Automated analysis and visualization of preclinical whole-body microCT data

Baiker, M.

8

Summary and Future Work

8.1 Summary and Conclusions

In this thesis, several strategies are presented that aim to facilitate the analysis and visualization of whole-body in vivo data of small animals. Based on the particular challenges for image processing, when dealing with whole-body follow-up data, we addressed several aspects in this thesis. The developed methods are tailored to handle data of subjects with significantly varying posture and address the large tissue heterogeneity of entire animals. In addition, we aim to compensate for lacking tissue contrast by relying on approximation of organs based on an animal atlas. Beyond that, we provide a solution to automate the combination of multimodality, multidimensional data.

Chapter 1 gives a general introduction to molecular imaging and identifies the challenges for image processing. Furthermore it presents a review of published work that might contribute to face the challenges. It is pointed out that crucial aspects like dealing with large postural variation and tissue heterogeneity as well as handling missing tissue contrast in whole-body data are still problems to be solved and these issues are defined as the goals of this research. The chapter concludes with an outline of the thesis.

Chapter 2 describes the process of constructing articulated skeleton atlases, based on three publicly available whole-body animal atlases (MOBY mouse [59], Digimouse [36], SD Rat [60, 61]). The process includes manual skeleton segmentation, joint localization and definition of anatomically realistic kinematic joint models. The chapter also contains some application examples for the usage of these articulated atlases.

The articulated MOBY atlas is used in Chapter 3 for whole-body segmentation of mice in in vivo MicroCT data. The method is based on registration of the articulated MOBY skeleton to the skeleton extracted from the data. A hierarchical model, combined with restrictions on the amount of DoFs for the registration of the individual bones, renders the method robust to postural variation and pathological bone tissue modifications like osteolytic lesions. Subsequently, the lungs and the skin of the MOBY are registered to the data and major organs are approximated using TPS interpolation. We are able to obtain registration errors (Euclidean surface distances) of less than two voxel dimensions

for the skeleton and the lungs respectively and less than one voxel dimension for the skin. We obtain Dice coefficients of organ volume overlap between manually segmented and interpolated skeleton and organs that vary between 0.47 ± 0.08 for the kidneys and 0.73 ± 0.04 for the brain. Given the very large postural variability of the data that we used for evaluation, the results are very promising and comparable to previously published results. If subvoxel accuracy for the skeleton and the lungs is needed for a specific application, using this method alone does not suffice. However, the results can be used for initialization of a more accurate, intensity-based registration approach. How this can be done is presented in Chapter 6. The most suitable modality for applying the atlas-based skeleton registration presented in this chapter is MicroCT, because the skeleton can easily be extracted from the data. However, the method can be applied to other modalities, such as MicroSPECT [165], as well, as long as an approximation of the skeleton is available.

In Chapter 3 we showed that the skeleton of an animal can be coarsely segmented into individual bones. The true potential of being able to do this automatically is demonstrated in Chapter 4. In this chapter, a method called Articulated Planar Reformation (APR) is presented that maps subvolumes of the original MicroCT dataset to a standardized reference frame. With this framework, a structure of interest, like a particular bone, in multiple MicroCT datasets of different animals or of the same animal at different time points, can easily be related to each other in a side-by-side visualization. Usage of a common reference assures that the structures of interest from all input datasets are aligned to each other. In this way, morphological changes, which can e.g. be disease or treatment related, become apparent. Thus the method allows to easily and automatically navigate through whole-body data and to intuitively assess follow-up datasets. To be able to compare the datasets in a quantitative manner as well, several measurement and visualization strategies are introduced, including bone change visualization, overlays and checkerboard visualizations. Color-coded mapping of the registration error onto the structure of interest allows to judge the reliability of the results. Furthermore, definition of a structure of interest beforehand has the advantage, that not the entire whole-body datasets have to be loaded into memory. This is of great importance especially if multiple time points have to be compared. The data can be loaded with higher resolution because of the restricted Volumes of Interest. The approaches were evaluated by two field scientists who found it very helpful in assessing whole-body follow-up data.

