

Guided Visualization of Ultrasound Image Sequences

Paolo Angelelli[†], Ivan Viola[†], Kim Nylund^{‡¶}, Odd Helge Gilja^{‡¶}, and Helwig Hauser[†]

[†]Department of Informatics, University of Bergen, Norway

[‡]National Centre for Gastroenterologic Ultrasonography, Haukeland University Hospital, Bergen, Norway

[¶]Institute of Medicine, University of Bergen, Norway

Abstract

Ultrasonography allows informative and expressive real time examinations of patients. Findings are usually reported as printouts, screen shots and video sequences. However, in certain scenarios, the amount of imaged ultrasound data is considerable or it is challenging to detect the anatomical features of interest. Post-examination access to the information present in the data is, therefore, cumbersome. The examiner must, in fact, review entire video sequences or risk to lose relevant information by reducing the examination to single screen shot and printouts. In this paper we propose a novel post-processing pipeline for guided visual exploration of ultrasound video sequences, to allow easier and richer exploration and analysis of the data. We demonstrate the usefulness of this approach by applying it to a liver examination case, showing easier and quicker ultrasound image selection and data exploration.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation—

1. Introduction

Ultrasonography (US) is a powerful and inexpensive imaging modality appreciated by many physicians. It causes little to no patient discomfort and is non invasive, providing high safety with neither contraindications nor radiation exposure. Ultrasonography has also good spatial resolution, combined with very high temporal resolution. This makes it an invaluable tool for examinations, where both anatomic and dynamic information is of interest. In the clinical practice it is successfully used for examination, diagnosis and intra-operative guidance.

However, US also suffers from certain limitations when it comes to interpretation and retrieval of image information. US waves are heavily attenuated by air and bones, and fatty tissue causes artifacts in the images. Furthermore, the acquisition process is dependent on the examiner, as the image acquisition is done by free hand and the interpretation is done in real time. US has also drawbacks regarding

the data storage and reviewing modalities: the typical ultrasonographic examination work flow, in fact, consists of live diagnosis during the examination. When there is the need to communicate the examination further, the acquired data can be saved for later reviewing. Data exported by 2D US scanners consists of annotated, and often printed, still images, and video sequences containing all the acquired US images, captured during the examinations at a certain frame rate. During certain kinds of US examinations, the physician scans several different anatomical structures, without focusing exclusively on one part of the anatomy. In such situations, as, for example, during abdominal examinations, simple snapshots of US slices may lack contextual information. They may also miss some important information that the examiner may have scanned, but not recognized and thus stored, in the first place. Stored video sequences contain all the imaged data, and, to a certain extent, prevent the loss of important information about structures of interest or their context. Unfortunately, such video data lacks higher semantic information, present during the live examination, such as the 3D position and orientation of the US planes, knowledge of which anatomical structures are imaged, neighbor-

[†] Email: paolo.angelelli,ivan.viola,helwig.hauser@uib.no

[‡] Email: kim.nylund,odd.gilja@helse-bergen.no

ing anatomical structures, scanning direction, and so on. Therefore it can be challenging and time-consuming to review ultrasonographic sequences *after* the examination, especially if the review is performed by another physician, without first hand knowledge of the examination. Considerable efforts are required to mentally reconstruct the spatial position and orientation of the US images, and sometimes to understand which anatomical structures the displayed images refer to. Moreover, videos of US data, lacking semantic annotations such as the imaged anatomical structures, require the examiner to go through all the video sequences, during the reviewing process, to find the images with the structures of interest, taking (potentially) long time.

To solve these limitations, in this paper we present a pipelined approach that enriches the US data with semantic information. By using the added semantics, we want to guide the examiner during the data exploration and reviewing tasks. The key concept of the presented solution is to compute the *degree of interest* (DOI) for each US image with respect to which anatomical structures the examiner wants to see. We do this by considering a so called *DOI volume*, that describes how much each voxel belongs to the structure(s) of interest. In the following step we combine this user defined annotation of the space with the US images. This way we derive an aggregated measure of the DOI for each image in the US sequences. This measure is a semantic information that represents how much of the selected anatomical regions is present, or visible, in the US images, and can therefore be seen as a measure of the *importance* of the US images. It is used through proposed visualization methods to guide the examiner to the relevant images or video subsequences.

