial images, CPR or MIP alone is not sufficiently accurate to judge many pathologies. The outcome of the study is that MIP together with an axial slice view is the most accurate procedure. Our system is able to provide this by adapting the rule base accordingly.

The domain experts also liked the possibility to additionally blend specific views into their workflow. Their workplace usually includes four displays, where three of them are used for the PACS workstation. The fourth and left most screen remains mostly unused. Here they see some potential for our proposed system as an extension of their current workstation. Having additional information on suspicious regions and getting only relevant views, was considered to be a promising extension. Concerning the discrimination of the various views, the domain experts mentioned that it is intuitive for a user with some basic knowledge about the underlying algorithms. Furthermore, the interaction is seen as fairly easy.

The domain experts see further potential of our proposed system in interdisciplinary medical discussion rounds between radiologists, clinicians and surgeons. Our system can be adapted to their different requirements by providing individual rule bases. By changing the rules, the system will immediately adapt to the desired protocol. For example, a surgeon wants to see a vessel with DVR, whereas a radiologist prefers CPR.

According to the study presented by Northam et al. [19] and as a future extension, our system might use a high resolution but limited field-of-view for the overview visualization, while additionally providing smart super views using the full field-of-view of the underlying data. This would reduce the potential risk of missing pathologies along with showing the ROI in a high resolution.

The study by Schertler et al. [30] investigates the impact of axial slices on the localization of possible pathologies beyond the vessels in the chest region. Our system can account for such cases by defining an appropriate rule base. The potential risk to overlook pathologies is reduced by guiding the radiologist to certain views in specific regions. These views use only well-known and established visualization techniques according to the application scenario.

Summarizing the evaluation, the domain experts see our system as a promising addition to their current clinical workflow. They especially see a potential in the possibility to guide the user to important views in order to reduce risk factors. A comprehensive user study with the aim to investigate the impact of our proposed system on the daily clinical routine is still open.

9 DISCUSSION

Among several possibilities of fuzzy membership functions, we chose triangular ones for every linguistic variable. All linguistic input variables exhibit the structure shown in Fig. 10(a) and consist of the terms *low*, *medium* and *high*. These terms have triangular membership functions with the same area, but *low* and *high* are shoulder

Table 1: Outline of the *if-then* rules used in our system to determine the relevance for the vessel, bone and tissue smart super views.

if V_{Vessel} is Low and V_{Bone} is High and $V_{VesselTree}$ is Low then Suggestion is Bone
if V_{Vessel} is Medium and V_{Bone} is Low and $V_{VesselTree}$ is Low then Suggestion is Tissue
if V_{Vessel} is High and V_{Bone} is High and $V_{VesselTree}$ is High then Suggestion is Vessel

terms. The linguistic output variable for the vessel, bone and tissue views has the corresponding visualizations as terms, as illustrated in Fig. 10(b). The linguistic output variable for the slices consists of four terms and is shown in Fig. 10(c). The colors are the same as in the user interface of our system. Table 1 shows an outline of the *if-then* rules used to create the result images throughout this work. The values V_{Vessel} , V_{Bone} and $V_{VesselTree}$ are obtained by Equation 1.

Providing the FIS with meaningful information often requires a considerable amount of information or even specific information, such as the bone and vessel masks. In general, we cannot assume that we have such annotated data available. However, our system is implemented in a framework, which is used in the clinical routine of two hospitals. Radiological assistants create such annotations in their daily work. Hence, in our scenario, we have such data available. The tracking of the vessel tree is also done for every patient, and we can use it in our system. This vessel tree does not contain all vessels, but the clinical relevant ones are present. Concluding, the required data, even the specific one, is usually available, without the necessity of any user adaptation or changes to the workflow.

10 CONCLUSION

We proposed a knowledge-assisted user interface for medical visualization. The user can hover over a specific region and gets only relevant visualizations presented in the form of smart super views. These views are integrated into the overview visualization and visually convey relevance by shape and color. In order to inspect a view in detail, the user just needs to hover over it. He or she can then interact in the usual manner. Knowledge is provided by domain experts in the form of *if-then* rules. These are evaluated using fuzzy logic to obtain a ranking of suitable visualizations.

