



Universiteit
Leiden
The Netherlands

Automated analysis and visualization of preclinical whole-body microCT data

Baiker, M.

Citation

Baiker, M. (2011, November 17). *Automated analysis and visualization of preclinical whole-body microCT data*. Retrieved from <https://hdl.handle.net/1887/18101>

Version: Corrected Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/18101>

Note: To cite this publication please use the final published version (if applicable).

4

Articulated Planar Reformation for Change Visualization in Small Animal Imaging

This chapter is based on:

Articulated Planar Reformation for Change Visualization in
Small Animal Imaging

Peter Kok, Martin Baiker, Emile A. Hendriks, Frits H. Post,
Jouke Dijkstra, Clemens W. G. M. Löwik, Boudewijn P. F. Lelieveldt,
Charl P. Botha

IEEE Transactions on Visualization and Computer Graphics, 2010, 16(6):1396-1404

Chapter 4

Abstract

The analysis of multi-timepoint whole-body small animal CT data is greatly complicated by the varying posture of the subject at different timepoints. Due to these variations, correctly relating and comparing corresponding regions of interest is challenging. In addition, occlusion may prevent effective visualization of these regions of interest. To address these problems, we have developed a method that fully automatically maps the data to a standardized layout of subvolumes, based on an articulated atlas registration. We have dubbed this process articulated planar reformation, or APR. A subvolume can be interactively selected for closer inspection and can be compared with the corresponding subvolume at the other timepoints, employing a number of different comparative visualization approaches. We provide an additional tool that highlights possibly interesting areas based on the change of bone density between timepoints. Furthermore we allow visualization of the local registration error, to give an indication of the accuracy of the registration. We have evaluated our approach on a case that exhibits cancer-induced bone resorption.

4.1 Introduction

MOLECULAR IMAGING offers new possibilities in lifescience research, because it allows *in vivo* imaging of biochemical processes at a macroscopic level. This means that the onset of pathological processes can be studied at an early stage, long before large-scale anatomical changes occur [3], and followed over time in the same subject [2]. Nowadays a large range of small animal scanners is available that enable acquisition of data of the entire animal body and thus to assess disease progression globally. This is of high interest for example in oncology research because it enables monitoring metastatic behavior of different types of cancer [5, 6].

Typical molecular imaging modalities are Bioluminescence Imaging (BLI) and Fluorescence Imaging (FLI), which enable *in vivo* imaging of gene expression. Both have a very high sensitivity but their spatial resolution is low, which complicates linking the signal to the corresponding anatomy [7]. Therefore, the data is often fused with a modality that provides high anatomical detail, such as microCT. An example is given in [5], where BLI is used to monitor metastatic activity of breast cancer cells in the mouse skeleton over time. The BLI data is combined with microCT, to assess cancer-induced bone resorption. However, since light has a very limited penetration depth in biological tissue, the positioning of the animal in the scanner during the acquisition is critical to achieve maximum optical signal. As a result, it is difficult to standardize the data acquisition protocol and animal positioning can vary to a great extent among different animals at the same time point or the same animal in a follow-up study (Fig. 4.1). Lacking anatomical structure in optical imaging except for contours, the subject is placed in the BLI and the CT scanner in the same holder, in order to minimize posture differences and to facilitate the registration and fusion of BLI and microCT data [7]. Thus, the posture variability issue transfers to microCT as well. Also, in follow-up studies with only microCT, there is no standardized positioning protocol and the animal posture will greatly vary.

In the current workflow, the biologists resolve this postural variability by visually comparing subsequent image sets to find differences with regard to their hypotheses on drug effects and disease progression. In the case of an experiment on cancer-induced bone resorption, this would result in a search for local intensity decreases in the bones. Researchers that aim at comparing the same structures, e.g. a particular bone, face the problem that they have to localize the structure of interest in each dataset separately, which is complicated by the fact that these structures of interest are not aligned due to the large differences in animal posture between time points. To a large extent, the biologists are able to form a mental picture of the entire subject and the relative position of a region of interest, but we aim to standardize this into a fully automatic procedure that facilitates comparison and data interpretation.

In a first step to address the problem of varying animal posture, we earlier presented an approach for fully automated skeleton registration [17]. The method registers the skeleton of an anatomical animal atlas bone by bone to a skeleton derived from whole-body microCT data. Subsequently, it is possible to interpolate low-contrast areas (most major organs), yielding a segmentation into these organs. The resulting registration and segmentation allow to highlight the same structure in several datasets and significantly

Chapter 4




Figure 4.1: Three examples that demonstrate the large variability in animal posture between different scans. The subjects were scanned in supine (left image) and prone (right two images) position. Shown are isosurfaces of the CT data corresponding to the skin and the skeleton.

facilitate data navigation, but still require multiple, unaligned datasets to be analyzed side by side.

Therefore there is a great need for visualization methods that present one or multiple whole-body datasets in an integrated and interactive fashion: This would greatly facilitate the assessment of differences between datasets qualitatively and quantitatively.

