Published in final edited form as:

Med Eng Phys. 2013 July ; 35(7): 994–1004. doi:10.1016/j.medengphy.2012.09.012.

Hierarchical Model-based Tracking of Cervical Vertebrae from Dynamic Biplane Radiographs

Md. Abedul Haque^a, William Anderst^b, Scott Tashman^b, and G. Elisabeta Marai^a

^aUniversity of Pittsburgh, Dept. of Computer Science

^bUniversity of Pittsburgh, Dept. of Orthopaedic Surgery

Abstract

We present a novel approach for automatically, accurately and reliably de-termining the 3D motion of the cervical spine from a series of stereo or biplane radiographic images. These images could be acquired through a variety of different imaging hardware configurations. We follow a hierarchical, anatomically-aware, multi-bone approach that takes into account the complex structure of cervical vertebrae and inter-vertebrae overlapping, as well as the temporal coherence in the imaging series. These significant innovations improve the speed, accuracy, reliability and flexibility of the tracking process. Evaluation on cervical data shows that the approach is as accurate (average precision 0.3 mm and 1°) as the expert human-operator driven method that was previously state of the art. However, unlike the previously used method, the hierarchical approach is automatic and robust; even in the presence of implanted hardware. Therefore, the method has solid potential for clinical use to evaluate the effectiveness of surgical interventions.

Keywords

musculoskeletal; radiography; model-based registration; dynamic biplane imaging; kinematics; joints; cervical spine

1. Introduction

Accurate in-vivo motion tracking is important for understanding articulation kinematics [1, 2], musculoskeletal related diseases and the effectiveness of different treatments. For example, to correlate abnormal motion with morphological features such as inter-vertebral disc height (< 3mm in posterior space [3]), sub-millimeter accuracy is needed to avoid errors as large as 30% in disc-deformation measurements.

Model-based methods have been developed to measure 3D bone motion with high accuracy at knee or shoulder joints [4, 5, 6]; such methods employ 3D models of the bones, which they track through a sequence of dynamic x-ray images. Model-based methods are more accurate than skin marker-based methods, which suffer from errors as large as 10mm in translation and 8° in rotation [7]. They can also capture dynamic motion, unlike existing three-dimensional techniques such as Computed Tomography (CT, which also features

© 2012 Institute of Physics and Engineering in Medicine. All rights reserved

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Competing interests: None declared

higher radiation exposure, depending on the anatomic location) or Magnetic Resonance Imaging (MRI). Finally, unlike dynamic three-dimensional techniques (Cine-PC MRI [8]), model-based methods do not require continuous movement for long periods of time during data collection, support in general large ranges of motion, and pose fewer restrictions during imaging, thus leading to loadings more similar to most everyday movements. Given the advantages of model-based tracking, systems implementing model-based methods are utilized in various forms at several different academic institutions and medical research centers; the basic imaging hardware required for biplane radiography setup costs less than one third of what a modern 3T MRI scanner costs.

However, existing implementations of 3D model-based tracking methods suffer from the same critical issues. The existing tracking processes are extremely labor-intensive, requiring many (up to 30) hours of labor for every hour spent collecting data. Furthermore, a high level of expertise is required to generate trustworthy results. For this reason the tracking task cannot be reliably delegated to a crowd-sourcing approach such as the mechanical turk [9]. Accuracy and reliability, especially for the more automated algorithms, are inconsistent. Simultaneous acquisition of a pair of radiographic images is a prerequisite for all systems claiming high 3D accuracy, but this requirement creates significant image quality problems due to scatter radiation (a widely known issue for biplane radiographic imaging, which can become intractable for imaging the thicker parts of the body such as hips or the lumbar spine). It is also often difficult or impossible to obtain two radiographic views that avoid bone overlap in the images, which also degrades imaging matching performance using conventional tracking approaches. Surgically inserted hardware further decreases tracking accuracy and robustness. These limitations have thus far limited application of this technology to research studies, since the time and cost for data analysis is prohibitive for clinical use.

In this paper, we propose and validate an *automated* intelligent, hierarchical model-based method to track with sub-millimeter accuracy the 3D motion of cervical vertebrae from dynamic biplane radiographs. The specific goal of this work is automation while matching the accuracy of human expert operators on difficult, clinical cervical spine data. However, we note that similar challenges exist in accurate tracking of all multi-articular joints, many of which feature significant bony or hardware overlap (e.g. the hip, shoulder, wrist or ankle).

2. Methods

The basic premise for tracking bone motion is a model-based tracking approach that matches radiographic images to a known bone shape. 3D models of the bones of interest are obtained using conventional imaging (CT or MRI scans). A virtual model of the stereo-radiographic imaging system is generated using the precise locations of the radiographic sources and image detectors (which can be determined automatically by imaging a calibration object of known dimensions). Simulated x-rays are passed through the bone model to produce a pair of digitally reconstructed radiographs (DRRs) on the image plane. By manipulating the bone model within the virtual radiographic system, pairs of DRRs can be generated for any bone position. By calculating image similarity measures [10, 11] between the actual radiographic image pairs and the DRRs, the virtual bone position and orientation can be adjusted (manually or by an optimization algorithm) to identify the position that provides the greatest match, thus determining the position of the actual bone in space. This process is repeated for each pair of the images in the motion sequence, and repeated again for each bone of interest to yield the 3D position of the joint for the entire movement.

In conventional model-based tracking, the presence of overlapping bones (a common occurrence) reduces the quality of image matching and degrades tracking performance.

Also, tracking each bone independently ignores the known characteristics of joints that constrain the relative bone movements. We introduce a hierarchical multi-bone model approach, in which multiple bones are combined in a single 2D projection and simultaneously matched with the radiograph images. This approach takes advantage of the rich detail present in regions of radiographic bone overlap, which can therefore enhance, rather than degrade, tracking performance. By incorporating hierarchical, anatomically aware, multi-articular models of joints, as well as temporal coherence, tracking reliability can be further enhanced by exploiting known constraints that are defined across space and time.

The method uses a new digitally reconstructed radiograph generation procedure which takes into account multiple bones and results in a Multibone DRR (i.e., a MDRR), a standard image processing step, and a novel hierarchical optimization procedure. For each X-ray frame in the motion sequence, a 2D projection (the MDRR) is generated from the multiple reconstructed bone models. Next, both the X-ray and the 2D projection images are processed to reduce noise and enhance edges. Finally, an optimization method searches through different positions and orientations of the bone models to find the best match between the MDRR and the X-ray image. The process is repeated for all frames of a motion sequence (Fig. 1).

While a variety of imaging hardware setups can be used in practice to acquire dynamic radiograph images, for our experiments high-resolution X-ray images were captured using a dynamic stereo X-ray (DSX) system, and 3D volumetric images of the bones of interest were acquired with a high-resolution static computed tomography scanner (see subsection 2.5).