ACKNOWLEDGEMENTS

The work presented in this paper is part of the Knowledge Assisted Sparse Interaction for Peripheral CT-Angiography (KASI) project, supported by the Austrian Science Fund (FWF) grant no. TRP 67-N23. The data sets are courtesy of the Kaiser-Franz-Josef Hospital and the General Hospital of Vienna.

REFERENCES

- [1] J.-P. Balabanian, I. Viola, and M. E. Gröller. Interactive illustrative visualization of hierarchical volume data. In *Proceedings of Graphics Interface 2010*, pages 137–144, 2010.
- [2] E. Bier, M. Stone, K. Pier, W. Buxton, and T. DeRose. Toolglass and magic lenses: the see-through interface. In *Proceedings of ACM SIGGRAPH 1993*, pages 73–80, 1993.
- [3] S. Bruckner and M. E. Gröller. VolumeShop: an interactive system for direct volume illustration. In *Proceedings of IEEE Visualization 2005*, pages 671–678, 2005.
- [4] S. Bruckner and M. E. Gröller. Instant volume visualization using maximum intensity difference accumulation. *Computer Graphics Forum*, 28(3):775–782, 2009.
- [5] M. Chen, D. Ebert, H. Hagen, R. Laramée, R. van Liere, K.-L. Ma, W. Ribarsky, G. Scheuermann, and D. Silver. Data, information, and knowledge in visualization. *IEEE Computer Graphics and Applications*, 29(1):12–19, 2009.
- [6] S. Diepenbrock, T. Ropinski, and K. Hinrichs. Context-aware volume navigation. In *Proceedings of IEEE Pacific Visualization 2011*, pages 11–18, 2011.
- [7] B. Kadlec, H. Tufo, and G. Dorn. Knowledge-assisted visualization and segmentation of geologic features. *IEEE Computer Graphics and Applications*, 30(1):30–90, 2010.
- [8] A. Kanitsar, D. Fleischmann, R. Wegenkittl, P. Felkel, and M. E. Gröller. CPR - curved planar reformation. In *Proceedings IEEE Visualization 2002*, pages 37–44, 2002.
- [9] A. Kanitsar, D. Fleischmann, R. Wegenkittl, and M. E. Gröller. Diagnostic relevant visualization of vascular structures. In G.-P. Bonneau, T. Ertl, and G. Nielson, editors, *Scientific Visualization: The Visual Extraction of Knowledge from Data*, Mathematics and Visualization, pages 207–228. Springer Berlin Heidelberg, 2006.
- [10] G. Kindlmann. Superquadric tensor glyphs. In *Proceeding of VisSym 2004*, pages 147–154, 2004.
- [11] 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, Oct. 2007.
- [12] 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.
- [13] P. Kohlmann, S. Bruckner, A. Kanitsar, and M. E. Gröller. Contextual picking of volumetric structures. In *Proceedings of IEEE Pacific Visualization 2009*, pages 185–192, 5 2009.
- [14] E. LaMar, B. Hamann, and K. Joy. A magnification lens for interactive volume visualization. In *Proceedings of the Pacific Conference on Computer Graphics and Applications 2001*, volume Vi, pages 223–233, 2001.
- [15] E. Mamdani. Application of fuzzy logic to approximate reasoning using linguistic synthesis. *IEEE Transactions on Computers*, C-26(12):1182–1191, dec. 1977.
- [16] M. J. McGuffin, L. Tancau, and R. Balakrishnan. Using deformations for browsing volumetric data. In *Proceedings of IEEE Visualization 2003*, pages 401–408, 2003.
- [17] A. Motro. Superviews: Virtual integration of multiple databases. *IEEE Transactions on Software Engineering*, SE-13(7):785–798, July 1987.
- [18] J. Nam, M. Maurer, and K. Mueller. A high-dimensional feature clustering approach to support knowledge-assisted visualization. *Computers & Graphics*, 33(5):607–615, Oct. 2009.
- [19] M. Northam, J. Koonce, and J. G. Ravenel. Pulmonary nodules detected at cardiac CT: comparison of images in limited and full fields of view. *American Journal of Roentgenology*, 191(3):878–881, Sep 2008.
- [20] H. Ota, K. Takase, K. Igarashi, Y. Chiba, K. Haga, H. Saito, and S. Takahashi. MDCT compared with digital subtraction angiography for assessment of lower extremity arterial occlusive disease: Importance of reviewing cross-sectional images. *American Journal of Roentgenology*, 182(1):201–209, Jan 2004.