In this work, we present several methods to facilitate the comparison of whole-body data from cross-sectional or follow-up studies. Our contributions are threefold:

- We describe a method to fully automatically transform anatomically relevant regions of interest from a whole-body small animal dataset into a standardized layout in order to facilitate comparison despite large differences in animal posture.
- Based on this standardized layout, we present a comprehensive coupled view approach for the visualization of change in whole-body datasets. Our approach is based on a two-level localization strategy that enables researchers to quickly identify areas of interest in a global whole-body view, and then to zoom in on and study the changes in a linked local per-segment view. This process is enhanced with the integration of derived features that further highlight areas of change.
- Finally, we investigate the role and utility of our application in molecular imaging research by means of a case study with two experienced domain scientists.

In this paper we focus on the analysis of microCT data because our whole-body segmentation method is based on this type of data. However, the presented methods can be used for other modalities as well, given an atlas and a registration of it to the skeletal structure of the animal.

In the next section, we discuss related work and in Section 4.3 we describe our approach in detail. Section 4.4 addresses implementation details and in Section 4.5 we look at the performance and evaluate our work. Finally, we present conclusions and future work in Section 4.6.

4.2 Related work

Visualization of three dimensional data often suffers from effects such as occlusion and the lack of depth cues. In order to gain a better understanding of three dimensional data in this respect, it is useful to map it to a standardized frame or onto a 2D plane. This becomes even more relevant when different three-dimensional datasets are to be visually compared.

Mapping to a standardized frame can be done for elongated structures such as vessels or the colon. Methods to achieve this are often based on curved planar reformation (CPR), which can be used to map tubular structures onto a plane. Kanitsar *et al.* [115] describe three methods to generate CPR images: projection, stretching and straightening. An application based on CPR is presented by Ochi *et al.* [116], who describe how CPR can be useful for the evaluation of aneurysms in CT angiography data. Lee *et al.* [117] present an approach that results in a topological and orientation invariant visualization of vascular trees. Finally, Ropinski *et al.* [118] describe a method to create a standardized visualization of mouse aorta PET and CT scans by applying CPR. Several visualizations are applied to enable comparisons between data.

Similar approaches are available to visualize the colon, which can be virtually unfolded or flattened to a 2D representation. Two of these methods are described by Hong *et al.* [119] and Lee *et al.* [120]. The resulting visualization enables showing the entire inner colon surface in a single view. An approach that combines CPR with volume rendering of the colon is described by Williams *et al.* [121].

The visualization of three-dimensional objects with a complex structure is facilitated by decomposition of the structure into distinct elements. This decomposition enables the application of exploded views, which is a commonly applied technique for depiction in technical drawings. It can clarify a visualization because it shows elements that would otherwise be hidden behind others. In Li *et al.* [26], an interactive exploded view is described that can be automatically generated from a complex solid model. Sonnet *et al.* [122] present a technique that uses a probe to interactively explode a 3D model.

Exploded views can also be used for the direct visualization of volume data. An approach for this is described by Bruckner *et al.* [123], where user-indicated parts of volumetric data are arranged based on a force model. A different technique is presented by McGuffin *et al.* [124], who describe a number of volume deformation methods that for example spread apart or peel away the volume.

Our approach extends these ideas by integrating a priori knowledge of the subject under study, in this case the endo-skeleton and the known modes of articulation that it affords. By making use of the skeleton as a model for articulation, we are able to create a standardized visualization that facilitates navigation in complete whole-body datasets and enables effective comparative visualization between different timepoints or between different subjects.

There are certainly parallels between our work and the volume animation method of Gagvani *et al.* [125], where a derived skeleton is employed to enable the realistic animation of volume datasets such as the Visible Human. However, where the APR is completely automatic, their skeleton is derived from volume data by a thinning process, and joints

Chapter 4

are manually assigned by an animator. In addition, they make use of a number of sample points to transform spherical parts of the volume into new poses of the skeleton. Finally, the purpose of their method is the lifelike animation of volumes, whereas we focus on the standardized analysis of anatomically relevant subvolumes of whole-body datasets.

In this work, we focus on the role of articulated planar reformation in facilitating the visualization of change in small animal datasets. Within this context, we have also implemented a number of existing comparative visualization techniques [126,127]: Side-by-side views or small multiples, animation, red-green channel overlays and image checkerboarding. In addition, we present the integration of a domain-specific change metric, in our case bone change, that can be mapped onto the atlas geometry and acts to facilitate the rapid localization of interesting changes both in the whole dataset and also in separate elements that are being focused on. This is an extension of previous comparative visualization approaches where differences were explicitly represented by geometry [128] or where multiple volumes could be combined to form new features for comparison [129].

4.3 Method

The complete pipeline is illustrated in Fig. 4.2. We start by briefly summarizing the steps, after which the details are further explored in the following subsections.

A standard atlas is automatically registered to all of the CT datasets that are to be compared, either of multiple small animals or of multiple timepoints of the same animal, using the method of Baiker *et al.* [17]. The atlas is based on that published by Segars *et al.* [59], but has been further manually segmented into its constituent bones. For each segment of the atlas, we derive a linear transform using the articulated registration approach in Section 4.3.1 (Fig. 4.2b).

Then, for each bone in the atlas, an object-aligned bounding box is automatically determined based on the surface representation of the bone in the atlas and the corresponding linear transform that was determined during registration. Using the bounding box and the resampling transform, the volume data is resampled for each bone with the aim of obtaining the volume in a standard coordinate frame, which facilitates comparison (Fig. 4.2c). This is described in Section 4.3.2.

