



# HHS Public Access

Author manuscript

*Comput Methods Biomech Biomed Eng Imaging Vis.* Author manuscript; available in PMC  
2019 June 11.

Published in final edited form as:

*Comput Methods Biomech Biomed Eng Imaging Vis.* 2019 ; 7(4): 464–478. doi:  
10.1080/21681163.2018.1537859.

## A marker-free registration method for standing X-ray panorama reconstruction for hip-knee-ankle axis deformity assessment

Yehuda K. Ben-Zikri<sup>a</sup>, Ziv R. Yaniv<sup>b,c</sup>, Karl Baum<sup>d</sup>, and Cristian A. Linte<sup>a,e</sup>

<sup>a</sup>Center for Imaging Science, Rochester Institute of Technology, Rochester, NY, USA;

<sup>b</sup>National Library of Medicine, National Institutes of Health, Bethesda, MD, USA;

<sup>c</sup>TAJ Technologies Inc, Bloomington, MN, USA;

<sup>d</sup>Qmetrics Technologies, Rochester, NY, USA;

<sup>e</sup>Department of Biomedical Engineering, Rochester Institute of Technology, Rochester, NY, USA

### Abstract

Accurate measurement of knee alignment, quantified by the hip-knee-ankle (HKA) angle (varus-valgus), serves as an essential biomarker in the diagnosis of various orthopaedic conditions and selection of appropriate therapies. Such angular deformities are assessed from standing X-ray panoramas. However, the limited field-of-view of traditional X-ray imaging systems necessitates the acquisition of several sector images to capture an individual's standing posture, and their subsequent 'stitching' to reconstruct a panoramic image. Such panoramas are typically constructed manually by an X-ray imaging technician, often using various external markers attached to the individual's clothing and visible in two adjacent sector images. To eliminate human error, user-induced variability, improve consistency and reproducibility, and reduce the time associated with the traditional manual 'stitching' protocol, here we propose an automatic panorama construction method that only relies on anatomical features reliably detected in the images, eliminating the need for any external markers or manual input from the technician. The method first performs a rough segmentation of the femur and the tibia, then the sector images are registered by evaluating a distance metric between the corresponding bones along their medial edge. The identified translations are then used to generate the standing panorama image. The method was evaluated on 95 patient image datasets from a database of X-ray images acquired across 10 clinical sites as part of the screening process for a multi-site clinical trial. The panorama reconstruction parameters

**CONTACT** Cristian A. Linte [calbme@rit.edu](mailto:calbme@rit.edu) Center for Imaging Science, Rochester Institute of Technology, Rochester, NY, USA © 2018 Informa UK Limited, trading as Taylor & Francis Group.

Notes on contributors

*Yehuda K. Ben-Zikri* completed his PhD in Imaging Science at Rochester Institute of Technology in 2017. He conducted an internship at Qmetrics, where he commenced working on this project, and is now with Applied Materials in Tel Avia, Israel.

*Ziv R. Yaniv* is a senior scientists in the Office of High Performance Computing at the National Library of Medicine at the National Institutes of Health and also a co-founder of TAJ Technologies.

*Karl Baum* has a doctoral degree in Imaging Science from Rochester Institute of Technology; he was involved with this research while at Qmetrics Technologies, and is now a senior software engineer with Sparton Medical.

*Cristian A. Linte* received his PhD in Biomedical Engineering from Western University, joined the Biomedical Imaging Resource at Mayo Clinic for a research fellowship, and is now a faculty member in Biomedical Engineering and Chester F. Carlson Center for Imaging Science at Rochester Institute of Technology.

Disclosure statement

No potential conflict of interest was reported by the authors.

yielded by the proposed method were compared to those used for the manual panorama construction, which served as gold-standard. The horizontal translation differences were  $0.43 \pm 1.95$  mm and  $0.26 \pm 1.43$  mm for the femur and tibia respectively, while the vertical translation differences were  $3.76 \pm 22.35$  mm and  $1.85 \pm 6.79$  mm for the femur and tibia, respectively. Our results showed no statistically significant differences between the HKA angles measured using the automated vs. the manually generated panoramas, and also led to similar decisions with regards to the patient inclusion/exclusion in the clinical trial. Thus, the proposed method was shown to provide comparable performance to manual panorama construction, with increased efficiency, consistency and robustness.

## Keywords

Long-limb X-ray; hip-knee-ankle angle; axial deformity at the knee; panorama reconstruction; segmentation; image registration

## 1. Introduction

The presence of knee misalignment or deformity is one of the first indicators of arthritic conditions that require further disease management or interventional treatment. Knee alignment depends on the geometries of the long bone and the surfaces of the femur and tibia. The measurements of the alignment are used for assessing the affect of the arthritic condition on the knee joint. In addition, these measurements are also fundamental to various aspects of musculoskeletal research, as there is significant interest in frontal plane alignment measures to assess the pathogenesis of knee osteoarthritis (OA). However, high angular deformity limits the effectiveness of the therapy and/or accelerates the progression of OA regardless of the treatment.

The standard of care for assessing axial deformity of the knee and mechanical lower limb axis entails the measurement of the hip-knee-ankle (HKA) angle, which provides a measure of the knee varus or valgus misalignment. Moderate knee varus or valgus misalignment is characterised by an HKA angle ranging from  $2^\circ$  to  $10^\circ$  (Jordan et al. 2011). Significant angular deformity will tend to overload one compartment of the knee, either limiting the effectiveness of the therapy, or accelerating the progression of OA regardless of the treatment (Felson et al. 2013). As such, the HKA angle is used in most studies as the threshold – typically between  $\pm 7^\circ$  to  $\pm 10^\circ$  – for including or excluding patients in clinical trials. The knee deformity threshold depends on the design of the trial and the nature of the therapy being studied.

The orientation of the knee is best described by the mechanical axes of the bones. The mechanical axis of the femur (FM) is the line from the centre of the femoral head running distally to the mid-condylar point between the cruciate ligaments ((Yoshioka et al. 1987)). Similarly, the mechanical axis of the tibia (TM) is the line that joins the centre of the tibial plateau to the distal centre of the tibial plafond (Yoshioka et al. 1989). The angle between the distal extension of the FM and the TM is the HKA angle (Cooke et al. 1994; Cooke and Scudamore 2003).

When the HKA angle approaches  $180^\circ$ , the FM and TM are collinear and coincident with the line of ground reaction force passing from the ankle to the hip (i.e. the load-bearing axis (LBA)), rendering a neutrally aligned limb, (Cooke et al. 2007) (Figure 1B). In varus cases, the knee centre is lateral to the LBA axis (Figure 1A), whereas in valgus the knee centre is medial to the LBA axis (Figure 1C). As a convention,  $HKA = 0^\circ$  for neutral alignment, varus deviations are negative and valgus deviations are positive. For a helpful schematic diagram illustrating these various conditions of knee malalignment, we point the reader to the work by Derek et al. (Cooke, Sled and Scudamore 2007). Based on general observations varus leads to more severe knee damage, hence the association of varus angles with the negative convention (Cooke et al. 2007).

The HKA angle is traditionally measured on a full length weight-bearing (FLWB) X-ray of an individual in standing position. Since the measurements are based on the mechanical axes, using only a knee view X-ray increases the risk of missed deformity; therefore a standing X-ray panorama is required for proper axis assessment (Cooke et al. 1991). This observation has been recently confirmed again in (Zampogna et al. 2015), which compared two methods for the HKA angle measurement using standard X-ray images of the knee and concluded that the correlation between the gold-standard and the best of the two proposed methods was only moderate, further supporting the need for a FLWB X-ray image.

The major limitation to using full standing X-ray panoramas is that standard X-ray scanners are unable to capture images of an individual's standing posture in a single exposure, primarily due to the limited size of the X-ray detector. Consequentially, FLWB panorama image creation typically entails manual stitching of three sector images (hip, knee and ankle) by aligning each of the four image pairs (i.e. two images per left and right limb). However, given that the images are acquired sequentially and the leg opening of the standing patient may vary for each sector acquisition, the stitching process is not only time consuming, but is also highly subjective to reader-induced error and susceptible to user variability. Once the FLWB panorama is available, HKA angle measurement can be carried out in a consistent manner using semi-automated approaches (Sled et al. 2011; Skyttä, et al. 2011). Therefore, here we focus on the development of an automated stitching application, after which one can utilise any semi-automated tools for HKA angle computation, reducing both intra and inter-observer measurement variability.

Existing solutions for automatically creating FLWB X-ray images can be classified into three categories: 1) dedicated approaches that are tied to specific X-ray acquisition devices, 2) approaches that rely on external markers specifically introduced for image stitching, and 3) approaches that only rely on the image content without introducing dedicated markers.

Commercial X-ray devices that also provide a dedicated application for FLWB panorama construction are available from Carestream Inc. (Carestream 2012) and Shimadzu (Toshihiro Minami and Okamoto 2014). Both of these applications are tied to the manufacturer's X-ray device and are therefore not applicable at clinical sites that do not use the specific equipment. Another, academic, method for panorama construction that requires dedicated hardware is the Carm-C system (Wang et al. 2004, 2010). This is a modified X-ray device that incorporates a video camera for tracking the location of the X-ray source via a dedicated

pattern placed under the patient bed. Similarly to the commercial systems, this approach to panorama construction is not likely to see widespread use due to the requirement for unique hardware.

