

Bone Segmentation in CT-Angiography Data Using a Probabilistic Atlas

Matúš Straka*
Austrian Academy of Sciences
Miloš Šrámek
Austrian Academy of Sciences

Alexandra LaCruz
Technical University Vienna
Dominik Fleischmann
University of Vienna

Leonid I. Dimitrov
Austrian Academy of Sciences
Eduard Gröller
Technical University Vienna

Abstract

Automatic segmentation of bony structures in CT angiography datasets is an essential pre-processing step necessary for most visualization and analysis tasks. Since traditional density and gradient operators fail in non-trivial cases (or at least require extensive operator work), we propose a new method for segmentation of CTA data based on a probabilistic atlas. Storing densities and masks of previously manually segmented tissues to the atlas can constitute a statistical information base for latter accurate segmentation. In order to eliminate dimensional and anatomic variability of the atlas input datasets, these have to be spatially normalized (registered) first by applying a non-rigid transformation. After this transformation, densities and tissue masks are statistically processed (e.g. averaged) within the atlas. Records in the atlas can be later evaluated for estimating the probability of bone tissue in a voxel of an unsegmented dataset.

Keywords: CT Angiography, Knowledge Based Segmentation, Probabilistic Atlas, Thin-Plate Spline, Distance Fields, Histogram Classification

1 Introduction

In this paper we present results of an ongoing, interdisciplinary research project [1], aimed at visualization and treatment planning of the peripheral arterial occlusive disease (PAOD) by means of CT angiography (CTA). PAOD is characterized by the formation of atherosclerotic plaque and vessel wall calcifications, which lead to luminal stenoses (narrowing) or complete occlusion of the involved arteries [3]. Diminished blood flow to the legs causes restricted mobility, pain, and even necrosis, eventually leading to death.

CTA of the peripheral arteries is performed using a multiple-detector row CT scanner, which allows to acquire a series of 1500-2000 transverse images of 512×512 12-bit pixels through the anatomic region of interest. While the scanning procedure requires only 30–40 seconds, the analysis of such huge datasets within acceptable time is impossible without the application of dedicated visualization techniques. Manual image editing and segmentation us-

ing currently available commercial medical visualization workstations is prohibitively long—in the order of two to four hours, even for a well trained operator.

First results from our early research led to promising techniques for segmentation and visualization of peripheral CTA datasets [12, 13]. For segmentation we took advantage of the fact that the voxel density of the enhanced vessels is above that of the surrounding soft tissue. Thus, density based operators (thresholding, gradients) were used to detect the vessels by a graph-theoretic approach (vessel tracking [20]). Since the vessel density range overlaps with that of the bones, it proved advantageous to identify the bones first by locally dependent thresholding and to exploit the fact that bones occupy a much larger volume than vessels.

This approach allowed a considerable decrease in postprocessing time for visualizing the peripheral arterial tree, in the range of 15 to 45 minutes of user interaction for each clinical case. This allowed us to evaluate the technique in a controlled clinical environment [15]. Currently, 1 to 2 patients with peripheral arterial occlusive disease are processed at the Department of Angiography and Interventional Radiology, University of Vienna General Hospital (Austria) every day. For routine clinical use, however, user interaction times of at most 15 minutes are desirable. The user interaction, which is still required is due to the complex anatomy and the overlapping density properties of the contrast-medium enhanced vessels and neighboring bones.

The goal of this paper is to introduce a knowledge-based approach for identification and labeling of vessels and vessel trees in CTA data. When fully implemented, will provide us with information precise enough for reliable visualization of vessel trees for diagnosis by a radiologist, as well as for automatic detection and quantification of the PAOD. The first step along the way to achieve this goal is to reliably identify bone tissues in the CTA data, since they can interfere in vessel identification due to their similar density properties. Once identified, the bone tissue can be excluded from consideration, thus leaving enhanced vessels and calcifications as the brightest tissues in the data, which would considerably simplify their identification.

Bone tissue labeling in CT data is often a very simple problem. Cortical bone has significantly higher X-ray attenuation than other tissues and simple thresholding can usually be safely used for its identification. However, this is not applicable for trabecular bone, which is the building material of a considerable amount of bones of the human skeleton. Its density is only slightly above that of the soft tissue, and—in addition—it overlaps with the density of the enhanced vessels in CTA. Moreover, there is a considerable variability of the density of trabecular bone between individuals. Therefore, we propose a knowledge-based approach for segmentation of bones, based on a probabilistic atlas of the skeletal anatomy, which will provide us with supplementary information guiding the data classification in bone/no-bone classes. Probabilistic atlases have been used by several authors to perform various tasks predominantly in human brain analysis [10, 18]. However, a similar ap-

* matus.straka@assoc.oeaw.ac.at

proach to the segmentation of large bones from CT data is currently unknown to us.

Thus, the bone identification task splits into atlas construction and usage. In the construction phase, a number of data sets is manually segmented by an expert and both the original data and segmented data are non-rigidly transformed to a common coordinate space. There, statistical properties of the transformed data are evaluated (mean density field, probability of bone presence, etc). This information is later used in the segmentation of an unknown data set, which is similarly transformed to the same coordinate space. Currently, the transformation is defined manually by specification of several landmark points, which is the only necessary user interaction.

While the procedure presented in this paper is designed specifically for bone extraction, we believe that it would be applicable to a wider class of segmentation problems. For example, we intend to apply it in a similar way also to the segmentation of the vessel tree after the bone removal, where we want to represent spatial information about the vessel segments and bifurcations.

The paper is organized as follows. In the second section we give an overview of related techniques known from the literature. In the third we summarize properties of CTA data from the point of view of segmentation and in the fourth we present the considerations which led us to the idea of an atlas-based segmentation technique. In the fifth we describe the sequence of operations required to build the atlas, whereas in the sixth we show how to use it for segmentation of unknown data sets. In the last but one section we give details about its implementation, and, finally, in the last section we summarize the results and sketch directions of our future work.

