

Matúš STRAKA*, Alexandra La CRUZ,
Arnold KÖCHL, Miloš ŠRÁMEK,
Eduard GRÖLLER, Dominik FLEISCHMANN

3D Watershed Transform Combined with a Probabilistic Atlas for Medical Image Segmentation

Recent advances in medical imaging technology using multiple detector-row computed tomography (CT) provide volumetric datasets with unprecedented spatial resolution. This has allowed for CT to evolve into an excellent non-invasive vascular imaging technology, commonly referred to as CT-angiography. Visualization of vascular structures from CT datasets is demanding, however, and identification of anatomic objects in CT-datasets is highly desirable. Density and/or gradient operators have been used most commonly to classify CT data. In CT angiography, simple density/gradient operators do not allow precise and reliable classification of tissues due to the fact that different tissues (e.g. bones and vessels) possess the same density range and may lie in close spatial vicinity.

We hypothesize, that anatomic classification can be achieved more accurately, if both spatial location and density properties of volume data are taken into account. We present a combination of two well-known methods for volume data processing to obtain accurate tissue classification. 3D watershed transform is used to partition the volume data in morphologically consistent blocks and a probabilistic anatomic atlas is used to distinguish between different kinds of tissues based on their density.

1. INTRODUCTION

Modern medical 3D imaging modalities, such as multiple detector-row computed tomography (CT) provide huge high-resolution volumetric datasets. For example, modern CT scanners have a spatial resolution well below 1 mm^3 . For a CT dataset covering the entire abdomen and lower extremities, this results in 2000 transverse images or more than 1GB of data per patient. This puts enormous demands on radiologists when inspecting these datasets. There is a growing need for computer aided pre-processing of such data, because usually just a small part of the whole dataset is relevant for establishing a patient's diagnosis.

For visualization and diagnostic assessment of vascular disease it is useful if tissues in images are segmented. Much work has been done in this field [1, 3], typically relying on the fact that voxels belonging to one object lie in the same density range. Areas with high gradient are

*<mailto:matus.straka@assoc.oeaw.ac.at>

Commission for Scientific Visualization, Austrian Academy of Sciences, Vienna, Austria

considered edges or borders between objects. When classifying CT-Angiography (CTA) datasets, however, one cannot simply rely on density/gradient values, because the density range of the diagnostically relevant tissues (vessels) substantially overlaps with less important structures, notably bone. Human operators do not decide which voxel belongs to which tissue based on density/gradient information alone, but they use anatomic knowledge and “see” the entire neighbourhood of a voxel for its classification. Therefore it is our hypothesis that the additional use of higher order spatial information will yield more accurate differentiation between vascular and bony structures when compared to techniques, which rely on density/gradient information alone.

This work is part of an interdisciplinary research project [8], aimed at visualization and treatment planning of peripheral arterial occlusive disease (PAOD) by means of CTA. PAOD is characterized by the formation of atherosclerotic plaque and vessel wall calcifications, which may lead to luminal stenoses (narrowing) or complete occlusion of the involved arteries [2]. Diminished blood flow to the legs causes restricted mobility, pain, and tissue loss, eventually leading to death.

CTA of the peripheral arteries is performed using multiple-detector row CT which allows the acquisition of a series of 1500-2000 transverse images of 512x512 12-bit pixels through the anatomic region of interest. While the scanning procedure requires only 20-40 seconds, the analysis of such huge datasets within an acceptable time period is impossible without the application of dedicated visualization techniques. Manual image editing, segmentation and vessel identification (labelling) using currently available commercial medical visualization workstations is prohibitively time consuming – in the order of two to four hours, even for a well trained operator.

In CTA, vessels are enhanced by the intravenous injection of iodinated contrast medium. Unfortunately, contrast medium enhanced vessels possess a density range that overlaps with the density range of bone. This leads to vessel identification errors and therefore pre-segmentation of bones is desirable. Once the bone tissue is identified, it can be masked out. Then vessel visualization (e.g. with Maximum Intensity Projection (MIP)) is straightforward, because they are the highest-attenuation structures remaining in the dataset.

