Real-time integration of 3-D multimodality data in interventional neuroangiography

Daniel Ruijters
Drazenko Babic
Robert Homan
Peter Mielekamp
Philips Healthcare, Cardio/Vascular Innovation
Veenpluis 6, 5680DA Best
The Netherlands
E-mail: danny.ruijters@philips.com

Bart M. ter Haar Romeny
Technische Universiteit Eindhoven
Biomedical Engineering
Image Analysis and Interpretation
Den Dolech 2, 5600MB Eindhoven
The Netherlands

Paul Suetens
Katholieke Universiteit Leuven
ESAT/PSI, Medical Imaging Research Center
Universitair Ziekenhuis Gasthuisberg
Herestraat 49, B-3000 Leuven
Belgium

Abstract. *We describe a novel approach to using soft-tissue data sets, such as computer tomography on magnetic resonance, in the minimally invasive image guidance of intra-arterial and intravenous endovascular devices in neuroangiography interventions. Minimally invasive x-ray angiography procedures rely on the navigation of endovascular devices, such as guide wires and catheters, through human vessels, using C-arm fluoroscopy. Although the bone structure may be visible and the injection of iodine contrast medium allows one to guide endovascular devices through the vasculature, the soft-tissue structures remain invisible in the fluoroscopic images. We intend to present a method for the combined visualization of soft-tissue data, a 3-D rotational angiography (3-DRA) reconstruction, and the live fluoroscopy data stream in a single fused image. Combining the fluoroscopic image with the 3-DRA vessel tree offers the advantage that endovascular devices can be located within the vasculature without additional contrast injection, while the position of the C-arm geometry can be altered freely. The additional visualization of the soft-tissue data adds contextual information to the position of endovascular devices. We address the clinical applications, the real-time aspects of the registration algorithms, and fast-fused visualization of the proposed method. © 2009 SPIE and IS&T.* [DOI: 10.1117/1.3222939]

1 Introduction

To the present date, the fluoroscopic image with the live information about endovascular interventional devices and soft-tissue images, such as computer tomography (CT) or magnetic resonance (MR), are visualized on separate displays. This means that the clinician has to perform a mental projection of the position of the endovascular device on the soft-tissue data. It may be clear that a combined display of this information is of great advantage because it relieves the clinician of performing this task. Furthermore, a fused image allows more precise navigation of the endovascular devices because these devices are visualized together with pathologies and contextual information present in the soft-tissue data. In order to provide the maximum benefit of such an augmented image, the live fluoroscopy data and the soft-tissue data have to be combined in real time, with low latency and a sufficient frame rate (15 or 30 fps, depending on the acquisition mode). Because the visualization is targeted at the usage during an intervention, it should not only be fast but also easy to interpret, and the manipulation of the image should be interactive and easy to use.

In this paper, we describe the steps that are necessary to achieve such a combined visualization. Prior to fusing a peri-interventionally acquired 3-D rotational angiography (3-DRA) and preinterventional soft-tissue data set with the live fluoroscopy image stream, a preprocessing step is performed. In this preprocessing step the 3-DRA and soft-tissue dataset are registered, using an image-based registra-

tion algorithm, and the vessels are segmented from the 3-DRA data. The preprocessing step is briefly touted in Sec. 3.1. During the visualization phase, an on-the-fly registration of the 2-D fluoroscopy images and the 3-D data must be performed. This is achieved by using a machine-based registration, which only depends on the geometry incidence angles, the x-ray source-to-detector distance, and the calibration data. The machine-based registration is described in Sec. 3.2. Section 3.3 discusses how a fast-fused visualization of all three data sets can be implemented, using off-the-shelf graphics hardware. We discuss the clinical applications that can benefit from the presented work in Sec. 4. Section 5 describes the data we measured in order to quantitatively tout the performance aspects of our methods, and the conclusions are presented in Sec. 6. However, first we start with a review of related work.

2 Related Work

Two fundamentally different approaches can be distinguished when coregistering the 2-D fluoroscopy data to 3-D volumetric data. In the first approach, called image-based registration, the registration process is driven by the image content. Angiographic image-based 2-D–3-D registration has received ample attention in the literature.[1–10] The image-based algorithms typically take a considerable amount of time to compute, ranging from a few seconds for methods that use a model of the anatomy of interest up to a few minutes for some intensity-driven approaches.[7] Because these algorithms use the image content, it should contain sufficient landmark features. In registration methods for angiographic applications, the features are usually provided by filling the vasculature with iodene contrast medium, which is harmful for the patient. Most registration methods are based on a single projection, which leads to a rather large registration error for the in-plane translation. As long as the projection angle does not change, this is not a big hurdle because it only leads to a slight mismatch in the magnification factor between the 2-D and 3-D images.[9] When the C-arm is rotated, however, the in-plane translation error leads to a large shift between the 2-D and 3-D images. This effect can be overcome by using two projection images at an angle of ~90 deg,[8] but then the amount of contrast medium doubles.