2 Related work

It is impossible to give an overview of 2D and 3D data segmentation techniques in a short paper, since normally it is the subject of voluminous textbooks and monographs [9, 2]. Therefore, we restrict ourselves here only to topics relevant to the presented paper: application of distance fields and deformable and statistical atlases in segmentation.

Distance transformations (DT) and distance fields (DF, arrays holding distances to certain objects of interest) [5] represent a powerful means to encode knowledge about spatial relationships in both planar and volumetric scenes. For example, DT stands in the background of a fully automatic scheme proposed for brain segmentation in MR scans by Brummer [6]. There, data voxels are classified first according to a probabilistic density model. Since in MR data different tissues occupy the same spectral space, in the second step only those regions are classified as brain, which have at least a certain distance from the skin surface. The necessary DF is obtained by skin identification by thresholding and a DT. Zeng *et al* [22] proposed a technique for identification of the brain cortex, which takes advantage of the fact that the cortex is a layer of approximately 3mm thickness. DTs are in this case used to measure the distance between the inner and outer cortex surface.

Segmentation by means of deformable models and atlases are usually applied when it is not possible to identify objects of interest based solely on the density data. In this sense, they represent a complement to the density based techniques, which in a preprocessing step provide the low-level features for the subsequent atlas matching [18, 17]. For example, Iosifescu *et al* [10] differentiate between brain white and gray matter and cerebrospinal fluid in a first stage by a classification algorithm based on interactive selection of representative samples for each tissue. In the second stage, the segmented data set is elastically deformed to an atlas [14], in order to measure the volumes of small brain structures. The deformation itself is defined by the maximization of a local similarity structure, leading to a multiresolution FEM solution.

The adaptive, template moderated, spatially varying statistical classification (ATM-SVC) technique proposed by Warfield *et al* [21] combines matching to an anatomical atlas with density based classification in an iterative loop. Its aim is to propose a general purpose technique for segmentation in different anatomy areas and was tested, among others, on the segmentation of knee cartilage from MRI data. First, a template of normal anatomy is build by an operator-driven manual or semiautomatic technique. Subsequently, it is converted to a set of distance maps, each indicating locations where bones and cartilage are likely to be found. For each unknown data set its statistical density model is formed by manual selection of representative voxels for each tissue type. The iterative loop starts by a rigid transformation of the unknown data to the template space. Then, coregistered distance and density probability data are classified by K-nearest neighbor classifiers. In the loop, both the unknown data set and the template are matched by an elastic matching procedure, which leads to improved tissue classification. The loop usually converges in three to five iterations.

3 Segmentation of Bones in CTA Data

In spite of the fact that bones are probably the most often visualized tissue in medical imaging (with the exception of the brain), the literature on its segmentation is very sparse even in the case of CT, which is the major 3D modality used for bone imaging. This is due to the fact that simple thresholding is usually sufficient to identify the bone tissue based on its higher attenuation of X-rays as compared to soft tissues. The most often discussed problem in segmentation and visualization of bones is the partial volume effect (PVE). The PVE manifests itself as data blurring and shifting densities of structures with thickness comparable to the scanner sampling rate to significantly lower values. This effect is usually accounted for in visualization by means of non-binary classification and semitransparent rendering [7, 16].

However, only cortical bone features the favorable high density. A large proportion of the human skeleton is build up of porous trabecular bone. Due to its own low X-ray attenuation and PVE it results in data density equal to or only a little higher than that of soft tissues. This becomes even more important with older patients, who are the major group affected by the lower extremity arterial disease, since they often suffer from osteoporosis, too.

Specifically, in the case of CT-angiography, we distinguish four relevant tissues: bone, soft tissue, contrast material enhanced blood (vessels) and vessel wall calcifications, which often occupy the same density space (Figure 1a):

- vessel plaque calcifications overlap with high density bone, and
- blood overlaps with low-density bone, bone areas with PVE and bone marrow.

The most important features which can significantly influence the reliability of a segmentation technique are contacts between vessels and bones, which, due to the overlapping density, cannot be identified by thresholding (Figure 1b), and interrupted thin vessels due to the inhomogeneous distribution of the contrast agent and PVE (mainly in areas, where the vessel is parallel to the scanning plane).

4 General Description of the Method

The implemented method is based on ideas formulated in discussions with skilled radiologists. It seems to be inevitable to incorporate some pre-knowledge that comprises information about patient

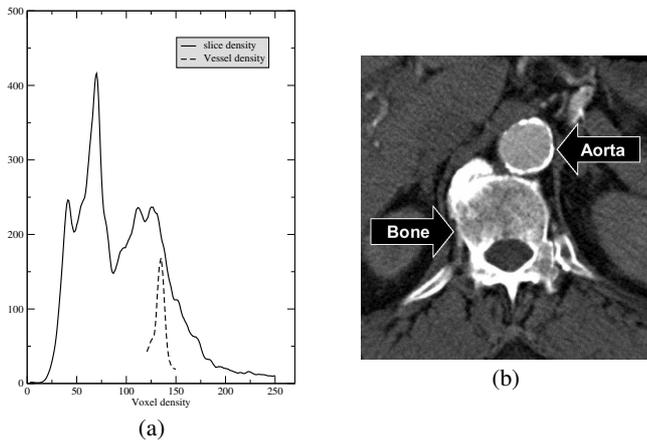


Figure 1: (a) CT-angiography data histogram, and (b) bone (vertebral body) and blood vessel (abdominal aorta) with overlapping densities. Close vicinity of both tissues may cause false interconnections in segmentation by thresholding.

anatomy in segmentation, because methods based only on density and gradient analysis simply cannot cope with the diversity of CTA data. Similar to a skilled human radiologist, the computer also must utilize spatial anatomic knowledge and be aware of what kind of tissue could be in the processed areas. This idea leads us to our goal to construct an atlas comprising spatial probabilistic existence information about given tissue types. This idea is also encouraged by the anatomic diversity of the processed data sets that may appear similar on the first sight, but are not on closer look.

