

## 3D reconstruction of the human rib cage from 2D projection images using a statistical shape model

Jalda Dworzak · Hans Lamecker · Jens von Berg · Tobias Klinder · Cristian Lorenz · Dagmar Kainmüller · Heiko Seim · Hans-Christian Hege · Stefan Zachow

Received: 18 December 2008 / Accepted: 21 June 2009 / Published online: 26 July 2009  
© CARS 2009

### Abstract

**Purpose** This paper describes an approach for the three-dimensional (3D) shape and pose reconstruction of the human rib cage from few segmented two-dimensional (2D) projection images. Our work is aimed at supporting temporal subtraction techniques of subsequently acquired radiographs by establishing a method for the assessment of pose differences in sequences of chest radiographs of the same patient.

**Methods** The reconstruction method is based on a 3D statistical shape model (SSM) of the rib cage, which is adapted to binary 2D projection images of an individual rib cage. To drive the adaptation we minimize a distance measure that quantifies the dissimilarities between 2D projections of the 3D SSM and the projection images of the individual rib cage. We propose different silhouette-based distance measures and evaluate their suitability for the adaptation of the SSM to the projection images.

**Results** An evaluation was performed on 29 sets of biplanar binary images (posterior–anterior and lateral). Depending on the chosen distance measure, our experiments on the combined reconstruction of shape and pose of the rib cages yield reconstruction errors from 2.2 to 4.7 mm average mean 3D surface distance. Given a geometry of an individual rib cage, the rotational errors for the pose reconstruction range from 0.1° to 0.9°.

**Conclusions** The results show that our method is suitable for the estimation of pose differences of the human rib cage in

binary projection images. Thus, it is able to provide crucial 3D information for registration during the generation of 2D subtraction images.

**Keywords** Geometry reconstruction · Biplanar projection images · Distance measure · 3D pose difference


### Introduction

In clinical routine, radiography is an inexpensive and frequently used imaging technique for screening and diagnosis of the chest region. However, interpreting radiographs is difficult, especially if conclusions about the 3D geometry need to be drawn. For this reason, computer-aided diagnosis (CAD) methods are developed and increasingly applied to support physicians [1].

A widely used diagnosis method is the comparison of chest radiographs with previously acquired radiographs of the same patient to detect changes in the health status. This approach is also commonly used in the context of interval studies, where radiographs are acquired intermittently to identify *interval changes* between subsequent images to observe the course of a disease, e.g. tumor growth. To detect interval changes by automated procedures, a CAD method known as *temporal subtraction* of subsequently acquired radiographs can be applied [2]. Subtraction images show the difference between previous and follow up images of a patient after a suitable image registration (c.f. Fig. 1). Ideally, unchanged anatomical structures are eliminated while any interval change caused by new opacities (e.g., tumors) appears strongly contrasted and stands out more clearly. Studies on the benefit of image subtraction indicate a significant improvement in the accuracy of detecting abnormalities in the chest region in case subtraction images are used [3]. One

J. Dworzak (✉) · H. Lamecker · D. Kainmüller · H. Seim · H.-C. Hege · S. Zachow  
Visualization and Data Analysis, Zuse Institute Berlin,  
Berlin, Germany  
e-mail: dworzak@zib.de

J. von Berg · T. Klinder · C. Lorenz  
Medical Imaging Systems, Philips Research Europe,  
Hamburg, Germany



**Fig. 1** Two different subtraction images of the same patient after image registration: **a** the interval change (the development of a pneumonia) is clearly visible (*arrow*); **b** no interval change is detectable due to artifacts that are caused by a strong anterior–posterior inclination

of the major problems for the generation of suitable subtraction images is the deviation of the patient's pose between subsequent images. For this reason, image registration prior to the image subtraction is indispensable, since the imaging geometry may differ for each radiograph. In case the registration fails (due to major pose differences), undesired artifacts emerge that might superimpose the interval change, which in consequence may remain undetected (see Fig. 1). In practice, especially anterior–posterior (AP) inclination and rotation around the longitudinal axis impair the quality of subtraction images [2].

The compensation of such *3D pose differences* using 2D to 2D image registration does not necessarily lead to sufficient results. A phantom study of von Berg et al. [4] showed that deformable 2D to 2D image registration, compared to non-deformable registration, better compensates for pose variations of up to  $2.3^\circ$ , still allowing faithful detection of interval changes from image pairs.

An accurate assessment of the 3D pose difference that causes the discrepancies between two images may yield valuable information that can be used in cases where deformable registration methods are likely to fail.

To quantify the 3D pose difference between subsequent images the 3D pose of a patient's thorax needs to be estimated from 2D radiographs. The ribs are suitable structures to serve for this task since they are rigid and do not notably change in shape over time. In comparison to scapulae and clavicle their pose is largely independent from the arm posture. Last but not least, ribs are reasonably contrasted in radiographs. Hence, they can be used as a reference system to define pose differences of a patient between time intervals.

We propose to use the rib cage to infer the 3D pose of a patient's thorax from 2D images. The general idea is to match the patient's individual 3D rib cage model to the image data such that the pose of the 3D model yields a good approximation of the patient's pose during image acquisition. For such an approach, the patient-specific rib cage geometry is required first. Therefore, this work addresses two problems:

1. *3D shape reconstruction*: The 3D geometry of the rib cage needs to be retrieved from patient-specific data. As an alternative to an expensive acquisition of tomographic image data with higher radiation exposure, the reconstruction shall be obtained from few 2D projections.
2. *3D pose reconstruction*: The patient's individual 3D rib cage, reconstructed via a 3D shape reconstruction, is used to recover the 3D pose from subsequent 2D images.

The general task of 3D reconstruction of anatomical structures from X-ray images is challenging, as only partial information is available. Our approach is based on an SSM of the rib cage, which aims at providing missing information in a reasonable and well-founded manner.

A factor that limits the information at hand is the small number of X-ray images to be used for a 3D reconstruction. While in clinical practice two radiographs (one posterior–anterior (PA) view and one lateral (LAT) view) are acquired in most examinations, there is no standardized imaging protocol for their acquisition. The patient commonly turns approximately  $90^\circ$  for the LAT-view. Consequently, PA and LAT X-ray images do not depict the same scene. This makes the problem of the 3D shape and pose reconstruction from two views more difficult. Although there are special X-ray imaging devices that can be configured to generate a calibrated, biplanar PA and LAT pair of radiographs, as described in [5] and used in [6], they are not widely used. Additionally, identifying the ribs in the LAT-view is even more difficult due to a higher degree of superposition by other thoracic structures. For the aforementioned reasons the 3D shape as well as pose reconstruction of a patient's individual rib cage should be preferably accomplished from one PA-projection.

Another difficulty involved in a 3D reconstruction from X-ray images is that the reconstruction process relies on the detection of features that represent the anatomy in the X-ray images. Extracting these features is a challenging problem in itself and strongly depends on the details of the imaging protocol. In this work, we assume that the ribs can be segmented from X-ray images. In most situations this assumption is met, i.e., rib boundaries can be identified visually in X-ray images and, at least, outlined manually. Work to extract the contours of posterior ribs from frontal X-ray images has been reported [7–9]. Park et al. address the problem of detecting and labeling the ribs [10]. Nevertheless, the automation of this task is still an open research problem not being within the scope of this work.

The objective of this work is to show that the proposed reconstruction method yields useful results within a well-defined experimental setup, and thus provides a solid basis for future clinical applications. Artificial binary projection images of the rib cage, for which the exact parameters to be determined are known beforehand, enable us to verify the

accuracy of our reconstruction results. Segmentations of the ribs can be easily obtained from such binary images. We show that an automated 3D reconstruction of shape and pose can be achieved from segmented binary projection images of the rib cage.

#### Related work

Matching objects is a fundamental problem that arises in various application fields of computer vision and medical image analysis. Ways and means are required to measure the similarity between objects to be matched. A common approach is to use distance measures, which rely either on geometrical features like edges (*feature based*) or on intensity information in images (*region based*), to quantify the similarity. The efficient computation of edge-based distance measures, as used in our work, has been investigated in the past, for example with edge-based template matching using the Hausdorff distance, which was used to quickly locate objects within images [11, 12].

