Precise 3D Angle Measurements
in CT Wrist Images

Johan Nysjö[1], Albert Christersson[2], Ida-Maria Sintorn[1], Ingela Nyström[1],
Sune Larsson[2], and Filip Malmberg[1]

[1] Centre for Image Analysis, Uppsala University and SLU, Sweden
johan.nysjo@it.uu.se
[2] Dept. of Orthopedics, Uppsala University Hospital, Sweden

Abstract. The clinically established method to assess the displacement
of a distal radius fracture is to manually measure two reference angles,
the dorsal angle and the radial angle, in consecutive 2D X-ray images
of the wrist. This approach has the disadvantage of being sensitive to
operator errors since the measurements are performed on 2D projections
of a 3D structure. In this paper, we present a semi-automatic system
for measuring relative changes in the dorsal angle in 3D computed to-
mography (CT) images of fractured wrists. We evaluate the proposed 3D
measurement method on 28 post-operative CT images of fractured wrists
and compare it with the radiographic 2D measurement method used in
clinical practice. The results show that our proposed 3D measurement
method has a high intra- and inter-operator precision and is more precise
and robust than the conventional 2D measurement method.

Keywords: Wrist fractures, CT, angle measurements, bone segmenta-
tion, interactive mesh segmentation, surface registration.

1 Introduction

Distal radius fractures occur when the radius bone in the wrist breaks between
the shaft and the joint surface. The clinically established method to assess the
displacement of such fractures is to manually measure two reference angles, the
dorsal angle and the radial angle, in 2D X-ray images of the wrist [4], as illus-
trated in Fig. 1. Although this approach usually works well enough for diagnosis
and treatment-guidance in clinical practice, it has the disadvantage of being
sensitive to operator errors since the angle measurements are performed on 2D
projections of 3D structures. An intra- and inter-operator variability of 3–4° has
been reported [4,5], which is too high for, e.g., orthopedic research studies where
more precise angle measurements are required to compare different methods of
treatment. The lack of depth in X-ray images makes it difficult to assess the 3D
positions of the bone structures. Moreover, consecutive X-ray images are seldom
acquired at exactly the same angle, and the contours of the bones are often
smooth, which makes it difficult to define reliable landmarks for the measure-
ments. In this paper, we aim to overcome these issues and precision limitations
by measuring the dorsal angle in 3D computed tomography (CT) images.
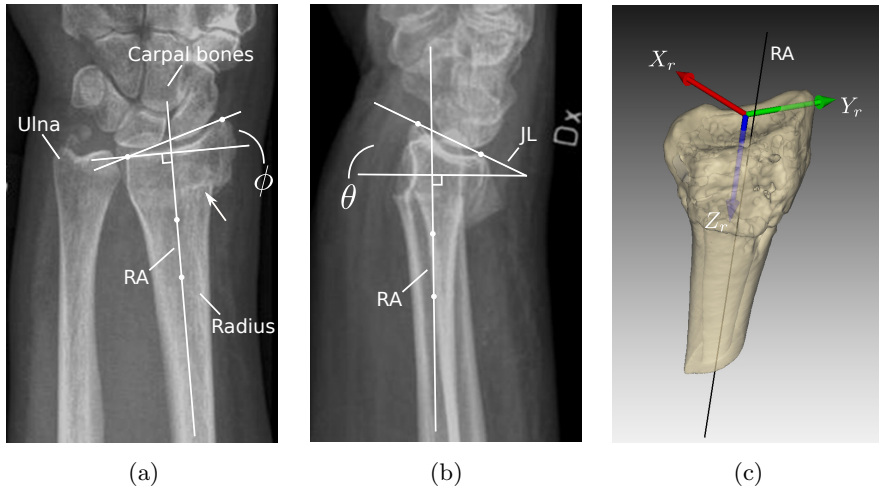
Fig. 1. (a) Frontal X-ray image of a distal radius fracture. The arrow marks the fracture location. ϕ denotes the radial angle. (b) The dorsal angle, θ, measured in 2D on a lateral X-ray image of the same wrist. θ is defined as the angle between the joint line JL and a line that is orthogonal to the long axis RA of the radius. (c) A 3D rendering of the radius bone and the reference axes we need to identify to measure the dorsal angle in 3D.