In Chapter 5, the APR framework is extended by a concrete example where accurate quantification is required to follow cancer induced osteolysis over time. First, a structure of interest is selected automatically, from datasets acquired at three different time points, based on the methods presented in the previous chapters. Since the articulated skeleton registration yields a coarse segmentation of the skeleton only (the DoFs of the individual registrations are restricted), we subsequently propose a method for accurate segmentation of different bones. In this chapter, the structure of interest is the tibia and we apply the segmentation method to separate it from the femur. However, the proposed strategy is suitable for segmentation of other bones or the entire skeleton as well. Based on the segmentation of the tibia, the bone volume is followed over time. We compare the automated bone volume measurements to those determined by two human observers. Thorough sta-

tistical analysis reveals that the automated method does not differ significantly from the human observers ($p = .10$). The measurements show high correlation between the observers ($r = 0.9996$) and between the observers and the automated method ($r = 0.9939$ and $r = 0.9937$). In addition, Bland-Altman plots reveal excellent agreement between the observers and the automated method (Interobserver variability: $0.59 \pm 0.64\%$, Obs1 vs. Auto: $0.26 \pm 2.53\%$ and Obs2 vs. Auto: $-0.33 \pm 2.61\%$). Therefore we can conclude that the method can substitute a human observer in monitoring osteolysis over time. Additionally, the chapter presents automated cortical bone thickness measurements. The determined values can be color-coded and mapped onto a surface representation of the tibia to follow thickness variations over time.

It was mentioned above that highly accurate registration of the skeleton cannot be achieved with the articulated skeleton registration presented in Chapter 3. However, intensity-based registration approaches that are capable of achieving high accuracy are not robust enough to deal with data of animals with greatly varying posture. In Chapter 6, we propose to combine the robustness of the articulated registration with the accuracy of intensity-based registration. To this end we formulate the registration between two datasets as an optimization problem. The registration criterion is based on intensity similarity on the one hand and on the Euclidean distance between corresponding sets of landmarks on the other hand. The sets of landmarks in the two datasets are determined by registration of the MOBY atlas to both datasets according to Chapter 6. The optimization can be performed very effectively and fast by the software package `elastix` [154] and we could achieve a significant speedup with an accuracy comparable to previously published work. The combination of the two terms in the optimization outperforms the registrations based on either of the terms used exclusively and we obtain subvoxel accuracy for the skeleton and voxel accuracy for the skin, based on the Euclidean surface to surface distance.

In Chapter 7 we focus on a very common problem in preclinical molecular imaging practice: the fusion of 2D bioluminescence imaging data and 3D MicroCT data. Our approach is to reconstruct an approximation of the animal skin surface in 3D, based on two or more skin silhouettes in 2D photographs, and the calculation of a 3D distance map from the reconstruction. Subsequently, we use the distance map to register a MicroCT-based animal skin surface. Several penalty terms are introduced to deal with the shape differences between the two representations of the animal skin. The method yields subvoxel accuracy for the Euclidean skin surface distance for synthetic data, demonstrating the potential of the method. The experiments for real data are less accurate, with a mean surface distance around two voxel dimensions. We argue that shape differences between the reconstructed skin surface and the MicroCT based skin surfaces are probably the main reason for this. For example the animal ears are clearly visible in the photographs but because they are very thin, they are not present after skin extraction from MicroCT. Despite the limitations, the results are comparable to the results that can be obtained, by registration based on manually determined anatomical landmarks.

Based on the research goals formulated in section. 1.4 and the considerations in this chapter, we think that we achieved our targets to a large extent. The issue of automated analysis of animal data including large postural variability between animals, like follow-

up data, is covered by the articulated registration (Chapter 3) as well as subsequent reformatted visualization and quantification (Chapter 4) and accurate whole-body registration (Chapter 6). Dealing with heterogeneous tissue properties is partly covered since we pay special attention to bone and the lungs in Chapter 3 and by regularization of the registration in Chapter 6. But other tissues (major organs) are treated as being equally stiff during the interpolation and unrealistic deformations may occur. The problem of missing tissue contrast in whole-body data is approached by the strategy to rely on an anatomical animal atlas. We demonstrated this for MicroCT, but approaches for the analysis of data from other modalities could benefit from using the animal atlas as well. Because of the time constraints of this research, the issue of combining data from multiple modalities could only be addressed partly. However, initial results are promising.