Approaches that rely on external markers visible in all images were proposed by several groups (Yaniv and Joskowicz 2004; Messmer et al. 2006; Gooßen et al. 2008; Apivatthakakul et al. 2013; Chen et al. 2015). These either use radio-opaque rulers or patterns that are visible across all sector images. This simplifies the alignment process at the cost of requiring additional hardware. An inherent challenge associated with marker based alignment is that when employing a marker which is not at the same distance from the X-ray source as the bone, a method must explicitly account for parallax. As such, the translation estimated using the marker needs to be scaled to correctly align the bone. In (Yaniv and Joskowicz 2004), the estimation of the scaling factor required that the user manually identify corresponding bone contours, an approach which we prefer to avoid. Additionally, these approaches assume no difference in translation between left and right limbs. A related approach that is often used in clinical practice is to place ad-hoc external markers for manual alignment of the sector images. These are usually ad-hoc markers that span the expected overlap between sector images and are placed on the patient (e.g. paper clips). Unfortunately, these markers often move between image acquisitions in an independent manner with respect to the patient anatomy, rendering them unreliable. Consequentially, an approach that only relies on anatomical structures is preferable and will better fit within clinical practice.

Approaches that do not require dedicated markers were described in (Wang et al. 2004) and (Lalys et al. 2018). The former focused on aligning the visible edges of the cassettes, a process which requires prior planning. This method can be applied only when the sub-images are exposed simultaneously and all the screen edges are completely visible in the acquired sub-images. The latter used standard image registration with a similarity metric combining gradient intensity and mutual information to obtain sufficiently robust results (81% success rate for automatic registration). This performance also required a minimal overlap between sector images of more than 20%.

Our goal in this work was to develop a fully automated method for panorama construction that does not require dedicated imaging or external markers and thus more readily adopted into clinical practice. Our approach relies on the detection of bone edges within the sector image overlap area, coupled with a one dimensional discrete search approach and a training-free registration technique to reliably align sector images to build panoramas. Moreover, the proposed method is fully automatic and does not need to deal with parallax issues as it directly aligns the anatomy of interest. As an overview, our method focuses on the extraction of continuous medial and lateral edges of the tibia and femur from the overlapping region, followed by a distance map-based (Maurer et al. 2003) implementation for edge alignment instead of a traditional optimisation-based registration approach. Moreover, while any of the traditional edge detection techniques, including Sobel, Prewitt, Roberts or Canny edge detection (Canny 1986), or even a monogenic filter (local-phase based boundary detection) (Rajpoot et al. 2009) may be used, we employed a horizontal second derivative filter to segment the bone shaft and its edges and retain only the lateral and medial bones edges.

We next describe the proposed approach in detail and its evaluation on a set of 95 highly variable patient X-ray image datasets acquired at 10 different clinical sites using different scanning protocols and energy levels. To assess the performance of our proposed technique, we compare the automatically generated panoramas to the ground truth panoramas manually generated by an expert X-ray technician. In addition, since these panoramas are used to measure the HKA angle and subsequently determine whether patients should be included or excluded in the clinical trial based on their extent of ankle deformity, we compared the HKA angles measured from the automated panoramas and those measured from the manual panoramas and used the estimated angles to establish patient eligibility for participation in the clinical trial. As such, we evaluate both the automated X-ray panorama generation and demonstrate its clinical utility vis-à-vis clinical standard of care.

## 2. Methodology

### 2.1 Imaging data

Our proposed technique was implemented and demonstrated on a population of long-limb X-ray datasets provided by Qmetrics Technologies, LLC, Pittsford, NY. The selected population features a wide variety of image resolutions, marker localisation, imaging artefacts, image quality and clinical conditions, hence providing a heterogeneous mix of scanner types and imaging parameters consistent with challenging clinical cases.

### 2.2 Methodology overview

In order to generate a panoramic image from multiple 2D sector images, a rigid registration is required to align the homologous image pairs. Since image rotation is deemed to only minimally affect the panoramas for patients in standing position, for the purpose of our application, this rotation is neglected. Therefore, only the horizontal (x-) and vertical (y-) translation pairs for each stitch was needed to construct the panoramic image.

Since the right and left limbs can also move as the patient slightly changes his/her leg opening between successive acquisitions of the image sectors, the left and right image stitches needed to be handled separately. Nevertheless, since the patient remains standing, then the left and right limb vertical translations (height) must be the same. Figure 2 demonstrates the stitching process which will be discussed next.

Each data set includes three individual X-ray images – the hip, knee and foot. The generated panoramas are constructed from the left and right image alignments of the upper limb, in which the lower 2/3 of the hip image is the fixed image and the upper 2/3 of the knee image is the moving image. Correspondingly, for the lower limb, the lower 2/3 of the knee image is used as the fixed image and the upper 2/3 of the foot image is the moving image, as shown in Figure 2b.

To construct the panorama image that optimally aligns the femur and the tibia, our method consists of four stages: 1) edge identification and extraction from all image sectors; 2) approximate segmentation of the bone shaft, followed by the identification and retention of the lateral and medial bones edges only based on the rough segmentation and previous edge

extraction; 3) discrete iterative registration of bone edges via DICE maximisation; and 4) panorama generation and HKA angle measurement.

## 2.3 Stage 1: edge detection and extraction technique

**2.3.1. Detection of the edges in the images**—To enhance the cortical bone edges, we filtered the images with a second derivative filter, a kernel consisting of three 20 mm H × 6 mm W intensity bands (−1, + 2, and −1) (Figure 3b).

As a result of the second derivative filtering, the intensity differences between dark-to-bright and bright-to-dark transitions were enhanced (Figure 3c). Based on the normalised cumulative histogram of the convolved images, we identified the threshold required to only maintain the bright regions of the filtered image, which include the bone edges. The optimal threshold was empirically chosen according to the mean intensity of the pixels within 65% – 95% of the highest intensities. Note that small variations of ± 5% around the threshold did not impact the results. An alternative approach to the second derivative filtering method is to use the Canny edge detection technique. While this approach may lead to the detection of finer, yet disconnected edges, depending on the standard deviation of the filter, alternatively, it could also yield continuous, but less accurate edges, attached to the surrounding edges. The resulting binary image shows the edge masks overlaid in red (Figure 3d) and the last image shows the acquired medial and lateral femur edges (Figure 3e).

## 2.4 Stage 2: identification and retention of tibial and femoral medial and lateral edges

**2.4.1. Bone mask initialization**—In order to enhance the cortical bone region, we filtered the images with a step function filter, a kernel consisting of three binary bands with values + 1, 0, + 1, each is 15 mm vertical and 1 mm horizontal (Figure 4b). Note that since the sum of the filter is not zero, it is not a derivative operation for edge detection. This filter rather enhances the brightness of the bright regions (Figure 4c). Similar to Stage 1, based on the normalised cumulative histogram of the convolved image, we identified the threshold required only to maintain the bright regions of the filtered image as shown in Figure 4d. Note that the mask might not include the bone marrow, because of its darker appearance, and might include flesh outside the bone, depending on the X-ray energy level and bone attenuation.

**Partial Segmentation of the Bone Shaft Region:** To include the bone marrow, we filled the gap in each row – wherever that gap exists and is larger than one cm, under the assumption that it is bone marrow. Additionally, based on the length of the shorter side of the cortical bone mask, we truncated the mask as shown in Figure 4e.

The approximately linear straight shape of the shaft region, which is the region of the mask at minimum width, assumed to be the femur width, is used to linearly extend the mask and remove the outside regions (Figure 4f). For the tibia, since the minimum mask width can be either the tibia, fibula or both, and is therefore inconsistent, we used the largest region with continuous boundaries to linearly extend and truncate the sides.

Figure 5 illustrates the summary of the first two stages – the overall rough segmentation of the femur and the tibia shafts, the extraction of the vertical edges, and the selection and

retention of the lateral and medial edges – and the final edges featuring constant width (Figure 5g,n).

In order to retain only the medial and lateral edges of the bone, we kept only the edges that overlapped with the shaft segmentation, then we retained only the edges that were closest to the medial and lateral boundaries of the shaft segmentation, as illustrated in Figure 5d–f,k–m. The medial edge was identified as the first edge detected when searching the image laterally, starting from the mid-line. The subsequent edges detected laterally to the medial edge were first interrogated in terms of their length and distance from the boundary of the segmented bone mask; as such, the lateral bone edge was identified as the longest edge lateral to the medial edge that was also close to the bone mask boundary. Additionally, where the edges were disconnected, we connected them based on the curvature direction of the femur and tibia in the given image. Specifically, we used the coordinates of the centre of each disconnected edge relative to the longest edge to connect the edge regions to match the generic curve on both sides of the bone and achieve a long and continuous medial and lateral edge masks for each bone in each of the fixed and moving images (Figure 5g,n). Similarly, in the lower leg, since the fibula is always lateral to the tibia, we only retained the medial and lateral tibia edges, which were identified as the first two continuous edges detected when searching the image starting from the mid-line towards the lateral direction. Lastly, although both the medial and lateral edges are detected, the alignment of the medial edges is both a necessary and sufficient condition to ensure bone alignment, and hence the alignment of the lateral edge in addition to the medial edge is implicit and neither necessary nor required to ensure overall bone alignment.