Per CT dataset, the collection of resampled data volumes can be visualized in a standardized layout to give a global overview of the data (Fig. 4.2d). This is explained in detail in Section 4.3.3. Importantly, having multiple datasets in the same reference frame, in this case the atlas, allows element-wise visual inspection (Fig. 4.2e). See Section 4.3.4.

Furthermore, features can be visualized that pinpoint interesting areas in the global view. This further facilitates data exploration. One such feature is the bone density change metric, which is discussed in Section 4.3.5. Aside from this, we enable visualization of the registration error, giving the user an impression of the registration accuracy. See Section 4.3.6.

The resulting user interface consists of the standardized layout, henceforth referred to as the APR view, at the top, with a number of focus views of the multi-timepoint data side-by-side at the bottom. By clicking on any segment of the APR view at the top, the focus views at the bottom switch to that subvolume. The APR views and

APR for Change Visualization in Small Animal Imaging




Figure 4.2: An overview of our method. For each timepoint, the input data is the CT data and the surface representation of the atlas (a). The atlas is registered to the CT data using articulated skeleton registration (b). Based on bounding boxes of the atlas segments, we determine the volumes of interest (VOI) from which the data will be sampled (c). Then we perform the articulated planar reformation and visualize the resulting VOIs in a standardized layout (d). After this, any of the segments can be interactively selected for closer inspection. This procedure is performed similarly for the other timepoints, which enables comparison (e).

Chapter 4




Figure 4.3: An articulated planar reformatted visualization of a mouse, where the skull has been selected and is shown at three different timepoints in the focus views. The visualization shows a slice of the volume with a corresponding bounding box. Additionally, the global overview shows a semitransparent surface rendering of the atlas surfaces.

focus views can be configured for any combination of visual representations, for example surfaces with mapped derived features, subvolume slices, and so forth. Fig. 4.3 shows an example where all views have been configured to show volume slices, whereas Fig. 4.5 shows surfaces with mapped features.

4.3.1 Whole-body registration

In order to register a common atlas to all whole-body datasets under investigation, we employ the automated approach described in [17]. For completeness, we briefly explain the method in this subsection. It is based on the publicly available MOBY whole-body atlas of a mouse [59], which we modified by segmenting the skeleton into individual bones, identifying joints and by adding anatomically realistic joint rotation models [17]. Using this modified atlas, we first register the elements that show high-contrast in non-contrast-enhanced microCT, namely the skeleton, the lungs and the skin. To deal with possibly large articulations between bones, we perform the registration based on a hierarchical anatomical model of the animal. After a coarse alignment of the entire skeleton, the individual bones are registered step by step, starting at the skull, proceeding with the spine and then moving down the front and hind limbs separately. The used transformation models for the individual bones are dependent on the joint type. Therefore, the number of degrees of freedom (DOFs) varies between seven for a hinge joint (translation, non-isotropic scaling, one rotation) and nine for a ball joint. The lungs and the skin are registered subsequently, based on the skeleton registration result. Finally, we derive a set of correspondences from the registered elements to define a Thin-Plate-Spline (TPS) interpolation of major organs. An example of a segmented animal is given in Fig. 4.4. For the work described in this paper, we made use of the skeleton only.

APR for Change Visualization in Small Animal Imaging




Figure 4.4: An example of a segmented mouse body. Shown are the original skeleton (beige), the atlas skeleton (red), major organs and the skin.

The algorithm is fully automatic and only needs the microCT dataset as input. The restricted number of DOFs for the individual bones renders the method highly robust to large postural variations but also to moderate pathological bone malformations (e.g. bone resorption as a result of metastatic activity).

4.3.2 Articulated Planar Reformation

In order to obtain the data for each atlas segment in a standard coordinate frame, we perform Articulated Planar Reformation (APR). A bounding box to resample from is acquired for each segment by performing principal component analysis (PCA) on the surface representation of the bone. This results in three principal components. The direction of the principal components are used to determine the edges of the bounding box and the magnitude of the components maps to the size. By using the direction of the principal components, the volumes of interest are automatically aligned with the bones. An additional margin can be added to the bounding box, partly for aesthetic purposes, but this also ensures that the entire bone falls inside the bounding box. The margin is a parameter that can be changed by the user. The resulting bounding boxes are shown in Fig. 4.2c.

Using the bounding box and the registration transform, a new data volume can be created for each bone, the size of which is equal to the bounding box. Using the registration transform and the necessary correction for the direction from the PCA, a cubic resampling is applied to the original CT data to obtain the new volume. The result is a series of small image volumes that each contains the resampled image data of the bone aligned with the principal components of the corresponding atlas bone.

4.3.3 Articulated Planar Reformation View

For each resulting image volume, we can now apply various existing visualization techniques to visualize its contents. The appropriate visualization method entirely depends on the kind of data and the user's requirements. We provide a range of commonly used techniques, such as image planes that can be moved through the volume interactively, isosurfaces and volume renderings. The image planes, such as those shown in Fig. 4.2, 4.3, 4.5 and 4.6 are positioned in three dimensional space in order to provide the

Chapter 4

user with a cue concerning their local context, which is anatomically relevant. Parameters for all the visualizations, such as brightness and contrast, color maps and isovalues can be interactively adjusted.