In order to build and use such an atlas it is necessary to deal with the following issues:

- the spatial variability of the processed datasets, caused mainly by the different positions of the patients during scanning,
- the anatomic differences between patients, mainly due to body proportions, sex and age,
- the different scanners and scanning protocols used, resulting in differences in data sampling, and
- to extract, store and interpret the appropriate information of a given tissue type.

Usage of non-rigid transformations (e.g. thin-plate spline warping) can help to solve the first three issues. A solution to the last one can be based on the idea of having a statistics of the possible densities in the given spatial range, as well as representing the probability of presence in this range. This idea is based on the observation that the densities change heavily in the global extent, but just slightly in the local range, and there exists a statistic measure of presence of the given type of tissue in the given spatial range. This spatial statistic measure can be transformed to a distance field which better represents the probability of tissue presence also in the areas where there is no relevant information.

The presented method consists of two independent steps: first, we have to construct the atlas (Section 5) and second, we have to use the information contained in the atlas for segmentation goals (Section 6).

5 Atlas Construction

The probabilistic atlas (Figure 2) stores for each voxel information regarding the *Average bone density*, *Average distance* to bone sur-

face and the distance variance. As an auxiliary information, necessary when adding new data sets (records), it also stores the total number of added records (*Record count*) and the number of records, for which the voxel was identified as a bone voxel (*Bone count*).

The atlas's *Reference frame*, which defines its dimensions, sampling rate and a set of landmarks for registration is given by the first inserted record. Having at disposal this reference frame, one can easily add new records and compute the necessary information. An input record for each new data set to be added to the atlas consists of the original density volume, a manually segmented bone mask and a set of landmarks, which define its correspondence with the reference frame. Currently, the landmarks are set manually by a simple interactive utility and pairwise correspond to those set in the reference frame. The density data is segmented by means of the semiautomatic utility ISEG [19], which is based on interactive thresholding, connected component analysis and application of morphologic and logic operations.

When adding a new record to the atlas, both the density and mask volumes are warped by means of the TPS transform (Section 5.1) to the reference space. However, since the manual landmark setting proved to be not precise enough, their position is optimized first by means of a mutual information based cost function (Section 5.2). At the next step, the atlas average bone density and average distance volumes are updated, using the warped density and mask volumes and the atlas record and bone count volumes. The distance volume of the processed record is computed from the transformed mask volume by the 345 chamfer distance transform [5]. Prior to the distance transform, the bone mask is subjected to morphologic closing in order to fill in-bone holes which belong to areas filled by bone marrow. A distance transform based closing [11] is used with closing distance equal to 16. Since in the segmentation phase we take advantage not only of the mean distance, but also of its variance, we store in the atlas also a sum of squared distances for each voxel.

5.1 The Thin-Plate Spline Transform

Thin-plate spline (TPS) is an elegant algebraic expression of the dependence of the physical bending energy of a thin metal plate on point constraints. Its extension to 3D space allows modeling of shape deformation changes in 3D datasets (Figure 3). The TPS transform is defined by [4]:

$$x = a_0 + a_1X + a_2Y + a_3Z + \sum_{i=1}^n F_i r_i^2 \ln r_i^2 \quad (1)$$

$$y = b_0 + b_1X + b_2Y + b_3Z + \sum_{i=1}^n G_i r_i^2 \ln r_i^2 \quad (2)$$

$$z = c_0 + c_1X + c_2Y + c_3Z + \sum_{i=1}^n H_i r_i^2 \ln r_i^2 \quad (3)$$

where $[X, Y, Z]$ and $[x, y, z]$ are points in source and target data space respectively, and $r_i^2 = (X - X_i)^2 + (Y - Y_i)^2 + (Z - Z_i)^2$. The TPS parameters can be determined if the coordinates of at least four corresponding points are known. When solving for the transform parameters a_i, b_i, c_i, F_i, G_i and H_i for the given point