## 2.5 Stage 3A: registration and corresponding edge alignment

### 2.5.1. Identification of optimal left and right limb translation parameters—

After identifying the medial and lateral tibia edges from both the fixed and moving images, we slide the moving image along the fixed image by incrementally ‘sliding’ a landmark on the fixed image (the blue dots in Figure 5g,n) along the medial bone edge in the moving image.

As explained earlier, although both the medial and lateral edges are detected, the alignment of the medial bone edge is a sufficient condition to ensure bone alignment. In terms of practicality, for the femur, the medial edge is also the first long edge encountered when searching the image space laterally starting from the mid-line; similarly, for the tibia, the medial tibia edge is easier to identify than the lateral tibia edge, the latter of which might also be confused with the fibula edge. Therefore, the alignment of the medial edge for both the femur and the tibia is a sufficient condition to ensure overall bone alignment.

Using this method, each vertical (y) translation is accompanied by only one horizontal (x) translation of the fixed landmark along the medial edge. Therefore, the registration problem is defined as a one dimensional (1D) alignment of masks along the medial bone edge, as opposed to a 2D image registration, enabling us to solve a single rather than a two degree of freedom problem. Additionally, this approach enables us to avoid the use of standard optimisation techniques which are dependent on the selection of initial parameter values. Instead, we use a discrete search evaluating all possible solutions.

We evaluated an objective function for all vertical (y) translation values, along with their matching horizontal (x) translations values. When re-sampling the images at 1 mm resolution, there are only 200–300 points along the medial bone edge at which the similarity metric is evaluated within the potential overlap region.

The chosen objective function is based on the Euclidean distance between the moving and fixed edges for a given translation pair. The employed distances were computed using the distance map function ((Maurer et al. 2003)), defined as follows: given a subset of metric space  $\Omega \subset \mathbb{R}^2$  with metric,  $d$ ,  $\partial\Omega$  denotes the boundary of  $\Omega$  and  $x \in \mathbb{R}^2$  then the distance function,  $f(x)$ , is defined by:

$$f(x) = \begin{cases} d(x, \partial\Omega) & \text{if } x \in \mathbb{R}^2 \setminus \partial\Omega \\ 0 & \text{if } x \in \partial\Omega \end{cases} \quad (1)$$

$$\text{when } d(x, \partial\Omega) = \inf_{y \in \partial\Omega} d(x, y),$$

where  $\inf_{y \in \Omega} d(x, y)$  denotes the *infimum* of the Euclidean distance between  $(x, y)$ . Using distance function, we implemented an efficient computation of the distances between the edge points in the fixed image to the closest edge points in the moving image and vice-versa. The translation pairs along the medial right and left edges that yield the minimum of the symmetric average distances correspond to the correct alignment. The objective distance function computes the average distance between the fixed and moving edges per image line:

$$Transferred = T[\partial\Omega_{Moving}] \quad (2)$$

$$Dist_{fromTransferred2Fix} = (\sum \sum f(Fix) \odot Transferred) / Nr,$$

$$Dist_{fromFix2Transferred} = (\sum \sum f(Transferred) \odot Fix) / Nr,$$

$$E_D = Dist_{fromTransferred2Fix} + Dist_{fromFix2Transferred}$$

where  $Nr$  is the number of rows,  $T[\partial\Omega_{Moving}]$  is the translated moving edge image, and  $\odot$  denotes the Hadamard Product (an element-wise product of matrices).

Additionally, we also tested an alternative similarity metric that is based on the DICE coefficient between the moving and fixed medial edges for a given translation pair. The employed DICE coefficient is defined by:

$$DICE = \frac{2(A \cap B)}{A + B}, \quad (3)$$

where  $A$  and  $B$  are the areas of the edges in the fixed and the moving images, respectively. The areas used are only within the overlap region of the two sector images after applying the given translation pair. Using the DICE coefficient, we calculated the similarity function, which is the sum of the DICE coefficients calculated along the lower and upper parts of the overlap region. Iterative alignment is achieved by sliding the medial bone edge from the moving image along the medial bone edge from the fixed image. The use of the DICE coefficient is based on the overlap and non-overlap areas (i.e., a binary measure), whereas the distance metric provides a distance weight. Thus, the distance metric was more robust than the DICE coefficient, especially when a slight rotation causes large non-overlap areas.

## 2.6 Stage 3B: selection of the optimal left and right vertical translation parameter

As mentioned previously, the left and right vertical translations for the upper and lower registrations need to be the same when reconstructing the standing panorama, as the height of the patient estimated from the left or the right limb must stay constant. In the event of different optimal translation parameters resulting from the left and right limb registration, we evaluate the distance function using the right and left limbs together as one image (Figure 6b), while assuming that the optimal overall left and right similarity (i.e., minimum distance) would provide the correct alignment for both.

Nevertheless, we also tried to use the intensities around the edges by comparing the Normalised Mutual Information (NMI) ((Studholme et al. 1999)) around the bone edges. However, since there are non-homologous markers placed on the patient and appearing near the bone edge, the NMI measure in some cases was low even when the bones were aligned. Figure 6a illustrates the evaluation of the objective function for the femur registration. Here, the left and right minima (marked with yellow X) were very close (i.e., 121 and 122 mm from the top of the fixed image, respectively), and the chosen one was 122 mm vertical translation marked with purple \*.

In order to compare the differences after applying the optimal translation pair to the moving edge image and to the moving original intensity image, we overlaid the translated images (magenta) onto their corresponding fixed images (green) with 50% transparency Figure 6b–c, respectively. Lastly, Figure 6d shows the upper and lower sections of the fix edges distance map (green) together with the translated moving edges, overlaid in magenta. Note that the distance map values at the pixels located at the translated edges are the distances from the fixed edges to the closest points in the translated edges.

## 2.7 Stage 4: automated panorama generation and HKA angle measurement

**2.7.1. Automated panorama generation**—To generate the full long limb panorama, the four identified translation pairs are applied to the right and left femur and tibia images, as shown in Figure 2. As a result, the upper knee sector image of the left and right limb was ‘stitched’ to the corresponding lower hip image, and similarly the corresponding upper foot

image was ‘stitched’ to the recently generated lower hip and upper knee partial panorama, in the end yielding a full panorama depicting both limbs in standing position

**2.7.2. HKA angle measurement**—Since the main purpose of reconstructing the panoramas is to measure the HKA angle and compare the angles estimated using the automated vs. manually constructed panoramas, we followed the standard procedure outlined in (Cooke et al. 2007) to select the required points for establishing the mechanical axes of the tibia and femur and calculated the axis angle as shown in Figure 7.

Following the development and implementation, we evaluated the method by comparing the vertical (y-) and horizontal (x-) translation parameters yielded by the automated panorama generation method to the translation values chosen by the expert technician during manual panorama construction. In addition, while this evaluation focuses on assessing the accuracy of the image processing technique, it does not speak to the clinical relevance or benefit provided by the proposed method. To this extent, we conducted an additional study in which we estimated the hip-knee-axis angle from all automated panoramas and compared them to the corresponding angles estimated from the manually generated panoramas. Lastly, we used the clinical trial patient eligibility criteria to compare the eligibility established according to both the manual and automated panoramas.

### 3. Results

#### 3.1 Panorama reconstruction evaluation

We first evaluated the accuracy of the image processing pipeline by comparing the horizontal and vertical translation parameters yielded by our proposed panorama generation method to the corresponding translation parameters chosen by the expert technician during the manual panorama creation for each of the 95 datasets.

Table 1 summarises the mean, standard deviation and range of these horizontal (x-) and vertical (y-) translation differences between the manual and the automated panoramas. Note that the horizontal differences are much smaller than the vertical differences since the horizontal translations are computed in response to the unique vertical translations along the medial femoral and tibial edges.

Figure 8 illustrates the differences between the horizontal and vertical translation parameters corresponding to the manual panoramas generated by the expert technician and the automated panoramas across all 95 patient datasets. The Bland-Altman plots show the mean  $\pm 2$  standard deviation of the differences between the automated and the expert manual translation parameters (horizontal (x), i.e. medial/lateral, and vertical (y), i.e. superior/inferior), for the right and left femur and tibia. The patients with high differences, larger than 2 standard deviations, were marked in red and two of them are further analysed in the discussion. We choose these two cases to show the causes responsible for the largest differences, either due to manual error of the expert or due to errors with our automated method associated with very low variations of the bone width in the shaft region of the femur.

### 3.2 Hip-knee-axis angle assessment and clinical trial eligibility assessment

In addition, since the panoramas are used to assess the HKA axis, we also conducted a study in which we compared the axis angle estimated based on the automated and manual panoramas, as well as asses patient eligibility for participation in the clinical trial according to the inclusion/exclusion criteria quantified based on the automated and manual panoramas.

Figure 9 illustrates the distributions, using box plots, of the resulting vertical translations and the HKA angles estimated using the automated and the manual panoramas. For the statistical comparison between these distributions, we inspected the histogram shapes to verify normal distributions and preformed pairwise t-tests. The two-sample t-test with 5% significance level showed no significant difference between the automated and the manual vertical translations ( $p = 0.15$  for the femoral stitch,  $p = 0.53$  for the tibial stitch, and  $p = 0.38$  for the overall femoral and tibial stitches together).