Considering the entire collection of image volumes, we can apply an articulated planar reformation view (APR-view), which puts all the image volumes in the scene together in any desired layout. To prevent the image volumes from occluding others, we have designed a planar arrangement so that subvolumes can be separated from each other using an “exploded view” layout (Fig. 4.2d). This makes all the elements separately visible at the same time. The layout in the figure has been created beforehand by manually positioning the atlas segments in space and saving it to a configuration file, combined with the atlas itself. The standardized layout was carefully designed, keeping in mind planarity, anatomical consistency and occlusion reduction. The layout is automatically applied to all datasets. For a different layout or a new atlas, a new configuration file will be required. Alternatively, the user has the option of generating the layout automatically, which moves all segments to the same plane (flattens) and then outward in that plane with respect to the center (explodes).

4.3.4 Focus view

An articulated layout of all image volumes may be useful for some purposes, but one of the strengths of our approach lies in being able to visualize each element of the atlas separately in its own reference frame. The user can interactively select a part of the atlas from the APR-view by clicking on it. This element is then shown in a separate focus view, which allows closer inspection of the data without any of the other elements being in the scene. To facilitate orientation, the current camera position and direction is visualized in the global APR-view by way of a cone. This is illustrated in Fig. 4.5a. The selected bone, in this case the femur, has been selected and is shown in the focus views (Fig. 4.5b and c).

A further advantage of the alignment between the segment and the axes of the reference frame is that the segment is now in an anatomically relevant orientation, which facilitates visualization using image planes by slicing through the volume in the direction of the principal components.

In the following subsections, we discuss the features that we implemented to enable comparison of corresponding bones in different datasets. These techniques could be applied for intertimepoint comparisons as well as intersubject or other types of comparisons.

As shown in [17], the automatic atlas registration is accurate enough for many comparative purposes. However, a further rigid per-segment registration can improve registration results and facilitate finer comparison, such as some of the methods described below. Therefore, we have integrated an optional additional registration, the results of which were used to create the visualizations in Fig. 4.6. This additional registration step applies an adaptive stochastic gradient descent method using voxel-based normalized correlation. This method is described in [130]. It is an optional step, which we applied manually in order to optimally demonstrate the comparative visualization methods. We are currently in the process of integrating this method into the pipeline.

APR for Change Visualization in Small Animal Imaging

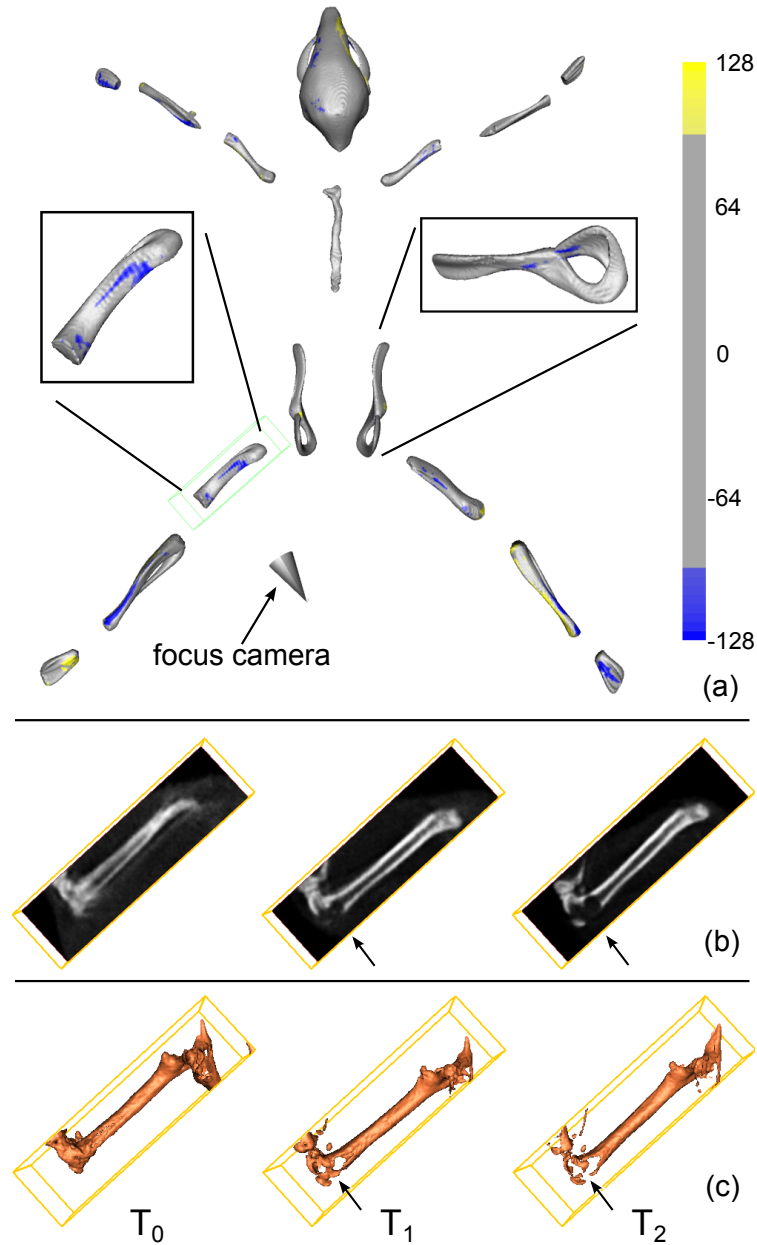


Figure 4.5: A global articulated planar reformatted visualization of the atlas with the bone change metric between the first two timepoints color coded onto the surface (a). The bone change metric highlights potentially interesting areas that the user may want to further inspect. The femur has been selected and is shown side-by-side in the focus views for the three available timepoints at 3, 5 and 6 weeks into the experiment (b) and (c). Shown for each timepoint are an image plane that can be interactively moved through the volume (b) and an isosurface (c). At 5 and 6 weeks, bone resorption can clearly be seen near the knee area (indicated by the arrows), even though the animal posture in the original data was highly variable. Note that between the first and second timepoint, with the subject at 10 and 12 weeks of age respectively, some growth is still taking place, which has to be taken into account when analyzing the images. More details about the data are presented in Section 4.5.