Bone tissue labelling in CT data is often a very simple problem. Cortical bone (Figure 1) has significantly higher X-ray attenuation than other tissues and the attenuation values (expressed in Hounsfield units (HU)) are rather constant across individuals. Thus, simple thresholding may be sufficient for its identification. However, simple thresholding will not identify trabecular bone, which constitutes a considerable proportion of bone volume. Its density is only slightly above that of soft tissue, and – in addition – it overlaps with the density of contrast-medium enhanced vessels. Moreover, there is a considerable variability of trabecular bone density between individuals. Partial volume effect (representing density of both bone and soft tissue in a single voxel) and image noise further aggravate the problem. Thus, there is a need to process the datasets in large blocks (if not in whole) and to incorporate some kind of anatomical knowledge. Any algorithm that lacks this information will fail in certain situations such as in the presence of noise, artifacts, and poor vessel opacification and vessel occlusions.

The basic idea of the method presented in this paper is reflected in the combined application of two techniques: a-priori knowledge derived from a density based probabilistic atlas is used to locate characteristic parts of bones and a watershed transform to identify spatially coherent subvolumes, regardless of their density.



Figure 1: Different types of bone tissue (femur area)

The fundamental idea of the proposed method of tissue classification, using information in both the spatial domain and the density domain is reflected in the combined use of density based spatial information (derived from a bone density probabilistic atlas) and shape based information obtained from a watershed transform.

In the first step of our method, a pre-computed atlas of bone density information was used to assign a bone-probability to each voxel. This constituted several regions of neighbouring voxels that possessed high probability. These regions typically follow objects in the datasets, mainly individual bones and vessels. To label voxels belonging only to selected object (e.g. femur), region-growing algorithm was used. Selection of object was provided by seeding start points with operator intervention.

The next step was partitioning of the whole volume in homogenous regions using the watershed transform, followed by classification of 3D watershed regions using the previously acquired bone mask. Finally we dilated the mask to cope with partial volume effect.

While the method presented in this paper is designed specifically for bone extraction, we believe that it is applicable to a much wider range of problems and for identification of any kind of object that might have variable densities throughout the dataset.

The paper is organized as follows. In the second section we describe the generation of the density based probability atlas and in the third section we describe the 3D watershed transform. In section four we propose how to use the results of both, the probabilistic atlas and a 3D watershed transform for segmentation of unknown data sets. Section five describes our method of implementation and results, and in the last section we summarize the paper and sketch directions for future work.

2. BONE DENSITY PROBABILITY ATLAS

Distribution of bone density in lower extremity CTA data is dependent on anatomic location, which allows the creation of an atlas of bone densities. We based our density atlas on several assumptions:

- Density of human bones varies substantially along the z-axis (longitudinal axis of a human body, Figure 2),
- Density of human bones varies much less in transverse slices (compared to inter-slice density differences),
- Human body is left-right symmetric, and
- Density variation between patients in corresponding body points is lower than the density variation between different kinds of bone tissue in the same patient (Table 1).

Such a density atlas [4] was constructed by manually segmenting bones [5] in four clinical CTA datasets. Then, using manual landmark specification and non-rigid

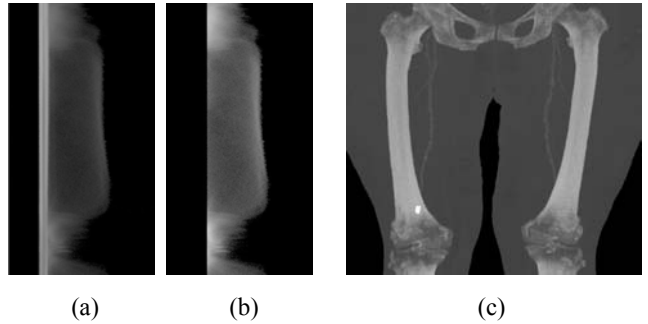


Figure 2: Dependency of density histograms on the slice position. Each row in (a) and (b) shows the histogram (-1000HU to 3095HU) of the corresponding slice in a dataset, the MIP of which is presented in (c). (a) shows histograms of all body tissues while (b) only those of segmented bones.