A possible use-case of our technique is when a physician wants to review previously acquired US scans of, for example, gastrointestinal examinations of a patient, and focus on regions that look suspicious. Normally, she would have to browse through each video sequence, trying to understand the position and orientation of each image with respect to the anatomy of the patient. Then she would identify the region of interest in each image of the sequences. Using our approach, the same physician could select the desired suspicious regions, specified in one of the proposed ways. The system computes a DOI profile, showing the DOI of each of the images for all the US scans she wants to review. Using this additional semantics, the system helps the physician to find quickly the images showing the desired anatomical regions in the videos. It can also superimpose the corresponding regions over the US images to allow for easier and quicker identification of the interesting anatomical regions. Our technique could also be useful for physicians who need to communicate examinations results to other doctors: video sequences together with relative regions of interest contain much more information as compared to single screen shots. Using our method, these videos are much easier to analyze by a physician without first hand knowledge of the examination.

Our implementation of the proposed technique has, at the moment, two pre-examination requirements. First, a volumetric dataset of the patient, such as MRI scan. This is used as anatomical context, as reference coordinate frame and as basis for the definition of DOI volumes along the anatomical regions of interest. Second, even if, in this work, we make the assumption that the multi-modal data are already spatially co-register, for the registration method adopted in our prototype the US videos must be recorded with US plane positioning information, acquired via any suitable tracking device. This provides registration of all the frames in the videos with the DOI volumes by just having to register one frame [BHW*07, VNØ*08]. Our pipeline is, in principle, also applicable without a volumetric dataset of the patient: it could be possible to specify the DOI information based solely on the tracked US data, as explained in Section 3.1. The presented technique is, however, useful even in presence of a pre-acquired volumetric scan of the patient: ultrasonography allows to re-examine patients in an effective, fast and inexpensive way, without having to let the patient undergo other complex, expensive, and potentially harmful examinations such as CT or MRI.

The paper is structured as follows: in the next section we discuss the related works regarding multimodal visualization focused on US with special attention on guidance, focus+context and importance driven visualization. In Section 3 we present the details of each stage of the pipeline. In Section 4 we present the results of this technique applied to a case of liver examination, showing the benefits achieved through our guided visualization system. We conclude discussing the presented work in Section 5.

2. Related work

Our work aims at improving the diagnosis and treatment planning, which is one of the main challenges in medical visualization research. The challenge is to enable a clear understanding of the medical conditions depicted, and to guide the US examiner to the most relevant information during the reviewing process. Previous research related to advanced US data visualization has been mainly focused on the development of techniques for noise-free image rendering, especially in the case of three-dimensional visualization. The direct volume rendering of 3D US data requires a filtering stage to improve the image quality [SSG95]. More recent approaches use probability metrics to evaluate a presence of an interface between tissues [HRH03]. Furthermore, redundant information from 3D US measurements, resulting from volume overlap of consecutive scans, can be exploited to improve the rendering by preserving the temporal coherence [PHHH05].

2D US data rendering has been previously combined with augmented reality hardware to blend US with the real environment. US images have been displayed in the context of the body of the patient to show where they intersect the

body [BFO92]. Currently, registration techniques are usually based on internal landmarks and external markers, visible in the US data and in the pre-interventional 3D acquisition modality. In clinical practice non-rigid registration techniques are occasionally used [LMPT07, NFN07]. US images have been fused with MRI for neurosurgical interventions [NHL*03, RHR*03]. Most of the commercially available techniques, such as fused visualization of PET-CT through image overlays, or linked CT-US slicing, or fused visualization of CT and US operate primarily on the data level [Son]. Very recently, GE added a point tracking feature in their last generation of Logiq scanners (Logic E9), to aid the physician during the examination to find previously analysed areas.

To provide better 3D orientation, an integration of a 3D CT visualization with 2D interventional US has been recently proposed for CT-US guided intervention incorporating cutaway views [BHW*07], or superimposed information on imaged liver segments on US images [VNØ*08]. These works originate from importance-driven visualization techniques, which use data segmentation and relevance information to automatically generate expressive visualizations [VFSG06]. Approaches to visually emphasize features in volume renderings have also been discussed in different contexts. In the visualization of volumetric scalar data, two-level volume rendering uses segmentation information to render objects in the data with different composition and rendering techniques [HMBG01]. In the visualization of 3D flow data, a user-specified DOI function affecting optical properties is shown to visualize important flow features [HM03]. A more comprehensive overview about focus+context visualization is given by Hauser [Hau05].

Our work is also related to the context of video visualization as we deal with video sequences. Chen et al. [CBH*06] propose a technique to extract features from video sequences. 2D US videos have also been automatically classified using a machine learning approach [PZS*07]. With respect to the visualization of time-varying data, interesting approaches to visualize the changes in time are also available [JR05, WS03].