Chapter 4




Figure 4.6: Using the same data as in Fig. 4.5, different modes of comparative visualization are applied to the first two timepoints. We have performed an additional rigid registration in order to maximally illustrate the capabilities of these techniques. A red-green overlay is shown in (a). In (b), a checkerboard visualization is shown that can be used to assess the quality of the registration. In (c), the bone change metric has been applied and mapped to the surface representation of the atlas.

Side-by-side visualization

The first technique is a side-by-side comparison and involves showing the data from different timepoints in separate viewer frames. The cameras of these viewer frames can be synchronized, as well as any image planes that are used to slice through the volumes. This way, the user has exactly the same view for the different timepoints. This is demonstrated in Fig. 4.5. There are several advantages to this approach. First, side-by-side visualization closely resembles the traditional way in which biologists compare datasets. Second, even when there is a relatively large registration error, the human visual system can correct for this discrepancy and link corresponding regions of interest. Finally, this approach allows for all kinds of visualization, including slicing image planes, isosurfaces and volume renderings, without the rest of the mouse body obstructing the view.

Comparing by switching

In this approach, the visualizations of all timepoints are shown in the same viewer frame and a slider can be used to quickly switch between the datasets. Assuming the registration between the timepoints is satisfactory, this approach allows the user to spot minute changes by flipping back and forth between timepoints. As in the side-by-side approach, this visualization allows all available types of visualization.

Overlays

In this approach, the visualizations of all timepoints are shown in the same viewer frame and a slider can be used to quickly switch between the datasets. Assuming the registration between the timepoints is satisfactory, this approach allows the user to spot minute changes by flipping back and forth between timepoints. As in the side-by-side approach, this visualization allows all available types of visualization.

APR for Change Visualization in Small Animal Imaging

Checkerboard

A final option is to use a checkerboard combination of two timepoints (Fig. 4.6b), which is useful for assessing whether two datasets are properly aligned. When this is not the case, this will be clear by the appearance of jagged edges in the visualization.

4.3.5 Feature visualization

Our approach is well suited for integration with feature visualization. For the case of research into cancer-induced bone resorption, which we will elaborate on in the evaluation (Sec. 4.5), we provide a bone density change metric. This tool can be used to estimate the change in bone density, which is directly related to the intensity values in the CT data. The bone density change is mapped onto the atlas surfaces and can be visualized using a color map as is shown in Fig. 4.6c.

The metric is calculated using one of two variations of the same method: With *simple bone change*, values are first interpolated at all vertices of the atlas mesh, from both of the atlas-registered volume datasets that are to be compared. At each vertex position, the value interpolated from the second dataset, referred to as the followup, is subtracted from that of the first dataset, referred to as the baseline. The end result is a new dataset, defined over all vertices of the atlas, where each value describes the signed change in density. If the value is negative, bone density has decreased, and vice versa. With *locally aggregated bone change*, we additionally sample values at a configurable distance, measured along the surface normal, both on the outside and the inside of each vertex. For each vertex, the average of the three samples (inside, outside and at the vertex itself) from the followup dataset is subtracted from the average of the three samples from the baseline dataset. This variation is more robust but less sensitive than the simple bone change.

The integration of this metric into both the global APR view and the focus views enables a two-level change localization approach. When we apply the visualization of this metric in the global overview, such as in Fig. 4.5a, the user can directly spot potentially interesting regions. The corresponding sub volume can be selected and brought into focus for a closer inspection of the bone change metric. This process assists the user in rapidly pinpointing specific regions where high bone change indicates that further analysis could be fruitful.

To further facilitate the visual pinpointing of interesting regions, we have implemented an interactively configurable significance threshold for all colormaps. With this functionality, the user determines above which threshold a feature is considered interesting. An example is shown in Fig. 4.5a, where areas below the significance threshold are colored gray, and regions outside are assigned a color, usually from a perceptually linearized colormap. The surface rendering is interactively updated during changes to the significance threshold, further helping to understand the mapped feature.