A similar statistical analysis conducted for the estimated HKA angles computed using the automated and manual panoramas which showed no statistically significant difference between the two datasets. Specifically,  $p = 0.99$  for the right limb,  $p = 0.65$  for the left limb, and  $p = 0.77$  for the overall (right and left) HKA comparison.

Furthermore to assess intra-user variability, we repeatedly selected the points defining the mechanical axes 6 times, in random order on both the automated and manual panoramas of five patients. The maximum variance of the 4 HKA angles (estimated for the left and right limb using the manual and automated panoramas) of the five patients was  $0.18^\circ$ . In addition, for a single patient dataset, the left HKA angle from the automated panorama had a maximum variability of  $1.18^\circ$ .

Figure 10a shows the HKA angle differences (in degrees) across 95 patient datasets for both right (blue diamond) and left (orange circle) limbs and the mean  $\pm 2$  standard deviation. The differences are between the HKA angle calculated from the selected points on the automated panorama and the HKA angle calculated from the selected points on the manual panorama generated by the expert. As mentioned previously, studies use varus/valgus malalignment as the threshold for excluding patients from the clinical trial. We divided the patients according to the absolute value of the HKA angle – less than  $7^\circ$  (included), greater than  $10^\circ$  (excluded) and in between  $7^\circ$ – $10^\circ$  (possibly included). Figure 10b shows the resulting absolute HKA angle and the common thresholds for excluding patients from clinical trial. Moreover, it also shows two cases with different clinical decisions, depending on which HKA angle is used – the angle computed from the automated panorama image or the manual panorama image stitched by the expert. Note that the clinical decisions for cases 25, 52 and 56, according to the automated or manual HKA angles are the same, although according to the vertical translation differences, they are considered as outliers.

Similarly, Table 2 reports the number of included/possible/excluded patients based on the automated vs. the manual HKA angles. The two cases that would result in different clinical decisions, Case 66 and Case 71, being both excluded based on the manually calculated angle, are analysed in the discussion. Note that a patient with an incorrect eligibility decision

might interrupt the results of the clinical trial or cause additional expenses to the clinical trial, as additional patients would need to be enrolled.

Figure 11a shows for each patient the HKA angle difference (blue left y axis and blue diamond) and the total (femoral plus tibial) vertical translation difference (orange right y axis and orange circle), while Figure 11b shows the HKA angle difference (y axis) as a function of the total vertical translation difference (x axis). These figures show that even for a large translation difference of 90 mm, the HKA angle difference was only 1°. Case 66 exhibited the maximum vertical translation difference between the automated and the manual panoramas of 98 mm, which, in turn, resulted in a HKA angle difference of 2.35°. This particular case featured the large difference between the manual and automated panorama translation parameters due to an error associated with the femoral stitch of the manually reconstructed panorama, as shown in Figure 12. All other cases featured HKA angle differences of less than 1.25° for vertical translation parameter differences of up to 88 mm.

The performance evaluation was conducted using non-optimised MATLAB code running on Windows 7 on a 64 bit PC, 3.6 GHz Intel CPU with 32 GB RAM. When the images were re-sampled into 1 mm isotropic spacing, the average processing time for the edge extraction was 6 s and for the optimal translation identification was another 6 s. Lastly, the panorama generation required 1 s, leading to an overall automated panorama reconstruction panorama in less than 15 s from start to finish. Because of our reduction of the ‘stitching problem’ into a single degree of freedom alignment along the medial edge, we avoided the use of an optimisation solver and consistently reached a global minimum convergence.

#### 4. Discussion

The task of creating an FLWB panoramic image is at its core a registration task. Our initial attempts at addressing this task explored the use of standard intensity based registration algorithms. Our first approach was described in (Miller et al. 2016). In that approach, we started by performing image de-noising using a Wiener filter followed by image enhancement with the Sobel filter. The resulting gradient magnitude images were used as input for rigid intensity based registration with the sum-of-squared differences as the similarity metric. This method performed reliably on less than 50% of the data. We then evaluated another intensity-based registration approach implemented via two registration steps, without image preprocessing, using the SimpleITK registration framework (Yaniv et al. 2017). The first step consisted of an initial registration by an exhaustive search of the maximum normalised cross correlation (NCC) on a coarse grid within the potential overlap region. The second step entailed the final registration using a multi-resolution pyramid, NCC with 20% sampling rate, and gradient descent optimizer. For both steps, we enforced the constraint that both the left and right limb vertical translations be equal. While this approach yielded acceptable results in terms of bone alignment, it performed reliably only 30% of the time (30/95 datasets). Based on these results, we did not pursue intensity based registration any further.

In this work, we presented a pipeline that enables us to detect sufficiently reliable bone edges that are subsequently used as features for sector image alignment. We first detect and enhance all intensity transitions in the images by using a vertical second derivative filter. We then identify the bone edges according to the partial segmentation of the shaft region. We identify the optimal sector image alignment by minimising the distance between the bone edges. This metric was evaluated for all iterative vertical translations along the medial edge of the bone. This approach allowed us to avoid the use of standard optimisation techniques that may converge to a local minimum when minimising the similarity metric. Lastly, we assessed the HKA angle for both the right and left limbs using both the automated and manually constructed panoramas.

In our prior work ((Ben-Zikri et al. 2017)) to keep only the cortical bone edges in the fixed and moving images, we measured the maximum intensity ( $I_{max}$ ) and the difference between the maximum and minimum ( $I_{min}$ ) intensities within a window of each edge. The chosen window has the same width as the edge mask and a height of 5–10 mm and consists of the edges in the lower half of the fixed image and the edges in the upper half of the moving image. We normalised the image intensities to a 0–1 range and only retained the edges that satisfied the relationship  $2I_{max} - I_{min} > 1:2$ . This selection method aimed to ensure that only the high intensities features corresponding to the cortical edges were featured in the window. The edges that satisfied this condition were overlaid on the fixed and moving images. However, this approach proved to not be as reliable as the shaft segmentation, since it depends highly on the chosen window, and hence yielded a high rate of false positives.

We summarise the results by first evaluating the accuracy of the automated panorama reconstruction method by comparing the yielded translation values to those selected by the expert technician during manual stitching. In addition, we also assessed and compared the HKA angle computed from both the automated and manual panoramas. Figure 9–10 show the distribution and the differences of the resulting translations and HKA angles between the manual stitching of the expert and the automated stitching.

For the femoral stitch, where the shape of the femur body is not changing width, the maximum vertical translation differences are higher (Table 1). Specifically, patient 66 showcased the largest vertical translation difference of  $-97.75$  mm. Figure 12a,b shows the two panorama images – automated and expert manual – side by side and illustrates an example where the manually constructed panorama is wrong, as the femur is shorter than the tibia.

We also discuss an additional case in detail – patient 71 featuring a femoral vertical translation difference of  $-61.62$  mm. As illustrated in Figure 13a,b, we can see that the femur length is too long in the automated panorama. Figure 13c,d show good alignment for the automated translations and Figure 13e,f show good alignment for the expert translations. Since the femur width is not changing much in this region, the average distance metric for both the automated method and expert are very close (Figure 13).

Unlike Patient 71, for Patient 66 the distance metric values for the expert translations were very high and therefore different from the values from the automated translations, which

suggested an error associated with the manual panorama generation, while for Patient 71, it suggests an error in the automated panorama reconstruction.

#### 4.1 Study limitations

The panorama reconstruction method relies on the identification of the optimal translation between the hip, knee and ankle image sectors based on the anatomical features associated with the bone in the overlap regions. Although the leg opening (i.e., distance between the patient's feet) may not be similar during the acquisition of all three image sectors, our method does not take into account any rotations for image sector alignment. Although perceived that leg opening may be better captured via rotation, the currently employed clinical approach for panorama reconstruction relies solely on manual translation of the image sectors. Since our method performed as well as the current clinical method, we don't believe the panorama reconstruction algorithm needs to incorporate rotation effects.

Also to account for potential changes in the leg opening in the acquisition of the subsequent image sectors, rather than handling the sector images simultaneously for the left and right limb, we split each sector in the left and right limb and essentially identified the optimal translational parameters for each leg independently, hence providing two 'recipes' for panorama reconstruction – one based on the left and one based on the right leg. Nevertheless, the overall height of the patient and hence the height of the standing panorama should be the same regardless of which side it is estimated. In the event of different optimal translation parameters resulting from the left and right limb registration, we evaluate the distance function using the right and left limbs together as one image, while assuming that the optimal overall left and right similarity (i.e., minimum distance) would provide the correct alignment for both.

We conducted this study and demonstrated and evaluated our method using a database of 95 patient image datasets, each consisting of three sector images and acquired across more than 20 clinical sites using different imaging protocols and scanners. Successful panorama generation typically requires a certain amount of overlap between the sector images. For example, the recent work cited earlier (Lalys et al. 2018) required a minimum of 20% overlap between the image sectors, as validated on 30 patient datasets. Similarly, most intensity-based registration techniques perform optimally given 40% or greater overlap. Nevertheless, our method does not require any specific overlap between the image sectors, as confirmed by the wide range of overlap features in the 95 patient datasets used for validation. The overlap range for the femur ranged from 3 to 43%, with a mean overlap of 21%, and for the tibia the overlap ranged from 10 to 58% with a mean overlap of 34%. Moreover, more than 35% of the datasets featured less than 20% overlap between the image sectors, and our proposed method still performed robustly despite the very limited overlap.