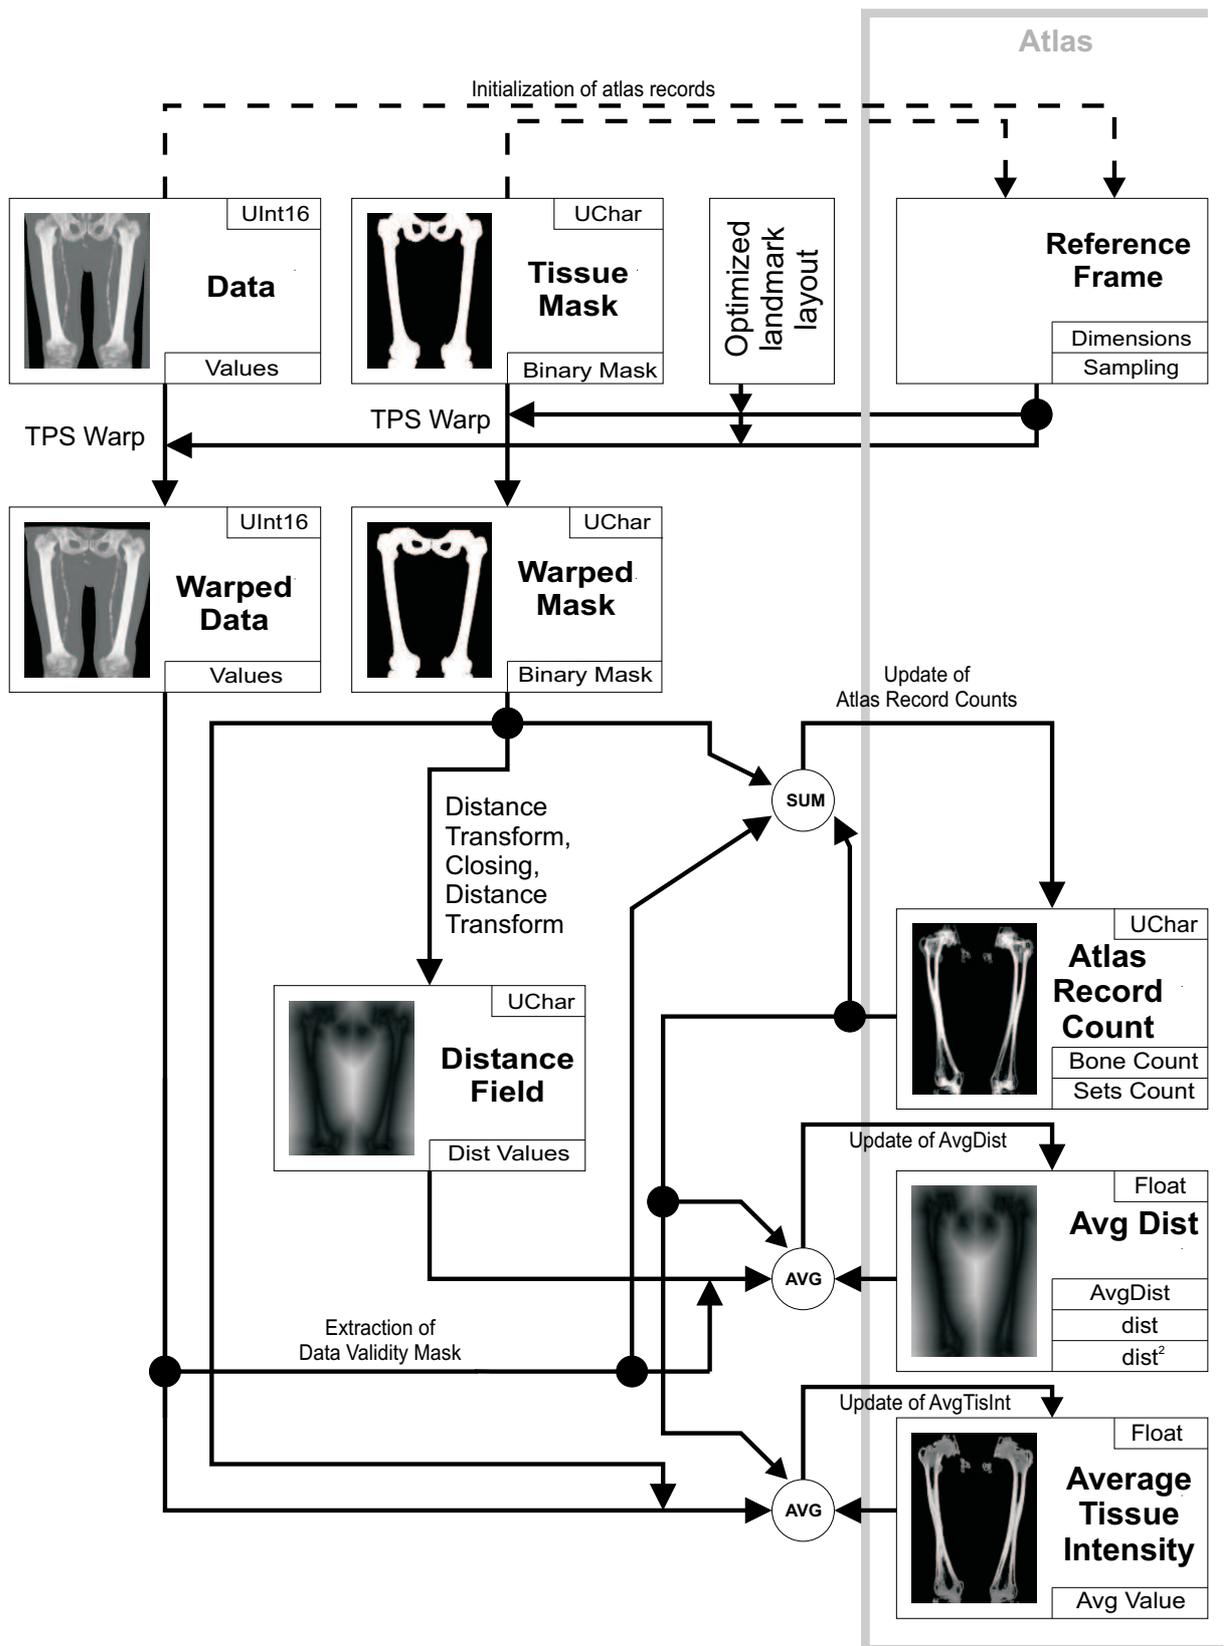


Figure 2: Atlas construction.

pairs, the following constraints must be also taken into account:

$$\begin{aligned}
\sum_{i=1}^n F_i &= 0 & \sum_{i=1}^n G_i &= 0 & \sum_{i=1}^n H_i &= 0 \\
\sum_{i=1}^n X_i F_i &= 0 & \sum_{i=1}^n Y_i F_i &= 0 & \sum_{i=1}^n Z_i F_i &= 0 \\
\sum_{i=1}^n X_i G_i &= 0 & \sum_{i=1}^n Y_i G_i &= 0 & \sum_{i=1}^n Z_i G_i &= 0 \\
\sum_{i=1}^n X_i H_i &= 0 & \sum_{i=1}^n Y_i H_i &= 0 & \sum_{i=1}^n Z_i H_i &= 0
\end{aligned} \tag{4}$$

These constraints ensure that the volume splines remain stable under application of point loads (Figure 3).