To measure the dorsal angle in 3D, we need to identify two components: (1) the long axis RA of the radius shaft; and (2) a local reference coordinate system (X_r, Y_r, Z_r) defining the orientation of the joint surface. These components are illustrated in Fig. 1c. The first component, the long axis, can be identified using surface normal information and random sampling consensus (RANSAC), as proposed in [9]. In this paper, we focus on the second component, presenting a precise semi-automatic method for identifying and tracking the orientation of the joint surface over time. By combining this method with the previously presented axis estimation method [9], we obtain a system that can be used to measure relative changes in the dorsal angle in 3D CT images of fractured wrists.

2   Image Data and Preprocessing

The image data used here consists of 28 post-operative CT images of fractured wrists in six patients. Each patient was scanned at 4–6 different occasions, 0–24 weeks after surgery. The CT images were acquired with a pixel spacing of 0.16–0.39 mm and a slice thickness of 0.4–0.8 mm. The original image dimensions were $512 \times 512 \times N_z$ voxels, where the number of slices, N_z, ranged from 72–346. We converted each of these CT images from a stack of DICOM images with graylevel values between -1024 and 3071 to an 8-bit VTK volume image with normalized graylevel values between 0 and 255. Thereafter, we cropped the converted CT images closely around the radius to reduce the amount of data and

speed up processing. We subsequently reflected left wrists into right wrists to simplify the construction of reference coordinate systems (our method assumes a right-handed coordinate system). Finally, we resampled the volume data from anisotropic to isotropic voxel size using cubic interpolation.

## 3   Identifying and Tracking the Joint Surface Orientation

This section describes our proposed semi-automatic method for identifying and tracking the orientation of the joint surface of a fractured radius over time. The method consists of three main steps: segmentation, template generation, and surface registration.

3.1   Segmentation

The first task is to create an accurate surface-mesh representation of the radius. The radius shaft is mainly composed of dense (cortical) bone, which appears significantly brighter than other tissue types in CT and is straight-forward to segment using, for instance, global thresholding. The joint surface, on the other hand, needs to be more flexible and is therefore composed of spongy (trabecular) bone, of which intensity distribution partly overlaps that of skin and soft-tissue. Because of this intensity overlap, simple intensity-based segmentation methods such as global thresholding cannot completely separate the radius from the rest of the image. The segmentation task is further complicated by partial volume effects (PVE) and the, in comparison with the image resolution, very narrow spacing between the articulated surfaces of the wrist bones. The PVE leads to blurring of the joint boundaries, making it difficult to separate the wrist bones from each other because the articulated surfaces appear to be in direct contact.

A common approach to deal with intensity inhomogeneities and image imprecisions is to use segmentation methods that not only consider the intensity of the voxels but also their spatial relationships. One such method is hysteresis thresholding [3], where the idea is to (1) define an upper threshold t_{high} and a lower threshold t_{low} for the object of interest, (2) threshold the image at t_{high} to generate a set of bright seed voxels that are assumed to belong completely to the object of interest, (3) threshold the image at t_{low} to extract darker candidate object voxels, and (4) apply connectivity analysis to identify and remove all candidate object voxels that are not connected to at least one seed voxel. This method lends itself well for segmenting bone in CT since the approximate Hounsfield unit (HU) ranges for cortical and trabecular bone are known.

We segmented the wrist bones using hysteresis thresholding with fixed thresholds $t_{low} = 76$ and $t_{high} = 89$, which are empirically selected and correspond to the lower intensity range for trabecular and cortical bone, respectively. We then used the marching cubes algorithm [8] to generate a triangular mesh representation of the segmented bones, which, to reduce staircase artifacts, had been postprocessed with a Gaussian smoothing kernel of size $\sigma = 0.7$ mm. The resulting mesh was simplified from on average 1,000K triangles to 340K triangles