2.1. MDRR Generation

Model-based tracking [5] is based on the idea that an X-ray image can be computationally produced using a simplified X-ray generation model [12]:

$$I(\mathbf{b}) = \int_L \mu(r) dr \quad (1)$$

where $I(\mathbf{p})$ is the intensity of the X-ray detector pixel $\mathbf{p}$, $\mu(r)$ is the X-ray attenuation coefficient, and L is the projection beam from source to point $\mathbf{p}$. This model assumes that the X-ray system corrects for beam divergence and that the sensors have logarithmic response.

$I(\mathbf{p})$ can be estimated using ray casting of the CT image and the resulting image is known as the Digitally Reconstructed Radiograph (DRR). The segmented CT image serves directly as the 3D bone model.

However, in reality the intensity of a pixel $I(\mathbf{p})$ not only depends on the bone of interest, but also on any other structure the beam L passes through. Thus, $I(\mathbf{p})$ can be thought of as integration of contributions from several different sources:

$$\begin{aligned} I(\mathbf{p}) = & I_c(\text{bone of interest}) \\ & + I_c(\text{neighboring bones}) \\ & + I_c(\text{metallic implants}) \\ & + I_c(\text{other noise}) \\ & + I_c(\text{soft tissue}) + I_c(\text{random noise}) \end{aligned}$$

where $I_c(O)$ represents contribution to pixel (x, y) due to object O .

Existing conventional bone-tracking methods assume that the contribution to a pixel intensity from all the sources except the bone itself is negligible or constant [5]. This assumption does not hold for spine data (and other joints) where neighboring bones and implants overlap significantly. A single-bone DRR (SDRR) is nonetheless generated using only the model of the bone being tracked, and accounts only for the first term, $I_c(\text{bone of interest})$.

To account for contributions from neighboring bones and surgical implants, we propose the following model for multi-bone DRR (MDRR) generation:

$$I_{MDRR}(\mathbf{p}) = I_c(\text{bone of interest}) + I_c(\text{neighboring bones}) + I_c(\text{metallic implants})$$

In this approximation, it is assumed that contribution due to non-neighboring bones (e.g. shoulder, skull) can usually be avoided by careful positioning of X-ray sources and detectors. Any contribution due to soft tissue and random noise is assumed negligible here. This approach also does not account for possible differences between kVp settings of the CT and the fluoroscopy system [13]. Although still an approximation, the resulting MDRRs have a more realistic signature in adjacent regions than the SDRRs (Fig. 2).

We use ray-casting volume rendering through multiple neighboring bones to generate MDRRs. On a single 2.0 GHz processor, each MDRR generation requires on average 100ms. However, during bone position and orientation optimization more than 500 MDRRs may need to be generated for each motion frame. A sequential implementation would require approximately 6 hours of computation time to track 6 vertebrae over 60 frames. To reduce the computation time, we use a parallel implementation of the MDRR generation process.

2.2. Image Processing

The image processing steps and the similarity metric are similar to earlier studies [4, 5] enabling us to isolate and investigate the effect of MDRRs and hierarchical optimization. The MDRRs and the X-ray images are passed through several standard image processing steps to reduce noise and enhance edges of the images. Summarizing [4], a 3×3 discrete Gaussian filter is first applied to reduce noise (Fig. 3a and b), then a 3×3 Sobel filter is applied to extract edges (Fig. 3c) from the smoothed images. Finally, an edge-enhanced image (Fig. 3d) is produced by taking the weighted average of the Gaussian-smoothed image g and the Sobel edge-extracted image h :

$$k(x, y) = w_0 * g(x, y) + w_1 * h(x, y) \quad (2)$$

w_0 and w_1 are empirically determined to be 0.1 and 0.9 for MDRRs and 0.15 and 0.85 for X-ray images.

2.3. Similarity Metric

The similarity between edge-enhanced X-ray and MDRR images is measured using normalized correlation (r) (previously used in [14]) which has been reported as one the best performing similarity metrics [15] for high-resolution, real clinical data.

$$r_{(j=1,2)} = \frac{\sum (I_{Xray_j}(x, y) - \bar{I}_{Xray_j}) (I_{MDRR_j}(x, y) - \bar{I}_{MDRR_j})}{\sqrt{\sum (I_{Xray_j}(x, y) - \bar{I}_{Xray_j})^2} \sqrt{\sum (I_{MDRR_j}(x, y) - \bar{I}_{MDRR_j})^2}}$$

for all pixels (x, y) such that $I_{MDRR}(x, y) > 0$. Here, j denotes the X-ray camera number of the bi-plane X-ray system. Following Bey et. al. [4], correlations for two sets of X-ray and MDRR images are multiplied to get the final matching score.

$$corr(I_{Xray_1}, I_{Xray_2}, I_{MDRR_1}, I_{MDRR_2}) = r_1 * r_2 \quad (3)$$

2.4. Hierarchical Optimization

While MDRRs have the potential to increase the accuracy and robustness of tracking, they also pose challenges in terms of the optimization of search space. We use a quasi-Newton optimization method for finding the point of the maximum matching score. Each bone i has 6 degrees of freedom, expressed in matrix form as: $M_i = (T_x, T_y, T_z, R_x, R_y, R_z)$, the concatenation of 3 translations and 3 rotations. Translations and rotations are specified with respect to a local coordinate system. The origin (C_{Tx}, C_{Ty}, C_{Tz}) of the local coordinate system (Fig. 4.a) is the center of mass of a CT object:

$$CT_{(X=x, y, z)} = \frac{\sum X * f(x, y, z)}{\sum f(x, y, z)}$$

where f denotes the CT image.

The objective function for the optimization method is:

$$\max_{T_x, T_y, T_z, R_x, R_y, R_z} \prod_{j=1,2} corr[I_{Xray_j}, I_{MDRR_j}(T(CTObject))] \quad (4)$$

where $T = T_x, T_y, T_z, R_x, R_y, R_z, T_{-CT_x}, T_{-CT_y}, T_{-CT_z}$ is the composite transformation matrix and j denotes the two cameras of a DSX system. The equation shows optimization of a single bone; however, it can be expanded for multiple bones as necessary. The search is initialized by a human operator for the first two frames. For the remaining frames, initialization is done by exploiting temporal coherence.

Ideally, all bones of the hierarchy would be present in the MDRR and the optimization algorithm would tackle all bones simultaneously. However, such an approach becomes computationally intractable due to the large, relatively unorganized search space and to the large number of degrees of freedom ($6n$ where n is the number of the bones). In contrast, employing temporal coherence, inter-frame and inter-bone motion priors, and a systematic hierarchical technique can reduce the search space significantly. At the same time, a hierarchical approach should allow sufficient degrees of freedom between bones to capture joint motion accurately.