A general method for aligning shapes is the iterative closest point (ICP) algorithm [13]. Starting with an initial transformation estimate of the shape to be matched to a reference shape, pairs of corresponding points are generated between the two shapes. Then, the initial transformation and the point correspondences are iteratively adjusted, such that an error metric is minimized. There are many different variants of the ICP method, and it was also adapted and modified to be used with 2D to 3D reconstruction problems [14–17]. A comparison of ICP variants can be found in [18].

In recent years, several methods for the *pose reconstruction* of 3D objects from 2D projections were introduced [14, 19–21]. The work of Lavallée et al. [19] deals with determining position and orientation of arbitrary surfaces based on contour points in 2D image data. The matching is achieved by minimizing signed distances from projection lines, which are defined between 2D contour points and the projection center, to the surface. A similar approach to [19] that uses an extended ICP method was introduced by Feldmar et al. [14]. Instead of using signed distances the approach exploits information of the surface normals. Since ICP-based methods highly depend on the initialization and may lead to local optimal solutions, Cyr et al. [20] propose a hierarchical iterative matching approach to avoid this problem. Here, the projection direction of vertebrae is recovered from one binary projection image. However, recovering translation parameters of the object with respect to the location of the projection source is not considered.

The aforementioned methods address the problem of recovering the pose of *known* shapes from their projections. With regard to medical applications, however, it is often required to assess the pose of an *unknown* patient-specific shape of a known anatomical class. The method presented

by Bhunre et al. [21] accomplishes this for the proximal and distal femur by fitting a generic model to one patient-specific X-ray image. For complex geometries, such as the rib cage, this approach is not likely to yield accurate results, because such structures are often subject to a high inter-patient variability and involve symmetry or even repetitive structures as the ribs. For this reason, it is also necessary to retrieve the specific 3D shape of a patient's anatomy, to estimate its pose from projection images correctly.

Several *shape reconstruction* methods have been proposed that use a priori knowledge, e.g., in form of SSMs [22], to incorporate information about reasonable variations in shape within a class of shapes [15–17, 23, 24]. Fleute and Lavallée [15] retrieve the shape of the femur from segmented contours in X-ray images with an SSM. They use a generalized version of the ICP algorithm to minimize the distance between apparent contours on the SSM's surface and a set of projection lines as already used in [19]. The shape reconstruction method presented by Lamecker et al. [16] minimizes the distances between 2D images and projection images of an SSM based on the silhouettes in both images. The pose of the shape to be reconstructed is assumed to be known in advance.

These previous methods deal with anatomies of the femur [15, 17, 21, 24], the pelvis [16] and individual vertebrae [19, 20]. Thus, they address only single bones. The reconstruction of bone ensembles with repetitive structures as the rib cage and the spine is considered much more challenging, because similarity and redundancy of sub-structures can lead to mismatches.

One of the first methods for the reconstruction of the human rib cage was proposed by Dansereau and Srokest [25] to assess geometric properties of the rib cage of living subjects. It uses direct linear transformation (DLT) [26] for the reconstruction of rib midlines from a pair of stereo-radiographs (one conventional PA-view and a second PA-view with 20° difference of the X-ray source's incidence angle). Delorme et al. [27] presented an approach to generate patient-specific 3D models of scoliotic spines, pelvises and rib cages. They used the method of Dansereau and Srokest [25] in combination with an additional LAT-view to obtain 3D coordinates of anatomical landmarks that need to be manually identified in the 2D radiographs. To obtain patient-specific surface models a generic 3D model of a scoliotic patient, reconstructed from computer tomography (CT), was adapted to these landmarks using free form-deformation. The work of Novosad et al. [28] addresses the problem of estimating the pose of vertebrae to reconstruct the spinal column for the analysis of the spine's flexibility. This work is similar to ours in that patient-specific 3D shape reconstructions from projection images are used to perform a subsequent pose reconstruction. For the prior 3D shape reconstruction Novosad et al. use the method described by Delorme et al. [27]. The pose is then reconstructed using only one PA X-ray image.

Recently, two works using a semi-automated framework for the reconstruction of the rib midlines have been presented by Mitton et al. [29] and Bertrand et al. [6]. Their approach depends on a prior 3D reconstruction of the spinal column and two calibrated, exactly perpendicular, and simultaneously acquired radiographs (PA and LAT) [5]. A generic model, fitted to previously reconstructed landmarks of the sternum and entry points of the ribs at each vertebra (using the same technique as in [27]), yields an initial estimate for the reconstruction. This estimate is iteratively improved, with the interaction of an operator who manually adapts projected rib midlines of the matched generic model to image information in the radiographs.

These rib cage reconstruction methods [6,25,27,29] depend on the identification of landmarks within the X-ray images. The correct identification of anatomically relevant landmarks of the ribs in X-ray images is difficult even when performed manually. A promising alternative is the use of features from X-ray images that are more likely to be detected automatically, e.g., rib contours or image gradients. Moreover, the methods do not take advantage of a priori knowledge about the rib cage's shape variability due to the use of the DLT technique. Benameur et al. [23] approach this problem using 3D SSMs of rib midlines and extracted contours from a calibrated pair of X-ray images (PA and LAT-view). The reconstruction is obtained by minimizing an energy function defined via edge potential fields between rib contours and projected rib midlines of an SSM. This method is used to classify pathological deformities of the spinal column in scoliotic patients.

### Contribution

Our work is based on the method of Lamecker et al. [16] that uses a 3D SSM for the reconstruction of complex 3D *shapes* from 2D projection images. We extend this method in two directions: (1) we additionally handle *pose* reconstructions with respect to a known image acquisition setup. This is crucial for supporting interval studies as described above; (2) we consider the distinct and complex geometry of the rib cage. Here, new problems arise which call for new solutions compared to the case of the pelvic bone geometry [16]. In this context, we compare different edge-based distance measures with regard to their application to the 3D reconstruction of the rib cage and thereby extend preliminary work presented in [30]. Additionally, we evaluate the accuracy of the pose reconstruction from one and from two calibrated binary projection images, which has not been done with other rib reconstruction methods.

Our method is similar to the work of Benameur et al. [23] inasmuch as it uses a 3D SSM and relies on edge-based features in the image data, instead of landmarks that have little prospect of being detected automatically from radiographs.

However, in contrast to [23], we use an SSM of the rib's *surfaces*. Our approach allows for a direct comparison of rib contours in both, the model projection and the image data, where Benameur et al. propose to compare projected midlines of the model with contours of the ribs in the image data. Furthermore, Benameur et al. [23] use 3D reconstructions obtained by the method of Dansereau and Srokest [25] as gold standard to evaluate the method, which limits the assessment of the method's accuracy, as already stated in [29]. We use 3D surface models of rib cages, extracted from CT-data of 29 different subjects, as gold standard for evaluation.

## Material and methods

### Overview

The present work addresses two problems: (1) the 3D shape reconstruction of an individual rib cage and (2) its 3D pose reconstruction, each from few binary 2D projection images. The goal of the 3D shape reconstruction is to estimate a rib cage's individual shape from the projection images. For this not only the 3D shape, but also the 3D position, orientation, and size must be determined. For the pose reconstruction only the 3D position, orientation, and size are recovered.

The aim of the reconstruction process is to infer these parameters from the projection data. The idea is to obtain a 3D model that serves as an estimate for the 3D pose (and 3D shape) of the individual rib cage. For this purpose a parameterized 3D model is fitted to the 2D projection images (*reference images*). In case of the 3D shape reconstruction, this model is deformable and has been derived from a set of training data. For a pose reconstruction we use a patient-specific, non-deformable model. Fitting such a model to the image data is accomplished as follows:

We generate projection images of the 3D model (*model images*). On the basis of images features, i.e., contours of the rib boundaries, these model images are compared to the reference images. The dissimilarities of the model images to the reference images are iteratively reduced by adapting parameters of the 3D model that control its 3D pose (and 3D shape) variations.

### Statistical shape model of the human rib cage

Since projection images can only provide limited information about a 3D anatomy it is beneficial to incorporate a priori 3D shape knowledge to infer missing information in an anatomically plausible way. Cootes and Taylor [22] proposed statistical shape models (SSM) that comprise such shape knowledge. An SSM is a deformable model that is based on statistical analysis of a set of training examples of the same



shape class. Only deformations according to the variability within the set of training shapes are permitted.