(a)   (b)   (c)
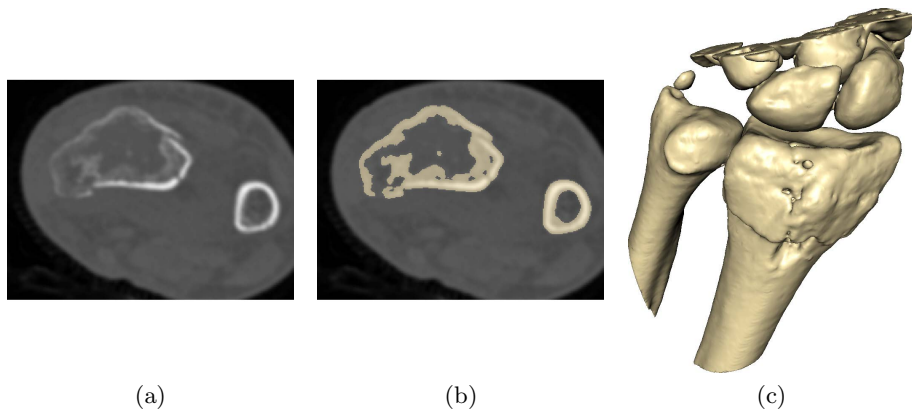
Fig. 2. Segmentation of the wrist bones: (a) preprocessed CT image; (b) segmentation result obtained with hysteresis thresholding; (c) surface mesh extracted from the segmented image with the marching cubes algorithm

using the vtkDecimatePro filter from the Visualization Toolkit (VTK) library[1] and moderately smoothed with a Laplacian filter (20 iterations, relaxation factor 0.1) to improve the mesh quality. Thereafter, we used interactive mesh-cutting by random walks [7] to separate the radius from the carpal bones and the ulna. Figures 2 and 3 illustrate the segmentation process. The sketch-based seeding interface allows the user to draw seeds directly on the surface mesh and was implemented using 3D ray-picking accelerated by an octree.

3.2   Template Generation

Having segmented the radius, our next task is to construct a template of the joint surface. We also need to derive a local reference coordinate system for this template that can be used to describe the orientation of the joint surface.

Building a statistical shape model of the joint surface is difficult due to the high shape-variability of the radius. Instead, we use the joint surface in the first postoperative CT image as template and construct a local reference coordinate system (similar to that in [6]) from user-defined landmarks. Figure 4 illustrates the idea. The user has to perform two interactive tasks: (1) select three landmarks corresponding to the small peaks located at the corners of the joint surface; and (2) position a 3D cutting plane so that it separates the joint surface from the radius shaft and the fracture. To keep the interaction task as simple as possible, the cutting plane is initially aligned to the landmarks and cannot be scaled or rotated but only translated along its normal direction. We define an initial reference coordinate system (X'_r, Y'_r, Z'_r) as follows: $Y'_r = l_1 - \frac{l_2 + l_3}{2}$; $Z'_r = (l_1 - l_2) \times (l_3 - l_2)$; and $X'_r = Z_r \times Y_r$. The actual reference coordinate system

---

[1] URL: http://www.vtk.org/

(a)   (b)   (c)
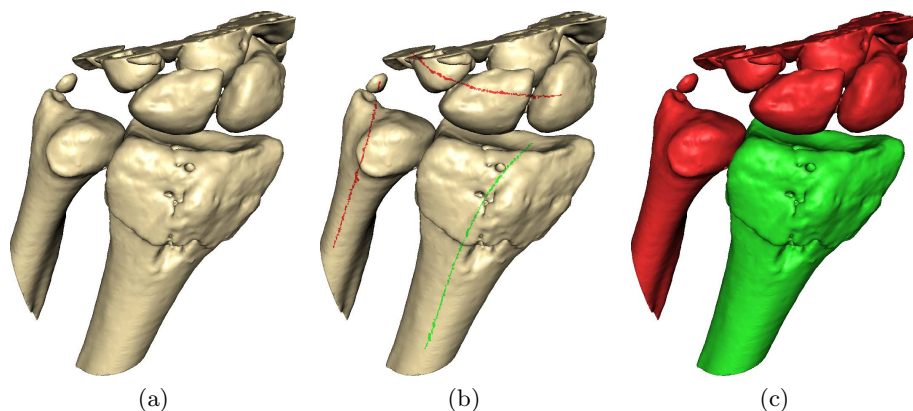
Fig. 3. Interactive separation of the radius from the carpal bones and the ulna: (a) original mesh; (b) user-defined seeds that have been drawn directly on the bone surface using our sketch-based segmentation interface; (c) mesh segmentation obtained with random walks

(X_r, Y_r, Z_r) is then obtained by applying a transform R that performs a 20-degree counterclockwise rotation of (X'_r, Y'_r, Z'_r) about X'_r, so that $X_r = X'_r$, $Y_r = RY'_r$, and $Z_r = RZ'_r$. The rotation corresponds roughly to the normal radial angle [5] in intact wrists and is required to align the axes correctly. X_r corresponds to the reference line JL shown in Fig. 1b. The landmarks positioning does not need to be very precise, but must be performed as indicated in Fig. 4.