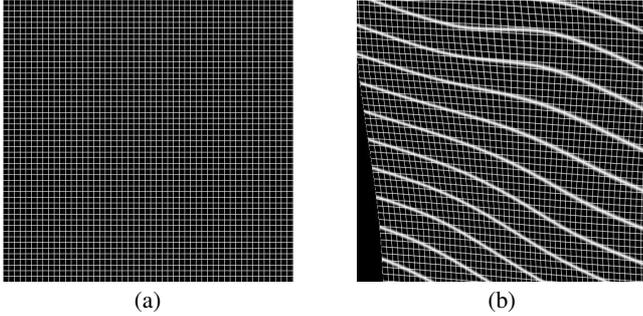


Figure 3: A typical result of the thin-plate-spline warp. (a) A 2D slice through a 3D grid filled with 3 orthogonal sets of planes (b) a 2D slice through the warped grid. The third set of planes is not visible in (a).

5.2 Optimization of the TPS transformation

According to our experience, it is impossible to define the optimal location of the control point pairs only by manual specification of the landmarks. Such manual specification usually results only in approximate data correspondence after the transformation (Figure 4a). Therefore there is a need to optimize the position of the control points in the datasets. The implemented optimization procedure is based on maximizing the mutual information of the transformed and reference data sets. The mutual information I_{AB} represents a measure of the mutual correspondence of two data sets A and B , by reflecting the number of relevant graphic elements that coincide in both volumes:

$$I_{AB} = \sum_i \sum_j p(a_i, b_j) \log \frac{p(a_i, b_j)}{p(a_i)p(b_j)}, \tag{5}$$

where $p(\cdot)$ is the density probability in images A and B respectively and $p(a_i, b_j)$ is the probability for the combination of given densities to appear in the corresponding locations in both datasets. All aforementioned probabilities are approximated by means of either 1D histograms or a 2D scatterplot.

By maximizing the value of I we maximize the above mentioned correspondence between the transformed and reference dataset. Due to the high variability of human body properties we apply Eq. 5 only to the subset of the data defined by the bone mask. This decision stems from the observation of larger "similarity" between human skeletons than between whole human bodies. A straightforward implementation of the optimization method leads to unsatisfactory results. The mutual information has many local minima because of the "nearly binary" character of both registered images which is caused by data masking through the bone selection

mask. To overcome this shortcoming, we filter the data by a low-pass Gaussian filter.

In cost optimization it is necessary to evaluate the optimized function many times. Since the computational complexity of the TPS transform is linear both in the number of voxels and the number of control points, one pass of the transformation for a full-sized data set can take a considerable amount of time and the duration of the optimization would become a problem. Fortunately, sub-sampled datasets and scaled coordinates of landmarks proved to be meaningful and delivered good and stable results in acceptable time.

The results acquired by the optimized TPS transform on sub-sampled datasets are, although not ideal (Figure 4b), still much better in comparison with unoptimized ones (Figure 4a).

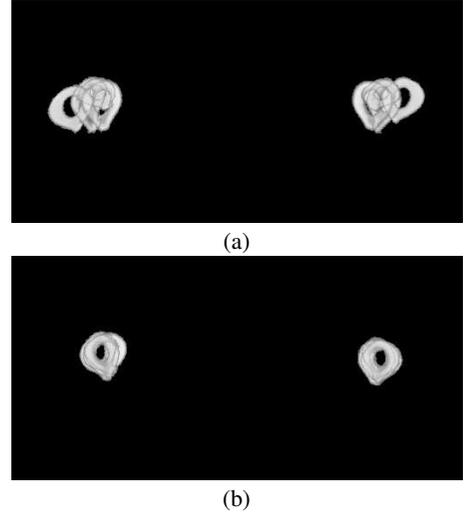


Figure 4: An overlaid cross-section of warped data sets before (a) and after (b) the optimization.

5.3 Probability Models Derived from the Atlas

The data gathered in the atlas and providing information on average tissue density and average distance to bone surface is not used for the segmentation of unknown data sets directly.

Firstly, the information derived from the atlas is specific to the CTA scans of lower extremities. Central parts of long bones are characterized by a high density. Therefore bone-tissue histograms of 2D transversal slices depend on the slice position (Figure 5). This leads to the idea to exploit this dependency for bone tissue segmentation by means of the analysis of spatially dependent histograms. For each slice of the average bone density volume a histogram is built (binned to 256 density levels) and approximated by a modified Gaussian distribution:

$$p_I = \begin{cases} e^{-\frac{(I-\bar{I})}{2\sigma}}, & \text{if } I \leq \bar{I} \\ 1, & \text{otherwise,} \end{cases} \tag{6}$$

where \bar{I} and σ are mean bone density and its standard deviation within the slice. I is the density of a voxel of the same slice. This formula assumes that bone is the most dense tissue and therefore all voxels with density above the bone average are assigned probability 1. Figure 6b shows this probability for all voxels in a transversal slice of one of the data sets used for atlas construction. In comparison to the unmodified image density (Figure 6a), we observe a significant enhancement of bone structures and their spatially homogeneous density distribution.

Probability p_I is, however, not sufficient for bone identification in an unknown data set, since it does not eliminate the problem of those no-bone tissues misclassified as bones, which share the same spectral space with bones.