Based on the above observations we chose a coarse-to-fine, multi-pass strategy. The first two passes correspond to a coarse search for the bone location and orientation. Two additional passes correspond to a fine tuning stage, and use fewer degrees of freedom and increasingly more constraints.

In terms of relative contribution, early experiments showed that phase 1 and 3 significantly impact the registration robustness. The exclusion of phase 1 in the search method produced off-track solutions and the exclusion of phase 3 produced very poor quality solutions. Phase 2 and 4 contribute primarily to tracking accuracy. Together, the four phases complement each other and lead to a robust and accurate solution, as demonstrated in Section 3.

Below we describe each phase in detail using an example linear hierarchy consisting of n bones, B_1 (first) through B_n (last). Let K_i be the set of bones which are used to generate MDPR at the i^{th} round of a given phase and $L_i = (B_j/F_k)$ be the subset of bones B_j in K_i which are optimized during the i^{th} round of the given phase, given the fixed position and orientation of the F_k bones.

Phase 1: Temporal Coherence—The first phase uses temporal coherence to find a good starting position for each bone (Fig. 5). This step is particularly important, given that the cost function is smooth across narrow ranges, and thus prone to local minima. To facilitate reliable convergence, it is important to start the search from a sufficiently large set of plausible initial placements. Using the position and orientation of a bone in the most recent two frames, we predict the most likely area in the current frame to search for the optimal position and orientation. We extrapolate the positions of the previous two solved frames (r_1 and r_2 in Fig. 5) to predict the position in the current frame. Then we generate intermediate positions (intermediate circles in Fig. 5) between the predicted position and the position of the closest solved frame. Finally, we find the best-match position among all these potential positions. The rationale of multiple intermediate points is that a bone can change its direction and speed of movement. Multiple intermediate points help to compensate for this variation in speed and direction. Subsequent phases thoroughly search the region surrounding a selected seed location (cross points in Fig. 5) in order to capture accelerated and decelerated motion. The intermediate seed points are generated uniformly so that the translation difference or the angle difference between two consecutive seed points are at most 0.1mm or 0.1°; the spacing was empirically determined by taking into account the image resolution (0.23mm CT resolution and 0.3mm pixel size, see Table 1), and the smooth kinematics of the data (under 2mm and 2° between-frame motion; the chosen spacing could capture ranges of up to 5 mm and 5° between-frame motion).

The sets K_i and L_i for this phase are: $K_i = \{B_i\}$ and $L_i = K_i$ where $i = 1 : n$.

Phase 2: Pairwise Optimization—The purpose of phase 2 is to refine the seed location from phase 1 with the help of adjacent bone location information. To this end, we optimize each pair of adjacent bones (e.g. $B_1 - B_2$, $B_2 - B_3$ and so on) in the hierarchy by optimizing over 12 degrees of freedom. Optimizing two adjacent bones simultaneously helps adjust their position and orientation with respect to each other. This optimization step is computationally tractable while providing the necessary and sufficient degrees of freedom between neighboring bones in order to adjust their relative position and orientation.

$$K_i = \{B_i, B_{i+1}\} \text{ and } L_i = K_i$$

where $i = 1 : n - 1$.

Since bone kinematics are temporally smooth, we further constrain the search space by including motion priors into the optimization. The search range is restricted within twice (empirically determined) the magnitude of the bone motion (under 2 mm and 2° in our datasets) in the most recent two frames. This motion prior allows sufficient freedom for bone movement but prevents sudden large movement. Theoretically, the search space could

be further reduced by incorporating the fact that different bones cannot penetrate each other. However, the computational cost of detecting 3D inter-penetration makes such an approach less desirable.

Phase 3: Biased Singleton Refinement—The purpose of this phase is to refine the position and orientation for each bone by searching in the region chosen by phase 2. We note that the bones of a hierarchy are unequally easy to track. For example, bones at the top (e.g. cervical vertebra C3 in a hierarchy of C3–C7) are easier to track due to less soft tissue and less interference from surrounding bones. The hierarchical search method takes advantage of this prior information by biasing the order in which individual bones are optimized. In the spine case, we optimize each bone of the chain sequentially (starting from the top of the chain and moving towards the bottom) in the presence of the previous bone in the hierarchy. Since we optimize one bone at a time, the optimization is over 6 degrees of freedom.

$$K_i = \{B_1, \dots, B_i\}$$

and

$$L_i = \{B_i | B_1, \dots, B_{i-1} \text{ are kept fixed}\}$$

where $i = 1 : n$.

Phase 4: Neighbor-Constrained Refinement—This is a refinement phase that works similarly to phase 3, but takes into account both predecessor and successor bones in the hierarchy. The movement of each bone during the optimization is restricted by the presence of all the surrounding bones of the hierarchy. This helps to find the optimal position and orientation for the whole hierarchy.

$$K_i = \{B_1, \dots, B_n\}$$

and

$$L_i = \{B_i | B_1, \dots, B_{i-1}, B_{i+1}, \dots, \text{are kept fixed}\}$$

where $i = 1 : n$.

2.5. Validation

Datasets—To validate and evaluate the hierarchical, multi-bone approach we used in vivo conditions and real clinical data. 3D volumetric images of the bones of interest were obtained from a high resolution static computed tomography (CT) scanner (LightSpeed 16, GE Medical Systems, Waukesha, WI). CT images were segmented using 3D medical imaging software (Mimics, Materialize Inc, Leuven, Belgium) to extract individual bone geometry.

A Dynamic Stereo X-ray (DSX) system was used to capture high resolution X-ray images at a high frame rate. DSX utilizes two frame-synchronized imaging systems, specifically designed for dynamic measurement and mounted in a custom-designed, flexible positioning system to optimize viewing angles and provide freedom of subject movement. Each imaging

system includes a 100 kW constant-potential high-frequency cardiac cine-radiographic generator (CPX-3100CV, EMD, Quebec, CA), a 0.3/0.6mm focal spot size X-ray tube (G-1582; Varian, Salt Lake City, UT), a 40 cm image intensifier (TH9447QX; Thales, France), and a high-speed camera providing 1800×2400 pixel resolution at up to 500 frames/sec with 14-bit dynamic range (Phantom V10; Vision Research, Inc., Wayne, NJ). The EMD X-ray generators included upgraded software (provided by the manufacturer) to provide 1 ms pulses at repetition rates up to 180Hz, providing blur-free images and a dose reduction of 4–16× (relative to continuous operation). A calibration object (a cube) is used to calibrate the camera system. Figure 4.c shows the system configuration for a flexion-extension trial; X-ray beams are approximately 50 degrees apart. For twist trials, the subject is rotated to face one X-ray source, and the X-ray sources are lowered and tilted up 15 degrees, while the image intensifiers are raised and tilted down 15 degrees.