The second approach is known as machine-based registration. With the introduction of motorized calibrated C-arm x-ray angiography, 3-D reconstruction of the vasculature came within reach. Because such 3-DRA data sets are obtained with the same apparatus as the 2-D fluoroscopy data, it is possible to calculate a registration based on the state of the geometry (viewing incidence angles, source-detector distance, detector size, etc) and calibration data, provided that there was no patient motion between the acquisition of the 3-DRA data and fluoroscopy data.[11–13] This method also allows one to obtain a registration when there are insufficient landmarks present in the images (e.g., due to the absence of contrast die in the fluoroscopy images). A further advantage of machine-based registration is the fact that it can be computed in real time. Machine-based registration and image-based 2-D–3-D registration have been compared by Baert et al.[14] A method for determining the incidence based on tracking a fiducial was proposed by

Video 1 A 3-DRA data set can be acquired and reconstructed peri-interventionally within a few seconds. To obtain such a data set, the X-ray C-arm geometry follows a circular trajectory around the anatomy of interest. The volumetric data are computed using a cone-beam reconstruction algorithm (QuickTime, 5.6 MB). [URL: http://dx.doi.org/10.1117/1.3222939.1].

Jain et al.[15] We, however, do not use any fiducials, but rather only use the information concerning the geometry state, as is provided by the C-arm system.

In earlier work,[16] we already proposed the combined visualization of soft-tissue data and vasculature, which was segmented in 3-DRA reconstructions. Here, we intend to augment these data with the live fluoroscopy image stream, which enables the clinician to real-time correlate, e.g., the guide wire or catheter position to the soft-tissue data.

3 Method

3.1 Preprocessing

Our approach relies on the acquisition of a 3-DRA data set at the beginning of the intervention (see Video 1). The 3-DRA data set is coregistered to a soft-tissue data set, such as CT or MR, which has been obtained prior to the intervention (e.g., for diagnostic purposes). Using 3-D image registration during interventional treatment poses a number of constraints on the registration algorithm. Especially, the calculation time of the algorithm has to be limited because the result of the registration process is to be used during the intervention. In order to reduce the calculation time, the graphics processing unit (GPU) is employed to accelerate the registration algorithm.[17,18]

Because we focus on cerebral applications, and there are only limited elastic transformations of the anatomical structures within the head, we can use a rigid registration (i.e., only a global translation and rotation). Rigid registration further has the property that it can be calculated relatively robustly and quickly. We use mutual information as similarity measure, as described by Maes et al.[19] because it performs very well on intermodality registration and does not demand any *a priori* knowledge of the data sets. The Powell algorithm[20] is used as an optimizer. Optionally, the image-based registration is preceded by a rough manual registration. Stancanello et al. have shown that the capture range of the registration is sufficient for usage during a clinical intervention.[21] Note that this preprocessing step must be performed only once.

A further preprocessing step forms the creation of a triangulated mesh, representing the vessel tree. In order to obtain such a mesh, the vessels are segmented in the
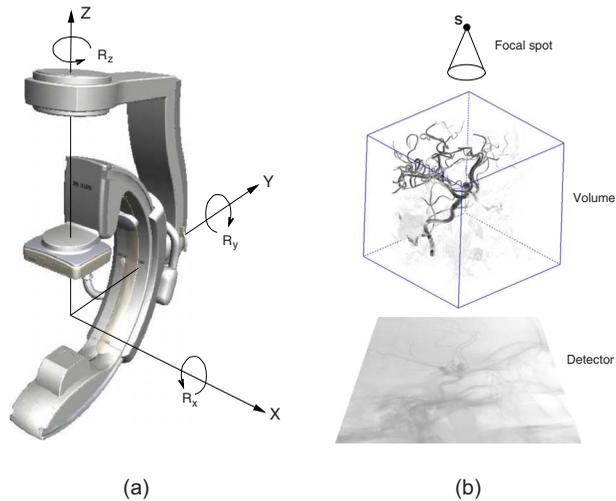
(a)   (b)

Fig. 1 (a) Degrees of freedom of the C-arm geometry and (b) the virtual projection of a 3-DRA dataset on a fluoroscopy image.

Video 2 Here the fused visualization of the 3-DRA vasculature (red) and a slab from a soft-tissue CT data set (gray) is shown. The CT slab is rendered semitransparent, and its position can be altered interactively by the user (QuickTime, 3.3 MB). [URL: http://dx.doi.org/10.1117/1.3222939.2].