A similar situation occurs in brain segmentation, where brain and non-brain tissues share the same density space in a similar way. As a solution, a probabilistic atlas has been build [8] by manual segmentation and rigid registration of brains of more than 300 healthy individuals. By averaging binary masks of tissues of interest a probability map was obtained, which was subsequently used in classification as the Bayesian prior. Our situation is different because we currently do not have a sufficient number of data sets in the atlas and also because of the higher variability of bone shapes. Therefore, instead of deriving the presence probability directly from the averaged binary masks, we derive it from average of distance maps. This significantly blurs the bone masks and thus makes the classification tool less sensitive to the insufficient number of data sets in the atlas (Figure 6c). From the averaged distance map \bar{D} and its standard deviation σ we derive the spatial presence probability (Figure 6d) in point X by

$$p_D(X) = e^{-\frac{\bar{D}(X)}{2\sigma(X)}}. \quad (7)$$

In classification of unknown data sets, we merge both probability measures p_I and p_D into a joint probability $p_{joint} = p_I p_D$, which, after thresholding, results in the final bone mask.

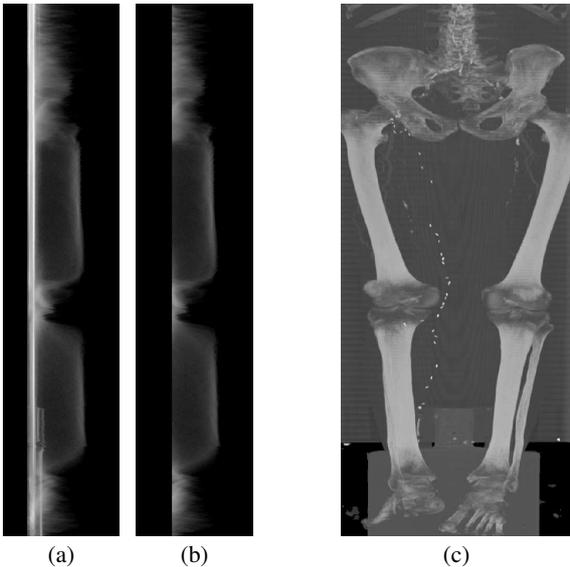


Figure 5: Dependency of density histograms on the slice position. Each row in (a) and (b) shows the histogram of the corresponding slice in a dataset, the MIP of which is presented in (c). (a) shows histograms of all body tissues while (b) only those of segmented bones.

6 Bone Segmentation by Means of the Atlas

Segmentation of an unknown dataset by means of the information provided by the atlas requires several steps (see Figure 7).

First, it is necessary to transform the unknown data to the reference atlas frame. Similarly to the atlas construction, we interactively define the set of landmarks, pairwise corresponding to those

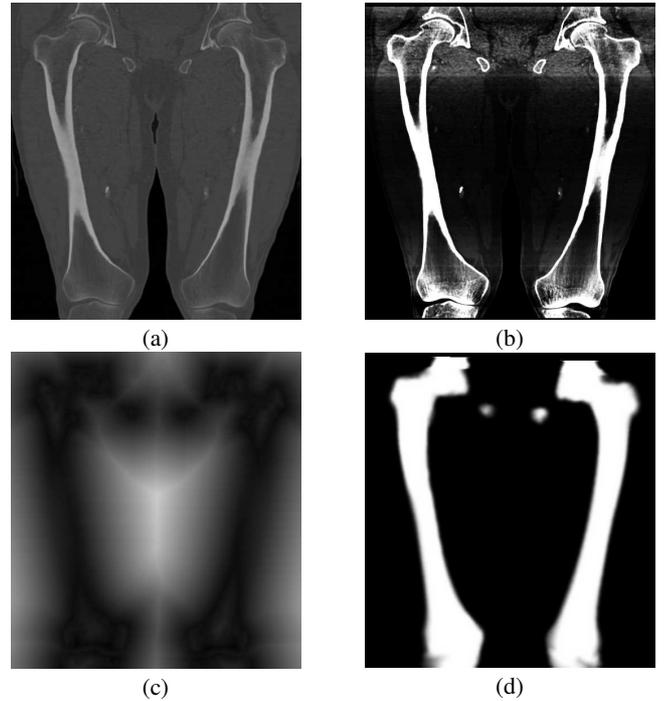


Figure 6: A transversal slice of (a) the original data set, (b) probability of bone presence based on the density, (c) distance fields showing distance to the bone surface and (d) spatial probability of the bone tissue.

used for construction of the atlas. Once the transformation is defined, the unknown data set is transformed to the atlas space and the probabilities p_I and p_D are computed according to Section 5.3. Finally, the bone mask is obtained by thresholding the joint probability p_{joint} at level 0.2.

Similarly to the atlas construction, the manually set landmarks result in an insufficient correlation of the data set to the atlas. However, in this case we do not have the necessary bone masks, which were obtained interactively in the case of atlas construction. Instead, we approximately identify the bone tissue by thresholding the density at a fixed threshold value. We use this mask for optimization of the transformation, in spite of the fact that also no-bone tissues are selected. Since these tissues occupy significantly less space than bones, their influence is suppressed by downscaling the mask and by Gaussian blurring.

7 Implementation and Results

The algorithm was implemented on a dual AMD Athlon 1.6 GHz CPU workstation with 2GB of RAM. The addition of a new dataset to the atlas takes about 7 minutes and the segmentation of an unknown patient dataset takes a comparable amount of time. A major problem of the current implementation is its very high memory requirements since for each data voxel ≈ 20 bytes of additional information have to be added. This causes an initially 300-600 MB dataset to grow easily over 2 GB, which results in memory management problems (the operation system limits us to 2GB address space per task) and rapidly deteriorates the performance of the whole application.