Following approval from our Institutional Review Board (Pitt IRB), 13 trials were acquired from 3 human subjects (1 male, 2 female, aged between 35 – 40) — 3 flexion/extension and 3 axial rotation trials from subject 1 and 2 and 1 flexion/extension trial from subject 3. Two of the test subjects have single-level anterior fusion in C5 and C6 and the remaining subject has single-level anterior fusion in C4 and C5. A fusion is performed by attaching a metal plate with two vertebral bodies using 4 screws. Trials from the subjects were taken between 6 and 7 months post surgery.

Tantalum beads were implanted into the fused cervical vertebrae and their adjacent bones during the fusion surgery so that a high accuracy ground truth solution could be produced by tracking the beads [5]. The subject with C4–5 fusion had beads into cervical vertebrae C3, C4, C5 and C6. Subjects with C5–6 fusion had beads into cervical vertebrae C4, C5, C6 and C7. Bead signatures were manually removed from the CT slices prior to MDRR generation by replacing voxels containing bead signatures with the average of the neighborhood voxels which do not belong to beads. In this way the beads did not influence the hierarchical and the operator-assisted tracking methods. For validation purposes, the center of each bead was also manually identified in the CT scans.

For these subjects and trials we tracked 5 cervical vertebrae (C3 – C7) and the fusion hardware. In the clinical study that generated this data cervical vertebrae C1 and C2 were not tracked (C1 and C2 are often completely overlapped by other anatomical structures), and so they did not have beads implanted into them to produce a ground truth solution. Table 1 shows additional specifications of our datasets.

Ground Truth—We obtain the ground truth by tracking the implanted beads in the distortion-corrected radiographs as described previously [5]. To validate this ground truth, we measure bias and precision [16]. Summarizing these references, first inter-bead distances (d_i where i is the frame number) are calculated per bone over an entire trial from the bead tracking (i.e. ground truth) results. True inter-bead distance (D) is measured by manually detecting the beads in the CT image. Next, differences between the inter-bead distance computed from the CT data and the bead tracking based solution ($D-d_i$) are calculated for each frame over an entire trial. Bias and precision are defined as the mean and standard deviation of the differences over the entire trial and are summarized over all trials to report in the form mean±standard deviation. Bead-based tracking is used as the “gold standard” to calculate the accuracy of the operator-assisted single-bone and hierarchical multi-bone tracking.

Analysis—Performance of the operator-assisted single-bone and the hierarchical multi-bone methods are compared for each bone and each axis in terms of bias, precision [5, 16] root-mean-squared (rms) error, and maximum error. These performance metrics are

computed with respect to the bead-based ground truth solution. For each method (operator-assisted, ground truth and hierarchical) the bead centroid locations are computed as the average of the three known bead coordinates of a bone over all frames of a trial. The differences between estimated bead centroid (hierarchical or operator-assisted method) and the ground truth bead centroid locations are computed for each bone and for each axis across all frames of a trial. Bias, precision, rms error and maximum error are defined as the mean, standard deviation, rms value and maximum value of this time-history of differences. Finally, bias, precision and rms error are summarized as the mean, standard deviation and maximum error over all trials. To compare the performance of the operator-assisted and the hierarchical method for a bone along all three axes simultaneously, Hotteling's T-squared tests ($\alpha = 0.05$) are performed for the three axes on bias, precision and also rms error of the bone.

To further compare the accuracy of both methods in clinically relevant terms, an anatomical coordinate system was defined at the center of each vertebra [17]. The axes of the anatomical coordinate system are Anterior-Posterior (AP) axis, Superior-Inferior (SI) axis and Medial-Lateral (ML) axis, and correspond to the main axes of the vertebral body (Fig. 4.b). Translations and rotations of each pair of vertebrae with respect to the three axes are computed from the bead-based, the assisted single-bone and the hierarchical multi-bone method solutions. Error (in terms of bias and precision) in the assisted single-bone method and the hierarchical method with respect to the bead-based ground truth is computed and then tested using a Hotteling's T-squared test ($\alpha = 0.05$).

In our experiments, the operator-assisted single-bone method was guided by an expert operator and the solution was checked and refined manually. The hierarchical method did not require any human assistance after initialization. We compare the solutions from these two methods in terms of accuracy, robustness and run time.

3. Experimental Results

There was no bias in the implanted bead tracking solution (Table 2) i.e. in the ground truth. Average precision over all bones is 0.11mm which is very similar to results published earlier [5, 18].

3.1. Accuracy

Figure 6 shows the bias in the hierarchical multi-bone method and the operator-assisted single-bone method. We did not find any statistically significant difference ($p > 0.01$) in bias between the hierarchical and the operator-assisted methods (Table 3) for any bone in any coordinate system direction. This finding indicates that the hierarchical multi-bone method has similar accuracy to the operator-assisted single-bone method. C5 and C6 have relatively higher bias than C3, C4 and C7 along the X and Y axes both for the operator-assisted single-bone and the hierarchical methods. Maximum bias of C3, C4 and C7 using the operator-assisted single-bone method along the X and Y axes are approximately 7 and 3 times higher than the maximum bias of C5 and C6 using the same method along X and Y axes. In contrast, maximum bias of C3, C4 and C7 using the hierarchical multi-bone method along the X and Y axes is approximately 3 and 4 times higher than the maximum bias of C5 and C6 using the same method along the X and Y axes. C5 and C6 tracking results show relatively higher bias than other bones because C5 was fused in all three of our subjects and C6 was fused in two of our three subjects. Fused vertebrae are typically harder to track due to relatively inaccurate CT-scan models of the fused bones (extracting a CT bone model from fused vertebrae is more difficult than extracting a bone model from non-fused vertebrae), interference from hardware during tracking etc. However, the bias for the hierarchical method for C5 and C6 along any axis was less than 0.62 mm; for the operator-

assisted single-bone method it was less than 1.07 mm which indicates that the hierarchical multi-bone method has sub-millimeter accuracy for all vertebrae, including the ones which have been affected by fusion surgery.

Figure 7 compares precision of the hierarchical multi-bone method and the operator-assisted single-bone method. Hierarchical multi-bone method precision ranged from 0.03mm to 0.34mm depending on axis direction. Operator-assisted single-bone method precision ranged from 0.04mm to 0.55mm depending on axis direction. We did not find any statistically significant difference ($p > 0.01$) in precision between the hierarchical and the operator-assisted methods (Table 4) along any coordinate system direction. Precision was bone-independent for both of the methods. The last bone to track (C7) shows slightly worse accuracy in the hierarchical method than in the operator-assisted single-bone method. A possible explanation is that C7-s have more soft tissue surrounding them compared to other bones in the dataset. Therefore, the assumption of negligible contribution in pixel intensity due to surrounding soft tissue for MDRR generation (Eq. 2) may no longer hold for C7, making human intervention particularly valuable. We note that for the operator-assisted single-bone method C5 and C6 show slightly lower precision (statistically not significant) than C3, C4 and C7.