3. Guided Ultrasound Visualization

3.1. DOI Volume Specification

The first stage of the pipeline, illustrated in Figure 1, consists of the acquisition of the DOI volumes. These volumes represent the anatomical regions that the examiner is interested to see in the US data, and are used as a guidance instrument in the following stages. The value of a voxel in a DOI volume represents to which degree the respective anatomical location is relevant, or interesting, to the examiner. Special attention is required not only for the placement of the boundaries of the region of interest, but also for their definition. Hard boundaries are characterized by a steep transition to the region of interest from the surrounding volume. DOI values

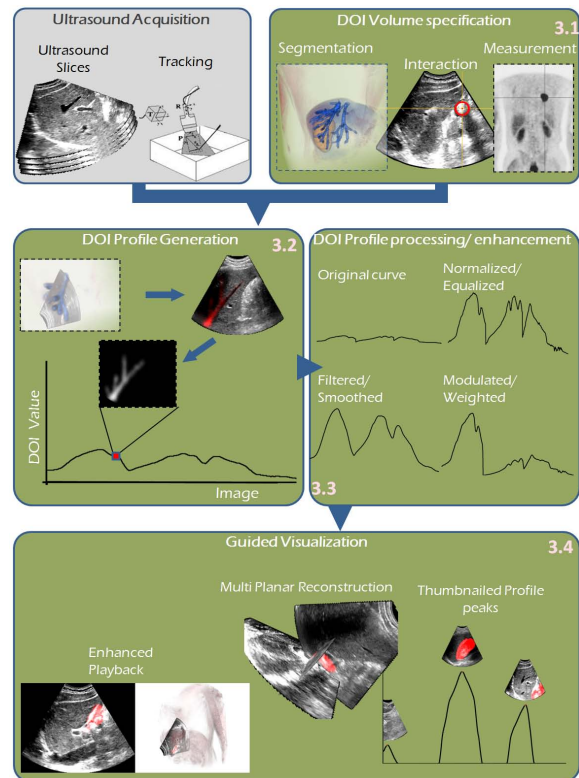


Figure 1: Schematic illustration of the proposed pipeline for guided visualization of US video sequences

change immediately from 0 to 1 when entering a region of interest with hard boundaries and thus defines precisely the region of interest. However, since there is no gradual transition, DOI volumes with hard boundaries cause the DOI profiles to be rougher. Moreover, hard boundaries amplify segmentation and registration errors present at the edge of segmented regions of interest. To attenuate this, we propose to allow for soft boundaries in the specification of DOI volumes [DH02], for example, by convolving the volume with a suitable smoothing kernel, such as an averaging or Gaussian 3D kernel. This way we attenuate small segmentation or registration errors by covering a slightly larger area with fading DOI values (Figure 3). By this, we also obtain a smoother transition of the DOI from one image to another. In case of DOI volumes with small interesting features, however, the smoothing operations should be carefully tuned to prevent information loss. A more formal definition of DOI volumes is then:

$$DOI(x,y,z) = \begin{cases} 1, & \text{structure in } (x,y,z) \text{ in focus.} \\ 0 < d < 1, & \text{structure in } (x,y,z) \text{ near focus.} \\ 0, & \text{structure in } (x,y,z) \text{ is context.} \end{cases} \quad (1)$$

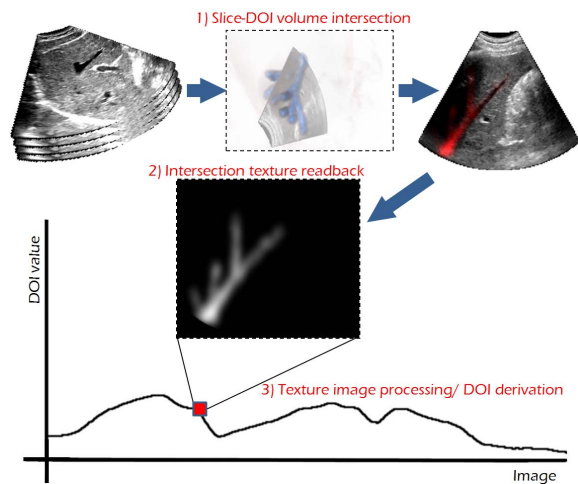


Figure 2: DOI profile generation through intersection between an US image plane and a DOI volume (segmented right hepatic vein tree), used as 3D texture. The result is a gray scale footprint of the DOI values in the 3D texture on the image plane. Summing each pixel of the intersection texture gives the DOI value for the US image