3.3   Surface Registration

The third and final task is to register the template mesh extracted from the first postoperative CT image against the radius meshes extracted from the remaining $n - 1$ follow-up images in the CT scan sequence, so that we can determine the orientation of the joint surface in each image. To accomplish this, we developed a semi-automatic surface registration interface based on the iterative closest point (ICP) [2] algorithm. The registration is performed in two steps:

1. Coarse alignment of the template by procrustes analysis of user-defined landmarks.
2. Precise surface registration using a modified ICP algorithm, with the template as source mesh and the segmented radius as target mesh.

Figure 5 illustrates the registration process. Step 1 is required because ICP needs a good starting guess to produce accurate registration results. After applying ICP, we obtain a rigid-body transformation that, together with the local reference axes of the template, defines the orientation of the target joint surface.
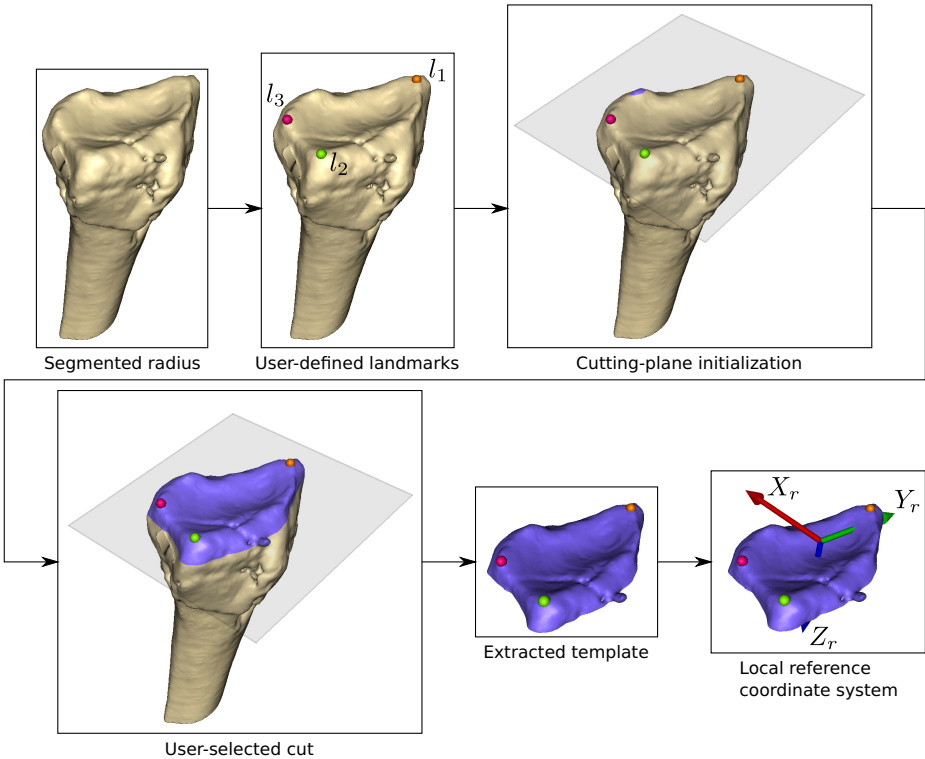
Fig. 4. Generation of a joint-surface template. The local reference coordinate system shown in the lower right image is derived from the three landmarks.

The original ICP algorithm described in [2] assumes that the input data is outlier-free and, further, that every point in the source mesh corresponds to a point in the target mesh. Unfortunately, neither of these assumptions holds for our data: due to segmentation errors and various CT image artifacts, the template might not overlap the target mesh completely, and even if perfect segmentation results were available, the joint surface might be fractured and have small disconnected fragments that move around during the healing. Furthermore, the original ICP algorithm does not take surface normal information into account, which means that there is no guarantee that it will actually fit the outer or the inner surface of the template to the corresponding surface of the radius. Our initial experiments with ICP showed that even with a good initialization, the algorithm can produce poor registration results by, for instance, fitting the outer surface of the template to the inner surface of the radius.