Our reconstruction process is based on such an SSM of the human ribs that was created from 29 triangulated surface models of different rib cages, automatically segmented from CT-data [31]. As not all ribs are displayed in each of the CT-images due to an incompleteness of the image data, the model we use is restricted to the ribs 2 to 8. The anatomical correspondences between the vertices of the individual surfaces are established during segmentation. The SSM is generated via a principle component analysis (PCA) on the set of the 29 rib cages. It captures the average shape of the training set as well as its variability in terms of  $n = 28$  modes of shape variation (*shape modes*) and is defined as

$$S(\mathbf{b}, T) = T \left( \bar{\mathbf{v}} + \sum_{i=1}^n b_i \mathbf{p}_i \right), \tag{1}$$

where  $\bar{\mathbf{v}}$  is the average shape and  $\mathbf{p}_i$  are the shape modes. The degrees of freedom of the SSM are its shape weights  $\mathbf{b} = (b_1, \dots, b_n)$ , which control the shape variation, and a linear transformation  $T$  describing the model’s pose. More precisely,  $T$  is composed of translation  $(t_x, t_y, t_z)$  and rotation parameters  $(\theta_x, \theta_y, \theta_z)$  as well as a uniform scaling  $s$ .

### Projection of the SSM

The reconstruction of a 3D rib cage from 2D projection images is based on the comparison of the projection of an SSM to the given reference images. For this purpose, instances  $S(\mathbf{b}, T)$  of the SSM with varying shape parameters  $\mathbf{b}$  and transformation  $T$  are projected.

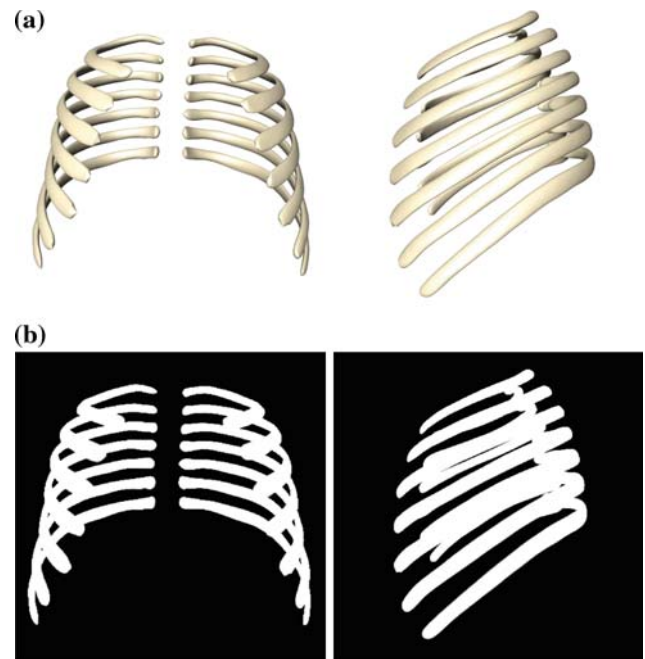
In our case, the projection image  $I_p$  of  $S(\mathbf{b}, T)$  is generated by means of rays that are emitted by a point source and propagated onto an image acquisition plane. The calibration of the point source with respect to the image plane is known and is consistent with the camera calibration that was used for the acquisition of the reference image. If the path of a ray to the image plane intersects the SSM then  $I_p(\mathbf{x}) = 1$ , where  $\mathbf{x}$  is the projected intersection. Otherwise,  $I_p(\mathbf{x})$  is set to zero. The projected SSM is then defined as

$$P = \{\mathbf{x} \in \mathbb{R}^2 : I_p(\mathbf{x}) > 0\}. \tag{2}$$

The mean shape  $\bar{\mathbf{v}}$  of the SSM and its projection images are shown in Fig. 2.

### Optimization

Our reconstruction process is defined in terms of an optimization problem. The goal is to determine values for both the shape weights  $\mathbf{b}$  and transformation parameters  $T$  of the SSM such that the projection of the 3D model  $S(\mathbf{b}, T)$  is an optimal approximation to the reference image. The model



**Fig. 2** The SSM and its projections: **a** shows the average shape  $\bar{\mathbf{v}}$  of the SSM and in **b** binary projection images of  $\bar{\mathbf{v}}$  are depicted

$S(\mathbf{b}, T)$  then serves as an estimate for the 3D shape and 3D pose of the individual rib cage’s anatomy. The quality of the approximation is measured via an objective function  $D$ , hence the goal is to compute

$$(\mathbf{b}^*, T^*) = \arg \min_{\mathbf{b}, T} D(\mathbf{b}, T). \tag{3}$$

By minimizing  $D$  the similarity between the reference images and the projections of  $S(\mathbf{b}, T)$  is maximized.

The objective function  $D$  depends on the specific data and thus may exhibit a highly complicated behavior with many local minima. There is no general rule how to treat such non-convex problems. Applying local minimization schemes requires additional suitable heuristics to guide the minimization into the right direction.

For the minimization of  $D$  a gradient-based optimization procedure was adapted [32]. Starting initially at  $\mathbf{v}_0 = (\mathbf{b}_0, T_0)$ , during step  $i$  of an iterative process, the gradient  $\nabla D$  is approximated at the current solution  $\mathbf{v}_i$ . Performing a line search along the direction of  $\nabla D$  an improved solution  $\mathbf{v}_{i+1}$  is computed. This is repeated until a termination criteria is reached: if  $\|\nabla D\|$  falls below a certain threshold at a step  $n$ , the iteration is stopped, and  $\mathbf{v}^* = \mathbf{v}_n$  is the final solution.

The line search is performed as follows: an initial large length parameter  $\delta_0$  and a minimal, final parameter  $\delta_f$  are defined by the user.  $\nabla D$  is scaled to  $\delta_0$  with regard to its maximal component,  $\mathbf{d} := \frac{\delta_0}{\|\nabla D\|_\infty} \nabla D$ . The step size is then defined by  $\|\mathbf{d}\|$ . As long as the value of  $D$  improves, another step is taken in this direction. Hereafter, the line search is refined with step sizes  $\delta_j$  iteratively reduced by a constant

factor, starting from the current solution: in each iteration, one step is taken in the direction of  $\pm \nabla D$ . This is repeated until step size  $\delta_f$  is reached.

Large steps in the beginning of the line search contribute to a faster convergence and prevent the optimization of getting stuck in small local minima, introduced by noise in the data. By narrowing the search interval with reduced step sizes the neighborhood of a minimum is sampled with increasing accuracy.

However, even with these heuristics, the optimization is still local in nature and cannot detect a global optimum. Therefore, a good initialization is of particular importance.

### Distance measures

A distance measure is used to assess dissimilarities between the projection of the SSM and the reference images. It serves as the objective function  $D$  for the optimization of the SSM’s parameters. In the following, we define different distance measures that are to be tested for their suitability for the 3D reconstruction of the rib cage. Different degrees of prior knowledge that could be extracted from X-ray images are considered. Reasons as to when to choose which specific distance measure are motivated and supported by qualitative results in the section on “Choice of distance measure”.

### Silhouettes

A silhouette  $S$  is a set of contours. The silhouette  $S_m$  of the projected SSM  $P$  in Eq. (2) is defined by the boundary  $\partial P$  and is referred to as the *model silhouette*. The *reference silhouette*  $S_r$  in a reference image is assumed to be known. We follow the approach of a silhouette-based distance measure as proposed by Lamecker et al. [16].


Given a reference and a model silhouette, the silhouette distance between a contour point  $\mathbf{x}_m \in S_m$  to  $S_r$  is defined as

$$d_e(\mathbf{x}_m, S_r) = \min_{\mathbf{x}_r \in S_r} (\|\mathbf{x}_m - \mathbf{x}_r\|), \tag{4}$$

where  $\|\mathbf{x}_m - \mathbf{x}_r\|$  is the Euclidean distance between  $\mathbf{x}_m$  and  $\mathbf{x}_r$ . The distance  $d_e(\mathbf{x}_r, S_m)$  from  $\mathbf{x}_r$  to  $S_m$  is defined analogously. In case the correspondences between individual rib silhouettes are given, i.e., the ribs are labeled within the images, the distance is defined between contours of these corresponding ribs only. Then,  $d_e(\mathbf{x}_m, S_r)$  for a point  $\mathbf{x}_m$  in a contour  $s_i \in S_m$  to another contour  $s_j \in S_r$  is defined with the constraint that  $i = j \in \{1, \dots, l\}$  applies, where  $l$  is the number of labeled ribs.