Table 1: Mean values and deviations of bone tissue in various locations in the atlas (atlas based on four manually segmented datasets)

Location	Mean value	Standard deviation
Distal femur (knee)	424, 62 HU	187, 01 HU
Middle of femur	1166,91 HU	415, 49 HU
Proximal femur (hip)	327,74 HU	151, 05 HU

transformation (thin-plate-spline warping (TPS)), all datasets were mapped to a “common atlas space”, which also solved problems regarding different sampling rates and patient positions. In the last step, the warped datasets were processed statistically and mean bone densities together with the variation for each transverse slice were computed.

When applying the information from the probabilistic atlas to an unknown dataset, the atlas space has to be warped to the space of the unknown dataset (by TPS warping). Then, by modelling the density distribution by the Gaussian distribution one can get the probability of a voxel belonging to the bone tissue for each slice in the dataset [4].

To identify voxels that should belong to one object (and being of bone tissue type), we apply the following technique: Starting in a specified (e.g. manually by operator intervention, or derived from spatial atlas) point in bone tissue, we grew a region using neighbourhood connections while checking the bone probability level. We grew the region only if the bone probability level of the actual voxel was above a specified probability value (in our case, 0.8). The fact that edges of bones/vessels affected by partial volume effects (PVE) usually do not belong to areas with high “bone probability” prevents merging with not-bone objects. Depending on patient anatomy, however, sometimes individual bones may be connected and therefore constitute one single object. This was not a problem at this stage because our goal was to identify any kind of bone tissue.

The result is typically a very coarse bone mask (Figure 3c). The mask is quite exact in areas of cortical bone, but in areas of trabecular bone many pixels are incorrectly labelled, which is an unsatisfying result. Therefore this intermediate result is combined with successive 3D watershed transform.

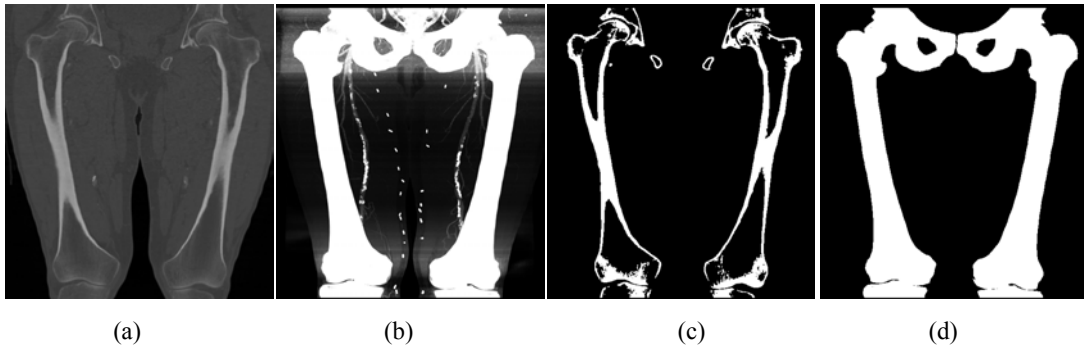


Figure 3: (a) Vertical slice of original the dataset, (b) MIP of the bone probability volume (vessels are also enhanced and still present), (c) bone mask of (a) (vertical slice) obtained by region growing with help of probability atlas (vessels suppressed), (d) a MIP of (c).