3-DRA volume, which is a fairly easy task because the iodine contrast medium absorbs more x-rays than any other substance present in the data set. The segmentation is achieved by applying two thresholds. Any voxel that has an intensity that is below the lower threshold is marked as background. Any voxel with an intensity higher than the upper threshold is marked as vessel. Intensities between the lower and upper thresholds are marked either as background or vessel, depending on a connectivity criterion.[22] The thresholds are automatically determined based on the histogram of the volumetric data. From the segmented data, a mesh is extracted by applying the marching cubes algorithm.[23]
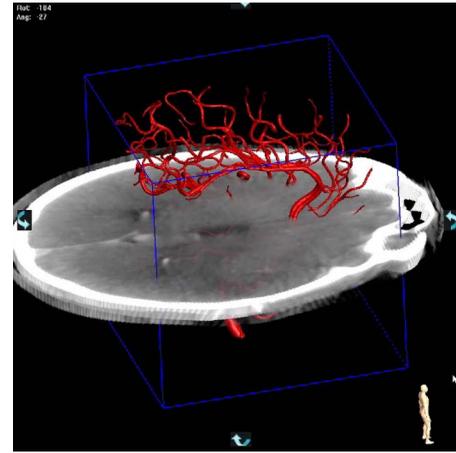
3.2 *2-D–3-D Registration*

The machine-based registration involves determining the transformation of the coordinate space of the 3-DRA data to the coordinate space of the fluoroscopy data. The x-ray C-arm system can rotate over three axes [see Fig. 1(a)]. The rotation of the detector coordinate system, with respect to the table, can be expressed as

$$M = R_x R_y R_z. \tag{1}$$

The C-arm system's isocenter serves as origin for the rotation matrices. The matrix M has to be corrected for deviations from the ideally calculated orientation, based on the calibration data. The calibration is performed by taking x-ray images from a known dodecahedron phantom from a large number of projections, equally distributed over the hemisphere of possible C-arm locations.[24,25] For any position in-between the calibrated positions, the deviations are linearly averaged from the neighboring calibrated data.

After the rotation of the 3-DRA data set into the appropriate orientation, and a translation of the origin from the system's isocenter to center of the detector, there still remains the task of projecting it with the proper perspective [see Fig. 1(b)]. The perspective depends on the x-ray source-to-detector distance (SID) and the detector dimensions. If the detector coordinate system uses the same metric as the coordinate system of the 3-DRA data set (e.g.,

millimeters), then the projection matrix will only depend on the SID. The projection matrix, which can be applied on homogenous coordinates (x, y, z, w), can then be written as

$$P = \begin{pmatrix} \text{SID} & 0 & 0 & 0 \\ 0 & \text{SID} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & \text{SID} \end{pmatrix}. \tag{2}$$

3.3 *Visualization*

To achieve interactive frame rates and a minimal latency, we seek to harvest the vast processing power of modern off-the-shelf graphics hardware. This power can be accessed by using the DirectX or OpenGL API. In order to render an image, first the triangulated mesh, representing the vessels, is rendered in the frame buffer. Simultaneously, the depths of the triangles are written in the z buffer. A stencil buffer operation is defined to write a constant S_1 to the stencil buffer for every pixel in the frame buffer that is filled by the mesh.

Consequently, a slab out of the soft-tissue data set is mixed into the scene using direct volume rendering (see Video 2). The position, orientation and thickness of the slab can be altered by the clinician. The slab is rendered by evaluating the direct volume-rendering equation for each pixel in the view port. The volume-rendering equation can be approximated by the following summation:[26,27]

$$i = \sum_{n=0}^{N} \left[\alpha_n c_n \prod_{n'=0}^{n} (1 - \alpha_{n'}) \right], \tag{3}$$

whereby i denotes the resulting color of a ray, α_n the opacity of the volume at a given sample n, and c_n the color at the respective sample.

This summation can be broken down in N iterations over the so-called overoperator,[28] whereby the rays are traversed in a back to front order
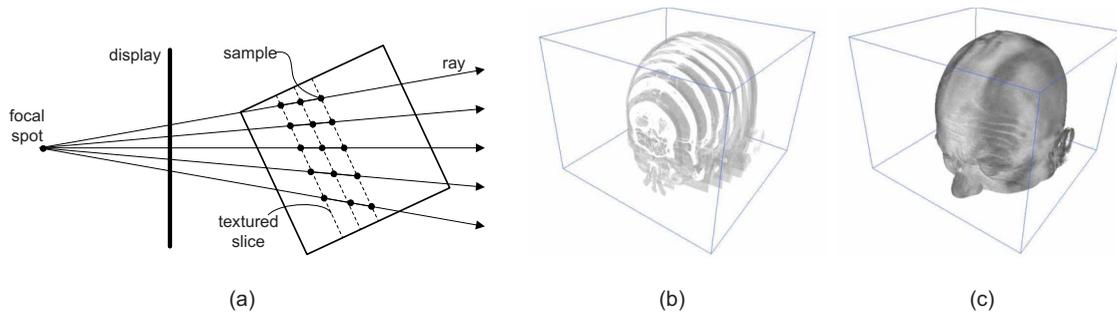
Fig. 2 (a) Volume rendering involves the evaluation of the volume render equation along the rays, passing through the pixels of the display. The usage of textured slices means that the rays are not evaluated sequentially. Rather for a single slice the contribution of the sample points to all rays is processed. (b) A volume rendered data set with large intervals between the textured slices. (c) The same volume-rendered data set with a small distance between the textured slices.