- [21] H. R. Portugaller, H. Schoellnast, K. A. Hausegger, K. Tiesenhäusen, W. Amann, and A. Berghold. Multislice spiral CT angiography in peripheral arterial occlusive disease: a valuable tool in detecting significant arterial lumen narrowing? *European Radiology*, 14(9):1681–1687, Sep 2004.
- [22] J. Rada-Vilela. Fuzzy-lite. <http://code.google.com/p/fuzzy-lite/>, 2012. [Accessed March 5, 2012].
- [23] A. Radloff, M. Luboschik, and H. Schumann. Smart views in smart environments. In *Proceedings of Smart Graphics 2011*, pages 1–12, 2011.
- [24] P. Rautek, S. Bruckner, and M. E. Gröller. Semantic layers for illustrative volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1336–1343, 2007.
- [25] P. Rautek, S. Bruckner, and M. E. Gröller. Interaction-dependent semantics for illustrative volume rendering. *Computer Graphics Forum*, 27(3):847–854, 2008.
- [26] C. Rezk-Salama, M. Keller, and P. Kohlmann. High-level user interfaces for transfer function design with semantics. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):1021–1028, 2006.
- [27] J. C. Roberts. State of the art: Coordinated & multiple views in exploratory visualization. In *Proceedings of the International Conference on Coordinated & Multiple Views in Exploratory Visualization 2007*, pages 61–71, 2007.
- [28] T. Ropinski and B. Preim. Taxonomy and usage guidelines for glyph-based medical visualization. In *Proceedings of SimVis 2008*, pages 121–138, 2008.
- [29] T. Ropinski, I. Viola, M. Biermann, H. Hauser, and K. Hinrichs. Multimodal visualization with interactive closeups. In *Proceedings of the Conference on Theory and Practice of Computer Graphics 2009*, pages 17–24, 2009.
- [30] T. Schertler, T. Frauenfelder, P. Stolzmann, H. Scheffel, L. Desbiolles, B. Marincek, V. Kaplan, N. Kucher, and H. Alkadhi. Triple rule-out CT in patients with suspicion of acute pulmonary embolism: Findings and accuracy. *Academic Radiology*, 16(6):708–717, Jun 2009.
- [31] M. Steinberger, M. Waldner, M. Streit, A. Lex, and D. Schmalstieg. Context-preserving visual links. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2249–2258, Dec. 2011.
- [32] T. Taerum, M. Sousa, F. Samavati, S. Chan, and J. Mitchell. Real-time super resolution contextual close-up of clinical volumetric data. In *Proceedings of EuroVis 2006*, pages 347–354, 2006.
- [33] C. Tietjen, B. Meyer, S. Schlechtweg, B. Preim, I. Hertel, and G. Strauß. Enhancing slice-based visualizations of medical volume data. In *Proceedings of EuroVis 2006*, pages 123–130, 2006.
- [34] F.-Y. Tzeng, E. Lum, and K.-L. Ma. An intelligent system approach to higher-dimensional classification of volume data. *IEEE Transactions on Visualization and Computer Graphics*, 11(3):273–284, 2005.
- [35] I. Viola, M. Feixas, M. Sbert, and M. E. Gröller. Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):933–940, 2006.
- [36] L. Wang, Y. Zhao, K. Mueller, and A. Kaufman. The magic volume lens: An interactive focus+context technique for volume rendering. In *Proceedings of IEEE Visualization 2005*, pages 367–374, 2005.
- [37] X. Wang, D. H. Jeong, W. Dou, S.-W. Lee, W. Ribarsky, and R. Chang. Defining and applying knowledge conversion processes to a visual analytics system. *Computers & Graphics*, 33(5):616–623, Oct. 2009.
- [38] L. Zadeh. Fuzzy sets. *Information and Control*, 8(3):338–353, 1965.
- [39] L. Zadeh. The concept of a linguistic variable and its application to approximate reasoning. *Information Sciences*, 8(3):199–249, 1975.
- [40] L. Zadeh. Fuzzy logic, neural networks, and soft computing. *Communications of the ACM*, 37(3):77–84, 1994.