### Symmetric and asymmetric distance measure

The objective function to be minimized can be defined symmetrically or asymmetrically. The symmetric distance measure is given by the integrated symmetric squared



**Fig. 3** Avoiding contour mismatches: **a** mismatches can be identified by different orientations of the silhouette normals, **b** whereas the orientation at corresponding contour points is similar

distance between two silhouettes  $S_m$  and  $S_r$ ,

$$D_S = \int_{\mathbf{x}_m \in S_m} d_e(\mathbf{x}_m, S_r)^2 d\mathbf{x}_m + \int_{\mathbf{x}_r \in S_r} d_e(\mathbf{x}_r, S_m)^2 d\mathbf{x}_r, \tag{5}$$

while the asymmetric distance measure is defined as the integrated squared distance from silhouette  $S_m$  to  $S_r$ ,

$$D_A = \int_{\mathbf{x}_m \in S_m} d_e(\mathbf{x}_m, S_r)^2 d\mathbf{x}_m. \tag{6}$$

### Contour normals

A feature that can assist in finding meaningful point correspondences between silhouette points are the contour normals of silhouettes. As illustrated in Fig. 3, orientation differences of the normals at two contour points indicate a mismatch. To avoid contour mismatches, we can extend the distance measure to take into account the orientation of the silhouette’s normals.

The modified distance is then defined as

$$d_n(\mathbf{x}_m, S_r) = \min_{\mathbf{x}_r \in S_r} ((1 + c \cdot e)(2 - n)), \tag{7}$$

where  $n = \mathbf{n}_r \cdot \mathbf{n}_m$  denotes the inner product of two contour normals at points  $\mathbf{x}_r$  and  $\mathbf{x}_m$ , and  $e = \|\mathbf{x}_m - \mathbf{x}_r\|$ .  $n$  serves as a measure of the orientation difference of both normals. Note that  $\mathbf{n}_r$  and  $\mathbf{n}_m$  are normalized. The purpose of the pre-defined constant factor  $c$  is to balance the impact of the Euclidean distance and orientation difference.

### Area difference

Another measure, which we propose to incorporate into the objective function, is the relative area difference of silhouettes. The distance measure is extended by a quotient  $a$  of the areas  $A_{S_r}$  and  $A_{S_m}$  enclosed by the silhouettes  $S_r$  and  $S_m$ ,

$$d_a(\mathbf{x}_m, S_r) = \min_{\mathbf{x}_r \in S_r} ((1 + c \cdot e)(2 - n)(2 - a)^2), \tag{8}$$

where  $a$  is given by  $A_{S_m}/A_{S_r}$  if  $A_{S_m} < A_{S_r}$  and  $A_{S_r}/A_{S_m}$  otherwise, with

$$A_{S_m} = \int_{\mathbf{x} \in P} 1 d\mathbf{x}. \tag{9}$$

Discrete distance computation

In the following details of the computation of the distances  $d_e$  and  $d_n$  defined in Eqs. (4) and (7) are described.

*Silhouette-based Euclidean distance*

Model silhouettes can be obtained automatically from  $I_P$ . They are represented as sets of discrete pixels, which define the outer rib boundaries in the rasterized projection image. The outer rib boundaries in the reference images are represented in the same manner. Consequently, the distance measures presented in Eqs. (5) and (6) are implemented as sums over centers of silhouette pixels. The distance  $d_e$  between two silhouettes in Eq. (4) is computed by means of 2D distance maps, which encode the Euclidean distance for an arbitrary image pixel to the closest point on a given silhouette. Therefore, for each point  $\mathbf{x}_m$  the distance can be retrieved from the distance map of  $S_r$  at  $\mathbf{x}_m$ .

*Extended distances*

The normal-extended distance  $d_n$  is computed as follows: for a point  $\mathbf{x}_m$  with a certain normal orientation  $\mathbf{n}_m$ , the point  $\mathbf{x}_r$  with normal  $\mathbf{n}_r$  must be determined that minimizes the combined distance in Eq. (7). This  $\mathbf{x}_r$  is not necessarily the point in  $S_r$  with the smallest Euclidean distance to  $\mathbf{x}_m$ . For this reason, the distance evaluation cannot be performed by calculating one distance map of  $S_r$  as previously described. Instead, we compute a set of Euclidean distance maps  $M_i (i = 1, \dots, k)$  as follows:

1. For all points in a reference silhouette and a model silhouette the normals are computed. A normal  $\mathbf{n}_m$  of  $\mathbf{x}_m \in S_m$  is given by the normalized image gradient of  $I_p$ .
2. All silhouette points are sorted into point sets  $p_i$  according to their normal orientation. Therefore, a set of  $k$  normalized vectors

$$\mathbf{n}_i = (\cos(\varphi_i), \sin(\varphi_i))^T \tag{10}$$

is built where  $\varphi_i = i \cdot 2\pi/k$  is the orientation angle. A point with a normal  $\mathbf{n}_r$  is assigned to the set  $p_i$  if  $\arccos(\mathbf{n}_i \cdot \mathbf{n}_r) \leq \pi/k$  holds true.

3. For each  $p_i$  a Euclidean distance map  $M_i$  is calculated.

The distance  $d_n(\mathbf{x}_m, S_r)$  for a point  $\mathbf{x}_m$  is then computed as follows: for each  $p_i$  the minimum Euclidean distance  $e_i$  of  $\mathbf{x}_m$  to  $p_i$  is retrieved from the distance map  $M_i$ . Then,

$$d_n(\mathbf{x}_m, S_r) = \min_i ((1 + c \cdot e_i)(2 - \mathbf{n}_m \cdot \mathbf{n}_i)). \tag{11}$$

For the area-extended distance  $d_a$  the area of a silhouette is given by the number of pixels enclosed by the silhouette contour.

**Choice of distance measure**


The distance measures defined in the previous section are based on silhouette contours. We choose edge-based features for the reconstruction of the rib cage over region-based measures for the following reason: in radiographs the cortical bone layer of the ribs appears reasonably contrasted against the background in regions where the X-rays penetrate the bone surface nearly in a tangential direction. Therefore, the contours of the ribs stand out more clearly and are more easily detectable. Region-based features, on the other hand, are almost non-existent, since the intensity of the inner regions of the ribs hardly differ from the background. Thus, with regard to future clinical applications, we choose to use edge-based features. Motives to choose a particular measure from the different silhouette-based distance measures are given below.

A silhouette-based distance measure (symmetric Eq. (5) or asymmetric Eq. (6)) is a promising choice for the reconstruction of the pose and the shape of a rib cage, provided that the silhouettes are correctly extracted and the ribs are *labeled* in the reference images. However, tests showed that a silhouette-based distance measure in combination with *unlabeled* rib silhouettes causes erroneous reconstructions. We identified two problems:

1. *Mismatches of rib contours* can emerge, i.e., ribs are fitted to the contours of two adjacent ribs in the reference images (see Fig. 4a). This is due to the spatial similarity of the ribs' width and the spacing between them.
2. *Mis-assignments of ribs* can occur, caused by the repetition of the ribs in the bone ensemble and the symmetric geometry of the rib cage (see Fig. 4b).

Contour mismatches can be avoided by using the normal-extended distance  $d_n$  (Eq. (7)). Mis-assignments of ribs are more difficult to cope with. Using a symmetric distance measure alleviates, but does not eliminate the effect. We solved this problem by proper automatic initialization that anticipates mis-assignments of the ribs, as described below.