$$C_{n+1} = \alpha_n c_n + (1 - \alpha_n) C_n. \tag{4}$$

Here, C_n denotes the intermediate value for a given ray. After N iterations, C_N represents the final color of that particular ray. N should be chosen such that every voxel is at least sampled once (we use two samples per voxel). Standard α blending, offered by DirectX or OpenGL, can be used to implement the overoperator. Equation (4) can be evaluated for all pixels in the frame buffer, simultaneously, by using a set of N textured slices containing the slab data (see Fig. 2). In iteration n, the textured slice n is then blended into the frame buffer, under the appropriate translation, rotation, and perspective, whereby the slices are processed in a back-to-front order, from the perspective of the viewer. After each iteration, every pixel in the frame buffer represents its respective C_{n+1} value[29]

The triangulated mesh is already present in the frame buffer when the textured slices are rendered. To mix the triangulated mesh and the direct volume rendering, we test the z buffer at each iteration of the overoperator. If the z-buffer test shows that, for a particular pixel, the position of the present sample of the ray is further away from the viewer than the triangle in the frame buffer, then the frame buffer remains unchanged. The first sample that lies closer to the viewer will take the present value of the frame buffer as input, which was written by rendering the triangulated mesh. In this way, the color of the mesh is blended into the volume-rendering equation at the appropriate place.

The registration matrix, which was calculated in the first preprocessing step, is applied to the position of the slices. This makes a resampling of the slab with the soft-tissue data to the grid of the 3-DRA data unnecessary, leading to a better image quality.[30] Also, while rendering the slab, a stencil buffer operation is defined to write constant S_2 to every pixel that receives a color value from the direct volume-rendering process, with $\alpha > 0$. The S_1 labels can be overwritten by this operation.

Finally, the current fluoroscopy image is blended into the frame buffer. This is done in two passes. The action that is performed on a given pixel in a certain pass is determined by the value in the stencil buffer. S_1 in the stencil buffer means that the vessel tree is depicted in that pixel, S_2 corresponds to the soft-tissue data. If the stencil buffer is empty at a certain pixel position, then that particular pixel has not been filled with any information yet (background).

Because the S_1, S_2, and empty regions are addressed individually, different blending and image processing operations can be performed to these regions [compare Figs. 3(a) and 3(b)]. For instance, a spatial sharpening to enhance small details and a temporal smoothing to reduce noise can be applied to the vessel region.

The fluoroscopy data that overlay the background can contain some anatomical landmarks that are relevant to the physician. The most important part of the fluoroscopy image, though, is to be found inside the vessel region, because the movement of the endovascular devices is supposed to be contained within this region. This hierarchy is reflected in the intensity and filtering of the fluoroscopy data stream. The fluoroscopic information that overlays the soft-tissue slab could be suppressed to reduce cluttering of information in this region (see Fig. 4).

4 Clinical Use

The availability of the live fluoroscopy image stream combined with the vasculature segmented from the 3-DRA data set and the registered soft-tissue (CT or MR) data set during the intervention is of great clinical relevance. The combination of the fluoroscopy image with the 3-DRA vessel tree provides the advantage that the guide wire and catheter
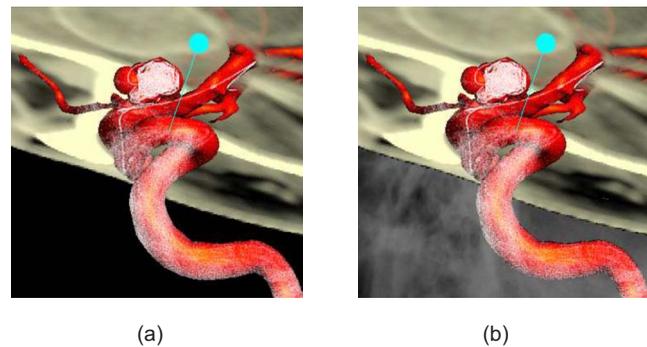