DOI volumes are scalar volumes and they must be co-registered with the US to match the desired anatomical structure in the US data with the related DOI value. There are several methods to specify the regions of interest. The most intuitive one is by segmentation of a volumetric scan of the patient, such as a CT or MRI. For example, if a PET scan of the patient is available, high tracer uptakes in the PET scan can be used as a DOI volume. Another possible method consists of using a transfer function to define structures of interest in a volume without the need of segmentation [BHW*07, KKH02, RS08]. In this case, modifications to the transfer function are equivalent to selecting a new DOI volume. One more way to specify DOI volumes consists of selecting interesting regions in US images, which can then be transferred into 3D space using the registration information. This selection can be as simple as a point, used to position a simple sphere of parameterizable radius, or to start a limited region growing process, following the gradients either from the volumetric data, if available, or from the US data.

3.2. DOI Profile Generation

The second stage of the pipeline consists of the derivation of the DOI for the US images in the video sequences the examiner wants to review. With the DOI of each image, we build a curve, or profile, of the relevance of each image in the video sequences, and we call it *DOI profile*. We compute the DOI of an US image with respect to a DOI volume as the amount of the volume intersected by the image. This

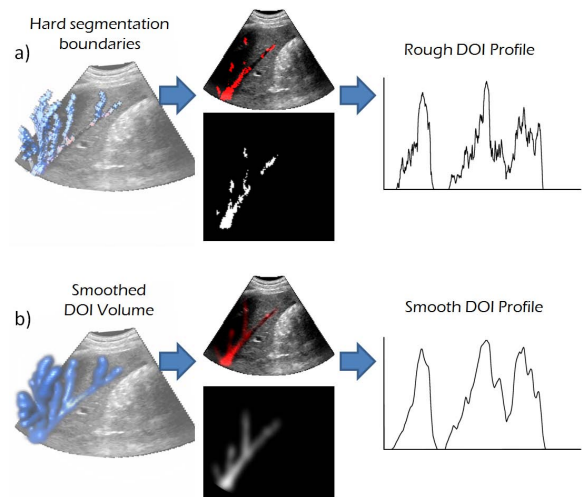


Figure 3: Original (a) and smoothed (b) DOI volume (segmented right hepatic vein tree). Small registration errors result attenuated and the DOI profile of the sequence results smoother

is done by using the registration information for the image, coming from a (suitable) registration method, and calculating the image spatial coordinates with respect to the DOI volume. Details on the registration method we employed in our prototype system are given in Section 4. We can now use the DOI volume as a 3D texture, and the image spatial coordinates as texture mapping for a rectangular polygon. The next step consists of texturing the polygon with the 3D DOI information from the DOI volume, and rendering it to a frame buffer object (FBO). The rendered image is a gray scale footprint of the values in the DOI volume on the polygon, representing the US image, as illustrated in Figure 2. Summing up the intensities of the pixels in the rendered image allows us to compute the aggregated DOI value of the US image. Storing the footprint enables us also to use it for visualization, as it highlights the interesting region in the US image. Keeping in mind that, by lowering the frequency of the US waves, the imaged area increases and vice-versa, a formal definition of the aggregated DOI, derived from the intersection footprint, is

$$DOI(image) = \frac{image\ area}{w \cdot h} \cdot \sum_{i=0}^{w \cdot h} pixel_i \quad (2)$$

The w and h parameters in Equation 2 represent the width and the height resolution of the intersection image. The $pixel_i$ parameter represents each pixel of the intersection image, while $image\ area$ is the physical area of the region in the US image. The DOI profile generated in this stage of the pipeline should be easy to modulate and use, and therefore needs to be normalized, or equalized, to be fitted into a

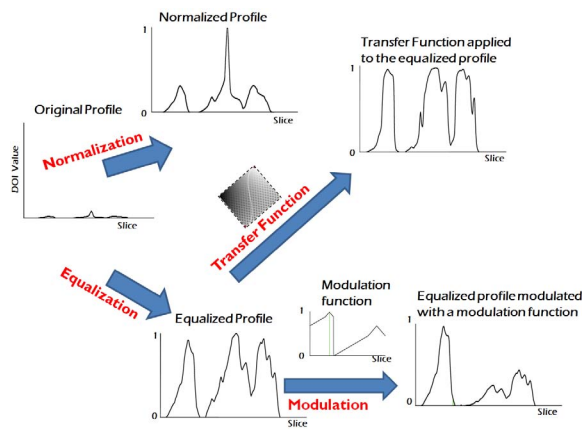


Figure 4: DOI Profile processing / enhancement operations

[0,1] range. These operations are going to be performed in the next stage of the pipeline.