The use of *asymmetric* distance measures in combination with unlabeled ribs introduce another problem: the incorrect determination of the scale factor  $s$  (c.f. Fig. 4c). The reason



**Fig. 4** Possible mis-adaptations: **a** with a distance measure  $D_S$  and a distance  $d_e$  contour mismatches occur, and **b** mis-assignments of ribs appear; **c** in case an asymmetric measure  $D_A$  is used, the problem of

incorrect scaling is introduced; **d** plot of the scale factor  $s$  in the neighborhood of the optimal reconstruction: a scale factor below the optimal factor has too low an impact on the objective function  $D_A$

for this is illustrated in Fig. 4d. The plot of the objective function  $D_A$  demonstrates that a scale factor below the optimal factor has a very low impact on the objective function, whereas too high factors are accounted for accordingly. Erroneous reconstructions due to too small scale factors are the consequence.

Nevertheless, asymmetric distance measures hold the power to better cope with incomplete and erroneous contours within the X-ray images.

In case the asymmetric distance measure is applied the adaptation of the scaling can be controlled by the distance  $d_a$ , which incorporates the relative area difference of the silhouettes.

## Quantitative results

We evaluated the accuracy of the method using the different distance measures to assess if the method is applicable to

1. *determine 3D pose differences* between images by estimating the 3D pose of a rib cage from projection images, given that the individual 3D rib cage model of the patient is available. We evaluated pose reconstructions from *two* biplanar, binary projection images (PA and LAT-view) as well as from *one* binary PA-projection.
2. *perform a 3D shape reconstruction* of an unknown rib cage from *two* biplanar, binary projection images (PA and LAT-view), i.e., the simultaneous reconstruction of the shape and the pose.

## Experiments setup

In order to evaluate the reconstruction quality, we used the 29 surface models of different rib cages (*reference surfaces*), which were extracted from CT-data and previously used for the creation of the SSM. These reference surfaces were used

to generate binary projection images to which the shape model was fitted.

All experiments involving shape reconstruction were performed with an SSM that excludes the shape to be reconstructed from the set of trainings shapes (leave-one-out test). The weights **b** of all remaining 27 shape modes were optimized.


For all experiments involving pose reconstructions, the rib cage models (SSM or non-deformable) were initialized with different, random pose parameters  $T_0$ , restricted to ranges of  $\pm 30$  mm for the translations  $t_x, t_y, t_z$  and  $\pm 15^\circ$  for the rotation parameters  $\theta_x, \theta_y, \theta_z$ . The initial scale was restricted to factors between 0.7 and 1.3. The ranges chosen for rotation and scaling follow reasonable pose variations observed during the acquisition of X-ray images in medical practice.

For reconstructions from *two* biplanar images, a PA and a LAT-view were used. The camera calibrations for both images were known. The orientation of the LAT-projection source was orthogonal to the PA-view. Its position is given by a  $90^\circ$  rotation of the PA-projection source around the longitudinal axis of the reference surface (see Fig. 5). In case a reconstruction was performed from *one* image, only the calibrated PA-view was used. In this virtual setup the distance from the projection sources to the image acquisition plane was  $1\text{ m}$ ; the angle of the field of view was  $22^\circ$ . In this well-defined setting the silhouettes of the reference images can be extracted from the binary images as described in the section on “Silhouette-based Euclidean distance”. The experiments were performed with  $k = 16$  distance maps. The weighting factor  $c$ , ranging from 1 to 65, was found empirically to yield satisfactory results for the different distance measures.

## Evaluation method

The virtual setup described above enables us to verify the accuracy of our reconstruction experiments in 3D, since the





**Fig. 5** Image acquisition setup: the projection sources of the PA-view and the LAT-view are calibrated to generate two orthogonal images of the reference surface; a possible initialization of the SSM within this setup is shown: the silhouettes to be matched are those of the projected SSM (*beige*) and the reference silhouettes (*red*)

exact parameters that need to be recovered from the reference images are known. The reference surfaces served as a gold standard for the evaluation of the reconstruction results. We measured the reconstruction quality in 3D via the mean surface distance between the surface of the adapted SSM and the reference surface. The surface distance was computed between corresponding vertices of the reference surface and the adapted model. The following reconstruction results are given as average and maximum mean surface distance across the individual reconstructions. The maximum mean distance represents the worst of the reconstructions (29 for experiments involving shape reconstructions, 20 for pose reconstruction experiments) in terms of its average distance to the reference surface and thereby yields evidence on the robustness of the method. Additionally, we determined the error for each transformation parameter to measure the quality of the pose reconstruction. Here, especially the rotational errors are of interest, as the quality of subtraction images is mainly influenced by AP-inclination ( $\theta_x$ ) and rotation around the longitudinal axis ( $\theta_z$ ) [2]. The results of the respective distances  $d_e$ ,  $d_n$ , and  $d_a$  are summarized in Tables 1, 2, 3 and are discussed in the remainder of this section.

### Pose reconstruction

We adapted patient-specific rib cage models with different, random pose initializations  $T_0$  to reference images of the respective rib cage geometry. Only parameters of  $T$  are to be determined. Ideally, the surface distance for such pose reconstruction would reduce to zero, since the shape of the rib cage is known. In terms of the reconstruction quality the

randomly initialized models deviated on average 21.8 mm (max. 37.4 mm) from the reference surface with the correct pose before our method was applied.

In case reconstructions were performed with *two* reference images, the distance  $d_e$  in combination with *labeled* ribs led to results with surface distances  $\leq 0.2$  mm and rotation errors  $\leq 0.2^\circ$  on average, regardless of whether an asymmetric or a symmetric distance measure was used (see Table 1). A symmetric measure with distance  $d_e$  and *unlabeled* ribs occasionally caused contour mismatches and mis-assignments of ribs, and it yielded an average reconstruction quality of  $2.0 \pm 3.0$  mm (max. 8.5 mm). With an asymmetric distance  $d_n$  applied to the reconstruction of unlabeled ribs the additional problem of incorrect scaling occurred ( $11.0 \pm 6.9$  mm (max. 27.4 mm)). Here, the scaling  $s$  highly deviates with a factor of 0.25 on average from the correct size. The effects of the contour mismatches and incorrect scaling were alleviated ( $4.6 \pm 3.9$  mm (max. 9.3 mm)) by using distance  $d_a$ . In this case the errors in the parameters  $T$  were  $\leq 2.9^\circ$  for the rotation and  $\leq 5.4$  mm for the translation. These errors are due to ribs mis-assignments that occurred in all the experiments with unlabeled ribs. In order to overcome the problem of these mis-assignments, we fit only a subset of ribs to the image data to initialize the pose. The ribs corresponding to this subset were labeled in the image data and were chosen to be spatially located as far apart from each other as possible in the projections (see Fig. 6). This avoids the periodicity of the ribs and prevents mutual occlusion in the LAT-projection image. After such an initialization, a subsequent pose and shape reconstruction was performed using the entire set of ribs. It yielded a surface distance of  $0.6 \pm 0.2$  mm (max. 1.2 mm) and rotation errors below  $0.9^\circ$  for distance measure  $D_A$  and distance  $d_a$ .

We additionally carried out pose reconstructions from *one* PA-view. In these experiments the scale parameter was not optimized, since depth, i.e., translation along  $\theta_y$ , and size are redundant. However, this is of little concern for a pose reconstruction, because the size of a rib cage is determined during the prior 3D shape reconstruction. Nevertheless, if two calibrated views are available, adapting the size of a rib cage during a pose reconstruction is possible and has its advantages: A possible scaling error that emerges during 3D shape reconstruction can be compensated for, which may result in an improved reconstruction of the pose. A drawback of this approach is that calibrated PA and LAT-views are needed that represent the same scene. For this reason, it is worthwhile to conduct pose reconstructions from one PA-image. Reconstructions from one PA-view led to an average accuracy of  $0.3 \pm 0.2$  mm (max. 1.0 mm) on labeled ribs and  $0.7 \pm 0.4$  mm (max. 1.8 mm) on unlabeled ribs after a proper initialization using  $D_S$  and  $d_e$ . The reconstruction qualities of these results are slightly lower in comparison to the respective reconstructions from two views. It is noteworthy that the rotation angles