Fig. 3 (a) In the first fluoroscopy overlay pass, the pixels that are labeled S_1 (vessel) in the stencil buffer are treated. In this case, a sharpening filter was applied to the fluoroscopy data before they were blended with the frame buffer content. (b) In the second pass, the pixels that were labeled as background in the stencil buffer are processed. The fluoroscopy data are written without being sharpened, and the intensity is reduced.
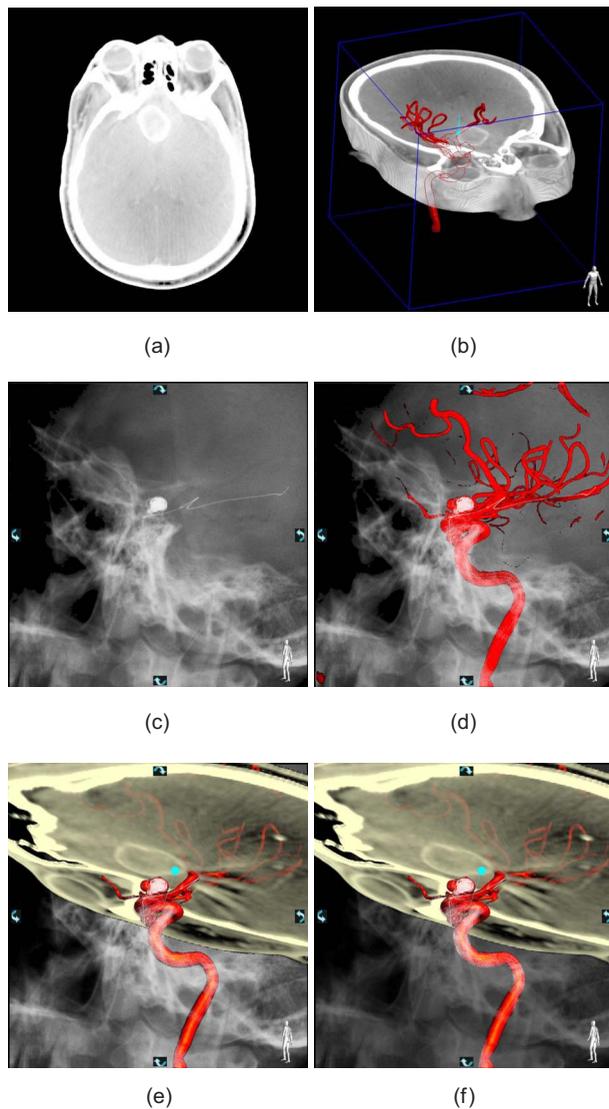
Fig. 4 (a) A CT image, clearly showing a tumor, (b) CT data set, registered with the 3-DRA data set, (c) a single frame from the fluoroscopy image stream, (d) the fluoroscopy image mixed with the vessel tree from the 3-DRA data set, and (e) the fluoroscopy image, the 3-DRA vasculature, and a slab from the CT data. (f) The fluoroscopy image outside the 3-DRA vessel tree is darkened.

position can be located with respect to the vessel tree, without additional contrast injection [see Fig. 4(d)], while the C-arm position and the x-ray SID can be altered freely.[13] Even during, e.g., rotations of the C-arm, the machine-based 2-D–3-D registration will always be up to date. The additional visualization of the soft-tissue data allows one to correlate the position of the guide wire and catheter to pathologies that are only visible in the soft-tissue data. Especially, the fact that this information is available in real time makes it very suitable for navigation.

The slab with the soft-tissue data can be moved, its width can be changed, and its orientation can be rotated freely to visualize different parts of the anatomical data set. In this way, the optimal view of a certain pathology can be determined. The implementation of the rendering running on the GPU offers interactive speed throughout.

The integration 3-D multimodality data can be used in the following treatments: (*i*) navigation to the optimal position for intra-arterial particle injection in endovascular embolization of intracranial neoplastic tissue, and arteriovenous malformation (AVM) treatment, prior to stereotactic radiation surgery, (*ii*) navigation to the optimal position for intracranial stenting in cases where aneurysms are pressing on surrounding eloquent and motoric brain tissue, (*iii*) navigation in the vessel portions to be embolized in, e.g., hemorrhagic stroke, (*iv*) navigation in the vessel segments where thrombolytic therapy should be applied in, e.g., ischemic stroke or vascular vasospasms.

Feedback from clinicians reported the presented approach to facilitate navigation in supra-aortic vessels from arch to skull base levels.[31] Less contrast medium was used than for traditional road mapping, while the hazard of thromboembolic events associated with direct catheterization was potentially reduced. The accuracy of registration was deemed satisfactory for clinical practice.
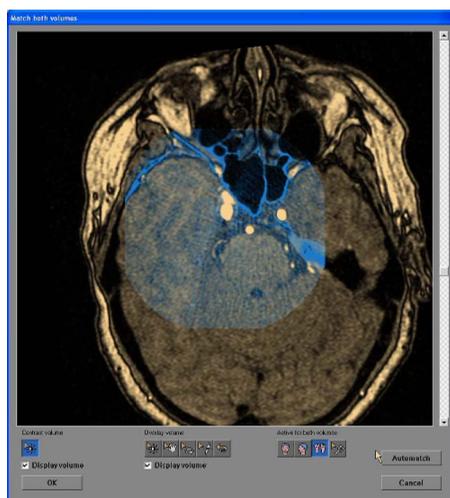
5 Results

The GPU implementation of the mutual information–based registration algorithm takes <8 s to register the 3-DRA data set and the soft-tissue data set in the preprocessing step. The extraction of the mesh that represents the vessels, the another preprocessing step, takes 300 ms. Overall, it can be concluded that these times are very acceptable and do not hinder the interventional procedure, especially because the preprocessing step has to be performed only once.