3.3. DOI Profile Processing / Enhancement

The third stage of the pipeline is dedicated to processing and enhancement operations on DOI profiles to enhance the usability of the profiles. As previously outlined, raw DOI profiles do not have any common value range, and the values can be very small if the region of interest is small. Therefore, the first step consist of bringing the values into the unit range. We propose two possible techniques: normalization and histogram equalization of the profiles [Jai89]. Both map the values onto the unit range, but normalization retains the proportions between values, while equalization evenly distributes the profile values in the value range. Equalization will enhance local variations of the DOI intensity instead of global variations, and it is useful when a small number of images have very different values compared to the others.

Once the profiles have been mapped into the unit range, they are ready for further processing, to improve the derivation of useful information for the visualization. The first operation we propose is temporal filtering of the DOI profiles. A filter kernel can be used for temporal smoothing, when the profile is too rough. This helps in the profile visualization, as it shows better the trends in the data. We then adopt a peak detection algorithm [Beu94, NH06], that identifies cluster peaks on the DOI profiles in a waterfall-like way, and thus finds a representative image for each cluster. To achieve good results, such algorithms are usually also applied on a relatively smooth curve, otherwise they lead to too many peaks along the profile.

The last operation we propose to process the profiles is profile modulation. This type of operation multiplies an input profile with one or more other profiles, to modulate the input profile values. If p_{DOI} is the DOI profile and p_m is the

modulation profile, a formal description of profile modulation is then

$$p'_{DOI}(x) = p_{DOI}(x) \cdot p_m(x) \quad (3)$$

Profile modulation is very useful as it can combine information from different profiles, and we use this operation to weight the DOI profile with a similarity profile of the scans. The reason for this is that the desired structures of interest may be poorly imaged or visualized from specific points of view, even when the image plane intersects well the structure of interest. So the image with the highest DOI value might not show the structures of interest in an optimal way. This can happen under certain circumstances like with air, or bones, attenuating the US waves. But the structures may become visible with a similar image plane position, or under different patient conditions. For this reason we want to provide the examiner with a similarity weighted DOI profile, different for each image in the videos, so that similar and important images can be found quickly. To provide this functionality, we needed a metric to define the similarity between two US images. We have chosen to use distance and orientation: closer images have closer image centers, while a small angle between image normals means similar orientation of images in space. So we compute the similarity between a selected image and another image by multiplying the distance between image centers with the angle between image normals. This way the system can build a similarity profile for a selected image, and use it to modulate the DOI profile (see Figure 4). Such augmentation is especially useful when there are many sequences to review and they are acquired with different transducers or under different examination setups or patient conditions. We also use DOI profile modulation with a profile of angles between image normals in the automatic generation of a multi planar reconstruction, as explained in the next section.

3.4. Guided Visualization

The first visualization technique we create using the DOI information is a plot of the DOI profile. In the plot we also highlight the position of the currently visualized US image (see Figure 7(d)). This becomes a navigation tool for the examiner during the data exploration process. We also apply a peak detection algorithm to the profile, as previously discussed. This enables us to find representatives for clusters of neighboring images. We use the peaks to place a selectable US image thumbnail on the corresponding profile location. The thumbnails work both as a high level overview of the data and as a bookmark for quick data browsing. The second technique we propose to enhance the exploration and playback of US images consists of using the intersection images already utilized in the DOI computation (Section 3.2), as an on-demand, semi-transparent layer to superimpose over the US images. This visualization highlights to the examiner the relevant regions in the image. We also integrate the classic 2D US image visualization, familiar to the examiner, in a




Figure 5: Image of a US image co-registered with the MRI volume

linked 3D view where the image is rendered together with the DOI volume or a volumetric scan of the patient. We do this by rendering a proxy geometry, textured with the US image, correctly positioned into a volume rendering of the volumetric data (Figure 5). This helps the examiner to better understand the position and orientation of the displayed US image.