3. 3D WATERSHED TRANSFORM

The watershed transform is a well-established image-processing technique for grey-scale images. It is based on topographical interpretation of the gradient image. The density magnitude is considered as a topographical relief, where the brightness value of each voxel corresponds to physical elevation. The water flowing down the elevation is always following the gradient direction (direction with maximal density change) to the nearest local minimum. Areas with common local minimum constitute watershed regions, and the borders between these regions constitute the watersheds. Vincent and Soille [7] showed that an alternative of the watershed transform as an immersion algorithm is also possible.

The procedure results in partitioning of the image in catchment basins, the borders of which define the watersheds. In this way, watershed algorithms create closed borderlines, which can be used for segmentation. A common problem of watershed algorithms is oversegmentation. Due to noise and randomness of natural medical data, the volume is partitioned into innumerable small regions. By low-pass filtering of the data (e.g. by the means of a Gaussian filter) it is possible to reduce the number of unimportant details and to reduce the number of regions. The disadvantage of this approach is that the borders of the regions are shifted (due to the nature of low-pass filtering) and become imprecise. Hierarchical merging of regions [6] has been suggested as a solution in this case. Alternatively, regions can be merged based on their statistical attributes (e.g. mean value, variation, etc.); eventually both techniques can be combined.

As it was already mentioned, the techniques results in heavily oversegmented data (Figure 4a). In our case, the simple hierarchical merging (based on Gaussian filtering) led either to omitting of important details (in case when the top of the hierarchy was heavily filtered), or the oversegmentation was not suppressed enough. Therefore, we included region merging based on the density mean value difference. We merged regions that differed no more than 10 HU in several




Figure 4: A vertical slices of a 3D dataset (hip joint area, cropped): (a) oversegmentation is a characteristic of the plain watershed transform (b) results of oversegmentation suppression achieved by merging regions according to their density mean values and (c) results of oversegmentation suppression achieved by both merging regions according to their density mean values and hierarchical merging

iterations (consecutively increasing the threshold from 1HU to 10HU). Next, we applied the hierarchical merging, which kept the region borders at their original (not-filtered) position but decreased the number of regions (Figure 4).

4. MERGING RESULTS OF PROBABILISTIC ATLAS AND WATERSHED TRANSFORM

To improve the results acquired using the probabilistic approach we incorporated edge information obtained with the 3D watershed transform into a final bone classification. In each watershed region we computed the number of voxels, which were labelled as bone in the probabilistic mask. If the number of labelled voxels in a region exceeded a specified threshold (e.g. 10%), we labelled the whole region as bone (Figure 6).

These results are satisfying if we need to identify objects for e.g. surface shaded visualization (Figure 7b). The fact that the borderlines follow the areas of the maximum of gradient magnitude area could cause that due to PVE there might be residual voxels not identified as bone in the datasets and in case of MIP this might cause artifacts (Figure 5c). Therefore we dilated the bone mask in the final stage (with 5x5x5 structuring element).