Given a certain set of viewing incidence angles, it takes a mere 1.5 μs to calculate the matrix, which expresses the 2-D–3-D registration between the 3-DRA data set and the fluoroscopy image. It is important that this part can be calculated in real time because it should be updated on the fly, when the geometry pose of the x-ray C-arm system changes. The augmented visualization, consisting of a mesh extracted from a 256^3 voxel 3-DRA dataset, a volume-rendered slab from a $256^2 \times 198$ voxel CT data-set and the fluoroscopy image stream, can be displayed at an average frame rate of 38 fps. All figures were measured on a Xeon 3.6-GHz machine with 2 GB of memory, and a nVidia QuadroFX 3400 graphics card with 256 MB of memory, using the data sets that are depicted in Fig. 4.

6 Conclusions

In this paper, a method for the combined visualization of the cerebral blood vessels segmented from 3-DRA data sets, data sets containing the surrounding anatomy such as CT or MR, and the live fluoroscopy data has been presented. The method is especially targeted for use in minimally invasive vascular procedures and distinguishes itself in the fact that it adds contextual information to the fluoroscopy images and 3-D vasculature.

The steps necessary to achieve this visualization have been described. First, an image-based registration of the 3-DRA data set and the soft-tissue data set has to be performed. We have demonstrated that the capture range is sufficient for interventional usage and that, due to the acceleration by the graphics hardware, the calculation time is very limited (Video 3). The machine-based registration between the fluoroscopy image and the 3-DRA data only depends on the geometry incidence angles, the x-ray SID, and the calibration data. It can be easily calculated in real time.
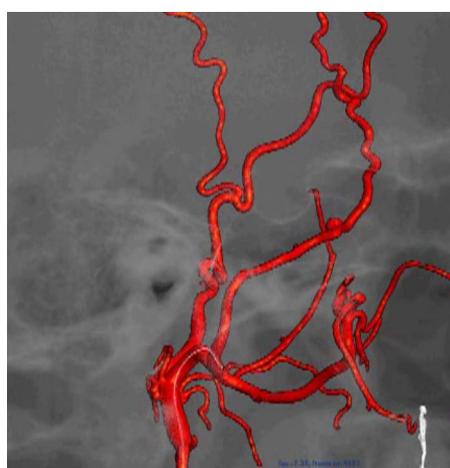
Video 3 During the preprocessing step, the 3-DRA data and the soft-tissue data are registered, using a GPU-accelerated mutual information registration method. The video shows the registration process in real time. At the beginning, the soft-tissue MR data (yellow) and the 3-DRA data (blue) are unregistered. At the end of the registration process, the boney landmarks in both data sets overlap and the brain tissue in the MR data are nicely contained within the skull in the 3-DRA data (QuickTime, 1.9 MB).
[URL: http://dx.doi.org/10.1117/1.3222939.3].

Also, we described how the visualization can be implemented to employ the possibilities of modern off-the-shelf graphic cards, allowing real-time display of the registered data with the live fluoroscopy image stream (Video 4). Further possible clinical applications have been identified, and it has been demonstrated how the presented method can be employed in those applications.

The strength of the described approach lies in its real-time nature, which is primarily achieved by the on-the-fly

Video 4 The machine-based 2-D–3-D registration allows one to overlay the 3-D vasculature (red) and the live X-ray fluoroscopy images (gray) in real time. The physician can navigate the guide wire (white line) without injecting contrast agent because the containing vessels and its bifurcations are clearly shown by the 3-DRA data. The video shows that any manipulations of the viewing incidence of the C-arm geometry are applied immediately to the 3-D vasculature (QuickTime, 1.9 MB). [URL: http://dx.doi.org/10.1117/1.3222939.4].

2-D–3-D registration, and the GPU-accelerated–fused visualization. The interactive real-time aspect contributes to the 3D perception of the anatomy and pathologies during an intervention. A possible disadvantage of the present method is the fact that patient motion will render the 2-D–3-D registration to be invalid. Therefore, future work could combine machine-based registration with image-based registration to correct for patient motion.