To remedy these problems, we implemented a modified ICP algorithm that, in every iteration, (1) identifies and rejects all closest point pairs of which normal directions differ more than 90 degrees and (2) rejects 10% of the remaining point pairs with the largest point-to-point distances. The second rejection criteria
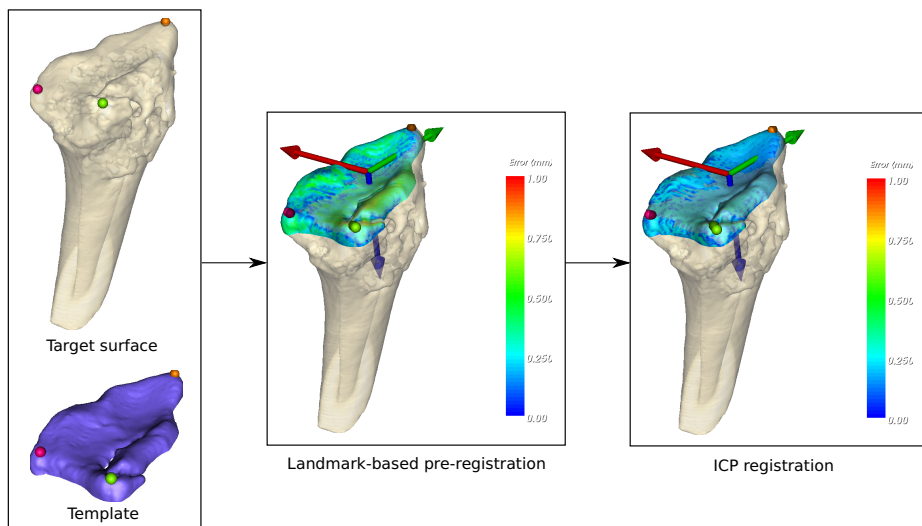
Fig. 5. Joint-surface registration. Left: user-defined landmarks. Middle: coarse registration obtained by aligning the landmarks. Right: registration result after refinement with ICP.

makes the algorithm more robust to segmentation errors. The implementation is CPU-based, sequential, and uses a kd-tree for accelerated closest-point search.

3.4 Implementation Details

We implemented the proposed segmentation and registration pipeline using VTK, Python[2], and NumPy/SciPy[3]. To solve the sparse linear system [7] generated in the mesh-segmentation step, we used a Ruge-Stüben-based algebraic multigrid solver obtained from the PyAMG library [1].

4 Computing the Dorsal Angle

Using the reference axes RA and X_r extracted with the methods described in [9] and in Section 3, respectively, we compute the dorsal angle θ as

$$\theta = \frac{\pi}{2} - \arccos(RA \cdot X_r). \tag{1}$$

5 Experiments and Results

Two test users performed repeated 3D angle measurements on the six CT scan sequences described in Section 2. The first user (U1) was an orthopedic surgeon with long experience of measuring wrist angles in 2D, whereas the second