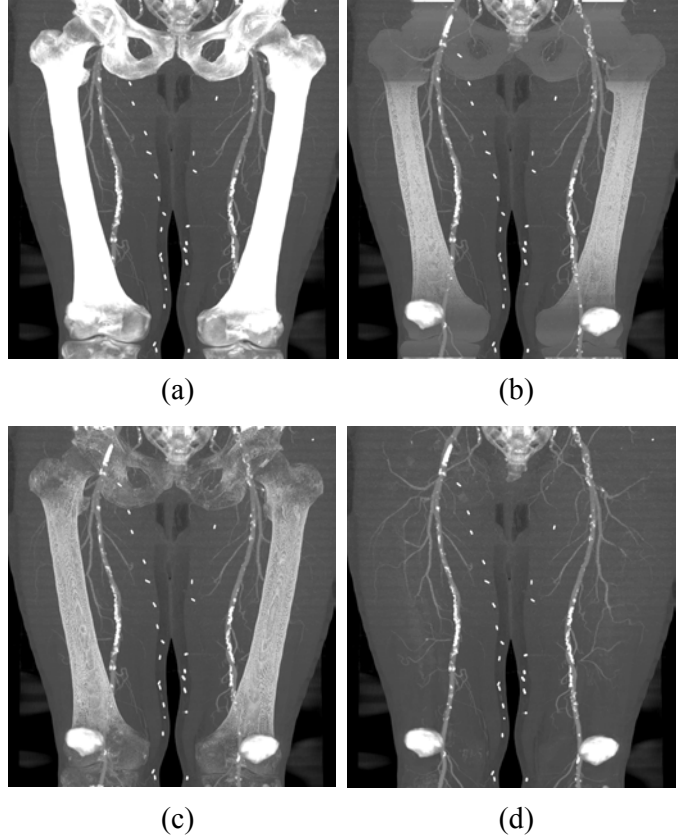


Figure 5: (a) MIP of the original dataset, (b) MIP, bone tissue partially removed using mask obtained with probabilistic atlas. (c) MIP, bone tissue removed using mask extended with 3D watershed transform, (d) MIP, bone tissue removed using the dilated mask. (Patella and sacro-coccygeal bone not modelled).

5. IMPLEMENTATION AND RESULTS

The presented technique shows a way to segment bone tissue in CTA datasets reliably. Compared to simple threshold/gradient algorithm it delivers much better results (Figure 7c).

To apply aforementioned technique, first building of a density probability atlas is needed. This is done by transforming manually segmented datasets into common atlas space and statistically processing them. Then, for segmentation, non-rigid transformation of the atlas data to the unknown dataset space is carried out and voxel bone-probabilities are evaluated. Region growing with threshold is used to separate different object in probability volume. In parallel, 3D watershed

transform is applied to the unknown dataset to obtain edge information. In the end, watershed regions are classified according to the mask obtained.

Determination of algorithm constants is quite self-explaining and straightforward. The threshold for region growing should express a probability in probability volume. The higher value, the less is the probability that the region growing will grow to incoherent objects, but also some voxels might be not labelled, therefore they will be missing in final mask. In our case, values approx. 0,7-0,8 proved to be suitable.

The threshold determination for watershed region classification according to the bone mask is even simpler. In a fact, even region with only one voxel labelled as bone can be considered as bone region. To be on a safe side, we used threshold about 10%. Need to be mentioned that both thresholds are tied together partially, so when choosing lower region-growing threshold one should consider higher value for region classification. But these values are not crucial for algorithm results.

Parameters of 3D watershed transform (kernel size of Gaussian filter, levels of hierarchy, threshold for watershed region mean value merging) are to be determined depending on the input data and required precision [6].

The aforementioned methods were implemented on a PC workstation with 1,6GHz CPU and 2GB of RAM. Processing of a 512x256x512 voxel dataset took 90 seconds for watershed labelling, 5-6 minutes to create an object map and 4-5 minutes to process density similarity merging. Other tasks were computed in a time comparable with data loading/saving time (4-5 seconds for loading, 15-20 seconds for saving.) The memory requirements were approx. 150MB for the dataset and additional 600 MB (gradient image, object map) for processing.

An example of achieved segmentation results is presented in Figure 7. Difference between mask obtained with probability atlas and mask extended with 3D watershed transform is not visible



Figure 6: Bone mask extended by the 3D watershed region classification (compare to Figure 3c)