References

1. E. B. van de Kraats, G. P. Penney, D. Tomaževič, T. van Walsum, and W. J. Niessen, "Standardized evaluation of 2D-3D registration," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'04)*, pp. 574–581, Springer, Berlin (2004).
2. A. Liu, E. Bullit, and S. M. Pizer, "3D/2D registration via skeletal near projective invariance in tubular objects," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, pp. 952–963, Springer, Berlin (1998).
3. G. P. Penney, P. G. Batchelor, D. L. G. Hill, D. J. Hawkes, and J. Weese, "Validation of a two- to three-dimensional registration algorithm for aligning preoperative CT images and intraoperative fluoroscopy images," *Med. Phys.* **28**(6), 1024–1031 (2001).
4. D. Tomaževič, B. Likar, and F. Pernuš, "3-D/2-D registration by integrating 2-D information in 3-D," *IEEE Trans. Med. Imaging* **25**(1), 17–27 (2006).
5. G.-A. Turgeon, G. Lehmann, M. Drangova, D. Holdsworth, and T. Peters, "2D-3D registration of coronary angiograms for cardiac procedure planning," *Med. Phys.* **32**(12), 3737–3749 (2005).
6. J. Weese, G. P. Penney, P. Desmedt, T. M. Buzug, D. L. G. Hill, and D. J. Hawkes, "Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery," *IEEE Trans. Inf. Technol. Biomed.* **1**(4), 284–293 (1997).
7. R. A. McLaughlin, J. Hipwell, D. J. Hawkes, J. A. Noble, J. V. Byrne, and T. C. Cox, "A comparison of a similarity-based and feature-based 2-D–3-D registration method for neurointerventional use," *IEEE Trans. Med. Imaging* **24**, 1058–1066 (2005).
8. J. Jomier, E. Bullitt, M. van Horn, C. Pathak, and S. R. Aylward, "3D/2D model-to-image registration applied to TIPS surgery," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'06)*, pp. 662–669, Springer, Berlin (2006).
9. M. Groher, T. F. Jakobs, N. Padoy, and N. Navab, "Planning and intraoperative visualization of liver catheterizations: new CTA protocol and 2D-3D registration method," *Acad. Radiol.* **14**, 1325–1340 (2007).
10. F. Bender, M. Groher, A. Khamene, W. Wein, T. H. Heibel, and N. Navab, "3D dynamic roadmapping for abdominal catheterizations," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'08)*, pp. 668–675 (2008).
11. S. Gorges, E. Kerrien, M.-O. Berger, Y. Trousset, J. Pescatore, R. Anxionnat, and L. Picard, "Model of a vascular C-Arm for 3D augmented fluoroscopy in interventional radiology," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'05)*, pp. 214–222, Springer, Berlin (2005).
12. J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Med. Image Anal.* **2**(1), 1–36 (1998).
13. M. Söderman, D. Babic, R. Homan, and T. Andersson, "3D roadmap in neuroangiography: technique and clinical interest," *Neuroradiology* **47**, 735–740 (2005).
14. S. A. M. Baert, G. P. Penney, T. van Walsum, and W. J. Niessen, "Precalibration versus 2D-3D registration for 3D guide wire display in endovascular interventions," in *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI'04)*, pp. 577–584, Springer, Berlin (2004).
15. A. K. Jain, T. Mustafa, Y. Zhou, G. S. Chirikjian, and G. Fichtinger, "FTRAC—a robust fluoroscope tracking fiducial," *Med. Phys.* **32**(10), 3185–3198 (2005).
16. D. Ruijters, D. Babic, B. M. ter Haar Romeny, and P. Suetens, "Silhouette fusion of vascular and anatomical data," in *Proc. Int. Symp. on Biomedical Imaging (ISBI'06)*, pp. 121–124, IEEE, Piscataway, NJ (2006).
17. R. Shams and N. Barnes, "Speeding up mutual information computation using NVIDIA CUDA hardware," in *Proc. Digital Image Computing: Techniques and Applications (DICTA)*, IEEE Computer Society, Washington, DC, pp. 555–560 (2007).
18. M. Teßmann, C. Eisenacher, F. Enders, M. Stamminger, and P. Has-

treiter, "GPU accelerated normalized mutual information and B-spline transformation," in *Proc. Eurographics Workshop on Visual Comput. Biomed. (EG VCBM)*, pp. 117–124, ACM Press, New York (2008).

19. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imaging* **16**(2), 187–198 (1997).
20. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, New York (1992).
21. J. Stancanello, C. Cavedon, P. Francescon, P. Cerveri, G. Ferrigno, F. Colombo, and S. Perini, "Development and validation of a CT-3D rotational angiography registration method for AVM radiosurgery," *Med. Phys.* **31**, 1363–1371 (2004).
22. J. Bruijns, "Segmentation of vessel voxel structures using gradient ridges," in *Proc. of Vision Modeling and Visualization Conf. (VMV)*, pp. 159–166, Max-Planck-Gesellschaft, Munich, Germany (2003).
23. W. E. Lorensen and H. E. Cline, "Marching cubes: a high resolution 3-D surface construction algorithm," *Comput. Graph.* **21**(4), 163–169 (1987).
24. S. Gorges, E. Kerrien, M.-O. Berger, J. Pescatore, Y. Trousset, R. Anxionnat, S. Bracard, and L. Picard, "3D augmented fluoroscopy in interventional neuroradiology: precision assessment and first evaluation on clinical cases," *Presented at MICCAI Workshop AMI-ARCS'06* (2006).
25. M. Grass, R. Koppe, E. Klotz, R. Proksa, M. H. Kuhn, H. Aerts, J. op de Beek, and R. Kempkers, "Three-dimensional reconstruction of high contrast objects using c-arm image intensifier projection data," *Comput. Med. Imaging Graph.* **23**(6), 311–313 (1999).
26. K. Engel, M. Kraus, and T. Ertl, "High-quality pre-integrated volume rendering using hardware-accelerated pixel shading," in *Proc. of Eurographics Workshop on Graphics Hardware*, pp. 9–16, ACM Press, New York (2001).
27. J. T. Kajiya, "The rendering equation," *Proc. SIGGRAPH'86*, *Comput. Graph.* **20**(4), 143–150 (1986).
28. T. Porter and T. Duff, "Compositing digital images," *Comput. Graph.* **18**(3), 253–259 (1984).
29. D. Ruijters and A. Vilanova, "Optimizing GPU volume rendering," *J. WSCG* **14**(1–3), 9–16 (2006).
30. K. J. Zuiderveld and M. A. Viergever, "Multi-modal volume visualization using object-oriented methods," in *SIGGRAPH Symp. on Volume Visualization*, pp. 59–66, ACM, New York (1994).
31. C.-J. Lin, R. Blanc, F. Clarençon, M. Piotin, L. Spelle, J. Guillermic, and J. Moret, "Overlaying fluoroscopy and pre-acquired computer tomography angiography for road mapping in cerebral angiography," *AJNR Am. J. Neuroradiol.* (in press).

Daniel Ruijters has been employed by Philips Medical Systems since 2001. Currently, he is working as senior scientist, 3-D imaging, at the Cardio/Vascular Innovation Department in Best, the Netherlands. He received his engineering degree at the University of Technology (RWTH) Aachen and performed his final project at ENST in Paris. Next to his work for Philips, he is currently conducting a PhD thesis supervised by Suetens and ter Haar Romeny. His primary research interest areas are medical image processing, 3-D visualization, image registration, fast algorithms, and hardware acceleration.

Drazenko Babic received his medical degree from the University of Zagreb, Croatia, and has been employed by Philips Medical Systems since 1996. Prior to his work for Philips, he worked at different medical universities on various research projects focusing on medical applications. Currently, he is a member of the Cardiovascular X-ray Department, working as a principal scientist responsible for new clinical developments in the neurovascular interventional field. His activities comprise establishing collaborations with different medical centers in North and South America, Europe, and Asia.

Robert Homan has been employed by Philips Medical Systems since 1997. Prior to that he worked for the research department of the University Hospital Nijmegen, developing software to simulate ultrasonic imaging in causal absorptive media. Within Philips, he worked on X-ray applications for a postprocessing workstation. Currently, he is working on prototyping and validation of applications, such as the 3-D roadmap for minimally invasive procedures.

Peter Mielekamp joined Philips in 1975 and worked on silicon waver optics in the late 1970s, on rasterizing vector fonts in the 1980s, and on virtual reality in the 1990s. After working for Philips Research for many years, he moved to Philips Medical Systems in 2001. There he worked on 3-D visualization, computer-aided analysis, and user interaction within the domain of interventional 3-D reconstructed X-ray data. He is author or coauthor of many patent applications in the field of computer graphics and medical image processing and visualization.

Bart M. ter Haar Romeny is full professor at the Faculty of Biomedical Engineering at the Eindhoven University of Technology. Before, he was associate professor at the Image Sciences Institute of Utrecht University (1986 to 2001). He received an MSc in applied physics from Delft University of Technology in 1978, and his PhD from Utrecht University in 1983. He has been chairman of the Dutch Society for Biophysics and Biomedical Engineering and the Dutch Society for Clinical Physics. He is head of the Biomedical Image Analysis (BMIA) group, with focus areas in biologically inspired multiscale image analysis, advanced 3-D visualization, and computer-aided diagnosis applications.

Paul Suetens is professor of medical imaging and image processing and head of the Center for Processing Speech and Images in the Department of Electrical Engineering (ESAT/PSI) at the Katholieke Universiteit Leuven. He is also chairman of the Medical Imaging Research Center at the University Hospital Leuven, which is a joint initiative of the Faculty of Engineering and the Faculty of Medicine. The focus of the research of the Medical Imaging Research Center lies on the clinical applications and the needs of a university hospital in the area of medical imaging and image processing.