Lastly, in this study, we used the manually reconstructed panoramas as ground-truth against which we assessed the panoramas reconstructed automatically using our method. However, not all manual panoramas were correctly reconstructed due to the user's inability to determine the optimal overlap. As a result, for a select few of the datasets, the difference between the parameters associated with the automated vs. manual translation were slightly large, rendering our method as sub-optimal, when in fact a large error contribution is caused

by the error in the ground truth. Furthermore, the manual panoramas used as ground truth in this study were not all reconstructed by a single technician, but rather by a single technician at each site. The inter-observer variability for different technicians is not known. An improved ground truth can potentially be obtained by having multiple technicians manually align the images. This would provide some knowledge with respect to the reproducibility of the manual panorama reconstruction techniques as it pertains to different technicians. Unfortunately, the manual alignment of the image sectors and panorama reconstruction by multiple technicians is currently not part of the clinical protocol at any of the sites – each manual panorama was reconstructed manually by only one technician. Thus, we can only say that the algorithm proposed in this work performs at a level that is close to that of a typical technician, but not a specific technician or the same technician, either at one site or across all sites. To mitigate this limitation, a potential avenue for future work could consist of having multiple technicians reconstruct all manual panoramas. This approach will help provide a measure of the variability associated with the manual panoramas that serve as gold standard to assess the automated panorama reconstruction method described and validated here.

## 5. Summary and conclusion

In this paper, we developed and evaluated a fully automated longlimb X-ray image stitching procedure for generating panoramas utilised to quantify the HKA axis deformity. This method is unique as it is based only on anatomical features and bone shape and the sector image registration does not require optimisation. The proposed method employs the medial and the lateral edges of the femur and the tibia captured in both the fixed and moving images and aligns the right and left sectors by minimising the Euclidean distance between the homologous bone edges. Moreover, by focusing on the registration of the bone edges, the achieved alignment is independent of the position of any markers, which may aid in the process, or mislead the process due to changes in their position during subsequent sector image acquisition.

We have implemented and evaluated our method using 95 patient datasets acquired as part of a multi-site clinical trial for patient screening. The images posed many of the traditional clinical challenges as a result of different acquisition protocols and imaging scanners. According to our results, the horizontal and vertical translation parameters recorded from the manually assembled panoramas were not significantly different from the automated registration parameters. There were also no significant visual differences between the manually and automatically generated panoramas images, except for five cases where the femoral vertical differences between the manual and the automated were greater than 40 mm. These cases featuring large differences are similar to the two cases in the discussion, with the large differences being primarily caused by the femur region featuring very small shaft width variations, human error, and possibly incomplete edges. Additionally, there were no significant differences between the HKA angles estimated from the manually and automatically generated panorama images or between patient eligibility for participation in the clinical established based on the measured HKA angles.

Lastly, in terms of performance, the average processing time for the edge extraction was 6 s and for the optimal translation identification was another 6 s. Panorama generation required 1 s, leading to an overall automated panorama reconstruction panorama in less than 15 s from start to finish, which is faster than the manual alignment workflow. In light of this study, as well as the simplicity of the method and its robustness across the highly variable testing dataset, our results are encouraging, rendering this method as suitable for the intended application.

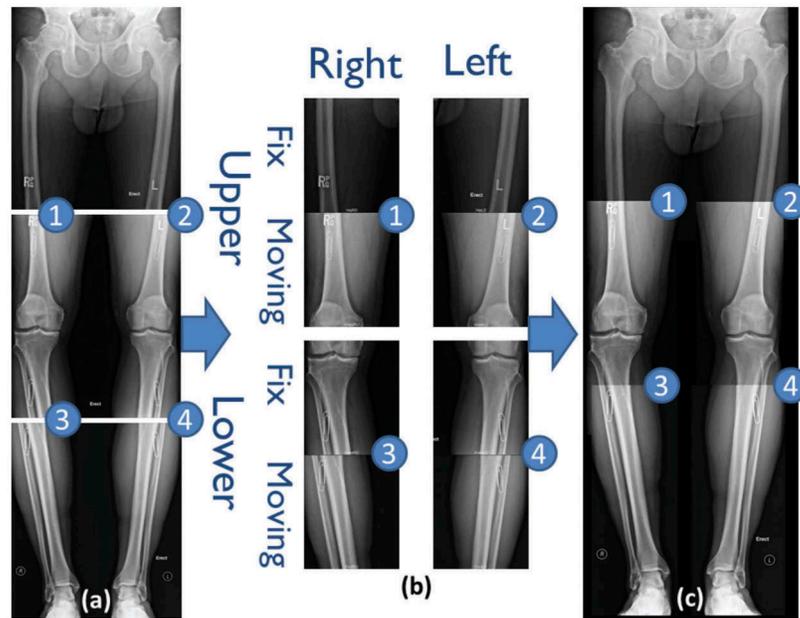
## Acknowledgments

The authors acknowledge Qmetrics Technologies for providing the X-ray images used in this manuscript. Also, Z. Yaniv was supported by the Intramural Research Program of the U.S. National Institutes of Health, National Library of Medicine. In addition, research reported in this publication was supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award No. R35GM128877.

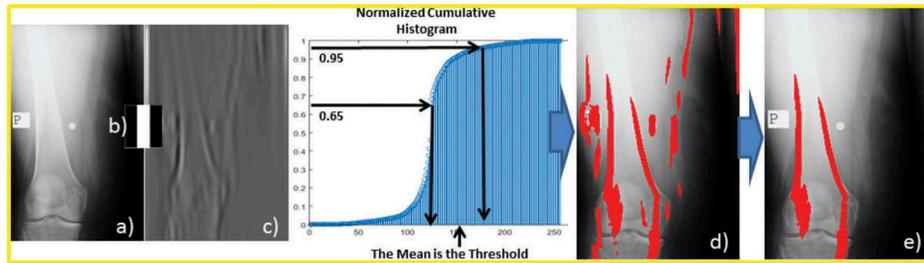
## References

- Apivaththakakul T, Duanghakrungs M, Luevitoonvechkit S, Patumasutra S. 2013 Intraoperative panoramic image using alignment grid, is it accurate? *Arch Orthop Trauma Surg.* 133(7):953–959. [PubMed: 23636319]
- Ben-Zikri YK, Mendez S, Linte CA. 2017 Anatomical-based registration of multi-sector x-ray images for panorama reconstruction. *Proc SPIE Medical Imaging: Biomedical Applications Molecular, Functional Structural Images.* 10137:101370H–101370H–11.
- Canny J 1986 A computational approach to edge detection. *IEEE Trans Pattern Anal Mach Intell.* 6:679–698.
- Caresream. 2012 White paper: advanced image processing with automatic stitching and manual adjustment capabilities produce a composite long-length image without visible seam-line artifacts CARESTREAM directview DR long-length imaging system with automatic and manual stitching. Carestream Health Inc. Report no.:
- Chen C, Kojcev R, Haschtmann D, Fe-Peterkete T, Nolte L, Zheng G. 2015 Ruler based automatic C-Arm image stitching without overlapping constraint. *J Digital Imaging.* 28(4):474–480.
- Cooke T, Li J, Scudamore RA. 1994 Radiographic assessment of bony contributions to knee deformity. *Orthop Clin North Am.* 25(3):387–393. [PubMed: 8028882]
- Cooke T, Scudamore R, Bryant J, Sorbie C, Siu D, Fisher B. 1991 A quantitative approach to radiography of the lower limb. principles and applications. *Bone Jt J.* 73(5):715–720.
- Cooke TDV, Scudamore A. 2003 Healthy knee alignment and mechanics The adult knee Philadelphia. Lippincott Williams & Wilkins; Philadelphia; p. 175–186.
- Cooke TDV, Sled EA, Scudamore RA. 2007 Frontal plane knee alignment: a call for standardized measurement. *J Rheumatol.* 34(9):1796–1801. [PubMed: 17787049]
- Felson DT, Niu J, Gross KD, Englund M, Sharma L, Cooke TDV, Guermazi A, Roemer FW, Segal N, Goggins JM, et al. 2013 Valgus malalignment is a risk factor for lateral knee osteoarthritis incidence and progression: findings from the multicenter osteoarthritis study and the osteoarthritis initiative. *Arthritis Rheum.* 65(2):355–362. [PubMed: 23203672]
- Gooßen A, Schlüter M, Pralow T, Grigat RR 2008 A stitching algorithm for automatic registration of digital radiographs *Proc. International Conference on Image Analysis and Recognition.* Póvoa de Varzim, Portugal p. 854–862.
- Jordan J, Sowers M, Messier S, Bradley J, Arangio G, Katz J, Losina E, Rovati L, Bachtell N, Cooper C, et al. 2011 Methodologic issues in clinical trials for prevention or risk reduction in osteoarthritis. *Osteoarthritis and Cartilage.* 19(5):500–508. [PubMed: 21396470]
- Lalys F, Favre K, Villena A, Durrmann V, Colleaux M, Lucas A, Kaladji A. 2018 A hybrid image fusion system for endovascular interventions of peripheral artery disease. *Int J Comput Assist Radiol Surg.* 13 (7):997–1007. [PubMed: 29549554]