Chapter 4

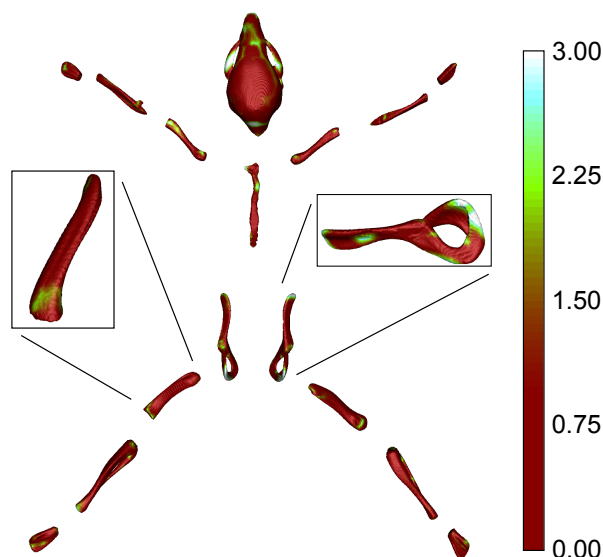


Figure 4.7: A global articulated planar reformatted visualization of the atlas with the registration error at the first timepoint color coded on the atlas surface. The registration error is the distance (in voxels) to the CT surface that was used in the registration (Sec. 4.3.6). We use a color map that is perceptually linear [131].

4.3.6 Confidence visualization

Residual error in the atlas registration process has an effect on the accuracy of the analysis. In order to give the user feedback on the localized registration error and hence help them to judge the results, we provide the option to color code the surface according to the registration error at each point. This is the distance between the surface representation of the atlas element and the CT surface that was generated and used in the registration process. Fig. 4.7 illustrates this idea.

Importantly, the registration error is also high in areas of high bone resorption, exactly the phenomenon under study in this case, an observation that should be used in the interpretation of both error and bone change. In areas of real bone change, we expect the confidence visualization to show focused low confidence surrounded by high confidence. When the error visualization shows a pattern that deviates from this, the bone change visualization should be verified by studying the direct comparative visualizations.

4.4 Implementation

We implemented our approach as an extension of the Cyttron Visualization Platform (CVP), which has been developed by us as a part of the Cyttron project [132]. The CVP is a software package for visualization of multimodal data and provides a number of tools for this purpose. It is mainly aimed to be used by biologists to visualize and explore their data. The CVP is written in Python and uses the Visualization Toolkit (VTK) for visualization, which in turn is based on OpenGL. Essentially, the CVP provides a

APR for Change Visualization in Small Animal Imaging

framework and a graphical user interface around visualization methods that are available in the VTK libraries.

We implemented the APR functionality as an extension package to the CVP. The articulated registration functionality was written in Matlab, the reformation and visualization functionality as plugins written in Python, using VTK libraries. We plan to release the APR-package as open-source at a later stage.

4.5 Evaluation

First, we will present some details about the performance of our implementation, after which the user evaluation will follow.

The articulated registration of the atlas to the CT data is a process that is only required once for each dataset. It takes a few minutes on a recent Intel Core2 PC and could be further optimized if required. The articulated planar reformation itself takes a second or two. When the entire pipeline is regarded, excluding the registration, but including loading the data, applying APR and setting up the preconfigured visualization, we are looking at about 10 to 20 seconds. After that, the visualizations can be inspected interactively. Due to memory constraints on the 32-bit machines that are in common use by the domain scientists, we applied cubic subsampling to reduce the size of the datasets from an original resolution of about 512x512x1200 voxels to about 256x256x600 voxels. On 64-bit PCs, to which our users will be gradually upgrading, this step would not be necessary.

In order to investigate the possible role of our application in the research pipeline of molecular imaging domain scientists, we performed a case study evaluation according to the guidelines set out by Yin [133]. We opted for this evaluation methodology, as the contributions of our work mainly target the outer levels of the nested model proposed by Munzner [134], and as we desired to study the real-world application of our pipeline. The main study question was defined as “How can the APR change visualization tool assist domain scientists in studying changes over time, pathology-related and otherwise, in small animal whole-body CT?” whereas the *case* was defined as the “guided application of the APR change visualization tool by two domain scientists referred to as CL, a widely published expert in molecular imaging and oncology research, and MH, a leading small animal and molecular imaging expert, to datasets of cancer metastasis induced bone resorption”. The application was guided, as at this stage our assistance was still required in operating the software. CL is also the sixth author of this paper. He contributed to this work with an extensive analysis of the problem domain and by giving feedback on our prototypes, but did not work on the implementation itself. In this context, his role as test user and case study subject was not compromised. MH was not involved in any stage of the project before taking part in the case study evaluation described in this section.

Together with each of the domain scientists, in two separate sessions, we used the software to analyze a dataset consisting of three small animal CT scans, made at 3, 5 and 6 weeks into the experiment. Breast cancer cells had been introduced when the subject was 7 weeks old (week 0 of the experiment). The purpose of the analysis was to study the metastasis of these cells into bony tissue. As per usual for case studies, we

Chapter 4

defined a number of propositions beforehand in order to better answer the main study question. During the analysis session, we gathered feedback on the application of the tool, structured according to the study propositions. In the following text, each proposition is stated, together with its related feedback, after which we discuss general comments and our preliminary conclusions.

The standardized APR view speeds up navigation in a dataset, i.e. selecting anatomical elements of interest. CL confirmed that this functionality was important in being able to select anatomical structures of interest, further adding that analysis is often based on *a priori* expectations of the pathology under study. For example, breast cancer is known to metastasize mostly to the long bones, favoring subregions of high bone turnover. MH agreed with the proposition, adding that he found the camera icon (section 4.3.4) intuitive in relating the focus view to the global view.