3.4.1. Multi-Planar Reconstruction from 2D US Data

The last visualization technique proposed here is an automatic multi-planar reconstruction (MPR) of the anatomy with intersecting 2D US images. It provides a “2¹/2D” visualization of the region of interest using the 2D US data (Figure 6). In the US domain such visualization is possible nowadays only with 3D US volume data, or has to be generated manually [KPKS05]. To create such an MPR visualization, we again employ a proxy geometry for each US image in the reconstruction, texture it with the image and then render each geometry in 3D space. For this technique we have also developed an algorithm to automatically compute a suitable selection of intersecting images showing the region of interest. It consists of a recursive modulation of the DOI profile of the US data with the angle between the images in order to favor images that are as normal to each other as possible. The algorithm works as follows: It takes the current DOI profile and selects the image with the highest DOI value. Then it calculates a modulation profile with $(1 - \cos(\text{angle}))$ of the angles between the selected image and all the others, to have a measure of the orthogonality of the all the images against the selected image. Finally, it modulates the current DOI profile with the computed modulation profile, and iterates. The second selected image will therefore be the image that conjugates best DOI value with orthogonality with the previously selected image. The third

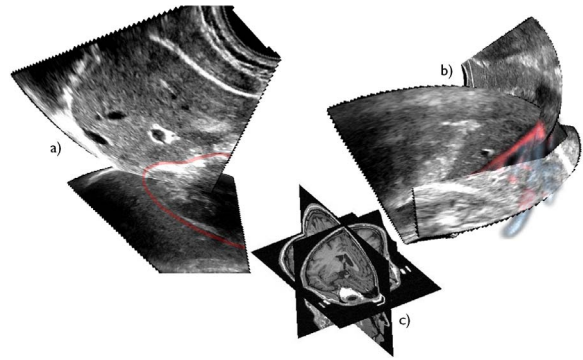


Figure 6: Multi planar reconstruction of the region of interest using 2D US images. a) MPR of the right kidney b) MPR of the right hepatic vein c) comparison image showing how MPR is currently applied to volume data

image will conjugate DOI value with orthogonality with the first and the second images. This algorithm is fast, taking only the normal of the US planes into consideration. However we also provide a manual tuning option to allow complete data exploration capability.

4. Results

To demonstrate the usefulness of the presented approach, and to obtain a first evaluation from the medical side, we built a prototype system and exemplified its capabilities in a proof-of-concept case study consisting of a trans-abdominal US examination with the focus on the liver. A 31 years old healthy male volunteer was examined after having first undergone an MRI scan of the abdomen. Several trans-abdominal US examinations were performed at different times. The US data (total of 7 scans) were obtained using different 2D transducers, in combination with a commercially available magnetometer-based tracking device (Flock of Birds, Ascension Technology) for image tracking during freehand US acquisitions. In our prototype system we have decided to adopt a landmark based rigid registration technique, well suited if tracking information is going to be employed. Our registration approach consist of identifying anatomical features visible in both the modalities, and indicating these through the placement of landmarks in the data directly on screen. These two sets of points are then used to compute a rigid transformation matrix from one dataset to the other [VNØ*08]. This matrix is then combined with the transformations recorded with the tracking system, to compute a suitable transformation for each US image in the video sequences. The presented visualization pipeline is, however, independent of the employed registration technique, as long as it can provide registration of each image in the video sequences to the DOI volumes. More advanced registration techniques are nowadays available. However, we




Figure 7: Screen shot of the system showing linked views of a US image, (a)with the intersection image superimposed, (b)forming a MPR with images crossing the DOI volume, (c)positioned in DVR of the MRI data to provide spatial orientation and context, (d)highlighted in the thumbnailled DOI profile (in green). The used DOI volume is the segmented and smoothed right hepatic vein tree. The yellow lines in the profile view (d) represent sequence boundaries. Here, three US sequences have been joined.

found that the employed registration method is quickly applicable to our type of data, a goal of the proposed technique, and provides sufficiently good results. Abdominal examination scans do not contain highly moving anatomy, and the deformation of the organs caused by the pressure of the probe proved to be minimal, so this method was able to register the two modalities sufficiently well. In case of changing the patient position, the physician simply started a new video sequence. The US data uptakes were stored in AVI format, then imported into our system. The MRI scan was acquired at a resolution of $256^2 \times 176$. Figure 5 shows both the modalities co-registered, with the intersection between the whole liver parenchyma and the US image highlighted in red. We have specified several DOI volumes, all of them through segmentation of the MRI scan. For the segmentation we used the ITK-SNAP tool. The DOI volumes we used define the liver parenchyma, the right kidney, the gallbladder, the right hepatic vein tree and the middle hepatic vein tree. With respect to the performances, processing a video sequence of ca. 3000 frames against a DOI volume requires around 5 seconds of computation, with our CPU based implementation, on an Intel Xeon 2.5GHz workstation. Figure 7 shows an examples of three US video sequences joined together, enriched with DOI information and visualized with our system. The development of the proposed pipeline into a prototype system benefited from our tight cooperation between technological and medical expertise. In the beginning we started