---

[2] URL: `http://www.python.org/`
[3] URL: `http://www.scipy.org/SciPy/`

user (U2) was a PhD student in image analysis who had no prior experience of evaluating wrist fractures. The first user had not used the system before and therefore received a short training session before the experiment started. Each user repeated the measurements five times on the four CT scan sequences that included most samples and two times on the remaining two CT scan sequences, measuring 111 angles in total. The experiments were performed on a laptop equipped with an Intel Core i7-3612QM 2.1 GHz CPU, 8 GB DDR3 RAM, an Intel HD Graphics 4000 GPU, and 64-bit Ubuntu Linux 12.04. In addition, one of the users (U1) performed conventional 2D angle measurements on plain X-ray images that had been acquired at the same occasions as the CT images.
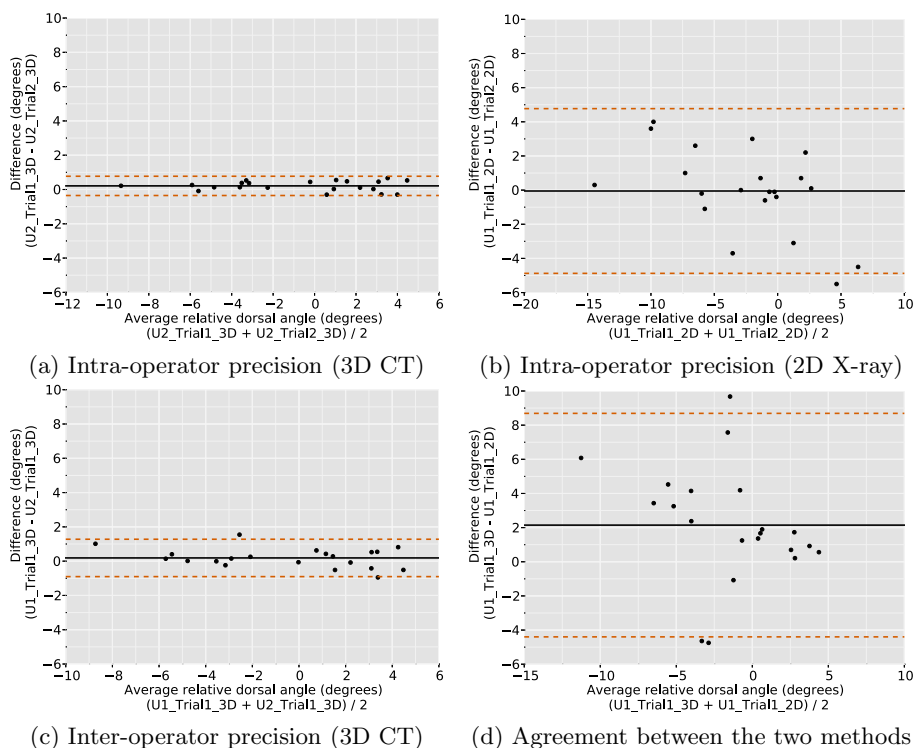
Figures 6 and 7 illustrate the angle measurement results (see caption descriptions). The intra-operator precision (mean angle difference ±1.96 SD) of the 3D
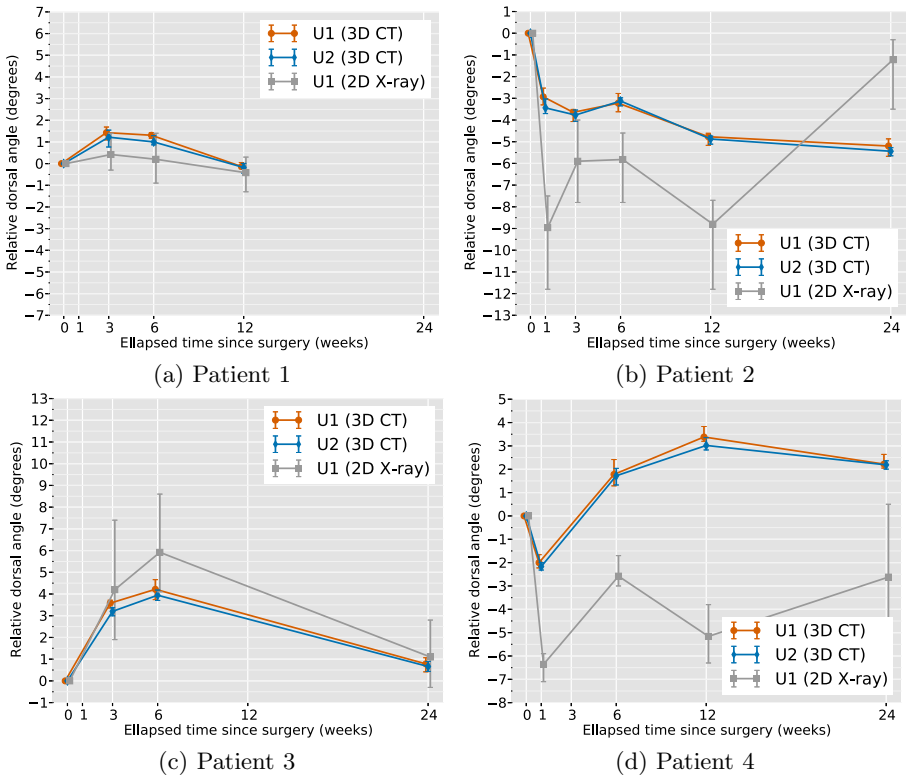
Fig. 6. Relative dorsal angles obtained in five trials with our proposed 3D measurement method (3D CT) and the conventional 2D measurement method (2D X-ray). U1 and U2 denote the two test users who performed the repeat measurements. The vertical bars show the mean and the range of the angles.