Automatic anatomical alignment of the subvolumes facilitates comparison of the data. CL confirmed this statement. Confirming the feedback of MH on the previous proposition, he found the camera orientation glyph useful in understanding the relation of a focus view to the global view. The user further indicated that aligning slices with the long bones, as is automatically done in our approach, was definitely preferred. MH commented that this type of alignment was in his view the most natural, further explaining that he would have requested it had it not been available. He added that he liked the way that every subvolume could be treated independently from the global slicing direction.

The switched view is the preferred view for qualitatively investigating bone change. Domain scientists usually eyeball datasets side-by-side in order to compare them. CL commented that with the switched view the registration error was more visible. He further remarked that being able to slice through the data was important in helping to distinguish between real bone change and registration errors. MH also noticed that there were remaining registration errors. He commented that switching was helpful to study geometric alterations, even in cases where registration can't be successfully performed due to gross changes, for example due to global growth or curvature.

The checkerboard view is primarily useful for further verifying the quality of the registration / matching. The utility of this view was not directly apparent to CL. It was finally remarked that it would be useful mainly for checking the quality of the alignment. MH was also not familiar with this representation and remarked that he preferred the switched view, also as they were used to paging through MR image stacks. He concluded by stating that it was an interesting / different approach, but would require training.

The overlay view is primarily useful for detecting minute changes between timepoints (when registration is accurate). Again the growth-related misalignment between the first two timesteps dominated. CL suggested that manual alignment functionality, with this view as feedback, would be very useful. In cases where the registration was better, for example between the second and third timepoints, he considered this view equally important to the switched and side-by-side views. MH was familiar with this representation, commenting that it was an ideal way to study edge differences and mentioning that the red-green mixing helped in interpreting differences. MH's lab already uses this type of representation to look for minute changes and to check the quality of registration.

The coordinated side-by-side view is the preferred view for qualitatively investigating

APR for Change Visualization in Small Animal Imaging

bone change. CL sees it as advantageous that this view extends and improves on the traditional method of manually comparing two datasets side-by-side. CL confirmed that it was possible to compensate for misregistration, as an internal 3D representation was mentally constructed and used to simultaneously judge change and registration error. MH explained that side-by-side comparisons were done routinely in his lab as part of a first screening of image data, albeit with side-by-side print-outs as they did not yet have suitable software tools for this. With side-by-side comparisons they are able to detect points of interest that should be further investigated. He commented that the 3D interaction possibilities made this representation more user-friendly, concluding that this would probably be the visualization that he would start with, followed by the overlay representation to further study suspected differences.

Global bone change metric facilitates localization of interesting areas in terms of bone resorption. and *The local bone change metric helps to further refine the location of areas of high bone resorption.* During the analysis sessions, the bone change metric visualization was affected by registration issues, but worked well enough to illustrate its use to the scientists. CL stated that in areas where bone cancer metastatic bone resorption occurs, there is a tell-tale pattern of bone density decrease surrounded by bone density increase. Other types of cancer have different tell-tale patterns. He continued that this visual feedback would be ideal for the fast location of interesting areas in the global view as well as further examination in the focus view. MH recognized that this served as an overview of all changes in the whole dataset and stated that, in his opinion, first screening the whole body for changes and then focusing on each of the changes for more in-depth study was indeed the most suitable approach. He agreed that the local bone change metric helped to further refine change locations.

The registration error visualization helps in the interpretation of the bone change feature as well as the comparison views (side-by-side, switching, overlay, checkerboard). CL considered this to be an important tool. He did make the suggestion that we should use distinctive color maps to differentiate between different types of metrics, especially registration error and change features. He further explained that he would check the data more carefully if the registration error is high. Interestingly, MH did not see a use for the registration error visualization in better judging change, stating that the domain expert was capable of differentiating between registration error and real change by local examination and that the error metric was in his opinion too implicit. He did see potential in using the error visualization as a way to study deviation from a standard atlas and hence to phenotyping individual small animals.

The tool will help guide the quantitative study of metastasis-related bone resorption, in that interesting areas can be localized and qualitatively studied before further analysis. CL agreed with this proposition and added that this correlated with the initial goals that were set for the analysis software. He further underlined that the software was to play a role at the start of the analysis pipeline, where domain scientists are exploratively analyzing acquired data to generate hypotheses for further analysis using traditional methods. MH strongly supported this proposition and explicitly underlined the potential value of the tool. He further commented that he believed that the tool could in some cases even answer research questions directly, i.e. without a further traditional analysis step.

Chapter 4

4.5.1 Study conclusions

CL was generally impressed with the system, and enthusiastic about its application to future studies by his research group. Other applications of the current functionality were mentioned, including the study of general growth-related deviations. He indicated that future functionality for group studies as well as for side-by-side comparisons of multi-timepoint datasets would be useful. MH was in general positive also, and saw potential for the APR approach also in the study of soft tissue changes in MR. MH independently suggested the future possibilities of projecting groups of datasets onto the standardized layout in order to study group statistics in the same space.

Further, we can derive a number of lessons from this case study. Importantly, it was confirmed that mimicking the existing analysis approach as far as possible with a new visualization application is a good strategy. It facilitates acceptance by the users, and also allows them to leverage their existing experience in applying the new tool. The case study was performed with the blue-yellow colormap for both the bone change and error. Since then, we implemented the suggestion to use intuitively differentiable colormaps for different types of quantities, as can be seen in Figs. 4.5a and 4.7, where respectively bone change and error are represented. Change metrics should be made more robust to residual registration error. The optional additional rigid registration, mentioned in section 4.3.4, is an important step in this direction. Finally, the suggestion to implement manual alignment with real-time comparative feedback will be added in a future version of the software.