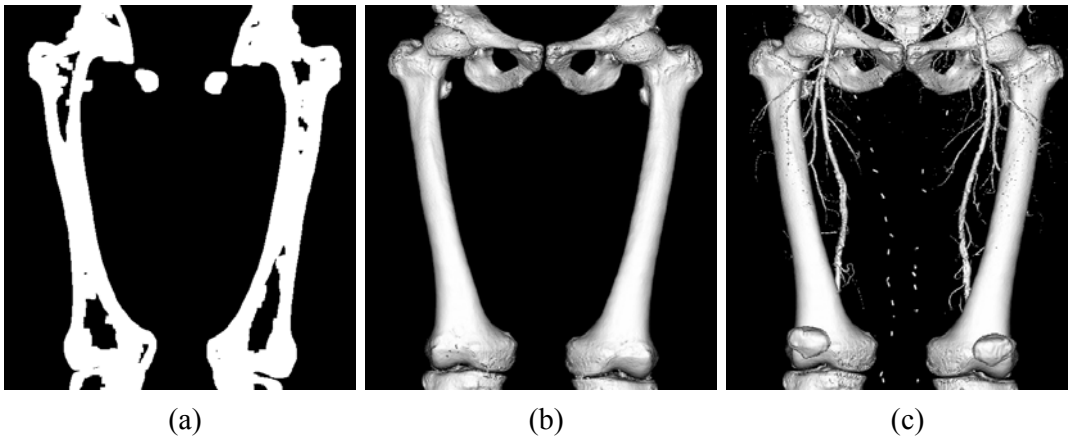


Figure 7: (a) Final bone mask after dilation (vertical slice from dataset), (b) final bone model (surface shaded), without dilation, (c) for comparison: object mask acquired only with simple thresholding (th = 200HU), surface shaded. (Patella and sacro-coccygeal bones were not modelled)

in MIPs (Figure 5). However, there is visible difference in the knee area in Figure 3c and Figure 6. (Compare areas of trabecular bone.)

6. SUMMARY AND FUTURE WORK

In this paper, we proposed a new method to identify bone tissue in 3D medical datasets. The proposed method delivers better results compared to methods based just on thresholding, region growing and methods currently used in the clinical environment.

There are still persisting problems, though. The most important is that bone marrow cannot be easily segmented as bone, mostly because of its density range, that overlaps with soft tissue density range. The spatial atlas in [4] was not yet precise enough to provide us with reliable information. Second problem is that this acquired mask might be not closed (or at least we cannot guarantee this); therefore identification of inside and outside will not work. Also, areas where bone is at least dense it might happen that no voxel of respective watershed region will be identified as bone and therefore the whole region will lack this information. This is not correct from our point of view, although it might not cause problems while using results of the bone classification for vessel tracking or creating MIP images, because density of these regions is below of that of the contrast medium enhanced vessels. This algorithm also needs further evaluation and testing, especially in areas where vessels and bones touch (vertebral area, tibia area, feet).

In the future, our work should further continue in investigation of reliable image segmentation techniques, mainly in shape recognition. We are convinced that there is no other way than to:

- Recognize partial outlines of objects (both in 2D and 3D)
- Compare them with known templates and
- Using these templates to complete the missing parts.

7. ACKNOWLEDGEMENTS

This work has been supported by the Austrian Science Fund (FWF) grant no. P15217 [8].

BIBLIOGRAPHY

- [1] Bankman I.N., editor, Handbook of Medical Imaging, Processing and Analysis. Academic Press, 2000.
- [2] Baum, S., Pentecost, M.J., Abrams' angiography: vascular and interventional radiology. Little Brown, Boston, 1997
- [3] Gonzales R.C., Woods R., Digital Image Processing. Addison Wesley, 1992.
- [4] Straka M., Šrámek M., Bone Segmentation in CT-Angiography Data Using a Probabilistic Atlas (Submitted paper). Vision, Modeling and Visualization Workshop, Munich, Germany, 2003
- [5] Šrámek M., 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.
- [6] Šrámek M., Dimitrov L.I., Segmentation of tomographic data by hierarchical watershed transform. Journal of Medical Informatics and Technologies, 3, pages MI-161—MI-169, 2002
- [7] Vincent L., and Soille P., Watersheds in digital spaces: an efficient algorithm based on immersion simulations, IEEE Transactions on Pattern Analysis and Machine Intelligence. 13:583-598, 1991.
- [8] Visualisierungsmethoden für die periphere CT-Angiographie (AngioVis), Project P15217, FWF Austria