- Maurer CR, Qi R, Raghavan V. 2003 A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Trans Pattern Anal Mach Intell.* 25(2): 265–270.
- Messmer P, Matthews F, Wullschleger C, Hugli R, Regazzoni P, Jacob AL. 2006 Image fusion for intraoperative control of axis in long bone fracture treatment. *Eur J Trauma.* 32(6):555–561.
- Miller L, Trier C, Ben-Zikri YK, Linte CA 2016 Automated reconstruction of standing posture panoramas from multi-sector long limb x-ray images *Proc. SPIE Medical Imaging: Physiology, Function and Structure from Medical Images.* San Diego, CA p. 97880B–01–97880B–08.
- Rajpoot K, Grau V, Noble JA 2009 Local-phase based 3D boundary detection using monogenic signal and its application to real-time 3-D echocardiography images *IEEE International Symposium on Biomedical Imaging: From Nano to Macro, (ISBI'09).* Boston, MA IEEE p. 783–786.
- Skyttä ET, Haapamäki V, Koivikko M, Huhtala H, Remes V. 2011 Reliability of the hip-to-ankle radiograph in determining the knee and implant alignment after total knee arthroplasty. *Acta Orthop Belg.* 77(3):329–335. [PubMed: 21846000]
- Sled EA, Sheehy LM, Felson DT, Costigan PA, Lam M, Cooke TDV. 2011 Reliability of lower limb alignment measures using an established landmark-based method with a customized computer software program. *Rheumatol Int.* 31(1):71–77. [PubMed: 19882339]
- Studholme C, Hill DL, Hawkes DJ. 1999 An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognit.* 32(1):71–86.
- Toshihiro Minami MY, Okamoto T 2014 Using the RADspeed Pro V4 general radiography system with FPD. Shimadzu Corporation Report no.:
- Wang L, Traub J, Weidert S, Heining SM, Euler E, Navab N. 2010 Parallax-free intra-operative x-ray image stitching. *Med Image Anal.* 14:674–686. [PubMed: 20561811]
- Wang X, Foos DH, Doran J, Rogers MK 2004 Fully automatic and reference-marker-free image stitching method for full-spine and full-leg imaging with computed radiography *Proc. SPIE Medical Imaging: Physics of Medical Imaging.* San Diego, CA p. 361–369.
- Yaniv Z, Joskowicz L. 2004 Long bone panoramas from fluoroscopic X-ray images. *IEEE Trans Med Imaging.* 23(1):26–35. [PubMed: 14719684]
- Yaniv Z, Lowekamp BC, Johnson HJ, Beare R. 2017 Simple ITK image-analysis notebooks: A collaborative environment for education and reproducible research. *Journal Digital Imaging.* 31(3): 290–303.
- Yoshioka Y, Siu D, Cooke T. 1987 The anatomy and functional axes of the femur. *J Bone Joint Surg Am.* 69(6):873–880. [PubMed: 3597501]
- Yoshioka Y, Siu DW, Scudamore RA, Cooke TDV. 1989 Tibial anatomy and functional axes. *J Orthopaedic Res.* 7(1):132–137.
- Zampogna B, Vasta S, Amendola A, Uribe-echevarria Marbach B, Gao Y, Papalia R, Denaro V. 2015 Assessing lower limb alignment: comparison of standard knee xray vs long leg view. *Iowa Orthop J.* 35:49–54. [PubMed: 26361444]

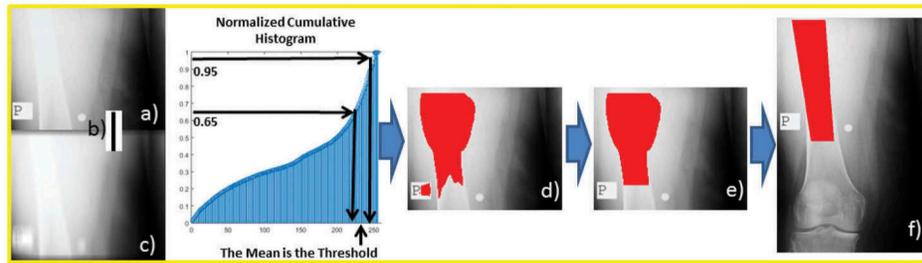


**Figure 1.** Illustration of the image stitching procedure: the three input images hip, knee and foot for one representative patient (a), the cropped regions from these images needed to construct the four subsequently registered image pairs (b) and the final long limb panorama image (c).



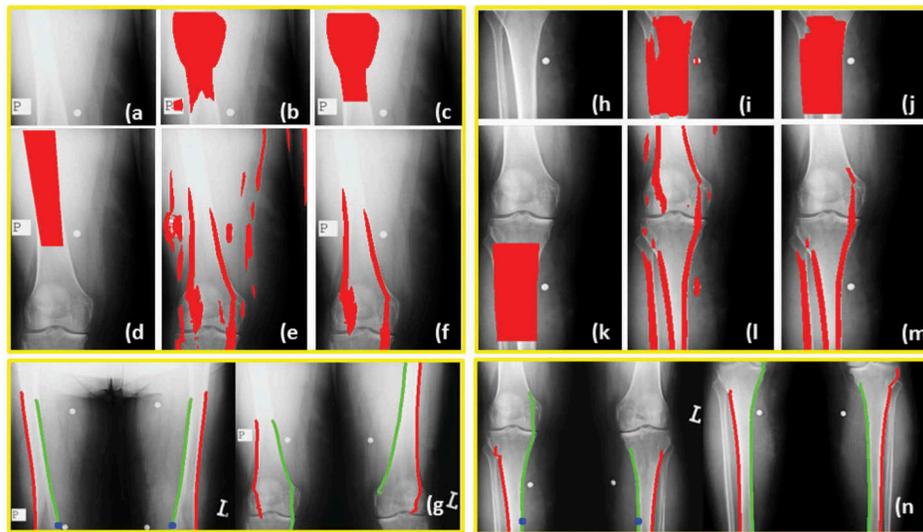
**Figure 2.**

a) right moving femur image (a); b) horizontal second derivative filter; c) output of image convolution with the aforementioned filter. The resulting image was thresholded using an intensity value selected according to the cumulative histogram, based on the mean intensity of all pixels within 65% – 95% of the highest intensity. The obtained edges are overlaid on the moving femur image shown in red (d). Sector image showing the retained/desired lateral and medial edges (e), obtained by partial segmentation of the bone shaft region, discussed in Stage 2.



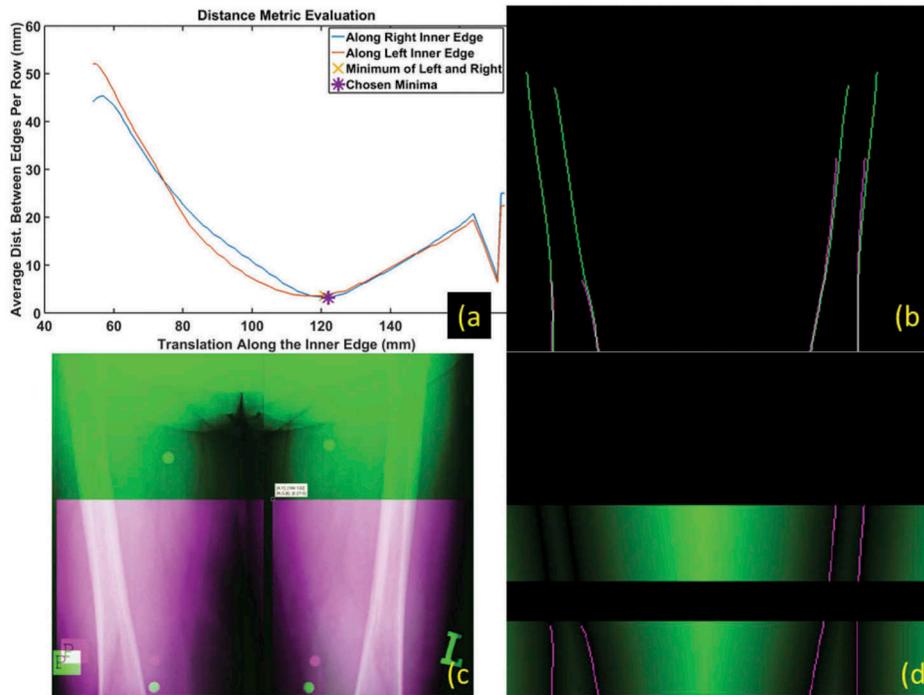
**Figure 3.**

Illustration of the partial segmentation of the bone shaft region: a) the region of interest (upper 2/3 of the upper-knee sector image); b) the vertical second derivative kernel; (c) resulting image with enhanced bright regions; (d) resulting bone mask overlaid onto the original image after thresholding, based on the mean intensity of the pixels within the 65% – 95% of the highest intensities; (e) filled-in and truncated bone mask and (f) final bone shaft region obtained by linearly extending the minimum width mask.