with a broader range of possible visualization techniques, and some of them did not prove to be promising, as, for instance, modulation of the playback speed according to the DOI values. The multi-planar reconstruction made out of 2D US images is a visualization which attracted particular interest, since it was not realized out of 2D data and from different sequences before, and it eventually turned out to be a useful visualization to inspect the data. The DOI profile was a handy tool for an active and quicker inspection of the US videos. The multimodal visualization of US images combined with DVR of the volumetric scan was useful for enabling a quicker understanding of the orientation of the images. The DOI region overlay was accepted as an interesting and useful method for examination training. After the demonstration of the test examinations with our system, it was acknowledged (on behalf of the medical side) that it combines assisted navigation and useful visualization techniques of US data in a novel way. It offers real time enhanced video playback and interaction for image selection and visualization customization, and changing the DOI volume takes also just a few seconds. The tool has been seen also especially interesting for doctor-to-doctor communication, as it enriches plain video streams with semantic information and allows to communicate the findings without leaving out part of the original data. Our system also potentially enables to find suspicious regions during the review, which was missed by the examiner during the live examination.

5. Summary and Conclusions

In this paper we have presented a pipelined approach for guided visualization during the review of US examinations. We have introduced the concept of *degree-of-interest* volumes in the context of US data visualization, to annotate the data with semantic information. We have presented a suite of visualization techniques that use the added semantic information to provide guidance and insight during the reviewing process, and aiming at improving the diagnosis and treatment planning process. We have implemented the proposed solution in a prototype system, and used it to review a case of a trans-abdominal US examination, achieving positive and useful feedbacks from our medical partners.

During the development of the prototype we tightly cooperated with our medical partners and addressed their needs. The presented prototype has been seen as possibly useful tool for post examination data exploration, to communicate examination results to other doctors, and for examination training. To the best of our knowledge, no mechanism has been previously presented to aid the examiner to focus on particular structures while reviewing 2D US examinations. Our approach extends the current ultrasonographic examination work flow during the live acquisition, since the data must be acquired with tracking, unless registration for the US data can be obtained by other means. We then add post processing steps currently non-existent in ultrasonographic

work flow to enrich the data with semantic information and thereby enable advanced data exploration. The presented method is also meant for examinations of anatomy that does not move particularly, or deform easily. When applied to cardiac data, for example, the high dynamic behavior of the imaged anatomy would represent a problem for the registration of the data to the DOI volumes. In such scenarios, alternative or additional solutions for the registration are needed.

Acknowledgements

This work has been supported by the IllustraSound research project, funded by the VERDIKT program of the Norwegian Research Council (Nr. 193180), and by the MedViz Initiative (medviz.uib.no, PK1760 Project 11). Furthermore, we are grateful to Stefan Bruckner for providing the VolumeShop visualization framework, and Timo Ropinski, for fruitful discussions enabled via the NRC's researcher exchange program DAADppp (Nr. 194673). A special mention goes to Yngve Devik Hammersland, for integrating the support for magnetic tracking devices into VolumeShop.