In the future, we will have to solve this problem by implementing some type of run-time data compression. A CTA dataset typically comprises substantial parts of the data representing uninteresting

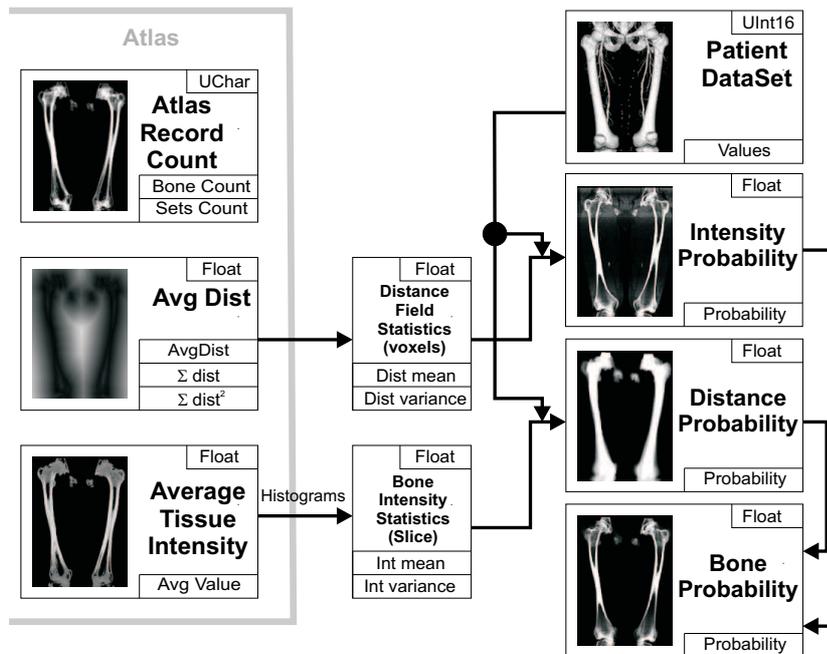


Figure 7: Atlas usage.

background (air and parts of the tomograph), which can be easily identified by a low threshold value. A modified version of the run-length encoding scheme, compressing only the background regions at both ends of each scanline, was used for this purpose in the aforementioned segmentation tool [19], saving up to 40% of memory space and could be applied here, too.

Figure 8 shows results of the actual test implementation. Here, only a part of the whole space captured by a CTA based PAOD study was used to ease and speed up the development. In areas where the bone anatomy does not vary significantly, we get very accurate results. Bone tissue is marked with very high probability, whereas vessels with contrast agent and calcifications are filtered out. On the other side, in regions with significant anatomical variability (Figure 8a), we do not get satisfactory segmentation yet. We partially attribute this to the currently insufficient number of records in the atlas and partially to the low number of control points used to define the TPS warping, rendering it to be not flexible enough.

8 Conclusion and Future Work

We proposed a new framework for knowledge-based segmentation of CT data sets by means of a probabilistic atlas. The atlas stores voxel density characteristics and spatial information about the anatomy by elastically matching manually segmented training data sets to a common reference space. In segmentation of an unknown data set, the information stored in the atlas is utilized by warping the data set to the atlas reference space and deriving probability information about the occupancy of a voxel from it.

This general concept was implemented for segmentation of bones from CTA data sets. The atlas was created by TPS warping of five manually segmented data sets to a common reference space and provides us with spatial distribution of the average bone tissue density and the average distance of bone/no-bone tissues. In segmentation, the dependency of the bone tissue density distribution along the longitudinal axis of the patient's body is extracted from the atlas, together with the spatial probability of bone tissue presence.

The technique is still in an early stage of development and is currently not reliable enough to be used in a daily routine. We see the highest potential for its improvement in the used elastic matching and classification techniques. Further investigation of the properties of the TPS transform is necessary, aimed at speed up and fit precision improvement. In this sense we intend to develop a procedure for automatic definition of the landmark set. Simultaneously, different elastic deformation techniques should be tested. The second source of possible improvements resides in the data classification technique. Obviously, the five data sets, which the atlas is currently based on, are not enough to build a reliable statistical model. Therefore, new data sets should be added, until atlas stability is achieved (i.e., adding new data sets should not influence the atlas significantly anymore). Further, new ways how to use the anatomic information of the atlas should be sought in order to make it more robust, namely in those cases, where the new anatomy significantly differs from the atlas. And finally, questions regarding the precision of the technique should be addressed, focused on limits of its application and ways how to measure its efficiency.

Acknowledgements

This work has been supported by grant No. P15217 of FWF Austria [1].

References

- [1] Visualisierungsmethoden für die periphere CT-Angiographie (AngioVis). Project P15217, FWF Austria.
- [2] Isaac N. Bankman, editor. *Handbook of Medical Imaging, Processing and Analysis*. Academic Press, 2000.
- [3] S Baum and MJ Pentecost. *Abrams' angiography: vascular and interventional radiology*. Little Brown, Boston, 1997.