Based on the case study, we answer the main study question with the observation that the tool can assist domain scientists in studying pathology-related changes over time by enabling the rapid localization and qualitative study of interesting areas, particularly by the combination of global view for navigation and focus view for more low-level investigation. More globally speaking, it has an important role in the early explorative stages of the analysis pipeline.

4.5.2 Limitations

During this evaluation, we focused on comparing multi-timepoint data of one mouse. However, our implementation also allows intersubject comparisons, which is required when comparing the subject to a control mouse. Judging differences will become more challenging, as apart from the pathological differences that we want to highlight, the two subjects will also have a different anatomical structure. Still, our technique constitutes an improvement over the existing approach of comparing regions of interest in a non-articulated and hence unstandardized manner.

In spite of the possible positive bias, it has been advantageous involving one of our collaborators in the case study, as it made possible far more extensive and in-depth discussion of the results and possible future directions. We did of course take reasonable precautions, for example not involving the user in the actual implementation and also explicitly requesting objectivity during the evaluation. Importantly, we involved an unbiased external domain expert at a later stage who, except for the role of the registration error visualization, confirmed the findings of our primary user.

APR for Change Visualization in Small Animal Imaging

Furthermore, due to the direct involvement of the collaborating research group, the tool will now find more widespread and longer term use. This will allow us to enter a new phase of evaluation during which we will be able to collect longer term feedback from a larger group of users.

4.6 Conclusion and Future Work

We have presented a method to fully automatically perform standardized articulated planar reformation for a whole-body CT dataset based on an atlas, in order to overcome the limitations caused by varying animal posture and occlusion effects that occur in traditional visualization approaches.

In an articulated registration step, linear transforms are acquired that map each atlas segment of the skeletal structure onto the CT data. Based on these transforms and a principal component analysis of the atlas surfaces, we set up a bounding box in the original CT data where we resample our data from. We called this method articulated planar reformation (APR). This results in a collection of sub volumes, each of which we can visualize separately in its own reference frame. When we apply this approach for each dataset of a cross-sectional or follow-up study, we can place the sub volumes side-by-side, which enables direct comparison of corresponding segments, even though in the original scans the animal posture was highly variable. An additional advantage of APR is that the segments, in our case bones, are automatically aligned with their principal axes, which is anatomically relevant and is more useful to our users.

We have shown how a two-level localization approach combined with an appropriate change metric, such as bone change, can be used to indicate interesting areas in the global whole-body view that can be further inspected in the focus view after interactively selecting the segment. In addition, we have implemented several approaches that allow for comparative visualization between multiple datasets, such as interactive switching, side-by-side, overlays and checkerboard visualizations. To further facilitate comparisons, we applied an additional rigid registration to each of the segments. Also, there is an option to show the local registration error mapped onto the atlas surface, which gives an indication of the accuracy of the registration and thus of the reliability of the comparison results.

Finally, we have presented an elaborate case study in which the potential role of our application was investigated in separate sessions with two experienced domain scientists. Although there were some questions about the usefulness of some of the visualizations, such as the checkerboard, the overall assessment was positive, with our primary user stating that he would certainly use it in his lab's research and the secondary user also expressing interest in applying the methods documented in this paper. We have thus shown that our approach is useful for comparing whole-body data in cross-sectional and follow-up studies.

For the time being, we decided to focus on the skeleton only, but our approach is also suitable for organs. Because of the lack of contrast in the data that was available, the location of the organs could only be estimated using interpolation based on the skeleton [17]. In the future, it will be interesting to use a contrast agent for imaging,

Chapter 4

which would allow proper registration of the organs and would increase the usefulness of visualizing them in an articulated manner. In addition, our approach can be applied for other modalities, such as MRI, as long as articulated registration transforms and the corresponding atlas segments are available.

Furthermore, with the resolution of scanners growing, datasets are rapidly increasing in size. Loading several gigabyte-size datasets will be problematic on many systems. We want to implement a level-of-detail approach where we load a subsampled version of the data for the global APR view, while on selecting a segment, we load a subset of the data in full resolution to show in the focus view. Also, in order to maintain the modular structure of our design, we currently apply a straightforward resampling method to obtain a collection of sub volumes. However, we would like to replace this with an approach that directly renders the transformed data, in order to eliminate any unnecessary interpolation steps and to reduce the amount of memory that is required to keep the sub volumes in.

We plan to develop more advanced comparative visualization techniques to assist in multi-timepoint and intersubject comparisons. In this context, we will investigate relevant change features other than bone density. Depending on the feature, we plan to apply different visualization methods. For example, it may be desirable to geometrically visualize the amount of change. It would also be interesting to try and discern between physiological changes and registration errors.

Globally speaking, we aim for our approach to become an invaluable and general tool for cross-sectional and follow-up studies in preclinical small animal research.

Acknowledgements

This research was supported in part by the European Network for Cell Imaging and Tracking Expertise (ENCITE), which was funded under the EU 7th framework program. We would like to thank Professor Mathias Hoehn for his willingness to serve as the external domain expert for our evaluation.

APR for Change Visualization in Small Animal Imaging