measurement method was $0.07 \pm 0.74°$ for user U1 and $0.21 \pm 0.56°$ for user U2, and the inter-operator precision was $0.19 \pm 1.09°$, indicating high repeatability. The intra-operator precision of the 2D measurement method was considerably lower, $-0.05 \pm 4.82°$, confirming the limitations of the method. Overall agreement between the 2D and 3D angles was poor, which was to be expected since the 2D angles were measured on 2D projections of the wrist. The mean computational time required to process a single CT image was 21.8 ± 7.0 seconds: 3.7 ± 2.2 s for hysteresis thresholding, 4.8 ± 3.2 s for surface extraction, 4.0 ± 1.1 s for random walks, 7.7 ± 2.4 s for ICP registration, and 1.6 ± 0.6 s for RANSAC axis estimation. The total time (interaction time plus computational time) required to process a sequence of five CT images was ∼10 minutes. There is no ground truth available to verify the obtained angles against, but the segmentation and registration results were considered successful in all trials, and the obtained axes RA and X_r have been visually inspected and deemed accurate enough.

(a) Intra-operator precision (3D CT)   (b) Intra-operator precision (2D X-ray)

(c) Inter-operator precision (3D CT)   (d) Agreement between the two methods

Fig. 7. Bland-Altman plots illustrating the intra- and inter-operator precision of the evaluated methods and the agreement between the 2D and 3D angle measurements. The solid line represents the mean angle difference, whereas the dashed lines represent ± 1.96 SD of the angle difference.

6   Conclusion

We have presented a semi-automatic method for identifying and tracking the orientation of the radius joint surface in 3D CT images of fractured wrists. By combining this method with a previously developed method for identifying the long axis of the radius, we have also developed a system that enables precise 3D measurements of relative changes in the dorsal angle, one of the reference angles orthopedic surgeons use to assess the displacement of wrist fractures. To the best of our knowledge, this is the first time such a system has been developed. The results presented in this paper show that our proposed 3D angle measurement method has a high intra- and inter-operator precision and is more precise than the conventional 2D measurement method used in clinical practice. The system is efficient enough for interactive usage and allows a user with no prior experience of evaluating wrist fractures to achieve similar results as an expert.

Next, we plan to deploy the presented system in orthopedic research studies where the objective is to compare different methods of fracture treatment. We also plan to extend the system so that it can be used to measure relative changes in other 3D rotation angles defined for the wrist, for instance, the radial angle [4]. Finally, we aim to further improve the robustness and accuracy of the angle measurement method. This could be achieved by, for instance, automatizing some of the interactive steps or replacing the underlying surface-based segmentation and registration methods with volumetric counterparts.

References

1. Bell, W.N., Olson, L.N., Schroder, J.: PyAMG: Algebraic Multigrid Solvers in Python (2008), http://www.pyamg.org (version 1.1)
2. Besl, P.J., McKay, N.D.: A Method for Registration of 3-D Shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence 14(2), 239–256 (1992)
3. Canny, J.: A computational approach to edge detection. IEEE Transactions on Pattern Analysis and Machine Intelligence 8(6), 679–698 (1986)
4. Friberg, S., Lundström, B.: Radiographic measurements of the radio-carpal joint in normal adults. Acta Radiologica Diagnosis 17(2), 249–256 (1976)
5. Johnson, P., Szabo, R.: Angle measurements of the distal radius: a cadaver study. Skeletal Radiology 22, 243–246 (1993)
6. Kobayashi, M., Berger, R.A., Nagy, L., Linscheid, R.L., Uchiyama, S., Ritt, M., An, K.N.: Normal kinematics of carpal bones: A three-dimensional analysis of carpal bone motion relative to the radius. Journal of Biomechanics 30(8), 787–793 (1997)
7. Lai, Y.K., Hu, S.M., Martin, R.R., Rosin, P.L.: Rapid and Effective Segmentation of 3D Models Using Random Walks. Computer Aided Geometric Design 26(6), 665–679 (2009)
8. Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3D surface construction algorithm. SIGGRAPH Computer Graphics 21(4), 163–169 (1987)
9. Nysjö, J., Christersson, A., Malmberg, F., Sintorn, I.M., Nyström, I.: Towards User-Guided Quantitative Evaluation of Wrist Fractures in CT Images. In: Bolc, L., Tadeusiewicz, R., Chmielewski, L.J., Wojciechowski, K. (eds.) ICCVG 2012. LNCS, vol. 7594, pp. 204–211. Springer, Heidelberg (2012)