References

- [Beu94] BEUCHER S.: Watershed, hierarchical segmentation and waterfall algorithm. In *Mathematical morphology and its applications to image processing* (1994), Serra J., Soille P., (Eds.), Kluwer Academic Publishers, pp. 69–76.
- [BFO92] BAJURA M., FUCHS H., OHBUCHI R.: Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. In *Proceedings of SIGGRAPH '92* (1992), pp. 203–210.
- [BHW*07] BURNS M., HAIDACHER M., WEIN W., VIOLA I., GRÖLLER E.: Feature emphasis and contextual cutaways for multimodal medical visualization. In *Proceedings of EuroVis '07* (2007), pp. 275–282.
- [CBH*06] CHEN M., BOTCHEN R., HASHIM R., WEISKOPF D., ERTL T.: Visual signatures in video visualization. In *IEEE Transactions on Visualization and Computer Graphics* (oct 2006), vol. 12, pp. 1093–1100.
- [DH02] DOLEISCH H., HAUSER H.: Smooth brushing for focus+context visualization of simulation data in 3D. *Journal of WSCG 10*, 1 (2002), 147–154.
- [Hau05] HAUSER H.: Generalizing focus+context visualization. In *Scientific Visualization: The Visual Extraction of Knowledge from Data* (2005), Springer, pp. 305–327.
- [HM03] HAUSER H., MLEJNEK M.: Interactive volume visualization of complex flow semantics. In *Proceedings of VMV '03* (2003), pp. 191–198.
- [HMBG01] HAUSER H., MROZ L., BISCHI G. I., GRÖLLER M. E.: Two-level volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 7, 3 (2001), 242–252.
- [HRH03] HÖNIGMANN D., RUISZ J., HAIDER C.: Adaptive design of a global opacity transfer function for direct volume rendering of ultrasound data. In *Proceedings of IEEE Visualization '03* (2003), pp. 489–496.
- [Jai89] JAIN A. K.: *Fundamentals of digital image processing*. Prentice Hall, 1989.
- [JR05] JOSHI A., RHEINGANS P.: Illustration-inspired techniques for visualizing time-varying data. In *Proceedings of Vis-Conference* (Los Alamitos, CA, USA, 2005), IEEE Computer Society, pp. 679–686.
- [KKH02] KNISS J., KINDLMANN G., HANSEN C.: Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics* 8, 3 (2002), 270–285.
- [KPKS05] KERN R., PERREN F., KREISEL S., SZABO K.: Multiplanar transcranial ultrasound imaging: standards, landmarks and correlation with magnetic resonance imaging. *Ultrasound in Medicine and Biology* 31, 3 (2005), 311–315.
- [LMPT07] LEROY A., MOZER P., PAYAN Y., TROCCAZ J.: Intensity-based registration of freehand 3D ultrasound and CT-scan images of the kidney. In *International journal of computer assisted radiology and surgery* (2007), vol. 2, pp. 31–41.
- [NFN07] NICULESCU G., FORAN D. J., NOSHER J.: Non-rigid registration of the liver in consecutive CT studies for assessment of tumor response to radiofrequency ablation. In *Proceedings of IEEE Engineering in Medicine and Biology* (2007), pp. 856–859.
- [NH06] NOVOTNY M., HAUSER H.: Outlier-preserving focus+context visualization in parallel coordinates. *IEEE Trans. on Visualization and Computer Graphics* 12 (2006), 893–900.
- [NHL*03] NIKAS D. C., HARTOV A., LUNN K., RICK K., PAULSEN K., ROBERTS D. W.: Coregistered intraoperative ultrasonography in resection of malignant glioma. *Neurosurg Focus* 14, 2 (2003), 338–343.
- [PHHH05] PETERSCH B., HADWIGER M., HAUSER H., HÖNIGMANN D.: Real time computation and temporal coherence of opacity transfer functions for direct volume rendering of ultrasound data. *Computerized Medical Imaging and Graphics* 29, 1 (2005), 53–63.
- [PZS*07] PARK J. H., ZHOU S. K., SIMOPOULOS C., OTSUKI J., COMANICIU D.: Automatic cardiac view classification of echocardiogram. *International Conference on Computer Vision* (2007), 1–8.
- [RHR*03] RICK K., HARTOV A., ROBERTS D. W., LUNN K. E., SUN H., PAULSEN K. D.: Graphical user interface for intraoperative neuroimage updating. In *Medical Imaging 2003: Visualization, Image-Guided Procedures, and Display* (2003), pp. 210–221.
- [RS08] REZK-SALAMA C.: Visual parameters and transfer functions. In *Trends in Interactive Visualization* (2008), Springer, pp. 99–116.
- [Son] SONOWAND: The sonowand system. www.sonowand.no.
- [SSG95] SAKAS G., SCHREYER L., GRIMM M.: Preprocessing and volume rendering of 3D ultrasonic data. *IEEE Computer Graphics and Applications* 15, 4 (1995), 47–54.
- [VFSG06] VIOLA I., FEIXAS M., SBERT M., GRÖLLER E.: Importance-driven focus of attention. *IEEE Transactions on Visualization and Computer Graphics* 12, 5 (Oct 2006), 933–940.
- [VNØ*08] VIOLA I., NYLUND K., ØYE O. K., ULVANG D. M., GILJA O. H., HAUSER H.: Illustrated ultrasound for multimodal data interpretation of liver examinations. In *Visual Computing in Biomedicine proceedings* (Oct 2008), pp. 125–133.
- [WS03] WOODRING J., SHEN H.-W.: Chronovolumes: A direct rendering technique for visualizing time-varying data. In *Volume Graphics* (2003), pp. 27–34.