**Table 1** Results for the silhouette-based distance  $d_e$ 

Experiment	Views	Pose surf. dist.	Transform. error ( $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, s$ )	Shape surf. dist.	Pose + shape surf. dist.	Transform. error ( $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, s$ )
$D_S, L$	PA/LAT	$0.16 \pm 0.06$ (0.28)	0.08, 0.18, 0.12, 0.05, 0.14, 0.21, 0.00	$2.08 \pm 0.96$ (4.11)	$2.18 \pm 0.82$ (3.91)	0.40, 0.56, 0.62, 0.61, 0.20, 0.30, 0.00
$D_A, L$	PA/LAT	$0.17 \pm 0.07$ (0.39)	0.06, 0.21, 0.14, 0.06, 0.13, 0.19, 0.00	$2.00 \pm 0.87$ (4.19)	$2.62 \pm 1.15$ (5.30)	0.60, 0.84, 0.90, 0.80, 0.19, 0.35, 0.00
$D_S, U$	PA/LAT	$1.97 \pm 2.98$ (8.45)	0.92, 2.31, 2.20, 0.44, 1.26, 1.66, 0.00	$2.34 \pm 0.98$ (4.26)	$4.53 \pm 2.00$ (8.98)	1.27, 3.68, 2.70, 1.85, 1.60, 2.16, 0.00
$D_A, U$	PA/LAT	$19.2 \pm 10.9$ (38.3)	9.06, 11.5, 35.7, 7.95, 3.64, 4.57, 0.25	$2.54 \pm 1.36$ (7.12)	–	–
$D_S, L$	PA	$0.29 \pm 0.21$ (0.97)	0.10, 0.57, 0.16, 0.06, 0.09, 0.14, –	–	–	–
$D_S, U, \text{Init.}$	PA	$0.72 \pm 0.37$ (1.75)	0.25, 1.41, 0.23, 0.13, 0.27, 0.41, –	–	–	–

The averaged mean surface distances with standard deviations (and maximum mean surface distances) are given in millimeter. Additionally, the average reconstruction error is listed for each transformation parameter. Rotation errors are specified by Euler angles in degree. Experiments using labeled ( $L$ ) or unlabeled ( $U$ ) silhouettes of the ribs as well as asymmetric ( $D_A$ ) or symmetric ( $D_S$ ) distance measures are listed. For the shape reconstructions leave-one-out tests were performed

**Table 2** Results for the extended distance  $d_n$ 

Experiment	Views	Pose surf. dist.	Transform. error ( $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, s$ )	Shape surf. dist.	Pose + shape surf. dist.	Transform. error ( $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, s$ )
$D_A, L$	PA/LAT	$0.37 \pm 0.20$ (0.81)	0.21, 0.62, 0.37, 0.15, 0.18, 0.26, 0.00	$2.00 \pm 0.87$ (4.34)	$2.48 \pm 1.04$ (5.26)	0.48, 0.99, 0.93, 0.92, 0.30, 0.44, 0.00
$D_A, U$	PA/LAT	$11.1 \pm 6.86$ (27.4)	4.95, 5.85, 25.8, 7.20, 4.28, 2.58, 0.11	$2.53 \pm 0.78$ (4.26)	–	–


**Table 3** Results for the extended distance  $d_a$ 

Experiment	Views	Pose surf. dist.	Transform. error ( $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, s$ )	Shape surf. dist.	Pose + shape surf. dist.	Transform. error ( $t_x, t_y, t_z, \theta_x, \theta_y, \theta_z, s$ )
$D_A, L$	PA/LAT	$0.28 \pm 0.13$ (0.52)	0.11, 0.47, 0.29, 0.10, 0.19, 0.30, 0.00	$2.11 \pm 0.82$ (3.96)	$2.85 \pm 1.09$ (4.99)	1.34, 1.44, 2.03, 1.13, 0.33, 0.52, 0.01
$D_A, U$	PA/LAT	$4.60 \pm 3.90$ (9.30)	2.03, 5.36, 4.63, 2.33, 2.88, 2.87, 0.00	–	–	–
$D_A, U, \text{Init.}$	PA/LAT	$0.55 \pm 0.22$ (1.15)	0.38, 0.80, 0.24, 0.13, 0.37, 0.86, 0.00	$2.53 \pm 0.78$ (4.26)	$4.69 \pm 2.35$ (12.0)	1.32, 2.85, 2.16, 2.02, 1.02, 2.32, 0.02

of interest can be recovered with a higher accuracy ( $\leq 0.1^\circ$  and  $\leq 0.4^\circ$ ) than from two views. From the translation error along  $\theta_y$  it is obvious that the error is mainly due to a slightly erroneous estimation of the rib cage's projection depth.

#### Shape and pose reconstruction

For the evaluation of a simultaneous reconstruction of the shape and the pose it is necessary to know to what extent



**Fig. 6** Pose initialization: a subset of the average shape model is fitted to the segmentation of the corresponding subset in the reference images

the SSM can approximate the shape of an unknown rib cage independently from the pose. Hence, the method's ability to assess only the shape of an unknown rib cage with a given correct pose is demonstrated at first. Then, the results of the simultaneous pose and shape reconstruction are presented.

#### Shape reconstruction

To assess the optimal value obtainable with our SSM as a reference value, we performed *direct surface optimizations* in 3D by minimizing the distance between the SSM and the reference surfaces. This leave-one-out test resulted in a mean surface distance of  $1.6 \pm 0.6$  mm (max. 2.9 mm).

For the evaluation of the 3D shape reconstruction from 2D image data, the SSM was adapted to 29 pairs of binary reference images with a given, correct pose. All the shape

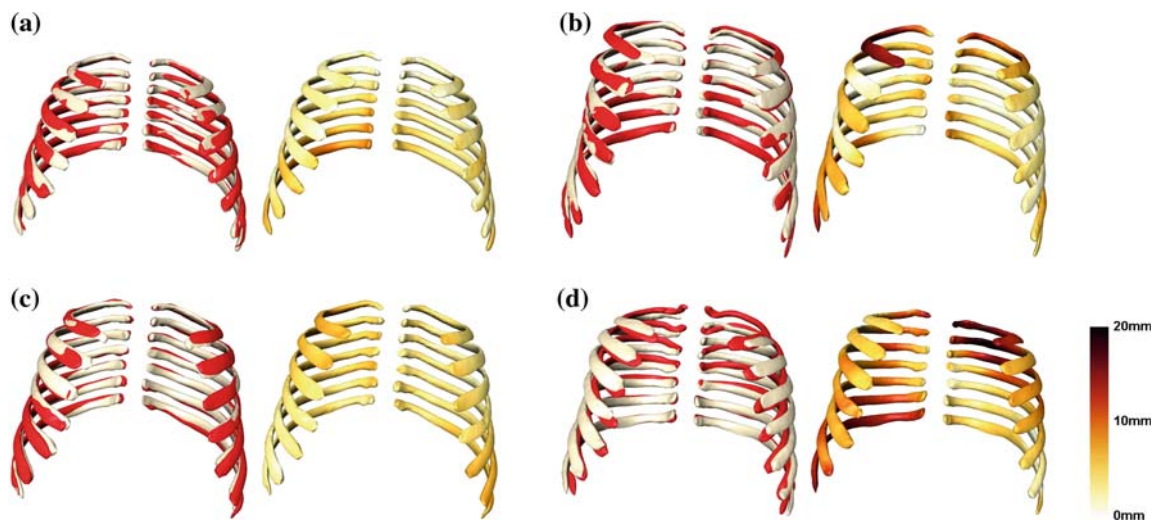
reconstruction experiments with labeled ribs yielded similar results of approximately 2 mm mean surface distance averaged across the individual 29 reconstructions. Using unlabeled ribs, the reconstruction quality slightly degrades for all distance measures with results ranging from 2.3 to 2.5 mm.

#### Combined shape and pose reconstruction

For the 3D reconstruction of an a priori unknown geometry from clinical data. The shape and the pose must be reconstructed simultaneously. To evaluate the method in this context, we fitted the SSM to pairs of reference images of an unknown rib cage using different random pose initializations.

With an average reconstruction quality of 2.2 mm a simultaneous pose and shape reconstruction on labeled ribs using a symmetric measure with distance  $d_e$  yielded a result that is comparable to the pose independent shape reconstruction (see section on “Shape reconstruction”). Reconstruction examples are shown in Fig. 7a and b. In case only unlabeled rib silhouettes were available the symmetric measure  $d_e$  resulted in  $4.5 \pm 2.0$  mm (max. 9.0 mm).