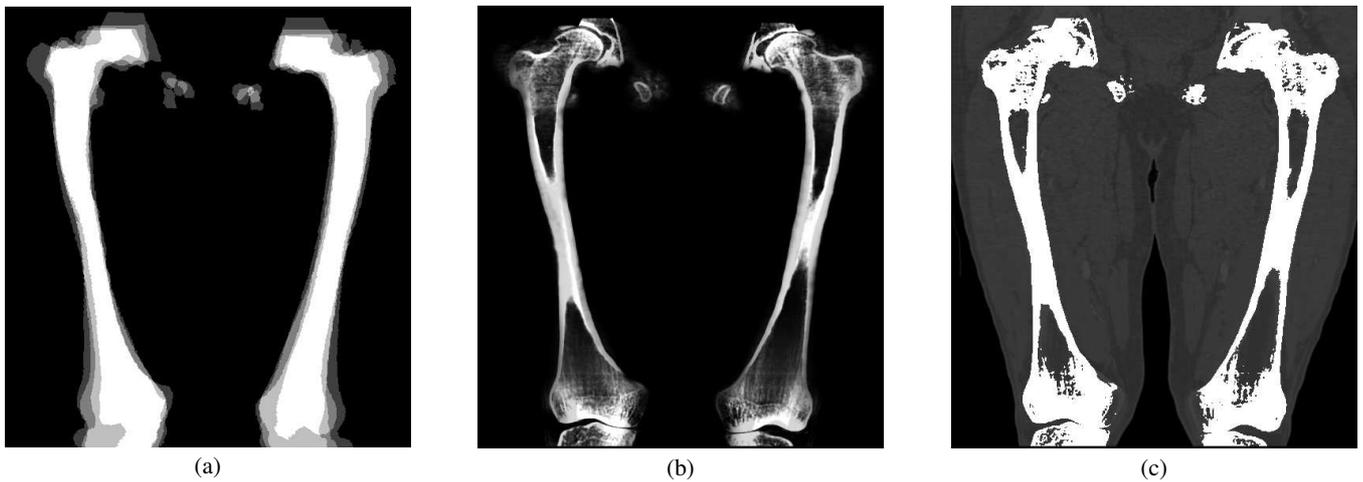


Figure 8: (a) Uncertainty caused by anatomic differences, which are not captured by the TPS transform. Compare the good overlap of the transformed bone masks in the central part of the bone and the not so good one at the ends. (b) the joint probability, (c) bone mask obtained from (b) by thresholding.

- [4] Fred L. Bookstein. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *T-PAMI*, 11(6):576–585, June 1989.
- [5] G. Borgefors. Distance transformations in digital images. *Computer Vision, Graphics, and Image Processing*, 34(3):344–371, 1986.
- [6] M. E. Brummer, R. M. Mersereau, R. L. Eisner, and R. J. R. Lewine. Automatic detection of brain contours in MRI data sets. *IEEE Transactions on Medical Imaging*, 12(2):153–166, June 1993.
- [7] R. A. Drebin, L. Carpenter, and P. Hanrahan. Volume rendering. *Computer Graphics (SIGGRAPH '88 Proceedings)*, 22(4):65–74, August 1988.
- [8] A.C. Evans, D.L. Collins, S.R. Mills, E.D. Brown, R.L. Kelly, and T.M. Peters. 3D statistical neuroanatomical models from 305 MRI volumes. In *Nuclear Science Symposium and Medical Imaging Conference*, volume 3 of *IEEE Conference Record*, pages 1813–1817, November 1993.
- [9] R. C. Gonzales and R. Woods. *Digital Image Processing*. Addison Wesley, 1992.
- [10] D.V. Iosifescu, M.E. Shenton, S.K. Warfield, R. Kikinis, J. Dengler, F.A. Jolesz, and R.W. McCarley. An Automated Registration Algorithm for Measuring MRI Subcortical Brain Structures. *NeuroImage*, 6:13–25, July 1997.
- [11] M. W. Jones and R. Satherley. Shape representation using space filled sub-voxel distance fields. In *International Conference on Shape Modelling and Applications*, pages 316–325. IEEE Computer Society Press, 2001.
- [12] A. Kanitsar, R. Wegenkittl, P. Felkel, D. Fleischmann, D. Sandner, and E. Groller. Computed tomography angiography: a case study of peripheral vessel investigation. In *Proceedings of IEEE Visualization 2001*, pages 477–480, October 2001.
- [13] Armin Kanitsar, Dominik Fleischmann, Rainer Wegenkittl, Petr Felkel, and Eduard Gröller. CPR: curved planar reformation. In *Proceedings of IEEE Visualization 2001*, pages 37–44, October 2002.
- [14] Ron Kikinis, Martha E. Shenton, Dan V. Iosifescu, Robert W. McCarley, Pairash Saiviroonporn, Hiroto H. Hokama, Andre Robatino, David Metcalf, Cynthia G. Wible, Chiara M. Portas, Robert M. Donnino, and Ferenc A. Jolesz. A Digital Brain Atlas for Surgical Planning, Model-Driven Segmentation, and Teaching. *IEEE Transactions on Visualization and Computer Graphics*, 2(3):232–241, September 1996.
- [15] A Koechl, A Kanitsar, Lomoshitz F, E Groeller, and D. Fleischmann. Comprehensive assessment of peripheral arteries using multi-path curved planar reformation of CTA datasets. *Eur Radiol*, 13:268–269, 2003.
- [16] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, May 1988.
- [17] J. B. Antoine Maintz and Max A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
- [18] John C. Mazziotta, Arthur W. Toga, Alan C. Evans, Peter T. Fox, and Jack L. Lancaster. A probabilistic atlas of the human brain: Theory and rationale for its development. the international consortium for brain mapping (ICBM). *NeuroImage*, 2(2a):89–101, June 1995.
- [19] Miloš Šrámek. ISEG - A system for interactive segmentation of 3D tomographic data sets. In J. Rozman, editor, *Proc. of 12-th international conference Biosignal '94*, pages 48–51, Czech Republic, 1994. Technical University Brno.
- [20] W. A. Barrett, W.A. and E.N. Mortensen. Interactive live-wire boundary detection. *Medical Image Analysis*, 1(4):331–341, 1996.
- [21] Simon K. Warfield, Michael Kaus, Ferenc A. A. Jolesz, and Ron Kikinis. Adaptive, Template Moderated, Spatially Varying Statistical Classification. *Medical Image Analysis*, 4(1), 2000.
- [22] X. Zeng, L. H. Staib, R. T. Schultz, and J. S. Duncan. Segmentation and measurement of the cortex from 3D MR using coupled surfaces propagation. *IEEE Trans. Med. Imaging*, 18(10), 1999.