8.2 Future work

There are several possibilities for future work, based on the obtained results in this thesis. Regarding the articulated skeleton registration it would be interesting to perform the registration of several bones simultaneously, subject to the motion constraints of the joints by which they are connected. In principle, the sequential nature of the bonewise registration bears the possibility that registration errors are propagated to registrations of distal bones. We compensate for that to a certain extent but it would be interesting to see how simultaneous registration performs in comparison. However, this would require an optimization strategy that can handle many more registration parameters. An improvement would be to include the limbs of the animal as well to derive corresponding landmarks on the skin. Currently, only the skin of the torso is used because the torso contains sufficient corresponding landmarks to define an interpolation of the major organs only. Furthermore, the approximation of the organs could be done in a more elaborate manner instead of taking TPS interpolation. To attach different stiffness properties to each organ, a Finite Element Model of the MOBY atlas could be created. The correspondences on the skeleton, lungs and the skin could be used to derive an estimate of external forces on the atlas, leading to realistic deformations. However, such an approach would be very calculation intensive.

A very interesting experiment would be to use the heterogeneous tissue model after segmentation of an animal with an implanted light source, and assign realistic optical tissue properties to the organs. This may be very useful for light source reconstruction by Bioluminescence Tomography because recently, these approaches were modified for in vivo applications [166, 167]. There is a great need for heterogeneous tissue models to ensure accurate results [90] and it seems, that the boundaries of the organs do not have to be known very accurately in order to still get good results [9].

The common reference of the Articulated Planar Reformation framework is currently tailored for MicroCT data and therefore only defined for the bones. However, one could take an organ approximation and define a VOI around it, similar to the automated scan planning strategy presented in [10]. Doing so for several timepoints it would be possible to map organs to the template as well. However, for MicroCT this is not very meaningful because of the missing contrast.

Summary and Future Work

The accurate bone segmentation technique presented in Chapter 5 was demonstrated for the femur, the tibia and the pelvis. However, the method could be applied to other bones of the skeleton as well, as long as the bone is included in the articulated MOBY skeleton, yielding an accurate whole-body segmentation of the individual bones. Furthermore, the bone volume is currently measured based on thresholded data. However, it might be a better choice to rely on bone mass rather than bone volume in case of bone reformation, based on the grayvalue distribution.

The whole-body registration was evaluated on non-contrast-enhanced MicroCT only. Normalized Cross Correlation is used as the similarity measure but other measures like Mutual Information are applicable as well. Therefore it would be interesting to test the method on contrast-enhanced data, eventually with varying organ contrast between different datasets. The more contrast there is, the better the registration should work.

An interesting extension of the 2D/3D registration method would be to make it independent of an accompanying MicroCT dataset. This could be achieved by relying on an articulated atlas that contains the skeleton and the skin. The skeleton could be articulated, subject to the joint constraints, with a subsequent approximation of the shape of the animal, thus the skin. This could be based on a weighted combination of the bone transformations [92] or paying special attention to the invertibility of the individual transformations [35] or by 'skinning' approaches from the computer animation literature (e.g. [168, 169]). The result would be an approximation of the animal interior, based on several photographs only.

Finally, the mentioned restrictions of the proposed 2D/3D registration method in Chapter 7 should be addressed by automation of the silhouette extraction from the photographs. This should include a priori knowledge of the shape, because not all features visible on the photographs (like the ears) are present in the extracted surface from MicroCT. The combination of such an approach with the heterogeneous tissue model of the animal, would be an important step towards a fully automated Bioluminescence Tomography system for in vivo applications.