For unlabeled ribs, the distance  $d_a$  performed best among the asymmetric measures, since incorrect scaling and contour mismatches were avoided. However, in this case an initialization of the pose using a less ambiguous subset of ribs is required before the actual reconstruction, as explained above. Experiments under these conditions yielded an average reconstruction quality of  $4.7 \pm 2.4$  mm (max. 12.0 mm) (see example reconstructions in Fig. 7c and d).



**Fig. 7** Exemplary results of combined pose and shape reconstructions for the two most suitable distance measures on labeled and unlabeled ribs: for each result both the reference surface (red) and the reconstructed surface (beige) are depicted on the left; the reconstructed surface with its distance to the reference surface is shown on the right

(additionally, the average surface distance is given). (a) A good result (labeled ribs,  $D_S, d_e$ ): 2.03 mm; (b) A more difficult case (labeled ribs,  $D_S, d_e$ ): 2.44 mm; (c) A good result (unlabeled ribs,  $D_A, d_a, \text{Init.}$ ): 1.81 mm; (d) Problematic case (unlabeled ribs,  $D_A, d_a, \text{Init.}$ ): 5.04 mm

## Discussion

### Accuracy of pose reconstruction

The requirements for the pose reconstruction are difficult to define as it is not clear yet which accuracy is needed to improve the quality of the image subtraction with respect to the detection of interval changes within a clinical setting. However, it would already be beneficial to know up to what 3D pose differences between previous and follow up image interval changes are still detectable using present 2D to 2D image registration approaches. To this end, von Berg et al. [4] performed a phantom study. They conclude that interval changes are detectable up to cut-off angles of  $4.1^\circ$  with non-deformable registration. While these cut-off angles are specific to the nature of the interval change, von Berg et al. reason that by using their deformable 2D to 2D registration approach the cut-off angle is increased by  $2.3^\circ$ .

With our method, rotation errors are below  $0.3^\circ$  on average for all distance measures on *labeled* ribs compared in Tables 1, 2, 3. On *unlabeled* ribs the pose reconstructions with distance measure  $D_A$  with  $d_a$  from two views as well as the distance measure  $D_S$  with  $d_e$  from one PA-view—both using prior initialization—yield an average accuracy of  $0.4^\circ$ . Hence, for these distance measures the accuracy is sufficient to detect undesirable pose deviations. Other distance measures on unlabeled ribs exhibit relatively high average reconstruction errors and may lack the required accuracy.

To the best of our knowledge there are no other methods that quantify pose reconstruction of the 3D rib cage from 2D projections. Bhunre et al. [21] recover the pose of the distal and proximal femur from one 2D image within an experimental setup that is similar to ours. They also measure the accuracy of their method by assessing the angular error using Euler angles and achieve results of  $0.8^\circ$  to  $2.3^\circ$  on average, where we obtain  $0.1^\circ$  to  $0.4^\circ$  for the rib cage from one PA-view.

### Accuracy of shape reconstruction

#### Causes of errors

The direct 3D surface optimization experiment of pose independent shape reconstructions (pose parameters are known in advance) shows that a considerable part of the reconstruction error is due to the limitation of the SSM's model space. The cause of the residual error is presumably the mutual occlusion of the ribs in the projection data. The accuracies of the pose independent shape reconstruction for all distance measures on labeled and unlabeled ribs are close to the results of the direct 3D surface optimization. Thus, all distance measures are applicable. The results are comparable to those achieved in the work of Lamecker et al. [16] for the reconstruction

of the pelvis (2.6 mm mean surface distance), despite of the geometrically more demanding shape of the rib cage.

The experiments on the pose independent recovery of the shape reveal the capability of the SSM to reconstruct an unknown shape from projection images using our method. In a realistic reconstruction of an unknown patient-specific rib cage the pose parameters need to be estimated simultaneously with the shape. For this combined shape and pose recovery all the tested distance measures are suitable in case the ribs are *labeled*. The reconstruction errors range from 2.2 to 2.9 mm, i.e., they are only 0.1 to 0.7 mm worse than the pose independent shape reconstructions. The error in the pose parameters ranges between  $0.2^\circ$  and  $1.1^\circ$ . Thus, the pose is also accurately estimated during combined shape and pose reconstructions on labeled ribs.

Robust reconstruction results on combined shape and pose reconstruction with *unlabeled* ribs can only be obtained with one of the distance measures tested ( $D_S$ ,  $d_a$ ). Due to the initialization procedure and the incorporation of the contour normals into the distance measure, global mis-assignments of ribs and contour mismatches are avoided. However, the accuracy of reconstructions with unlabeled ribs is lower than with labeled ribs. A possible reason for this inaccuracy is that for a combined reconstruction of the shape and pose we initially fit the average shape  $\bar{\mathbf{v}}$  to the image data. Thereby the final pose of the shape yet to be reconstructed is only roughly approximated. However, small deviations from the correct pose can cause mis-adaptations of the model driven by the adaptation of the shape weights  $\mathbf{b}$ , especially if point correspondences between silhouettes are not constrained by a labeling of the ribs. In some cases on unlabeled ribs, using  $D_S$  with  $d_a$ , this leads to local mis-adaptations, that diminish the overall reconstruction quality (c.f. Fig. 7d).

#### Comparison with other methods

Most existing rib cage reconstruction methods do not obtain a reconstruction of the rib cage surface, but 3D models of rib midlines. Moreover, they mainly assess their reconstruction qualities not directly but via derived quantities, such as chord length or maximal curvature [6, 25, 27, 29]. For our application these measures are not well suited, since the objective is to globally recover the shape of a patient-specific rib cage as accurately as possible, to be subsequently used for pose reconstructions.

Benameur et al. [23] validated their rib cage reconstruction method only against another method [25], and not against some reliable ground truth, as in our case. By comparison, our method uses surfaces extracted from CT-data, which serves as a more reliable gold standard. The mean 3D distance between reconstructed model to the reference midline model is 1.6 mm on average [23], where we yield surface distances from 2.2 to 4.7 mm, depending on the distance



measure chosen. Since we compute our 3D surface distance between *corresponding* points of the reconstructed model and the reference surface, the accuracies are comparable. However, measuring the 3D distance between midlines is less sensitive against surface deviations than measuring the surface distance directly, and hence is expected to generally provide better results.

The method of Dansereau et al. [25] achieves an accuracy of 1 mm. This value is determined from radiographs of 7 isolated ribs by means of radio-opaque markers.

## Conclusions

In this work, we have addressed the fundamental problem of estimating the 3D geometry of the rib cage from 2D projection images. Solving this problem is an important step towards improving CAD based on temporal subtraction of sequences of chest radiographs. Our approach is based on measuring the deviation between silhouettes of a projected 3D SSM and silhouettes of the ribs in the radiographs. The shape and pose parameters of the SSM are varied to minimize this deviation.

We have shown that with our approach the 3D shape and pose of the ribs can be recovered with an accuracy of 4.5 mm from two calibrated projection images. When the ribs in these images are additionally labeled, the accuracy increases to 2.2 mm. Furthermore, when the shape is known the pose alone can be accurately estimated even from a single projection image, given that only two ribs are labeled in the projections. The accuracy (rotation errors  $\leq 0.4^\circ$ ) is suitable to detect undesirable pose differences between image pairs.

There is still some room for improving the accuracy and robustness of our method with regard to the optimization approach as well as the optimal choice of the parameters involved in the reconstruction process. One idea is to use stochastic optimization methods to analyze the non-convexity of our distance measures. With this we may gain insight for developing multilevel optimization strategies, which are both efficient and able to overcome local minima. We have already performed first experiments with a sequential consideration of the shape parameters. This improves the accuracy of the 3D shape reconstruction from an average surface distance of  $4.7 \pm 2.4$  mm (max. 12 mm) to  $4.4 \pm 0.8$  mm (max. 5.5 mm). In future work we will systematically investigate such strategies.