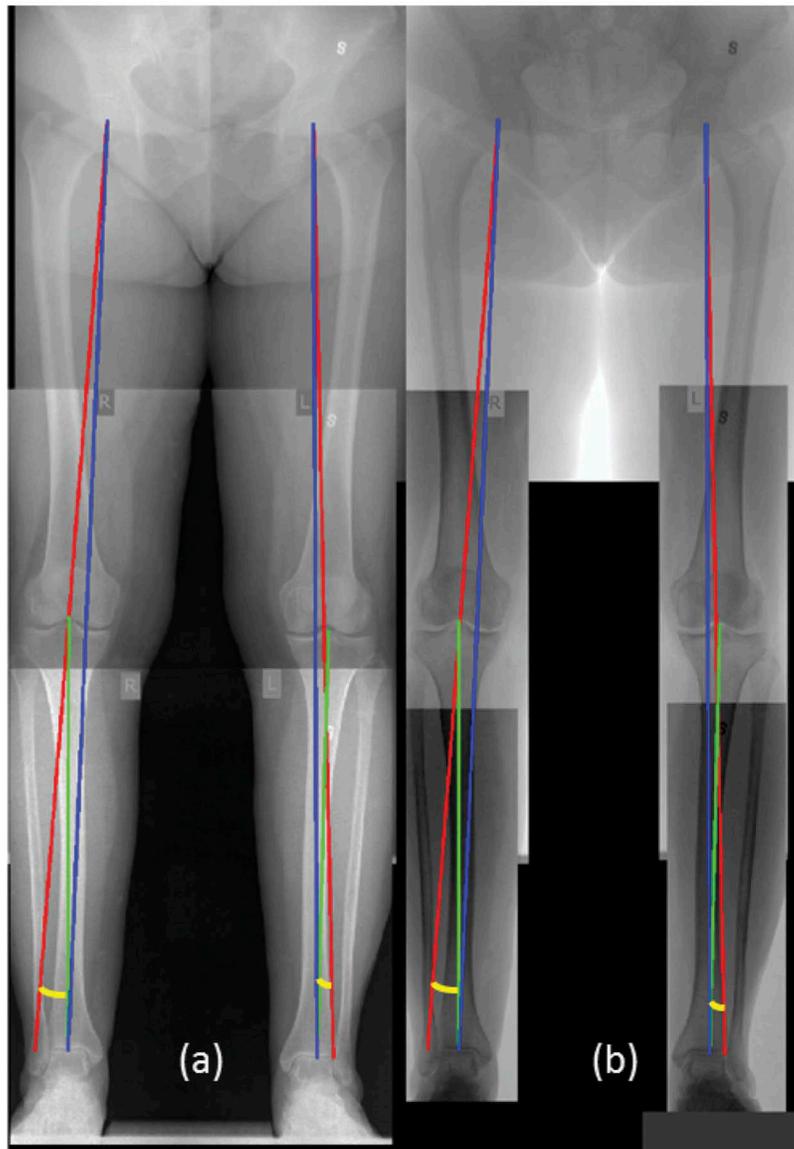
**Figure 4.**

Illustration of the overall rough segmentation of the femur and the tibia shafts, the vertical edges and the selection of the lateral and medial edges. Starting with the region of interest from the original right femur moving image (a); the outputs of processing steps are shown: the initial femur mask obtained by filtering and the thresholding (b); the filled-in and truncated bone mask (c); the linearly extended mask to exclude the non-bone regions (d); the edge masks obtained via the second derivative horizontal filter followed by thresholding (e); and the edge masks overlapping with the shaft segmentation (f). Image (g) shows the medial (green) and the lateral (red) edge masks of the right and left-sided, fixed and moving femurs. Note that the edges feature constant widths, enabling the use of a single horizontal translation increment for each vertical translation increment. The selected blue marker located in the lower region of the medial edge in the fixed image is used as an encoder for the vertical translation increments. The corresponding images (h)-(n) illustrate the same steps as shown in (a) – (g), respectively, however applied to the tibia from the lower-knee and upper-foot sector images.

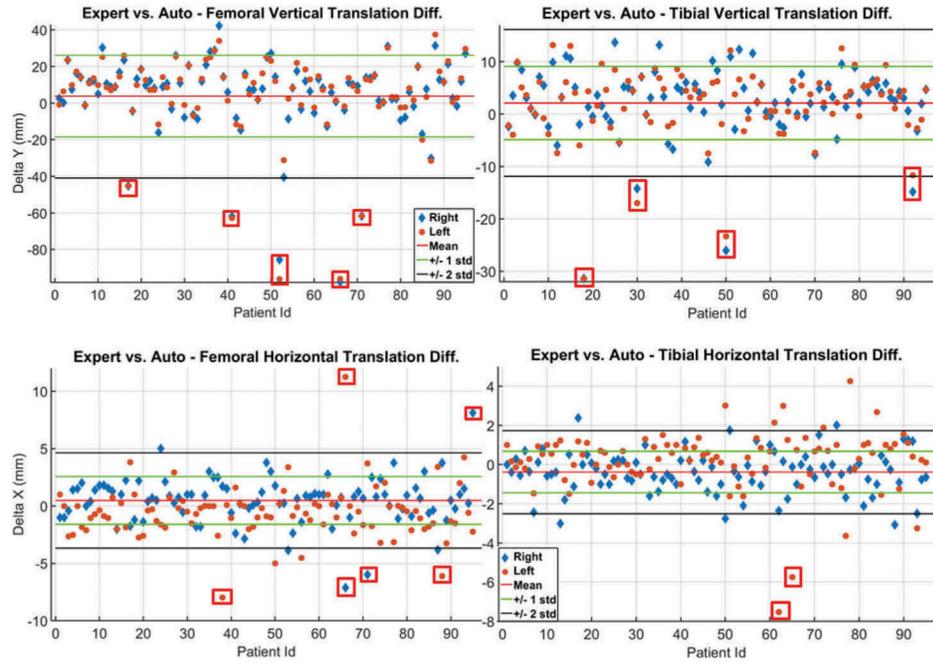


**Figure 5.**

Illustration of the right and left limb upper sector image registration: (a) the distance metric evaluation, the graphs show the average distance metric (y axis in mm) of the right (blue) and left (orange) limbs for each translation pair (x axis in mm), as the moving image slides along the medial bone edge. Note that in this example, the left and right minima (marked with a 'X') were very close (121 and 122 mm vertical translation from the top of the fixed image), and the chosen translation value was 122 mm, as marked with purple \*). In order to compare the differences after the chosen translation pair is applied to the moving edge images and to the moving original intensity images, the fixed (green) and the translated (magenta) images are overlaid with 50% transparency, as shown in (b) and (c), respectively. Finally, the distance map (green) of the fixed edges is shown in the upper and lower sections of the overlap region, together with the translated moving edges in magenta (d).

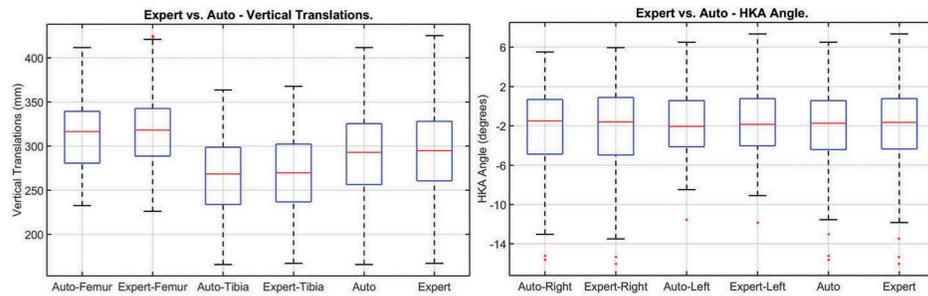


**Figure 6.** The calculated HKA angle based on the automated panorama (a) and the manual panorama (b).



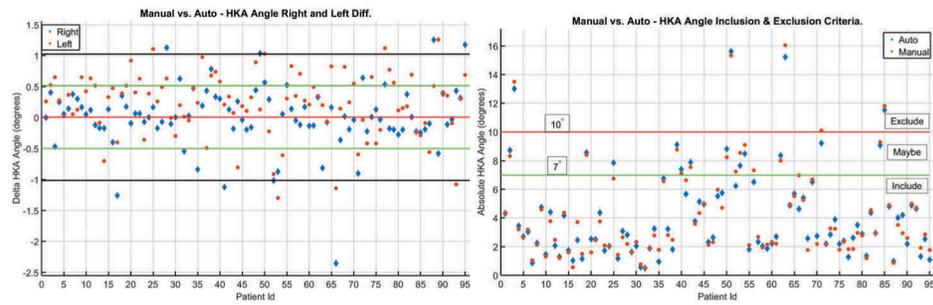
**Figure 7.**

Bland-Altman plots showing the mean  $\pm$  2 standard deviation of the vertical (delta-y) and horizontal (delta-x) differences (in mm) between the automated and the manually generated panoramas across 95 patient datasets for both right (blue diamond) and left (orange circle) limbs. The first column corresponds to the lower-hip and upper-knee sector images stitched based on the femoral edges, while the second column corresponds to the lower-knee and upper-foot sector images stitched based on the tibial edges. Note that patients exhibiting high translation differences, specifically differences larger than two standard deviations, are marked in red.