Figure 8 shows the root-mean-squared (rms) error of the hierarchical and the operator-assisted methods. Root-mean-squared error indicates no statistically significant difference ($p > 0.01$) between the hierarchical and the operator-assisted methods (Table 3) for any bone in any coordinate system direction. From Fig. 8, we can see that vertebrae C5 and C6 show slightly higher rms error than the vertebrae C3, C4 and C7 for the operator-assisted single-bone method. For the operator-assisted single-bone method, the maximum rms error is 1.11mm for C5 and C6, and 0.44mm for C3, C4 and C7. For the hierarchical multi-bone method, the maximum rms error is 0.70mm for C5 and C6, and 0.51mm for C3, C4 and C7. The rms error result again indicates that the hierarchical multi-bone method is as accurate as the operator-assisted single-bone method even for the vertebrae with fusion hardware.

Figure 9 shows the maximum error of the hierarchical and the operator-assisted method. It is interesting to see that for both of the methods the maximum error is often greater or close to 1mm (especially for C5 and C6). Considering the sub-millimeter level values found in bias, precision and rms error analysis, these high values in maximum error analysis indicate the presence of outliers X-ray frames where the image quality is relatively poorer than most other frames. According to our expert operators, this is often the case and the operator-assisted single-bone method fails mainly for these outliers frames.

Table 7 and Table 8 summarize the translation and rotation analysis with respect to anatomical coordinates. For the hierarchical approach, the average precision in measuring 3D joint kinematics is 0.27mm for translations (compared to 0.37 mm using the operator-assisted method) and 1.0 for rotations (compared to 1.10° for the operator-assisted method). Hotelling's T-Squared testing ($p > 0.01$) found no statistically significant difference between the two methods.

A quantitative analysis of the implant hardware motion could not be performed due to the lack of a bead-based ground truth solution. However, two human expert operators manually checked the hierarchical method tracking solution of the implant hardware. According to the expert operators, the hierarchical method tracking solution of implant hardware was accurate and those experts often commented that the hierarchical tracking solution seemed better than the operator-assisted method tracking solution.

To further analyze the effect of implant hardware on method accuracy, we used the hierarchical method to track a trial (subject 1, flexion-extension trial 1) without including the implant hardware in the MDRR generation process. This approach produced very poor quality solution for the fused vertebra C5, while the other fused vertebra C6 went off track around frame 15.

3.2. Robustness

To compare the robustness of the hierarchical approach against the single-bone method across trials, we tracked 13 trials from the 3 subjects with each method. Both methods were run without human operator assistance after initialization of the reference frames. The single-bone method failed to track 5 trials (3 flexion/extensions and 1 axial rotation from subject 1, 1 flexion/extension from subject 3). In all these trials, either vertebra C5 or C6 went significantly off track. For example, in the flexion/extension trial from subject 3, vertebra C6 went visibly off track around frame 6 and the bone remained off track for the rest of the frames. C6 also went off track in all 3 flexion/extension trials from subject 1 approximately around frame 10 (Fig. 10). C5 went off track in one axial rotation trial from subject 1 around frame 15. Although the single-bone method was able to complete tracking the remaining 8 trials, the tracking results required significant manual correction through human operator intervention. In contrast, the hierarchical approach completed successfully all trials, and the results required no manual correction.

3.3. Run Time

We used a Windows based cluster with 24 Intel Xeon (2.0 GHz) processors to run both the operator-assisted and the hierarchical tracking methods. The operator-assisted single-bone method required approximately 1 hour to track a single cervical vertebra in a single trial of 60 ~ 90 frames, leading to a total of 6 hours for the five cervical vertebrae and the metallic implant. 96% of the time was spent on human interaction during the tracking (to keep bones on track) and the manual refinement phase. The hierarchical method tracked five cervical vertebrae and the metallic implant on average in approximately 25 minutes, i.e., it attained a speedup factor of 12. While having a faster solution is not the main contribution of this work, these run times indicate the hierarchical method is cost-effective, which is essential for clinical application.

4. Discussion

The experimental results show that the hierarchical multi-bone method matches the sub-millimeter accuracy of the state-of-the-art operator-assisted single-bone method. At the same time, the hierarchical method is superior to the single-bone method in terms of robustness (100% compared to 61%) and run-time (1 compared to 12). Notably, the hierarchical approach dramatically reduces the labor required for imaging studies, while making the accuracy and robustness of the method operator-independent.

An advantage of the hierarchical method is that it is able to track bones that change both direction and speed within each camera view during a specific trial, a task significantly more difficult than tracking a bone that moves in the same direction at a constant speed. In our experiments, no restrictions were imposed on the subjects during motion. The maximum range of motion was approximately 20mm translation and 32° rotation; bone motion direction and speed had variations in all datasets.

In our experiments, we used post-operative CT scan data to extract a model of each hierarchy component to be tracked, be they vertebrae or additional hardware. However, a pre-operative CT scan can also be used with this approach. The pre-operative scan would be

used to extract models of the bones; while the implant model can be acquired either through scanning or directly thanks to CAD designs.

We followed an incremental approach in developing the four step hierarchical searching process. For example, the initial searching process included only two steps (Phase 1 and 3) and produced a poor solution for the first node of the hierarchy. We tried several other structures for the hierarchical searching e.g. different orderings of the phases and different orderings of the bones. These approaches did not produce any improvements over the approach described in this paper.

Furthermore, the hierarchical multi-bone method assumes that noise in X-ray images due to surrounding soft-tissue structure of cervical spine is negligible. However, it is often very challenging to avoid soft-tissue interference and get high quality noise-free X-ray images from both cameras of a DSX system due to restrictions imposed by a subject's motion and body structure. Current model-based techniques, including the one reported in this paper, frequently employ image filters as a pre-processing step to reduce noise due to soft tissue [19]. Many of these image processing techniques [20] require manual tuning even for a single dataset, primarily due to overall intensity variation from frame to frame. Adjusting these parameters for different datasets (or for different frames within the same dataset) requires significant human effort, outweighing the benefits of the pre-processing step. Robust and automatic image pre-processing remains a direction of future work.

In our experiments we compare the performance of the hierarchical multi-bone method against the operator-assisted single-bone method which is to the best of our knowledge the current state of the art in the field [17]. Biplane or stereo radiographic imaging similar to the one used in our experiments enables accurate quantitative 3D motion assessment for both static [5] and dynamic bone motion analysis. Bone location and orientation can be precisely measured by beads (Radiostereometric analysis or RSA) implanted into the bones [18]; bead-based tracking has shown good accuracy and has been used to produce ground truth to validate other tracking methods [5, 18]. Other model-based methods are primarily single-2D to 3D matching (i.e. they only collect one x-ray view [21, 22, 23]). Typically, motion is kept within a plane perpendicular to the direction of camera projection. A study on knee implant tracking [24] reported out of plane translation errors greater than 3 mm despite the fact that the movement was mainly in plane. Banks [21] compares the accuracy of his method to others. To summarize, all matching that is done using only one X-ray view has large errors perpendicular to the imaging plane. This class of methods is highly unlikely to produce higher accuracy results than the biplane radiographic approach. Several groups do collect 2 X-ray views [25, 26]. They build bone models from MRI/CT and either manually match the model to the 2 X-ray views or use invasively implanted beads for the matching process. Given the fundamentally manual, respectively invasive matching processes used, these approaches are not preferable to either the expert-assisted method or the hierarchical method described here.

While our experiments use data acquired through a stereo-radiographic imaging system and a CT scanner, a variety of imaging hardware setups, including single-plane radiography, can be used in practice to acquire dynamic radiograph images. Stereo-imaging is more likely than single-plane radiography to suffer from image quality problems due to scatter radiation and thus tracking motion from stereo images is likely to benefit more from our approach. However, radiographic bone overlap and temporal coherence are traits of both single-plane and stereo-imaging; our hierarchical tracking algorithm may enhance the accuracy and robustness of single-plane dynamic tracking. A major advantage of our approach is that in the long run it could eliminate the requirement for simultaneous image acquisition, thus leading to dramatic improvements in radiographic image quality.

5. Conclusion

In this paper, we have introduced and validated an intelligent, hierarchical algorithm which improves the accuracy, reliability, and/or flexibility of the dynamic radiograph tracking process. The two significant innovations proposed – 1) multibone projection (MDRR); and 2) temporally-aware constrained hierarchical optimization – can be applied in combination to enable rapid, automated, accurate bone motion tracking and to facilitate clinical application.

When applied to cervical spine data, the new algorithm matched the sub-millimeter accuracy of the expert-operator existing tracking process, while being automated and operator-independent. The approach was also more robust in the presence of implanted hardware than the single-bone tracking process. Finally, the approach sped up the total tracking time by a factor of 12. Preliminary evaluation indicates similar improvements on *in vitro* lumbar spine data [27].

Our automated process decreases the labor cost associated with human operators, which facilitates practical clinical application. Considering the relatively low levels of radiation involved by the imaging system (approximately half the amount of one cervical CT scan), the moderate hardware costs and the proliferation in recent years of biplane DSX systems, the automation of the tracking procedure shows promise for large-scale clinical application. Since the approach showed also good performance in the presence of implanted hardware, it can be used to study post-surgery cases and evaluate the effectiveness of a surgical intervention. Applications for this technology include (but are not limited to) assessment and diagnosis of musculoskeletal disorders, bone, ligament and joint injury, derangements of the spine and osteoarthritis.

Acknowledgments

This work was supported by NSF-IIS-0952720. Thanks are also due to Emma Bergeron for the manual tracking results, and to Brittany Frankowiack and Jamie Novak for their writing feedback.

Ethical approval: All subjects for the study signed Institutional Review Board (IRB)-approved informed consent forms before being enrolled in the study

Reference

1. Cerciello T, Romano M, Bifulco P, Cesarelli M, Allen R. Advanced template matching method for estimation of intervertebral kinematics of lumbar spine. *Medical Engineering & Physics*. 2011; 33(10):1293–302. [PubMed: 21764624]
2. McCane B, Abbott JH, King T. On calculating the finite centre of rotation for rigid planar motion. *Medical Engineering & Physics*. 2005; 27(1):75–9. [PubMed: 15604008]
3. Lu J, Ebraheim NA, Yang H, Rollins J, Yeasting RA. Anatomic bases for anterior spinal surgery: surgical anatomy of the cervical vertebral body and disc space. *Surgical and radiologic anatomy SRA*. 1999; 21(4):235–9. [PubMed: 10549078]
4. Bey MJ, Zauel R, Brock SK, Tashman S. Validation of a new model-based tracking technique for measuring three-dimensional, in vivo glenohumeral joint kinematics. *Journal of Biomechanical Engineering*. 2006; 128(4):604–9. [PubMed: 16813452]
5. Anderst W, Zauel R, Bishop J, Demps E, Tashman S. Validation of three-dimensional model-based tibio-femoral tracking during running. *Medical Engineering & Physics*. 2009; 31(1):10–6. [PubMed: 18434230]
6. Bingham J, Li G. An optimized image matching method for determining in-vivo tka kinematics with a dual-orthogonal fluoroscopic imaging system. *Journal of Biomechanical Engineering*. 2006; 128(4):588–95. [PubMed: 16813450]

7. Holden JP, Orsini JA, Siegel KL, Kepple TM, Gerber LH, Stanhope SJ. Surface movement errors in shank kinematics and knee kinetics during gait. *Gait & Posture*. 1997; 5(3):217–27.
8. Behnam AJ, Herzka DA, Sheehan FT. Assessing the accuracy and precision of musculoskeletal motion tracking using cine-pc mri on a 3.0t platform. *Journal of Biomechanics*. 2011; 44(1):193–7. [PubMed: 20863502]
9. 2010. <http://aws.amazon.com/documentation/mturk/>
10. Penney GP, Weese J, Little JA, Desmedt P, Hill DL, Hawkes DJ. A comparison of similarity measures for use in 2-d-3-d medical image registration. *IEEE Transactions on Medical Imaging*. 1998; 17(4):586–95. [PubMed: 9845314]
11. Wu J, Kim M, Peters J, Chung H, Samant SS. Evaluation of similarity measures for use in the intensity-based rigid 2d-3d registration for patient positioning in radiotherapy. *Medical Physics*. 2009; 36(12):53–91.
- [12]. Tomazevic D, Likar B, Slivnik T, Pernus F. 3-d/2-d registration of ct and mr to x-ray images. *IEEE Trans Med Imag*. 2003; 22(11):1407–16.
- [13]. Bifulco P, Cesarelli M, Allen R, Romano M, Fratini A, Pasquariello G. 2d-3d registration of ct vertebra volume to fluoroscopy projection: A calibration model assessment. *EURASIP Journal on Advances in Signal Processing*. 2010; 2010:1–9.
- [14]. Anderst WJ, Baillargeon E, Donaldson WF, Lee JY, Kang JD. Validation of a noninvasive technique to precisely measure in vivo three-dimensional cervical spine movement. *Spine*. 2011; 36(6):393–400. [PubMed: 21325933]
- [15]. Russakoff DB, Rohlfing T, Ho A, Kim DH, Shahidi R, Adler JR, et al. Evaluation of intensity-based 2d-3d spine image registration using clinical gold-standard data. *Image Registration*. 2003; 151:60.
- [16]. ASTM: Standard Practice for Use of the Terms Precision and Bias in ASTM Test Methods. West Conshohockent; PA: 1996.
- [17]. McDonald CP, Bachison CC, Chang V, Bartol SW, Bey MJ. Three-dimensional dynamic in vivo motion of the cervical spine: assessment of measurement accuracy and preliminary findings. *The spine journal official journal of the North American Spine Society*. 2010; 10(6):497–504. [PubMed: 20359957]
- [18]. Anderst WJ, Baillargeon E, Donaldson WF, Lee JY, Kang JD. Validation of a noninvasive technique to precisely measure in vivo three-dimensional cervical spine movement. *Spine*. 2011; 36(6):E393–400. [PubMed: 21372650]
- [19]. Cesarelli M, Bifulco P, Cerciello T, Romano M, Paura L. X-ray fluoroscopy noise modeling for filter design. *International journal of computer assisted radiology and surgery*. 2012
- [20]. Teyhen DS, Flynn TW, Bovik AC, Abraham LD. A new technique for digital fluoroscopic video assessment of sagittal plane lumbar spine motion. *Spine*. 2005; 30(14):E406–13. [PubMed: 16025017]
- [21]. Banks SA, Hodge WA. Accurate measurement of three-dimensional knee replacement kinematics using single-plane fluoroscopy. *IEEE Trans Biomed Eng*. 1996; 43(6):638–49. [PubMed: 8987268]
- [22]. Muggleton JM, Allen R. Automatic location of vertebrae in digitized videofluoroscopic images of the lumbar spine. *Medical Engineering & Physics*. 1997; 19(1):77–89. [PubMed: 9140876]
- [23]. Bifulco P, Cesarelli M, Cerciello T, Romano M. A continuous description of intervertebral motion by means of spline interpolation of kinematic data extracted by videofluoroscopy. *Journal of Biomechanics*. 2012; 45(4):634–41. [PubMed: 22277152]
- [24]. Mahfouz MR, Hoff WA, Komistek RD, Dennis DA. A robust method for registration of three-dimensional knee implant models to two-dimensional fluoroscopy images. *IEEE Trans Med Imag*. 2003; 22(12):1561–74.
- [25]. Brainerd EL, Baier DB, Gatesy SM, Hedrick TL, Metzger KA, Crisco JJ. X-ray reconstruction of moving morphology (xromm): Precision, accuracy and applications in comparative biomechanics research. *Journal of Experimental Zoology*. 2010; 313A(5):262–79. [PubMed: 20095029]
- [26]. Li G, Van De Velde SK, Bingham JT. Validation of a non-invasive fluoroscopic imaging technique for the measurement of dynamic knee joint motion. *Journal of Biomechanics*. 2008; 41(7):1616–22. [PubMed: 18394629]

- [27]. Haque, A.; Anderst, W.; Tashman, S.; Marai, G. Validation of a non-invasive automated hierarchical method to precisely measure lumbar spine movement. 2012 Annual Meeting of the Orthopaedic Research Society; 2012.

Figure 1.

Model-based tracking method overview. The method uses a three step process: 2D projection, image processing, and optimization. The process is repeated for all frames of a motion sequence.

Figure 2. Single-bone projection of C4 (a) and multi-bone projection of C4 (c). Notice how the presence of C3 and C5 (b) changes the signature of C4 (in (c)) in the adjacent region.

Figure 3.

X-ray image processing for both cameras of a random trial. (a) Raw images. (b) Gaussian-filtered images. (c) Gaussian and Sobel filtered images. (d) Gaussian, Sobel and weighted averaged images. Snapshots (c) and (d) have been contrast-enhanced for printing purposes.

Figure 4.

(a). Local coordinate system defined on a vertebra. (b) Anatomical coordinate system (used for analysis) defined on the same vertebra. (c) System configuration for cervical flexion-extension trials. In this image, X-ray sources (left) direct X-rays through the subject to image intensifiers (right).

Figure 5.

A conceptual 2D diagram of phase 1. $r1$ and $r2$ represent the bone in the previous two solved frames. p is the predicted point for current frame. Intermediate circles represent the additional intermediate points. The circle within the diamond represents the selected seed point. The cross points represent the surrounding regions which are going to be searched in the next phases.

Figure 6.

Bias of the hierarchical method and the operator-assisted method. X, Y and Z denote the axes of comparison. No statistically significant difference ($p > 0.01$) was found in bias between the hierarchical and the operator-assisted methods for any bone. However, note that the bias for C5 and C6 is relatively higher than the bias for C3, C4 and C7 due to the hardware implant in C5 and C6. Still, the hierarchical multi-bone method has achieved sub-millimeter accuracy for C5 and C6 (max bias 0.62 mm) along with the other vertebrae. Results are averaged over 13 trials from 3 subjects, 60 ~ 100 frames per trial.

Figure 7.

Precision for the hierarchical and the operator-assisted single-bone method. X, Y and Z denote the axes of comparison. We did not find any statistically significant difference ($p > 0.01$) in precision between the hierarchical and the operator-assisted method for any bone. Results are averaged over 13 trials from 3 subjects, 60 ~ 100 frames per trial.

Figure 8.

Root-mean-squared error (mm) for the hierarchical and the operator-assisted single-bone method. X, Y and Z denote the axes of comparison. We did not find any statistical significant difference ($p > 0.01$) in rms error between the hierarchical and the operator-assisted method for any bone. Results are averaged over 13 trials from 3 subjects, 60 ~ 100 frames per trial.

Figure 9.

Maximum error for the hierarchical and the operator-assisted single-bone method. X, Y and Z denote the axes of comparison. The automated hierarchical method is as accurate as the operator-assisted method. Please note that relatively higher values of maximum error (compared to rms error, bias and precision) indicate the presence of some outliers X-ray frames with poorer image quality than most other frames of a motion sequence. Results are summarized over 13 trials from 3 subjects, 60 ~ 100 frames per trial.

Figure 10.

Comparison of tracking results between the single-bone method without operator assistance (top row) against the hierarchical approach without operator assistance (bottom row). In each row, the two camera views are shown, with the vertebra (green) superimposed on the X-ray image (red); dotted lines show the target location of the vertebra. Note the single-bone solution is off-track, while the hierarchical one is on target.

Table 1

Experiment data set specification

Bone	Cervical Spine (<i>C3 - C7</i>)
Motion	Flexion/extension, axial rotation
Total trials	13
Original CT Resolution(mm)	0.23×0.23×1.25
Interpolated CT resolution (mm)	0.23×0.23×(0.23~0.5)
X-ray image resolution	1024×1024
Number of frames per trial	60 ~ 100
Frame capturing rate	30 frames/sec
Pixel size of x-ray image (mm)	0.30×0.30

Table 2

Bead-based tracking accuracy (bias and precision). All measurement units are in mm.

Bone	Bias	precision
C3	0.05 ± 0.03	0.11 ± 0.02
C4	0.04 ± 0.03	0.10 ± 0.03
C5	-0.04 ± 0.10	0.12 ± 0.02
C6	0.04 ± 0.07	0.12 ± 0.04
C7	0.02 ± 0.05	0.08 ± 0.03

Bias of the hierarchical method and the operator-assisted method. All measurement units are in mm. No statistically significant difference was found in bias ($p > 0.01$) for any bone between the hierarchical and the operator-assisted method.

Table 3

Mean $\pm$ standard deviation of bias over all trials						
Axis	Hierarchical multi-bone method			Operator-assisted single-bone method		
	X	Y	Z	X	Y	Z
C3	0.04 $\pm$ 0.12	0.01 $\pm$ 0.17	0.08 $\pm$ 0.05	-0.04 $\pm$ 0.17	-0.04 $\pm$ 0.22	0.03 $\pm$ 0.08
C4	0.07 $\pm$ 0.08	0.02 $\pm$ 0.15	0.18 $\pm$ 0.16	-0.03 $\pm$ 0.13	0.05 $\pm$ 0.09	0.21 $\pm$ 0.14
C5	0.19 $\pm$ 0.20	0.06 $\pm$ 0.39	-0.07 $\pm$ 0.31	0.22 $\pm$ 0.18	0.11 $\pm$ 0.29	-0.07 $\pm$ 0.42
C6	-0.08 $\pm$ 0.25	-0.04 $\pm$ 0.10	0.07 $\pm$ 0.14	0.00 $\pm$ 0.49	0.12 $\pm$ 0.22	0.04 $\pm$ 0.14
C7	-0.18 $\pm$ 0.14	-0.27 $\pm$ 0.06	-0.12 $\pm$ 0.18	-0.13 $\pm$ 0.06	-0.16 $\pm$ 0.10	0.02 $\pm$ 0.07

Precision of the hierarchical method and the operator-assisted method. All measurement units are in mm. We did not find any statistically significant difference ($p > 0.01$) in precision between the hierarchical and the operator-assisted method for any bone.

Table 4

Axis	Mean $\pm$ standard deviation of precision over all trials					
	Hierarchical multi-bone method			Operator-assisted single-bone method		
	X	Y	Z	X	Y	Z
C3	0.20 $\pm$ 0.05	0.16 $\pm$ 0.09	0.12 $\pm$ 0.02	0.21 $\pm$ 0.08	0.16 $\pm$ 0.04	0.14 $\pm$ 0.05
C4	0.09 $\pm$ 0.03	0.09 $\pm$ 0.03	0.07 $\pm$ 0.02	0.12 $\pm$ 0.07	0.12 $\pm$ 0.05	0.09 $\pm$ 0.04
C5	0.13 $\pm$ 0.04	0.14 $\pm$ 0.04	0.09 $\pm$ 0.02	0.22 $\pm$ 0.11	0.16 $\pm$ 0.05	0.12 $\pm$ 0.05
C6	0.14 $\pm$ 0.07	0.13 $\pm$ 0.07	0.07 $\pm$ 0.02	0.19 $\pm$ 0.10	0.19 $\pm$ 0.13	0.12 $\pm$ 0.08
C7	0.09 $\pm$ 0.02	0.11 $\pm$ 0.05	0.07 $\pm$ 0.03	0.10 $\pm$ 0.02	0.15 $\pm$ 0.09	0.07 $\pm$ 0.04

Root-mean-squared error of the hierarchical method and the operator-assisted method. All measurement units are in mm. We did not find any statistically significant difference ($p > 0.01$) in rms error between the hierarchical and the operator-assisted method for any bone.

Table 5

Mean $\pm$ standard deviation of root-mean-squared error over all trials						
Axis	Hierarchical multi-bone method			Operator-assisted single-bone method		
	X	Y	Z	X	Y	Z
C3	0.23 $\pm$ 0.05	0.23 $\pm$ 0.07	0.15 $\pm$ 0.03	0.26 $\pm$ 0.10	0.25 $\pm$ 0.08	0.16 $\pm$ 0.04
C4	0.14 $\pm$ 0.04	0.16 $\pm$ 0.08	0.23 $\pm$ 0.10	0.17 $\pm$ 0.09	0.15 $\pm$ 0.05	0.26 $\pm$ 0.10
C5	0.26 $\pm$ 0.16	0.35 $\pm$ 0.21	0.29 $\pm$ 0.13	0.37 $\pm$ 0.16	0.30 $\pm$ 0.16	0.32 $\pm$ 0.30
C6	0.24 $\pm$ 0.18	0.17 $\pm$ 0.07	0.15 $\pm$ 0.07	0.42 $\pm$ 0.31	0.26 $\pm$ 0.20	0.17 $\pm$ 0.11
C7	0.21 $\pm$ 0.12	0.30 $\pm$ 0.05	0.16 $\pm$ 0.16	0.17 $\pm$ 0.05	0.24 $\pm$ 0.08	0.10 $\pm$ 0.05

Table 6

Maximum error of the Hierarchical method and the operator-assisted method. All measurement units are in mm.

Axis	Maximum error						
	Hierarchical multi-bone method			Assisted single-bone method			
	X	Y	Z	X	Y	Z	
C3	1.01	0.55	0.53	0.56	0.50	0.43	
C4	0.48	0.43	0.50	0.43	0.45	0.59	
C5	1.28	0.93	0.55	1.34	0.97	0.56	
C6	1.28	0.72	0.43	1.68	1.57	0.99	
C7	0.27	0.11	0.18	0.24	0.44	0.30	

Translational (anatomical coordinates) bias and precision of the hierarchical method and the operator-assisted method. All measurement units are in mm. There was no statistically significant difference (Hotelling's T-Squared, $p > 0.01$) in bias and precision between the hierarchical and the operator-assisted method.

Table 7

Translation accuracy in anatomical coordinate system over all trials						
Axis	Hierarchical multi-bone method			Operator-assisted single-bone method		
	LM	SI	AP	LM	SI	AP
Bias	0.46±0.67	0.22±0.33	0.85±1.10	0.47±0.59	0.09±0.32	0.69±0.75
Precision	0.25±0.06	0.20±0.12	0.35±0.17	0.32±0.09	0.40±0.34	0.41±0.17

Rotational (anatomical coordinates) bias and precision of the hierarchical method and the operator-assisted method. All measurement units are in degrees. There was no statistically significant difference (Hotelling's T-Squared, $p > 0.01$) in bias and precision between the hierarchical and the operator-assisted method.

Table 8

Rotation accuracy in anatomical coordinate system over all trials						
Axis	Hierarchical multi-bone method			Operator-assisted single-bone method		
	LM	SI	AP	LM	SI	AP
Bias	-0.31±2.80	-0.17±1.59	0.31±1.49	-0.08±1.71	-0.14±1.65	0.32±1.55
Precision	1.26±0.46	0.95±0.26	1.08±0.47	1.37±0.51	0.94±0.24	0.97±0.19