From a clinical point of view our current approach has two major limitations. First, regarding the pose reconstruction, the exact shape is assumed to be known. With respect to image subtraction, however, the shape can only be estimated approximately at one time point and this estimation could then be used to determine the pose. Therefore, in future work it must be evaluated how the pose reconstruction can

cope with inexact shapes. The second limitation concerns the assumption that the silhouettes of the ribs are exactly known in the radiographs. In practice, however, it is difficult to automatically detect them in radiographs.

Therefore, we anticipate two directions for future research. One is to improve the detection of ribs in radiographs. The other one is to investigate how to extend the distance measure to deal with incomplete input data, which is likely to be produced by rib detection methods.

We have presented first experiments using asymmetric distance measures, which suggest that incomplete silhouettes, e.g., arising from edge detecting methods, may also lead to good results.

With this work, we have laid the solid foundation to explore these future research directions. We have shown that methods can be devised to estimate the 3D geometry of the rib cage from 2D rib silhouettes in radiographs with good accuracy. Our work suggests that it is possible to develop clinical solutions that are adequate for the recovery of 3D pose differences between subsequently acquired images to increase the detection rate of developing diseases in patients.

**Acknowledgments** Heiko Seim is supported by the European Commission under the FP6 IST Project DeSSOS (027252). Dagmar Kainmüller is supported by DFG Collaborative Research Center SFB 760. Hans Lamecker is funded by the German Research Center MATHEON in Berlin. The authors would like to thank Dr. Cornelia Schaefer-Prokop for providing the radiographs.

## References

1. van Ginneken B, Ter Haar Romeny B, Viergever M (2001) Computer-aided diagnosis in chest radiography: a survey. *IEEE Trans Med Imaging* 20(12):1228–1241
2. Kano A, Doi K, MacMahon H, Hassell DD, Giger ML (1994) Digital image subtraction of temporally sequential chest images for detection of interval change. *Med Phys* 21:453–461
3. Kakeda S, Kamada K, Hatakeyama Y, Aoki T, Korogi Y, Katsuragawa S, Doi K (2006) Effect of temporal subtraction technique on interpretation time and diagnostic accuracy of chest radiography. *Am J Roentgenol* 187(5):1253–1259
4. von Berg J, Manke D, Schaefer-Prokop C, Neitzel U (2008) Impact of patient pose differences on subtle interval change detection by temporal subtraction in chest radiographs. A phantom study. *Eur Radiol* 18, Suppl 1. In: Proceedings of ECR 2008, p 212
5. Dubouset J, Charpak G, Dorion I, Skalli W, Lavaste F, Deguise J, Kalifa G, Ferey S (2005) A new 2D and 3D imaging approach to musculoskeletal physiology and pathology with low-dose radiation and the standing position: the EOS system. *Bull Acad Natl Med* 189:287–297
6. Bertrand S, Laporte S, Parent S, Skalli W, Mitton D (2008) Three-dimensional reconstruction of the rib cage from biplanar radiography. *ITBM-RBM* 29(4):278–286
7. van Ginneken B, Ter Haar Romeny BM (2000) Automatic delineation of ribs in frontal chest radiograph. *Proc SPIE* 3979:825–836
8. Yue Z, Goshtasby A, Ackerman L (1995) Automatic detection of rib borders in chest radiographs. *IEEE Trans Med Imaging* 14(3):525–536

9. Loog M, van Ginneken B (2006) Segmentation of the posterior ribs in chest radiographs using iterated contextual pixel classification. *IEEE Trans Med Imaging* 25:602–611
10. Park M, Jin JS, Wilson LS (2003) Detection and labeling ribs on expiration chest radiographs. *Medical Imaging 2003: Physics of Medical Imaging*. In: *Proceedings of the SPIE*, vol 5030, pp 1021–1031
11. Rucklidge WJ (1997) Efficiently locating objects using the Hausdorff distance. *Int J Comput Vis* 24:251–270
12. Dubuisson M, Jain A (1994) A modified Hausdorff distance for object matching *Pattern Recognition*. In: *Proceedings of the 12th IAPR international conference on computer vision and imaging process*, vol 1, pp 566–568
13. Besl P, McKay H (1992) A method for registration of 3-D shapes. *IEEE Trans Pattern Anal Mach Intell* 14(2):239–256
14. Feldmar J, Ayache N, Betting F (1995) 3D-2D projective registration of free-form curves and surfaces. *ICCV'95*, p 549
15. Fleute M, Lavallée S (1999) Nonrigid 3-D/2-D registration of images using statistical models. *MICCAI 1999. LNCS*, vol 1679, pp 138–147
16. Lamecker H, Wenckebach TH, Hege HC (2006) Atlas-based 3D-shape reconstruction from X-ray images. In: *Proceedings of the international conference on pattern recognition (ICPR2006)*, Bd vol I. *IEEE Comput Society, Los Alamitos*, pp 371–374
17. Zheng G (2006) Reconstruction of patient-specific 3D bone model from biplanar X-ray images and point distribution models. *ICIP06*, pp 1197–1200
18. Rusinkiewicz S, Levoy M (2001) Efficient variants of the ICP algorithm. In: *Proceedings of the international conference on 3-D digital imaging and modeling*, pp 145–152
19. Lavallée S, Szeliski R (1995) Recovering the position and orientation of free-form objects from image contours using 3D distance maps. *IEEE Trans Pattern Anal Mach Intell* 17(4):378–390
20. Cyr C, Kamal A, Sebastian T, Kimia B (2000) 2D-3D registration based on shape matching. In: *Proceedings of the IEEE Workshop on mathematical methods in biomedical image analysis*, pp 198–203
21. Bhunre P, Leow WK, Howe TS (2007) Recovery of 3D Pose of Bones in Single 2D X-ray Images. In: *IEEE workshop on applications of computer vision, WACV '07*, p 48
22. Cootes TF, Taylor CJ, Cooper DH, Graham J (1995) Active shape models—their training and application. *Comput Vis Image Underst* 61(1):38–59
23. Benameur S, Mignotte M, Destrempe F, De Guise JA (2005) Three-dimensional biplanar reconstruction of scoliotic rib cage using the estimation of a mixture of probabilistic prior models. *IEEE Trans Biomed Eng* 52(10):1713–1728
24. Mahfouz M, Badawi A, Fatah EEA, Kuhn M, Merkl B (2006) Reconstruction of 3D patient-specific bone models from biplanar X-ray images utilizing morphometric measurements. *Proc Int Conf Image Process Comput Vis Pattern Recognit IPCV 2*:345–349
25. Danserau J, Strokest IAF (1988) Measurements of the three-dimensional shape of the rib cage. *J Biomech* 21:893–901
26. Marzan GT (1976) Rational design for close-range photogrammetry. PhD Thesis, Department of Civil Engineering, University of Illinois at Urbana-Champaign
27. Delorme S, Petit Y, de Guise J, Labelle H, Aubin C, Dansereau J (2003) Assessment of the 3-D reconstruction and high-resolution geometrical modeling of the human skeletal trunk from 2-D radiographic images. *IEEE Trans Biomed Eng* 50(8):989–998
28. Novosad J, Cheriet F, Petit Y, Labelle H (2004) Three-dimensional (3-D) reconstruction of the spine from a single X-ray image and prior vertebra models. *IEEE Trans Biomed Eng* 51(9):1628–1639
29. Mitton D, Zhao K, Bertrand S, Zhao CF, Laporte S, Yang C, An KN, Skallia W (2008) 3D reconstruction of the ribs from lateral and frontal X-rays in comparison to 3D CT-scan reconstruction. *J Biomech* 41:706–710
30. Dworzak J, Lamecker H, von Berg J, Klinder T, Lorenz H, Kainmüller D, Seim H, Hege HC, Zachow S (2008) Towards Model-based 3-D Reconstruction of the Human Rib Cage from Radiographs. In: *Proc 7. Jahrestag der Dtsch Ges für Computer- und Roboterassistierte Chirurgie (CURAC)*, pp 193–196
31. Klinder T, Lorenz C, von Berg J, Dries SPM, Bülow T, Ostermann J (2007) Automated model-based rib cage segmentation and labeling in CT images. *MICCAI 2007, Part II. LNCS*, vol 4792, pp 195–202
32. Rohlfing T (2000) Multimodale Datenfusion für die bildgesteuerte Neurochirurgie und Strahlentherapie. PhD Thesis, Technische Universität Berlin