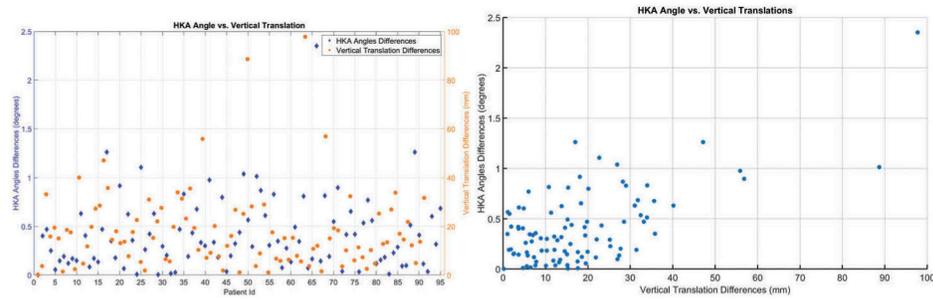
**Figure 8.**

Visual representation of the vertical translation (left) and the HKA angle (right) distributions using box plots, showing the medial, 25<sup>th</sup> and 75<sup>th</sup> percentiles, as well as the outliers. From left to right, the vertical translation plots correspond to the automated femoral stitch, manual femoral stitch, automated tibial stitch, manual tibial stitch, automated (femoral + tibial) stitch, and manual (femoral and tibial) stitch. Similarly, from left to right, the HKA angle plots correspond to the automated right limb HKA angle, manual right limb HKA angle, automated left limb angle, manual left limb angle, automated (left and right limb) angle, and manual (left and right limb) angle.



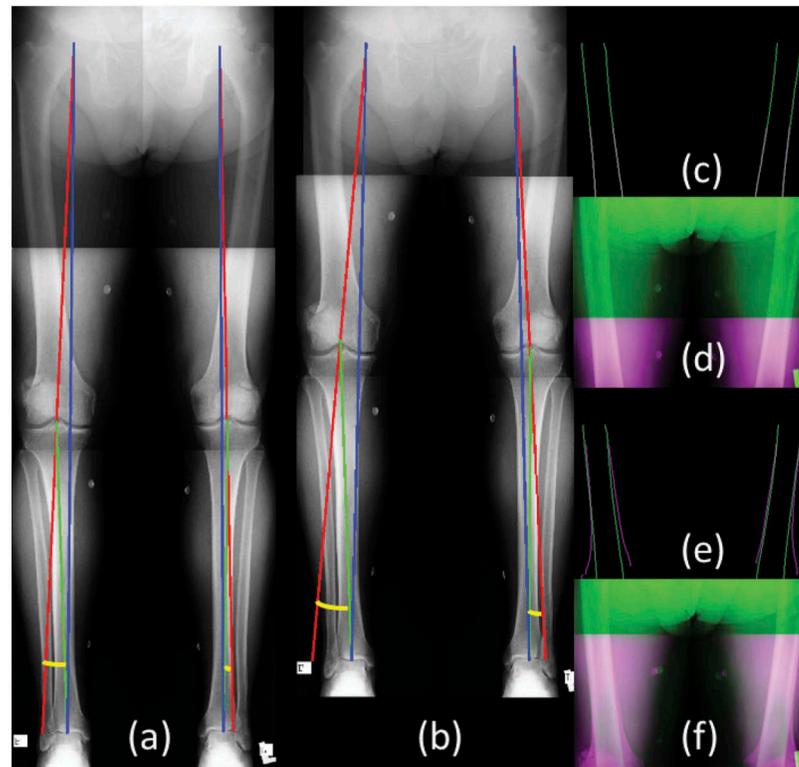
**Figure 9.**

**Left Panel:** Bland-Altman plots showing the mean  $\pm 2$  standard deviation of the HKA angle differences (in degrees) across 95 patient datasets between the HKA angle calculated from automated panoramas and the expert-generated manual panoramas for the right (blue diamond) and left (orange circle) limbs. **Right Panel:** Manual vs. automated HKA angle-based inclusion and exclusion criteria. As shown, the absolute value of the HKA angles of the 95 patients according to the automated and the manual panoramas. 7° (green line) is the upper threshold for including a patient in the clinical trial and 10° (red line) is the lower threshold for excluding the patient from the clinical trial.



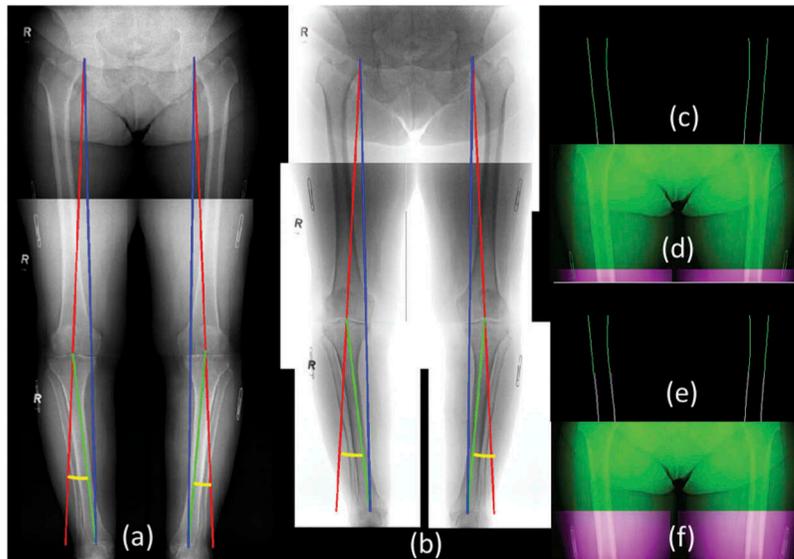
**Figure 10.**

**Left Panel:** HKA Angle vs. vertical translation differences between the automated and manual panoramas: left y axis (blue) shows the HKA angle difference (blue diamond), while the right y axis (orange) shows the total (femoral plus the tibial) vertical translation difference (orange circle) between the automated and the manual panoramas. **Right Panel:** Each patient is marked by a data point, where the x coordinate corresponds to the total vertical translation difference and the y coordinate corresponds to the HKA angle difference, both computed between the automated and manual panoramas.



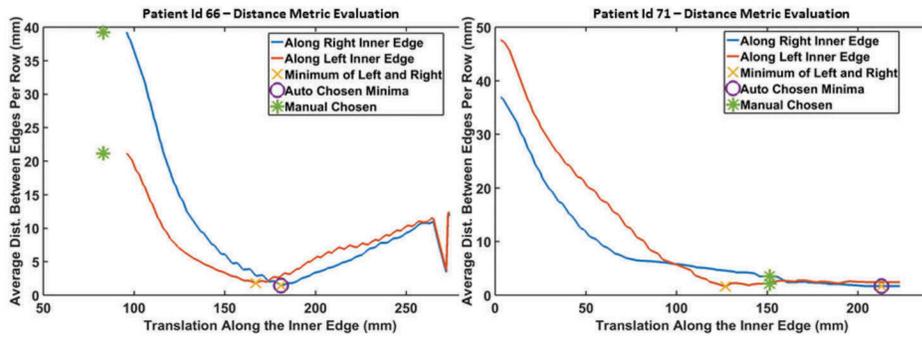
**Figure 11.**

Visualisation of patient No. 66 – automated and manual panoramas. From left to right: the automated long limb panorama (a); the expert manually generated long limb panorama (b); the differences between the edges and the original X-ray images in the upper stitch after applying the automated translations, with 50% transparency showing the fixed (green), moving (magenta) and overlap (white) regions (c) and (d), respectively. Similarly, the differences after applying the expert manual translations are also shown in panels (e) and (f). The manually constructed panorama is clearly erroneous given that in this image the femur is shorter than the tibia.



**Figure 12.**

Visualisation of Patient No. 71 – automated and manual panoramas. From left to right: the automated long limb panorama (a); the expert manual long limb panorama (b); the differences between the edges and the original X-rays in the upper stitch after applying the automated translations, showing the fixed (green), moving (magenta) and overlap (white) (c) and (d), respectively; similarly, the differences after applying the expert translations (e) and (f).



**Figure 13.**

Distance metric evaluation for cases 66 (left) and case 71 (right) upper stitch; the distance metric graphs for the right limb (blue) and left limb (orange) for each translation pair, as the moving image slides along the medial bone edge of the fixed image. The minimum of the left and right limbs are marked here with a yellow 'X' and the automated chosen minimum is marked with a purple circle, and correspondingly the expert manual translation is marked with green.

The femoral and tibial translation differences (horizontal – delta-x and vertical – delta-y) (in mm) between the expertly stitched and automatically generated panoramas: range, mean  $\pm$  standard deviation ( $\mu \pm \sigma$ ) for right and left limbs.

**Table 1.**

	Right Limb		Left Limb	
	Range	Mean $\pm$ St. Dev.	Range	Mean $\pm$ St. Dev.
Femur delta-x	[-5.00, 8.12]	0.43 $\pm$ 1.94	[-6.12, 11.25]	0.24 $\pm$ 2.31
Femur delta-y	[-97.75, 42.30]	3.76 $\pm$ 22.35	[-96.20, 37.49]	4.53 $\pm$ 22.47
Tibia delta-x	[-3.07, 3.00]	0.10 $\pm$ 1.09	[-7.52, 5.75]	0.26 $\pm$ 1.42
Tibia delta-y	[-31.40, 13.70]	2.13 $\pm$ 7.00	[-31.40, 13.17]	1.85 $\pm$ 6.78

**Table 2.**

The number of patients eligible for inclusion (HKA angle  $7^\circ$ ) in the clinical trial, exclusion, (HKA angle  $10^\circ$ ) and tentative inclusion assessed across all 95 patients based on the automated and the manual panoramas.

	Include		Tentative		Exclude		Total
	HKA	$7^\circ$	$7^\circ$ HKA	$10^\circ$	HKA	$10^\circ$	
Automated Panorama		79		12		4	95
Manual Panorama		77		13		5	95

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript