



Automatic annotation of hip anatomy in fluoroscopy for robust and efficient 2D/3D registration

Robert B. Grupp¹ · Mathias Unberath¹ · Cong Gao¹ · Rachel A. Hegeman² · Ryan J. Murphy³ · Clayton P. Alexander⁴ · Yoshito Otake⁵ · Benjamin A. McArthur^{6,7} · Mehran Armand^{2,4,8} · Russell H. Taylor¹

Received: 15 November 2019 / Accepted: 3 April 2020 / Published online: 24 April 2020
© CARS 2020

Abstract

Purpose Fluoroscopy is the standard imaging modality used to guide hip surgery and is therefore a natural sensor for computer-assisted navigation. In order to efficiently solve the complex registration problems presented during navigation, human-assisted annotations of the intraoperative image are typically required. This manual initialization interferes with the surgical workflow and diminishes any advantages gained from navigation. In this paper, we propose a method for fully automatic registration using anatomical annotations produced by a neural network.

Methods Neural networks are trained to simultaneously segment anatomy and identify landmarks in fluoroscopy. Training data are obtained using a computationally intensive, intraoperatively incompatible, 2D/3D registration of the pelvis and each femur. Ground truth 2D segmentation labels and anatomical landmark locations are established using projected 3D annotations. Intraoperative registration couples a traditional intensity-based strategy with annotations inferred by the network and requires no human assistance.

Results Ground truth segmentation labels and anatomical landmarks were obtained in 366 fluoroscopic images across 6 cadaveric specimens. In a leave-one-subject-out experiment, networks trained on these data obtained mean dice coefficients for left and right hemipelvises, left and right femurs of 0.86, 0.87, 0.90, and 0.84, respectively. The mean 2D landmark localization error was 5.0 mm. The pelvis was registered within 1° for 86% of the images when using the proposed intraoperative approach with an average runtime of 7 s. In comparison, an intensity-only approach without manual initialization registered the pelvis to 1° in 18% of images.

Conclusions We have created the first accurately annotated, non-synthetic, dataset of hip fluoroscopy. By using these annotations as training data for neural networks, state-of-the-art performance in fluoroscopic segmentation and landmark localization was achieved. Integrating these annotations allows for a robust, fully automatic, and efficient intraoperative registration during fluoroscopic navigation of the hip.

Keywords Landmark detection · Semantic segmentation · 2D/3D registration · X-ray navigation · Orthopedics

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s11548-020-02162-7>) contains supplementary material, which is available to authorized users.

✉ Robert B. Grupp
grupp@jhu.edu

¹ Department of Computer Science, Johns Hopkins University, Baltimore, MD, USA

² Research and Exploratory Development Department, Johns Hopkins University Applied Physics Laboratory, Laurel, MD, USA

³ Auris Health, Inc., Redwood City, CA, USA

⁴ Department of Orthopaedic Surgery, Johns Hopkins Medicine, Baltimore, MD, USA

⁵ Graduate School of Information Science, Nara Institute of Science and Technology, Ikoma, Nara, Japan

⁶ Department of Surgery and Perioperative Care, Dell Medical School, University of Texas, Austin, TX, USA

⁷ Texas Orthopedics, Austin, TX, USA

⁸ Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD, USA

Introduction

Minimally invasive surgical interventions of the hip manipulate, modify, or augment anatomical structures which are hidden or not reliably visible [1]. Clinicians commonly use intraoperative fluoroscopy in order to overcome this occlusion and ascertain the poses of anatomy, surgical instruments, or artificial implants. However, mental interpretation of these images is a difficult task and subject to an extensive learning curve [2,3]. Computer-assisted navigation systems ease this burden by tracking relevant objects and reporting their poses within the context of a surgical plan or scenario. Systems leveraging fluoroscopy have been developed for total hip arthroplasty [4], hip resurfacing [5], cement injection [6], and osteotomies of the acetabulum [7] or proximal femur [8].

In order to report object poses accurately, fluoroscopic navigation systems rely on 2D/3D registrations between intraoperative 2D images and the appropriate 3D models [9]. Large errors in pose estimates may occur when a registration is not initialized sufficiently close to the actual pose of an object. Quality initializations are derived from some manual human input, often through annotated landmark locations in the fluoroscopic image. Although these systems report favorable navigation-related results, the manual initialization of processing may interrupt surgical workflows and negatively affect patient outcomes due to increased operating time or blood loss.

Convolutional neural networks (CNNs) have excelled at detecting landmarks and performing semantic segmentation when sufficiently large annotated datasets are available for training [10,11]. However, since existing large-scale hip datasets have focused on 3D image modalities and pre- and postoperative radiography, rather than intraoperative fluoroscopy [12], applications of CNNs to fluoroscopy have been mostly limited to recognizing surgical instruments and tools [13–16].

Several authors have coupled image segmentation with landmark estimation using multi-task networks and achieved favorable results. By reusing segmentation features from an encoder–decoder style network for the computation of landmark heatmaps, Laina was able to automatically annotate segmentation labels and landmark locations of tools used in laparoscopy and retinal microsurgery [17]. Gao also leveraged this approach for the localization of a dexterous continuum manipulator in fluoroscopy [16]. Kordon demonstrated that a CNN, trained from 149 manually annotated preoperative radiographs, could successfully segment the four bones of the knee joint and locate two anatomical landmarks and a surgically relevant line [18].

Using a large collection of simulated fluoroscopy, Bier trained CNNs to annotate anatomical landmarks of the pelvis [19]. When evaluated on five sequences of actual fluoroscopy across two cadaveric specimens, mean annota-

tion errors of 12–24 mm in the detector plane were reported. Pelvis poses were estimated using these annotations, yielding reprojection errors of 14–34 mm for other landmarks not learned by the network. Their work was extended in [20], whereby each network was fine-tuned on simulated fluoroscopy for a specific patient of interest. The approach was evaluated by estimating landmark locations in previously unseen simulated images and using these estimates to produce quality initializations for 2D/3D registration. No analysis on actual fluoroscopy was conducted in [20].

In this paper, we propose a method for 2D/3D registration of hip anatomy that simultaneously combines image intensities with higher-level landmark and segmentation features, making it robust against large initial offsets from actual object poses. Segmentation labels and landmark annotations are produced by a CNN similar in architecture to those found in [16,17]. Contrary to [16,19,20], we train our networks using smaller datasets of *actual* fluoroscopy and achieve state-of-the-art results on clinically relevant data. Annotated fluoroscopy for *training* is semi-automatically obtained using a computationally expensive 2D/3D registration, with run-times on the order of several minutes per image.

The novel contributions of this paper are:

- A semi-automatic, *offline*, pipeline for creating the first annotated training dataset of semantically segmented individual bone structures and anatomical landmark locations in actual hip fluoroscopy,
- A demonstration that CNN models, trained using small datasets of less than 400 annotated images, can achieve state-of-the-art results for the tasks of semantic segmentation and landmark localization in actual hip fluoroscopy,
- An *online*, *intraoperative*, registration strategy, leveraging image intensities and CNN features, that is fully automatic, requires no initialization from a user, and completes in an order of seconds.


Methods

We now describe the details of the data preprocessing, the methods for creating an annotated, *training*, dataset of hip fluoroscopy, the CNN architecture, and the intraoperative registration strategy. The reader is referred to the supplementary document for details regarding the lower-level parameters used for the registration pipelines and network training.

Data preprocessing

Using the procedure described in [7], lower torso 3D CT scans are resampled to have 1 mm isotropic spacing. Seg-

Fig. 1 Three views of the 6 anatomical structures and 14 landmarks to be annotated in 2D fluoroscopy. All landmarks are bilateral with left (L.) and right (R.) denoted. The L. hemipelvis is shown in green, the R. hemipelvis in red, L. femur in cyan, R. femur in orange, vertebrae in blue, and upper sacrum in yellow. Each landmark is overlaid as a purple sphere



mentations of the pelvis, femurs, and vertebrae are obtained semi-automatically. A total of 14 landmarks are manually annotated in 3D: the left and right (L./R.) centers of the femoral head (FH), L./R. greater sciatic notches (GSN), L./R. inferior obturator foramen (IOF), L./R. medial obturator foramen (MOF), L./R. superior pubis symphysis (SPS), L./R. inferior pubis symphysis (IPS), and the L./R. anterior superior iliac spine (ASIS). These landmarks were previously identified as being useful for obtaining initial registration estimates of the pelvis [7]. The anterior pelvic plane (APP) coordinate system for each specimen is defined using the L./R. ASIS and L./R. SPS landmarks [21] and is later used to estimate nominal anterior/posterior (AP) poses and as a reference coordinate frame during registration. Segmentations of each hemipelvis and sacrum are separated from the full pelvis segmentation, and any sacrum labels inferior to the sacroiliac joint are discarded. Figure 1 shows an example 3D visualization of the individual bone surfaces and the anatomical landmarks.

Fluoroscopy is collected with a Siemens CIOS Fusion mobile C-arm with 30×30 cm² detector. Images are 1536×1536 pixels with 0.194 mm pixel spacings. Each image is cropped by 50 pixels along each border to remove collimator artifacts, and intensity values are log-corrected (“bone is bright”).

2D/3D registration

Our approach to 2D/3D registration of single-view fluoroscopy and CT builds upon the multiple resolution, multiple component, 2D/3D, intensity-based registration pipeline introduced in [7]. The registration problem of finding the rigid poses of the pelvis (θ_P), left femur (θ_{LF}), and right femur (θ_{RF}) with respect to a single fluoroscopic view, I , is defined by the optimization problem (1), where \mathcal{P} indicates a projection operator creating digitally reconstructed radiographs (DRRs), \mathcal{S} indicates a similarity measure between DRRs and fluoroscopy, \mathcal{R} is a regularization penalizing implausible poses, and $\lambda \in [0, 1]$ is a tuning parameter.

$$\min_{\theta_P, \theta_{LF}, \theta_{RF} \in SE(3)} \lambda \mathcal{S}(\mathcal{P}(\theta_P, \theta_{LF}, \theta_{RF}), I) + (1 - \lambda) \mathcal{R}(\theta_P, \theta_{LF}, \theta_{RF}) \quad (1)$$

In this paper, \mathcal{S} is defined as the weighted sum of normalized cross-correlations of 2D image gradients computed over image patches [22]. For all registrations using regularization, $\lambda = 0.9$.

Training dataset creation

The *training* dataset of annotated fluoroscopy images is constructed using an automatic 2D/3D registration of the pelvis and both femurs. Once anatomy is registered to each image, the 3D segmentation labels and landmarks are propagated to 2D. Since this registration is performed “offline,” we use a computationally expensive combination of global search strategies, followed by several local strategies. Manual inspection is performed so that images corresponding to failed registrations are pruned from the dataset. It should be emphasized that although this registration is automatic and global, the amount of computation precludes it from intraoperative application.

An attempt is first made to register the pelvis using a mixture of the differential evolution [23], exhaustive grid search, particle swarm [24], covariance matrix adaptation: evolutionary search (CMA-ES) [25], and bounded optimization by quadratic approximation (BOBYQA) [26] optimization strategies at multiple resolutions. Using a combination of the CMA-ES and BOBYQA strategies, the left and right femurs are registered once the pelvis is registered. The rotation components of the left femur and right femurs are independently estimated, keeping the pelvis fixed at its current pose estimate. Next, simultaneous optimization over the rigid poses of the pelvis and both femurs is performed. Multiple resolution levels are used throughout this process, with down-sampling factors along each 2D dimension ranging from $32 \times$

to $4\times$. For each of the preceding registrations, uniform patch weightings were applied for S .

The 2D location of each landmark is obtained by projecting the corresponding 3D landmark onto the detector. When a landmark projects outside the detector region, it is identified as not visible in the image.

Each 2D pixel is labeled as the anatomy for which the corresponding source to detector ray intersects. Discrete labels are used to assign a single class of anatomy to each pixel. Femurs are given the highest precedence in labeling: any ray/femur intersection yields a label of the corresponding femur. Hemipelves have the next highest precedence, with any rays intersecting both hemipelves assigned a label corresponding to the hemipelvis closer to the X-ray source. Vertebrae intersections are given next precedence, followed by the upper sacrum. All remaining pixels are assigned to background.

The 2D labels and landmarks for each projection are manually inspected and verified.

Network architecture and training

In constructing our network, we follow the approach described by [16] and [17], appending segmentation and landmark heatmap network paths after an encoder–decoder structure. Supp. Figures S-2, S-3, and S-4 describe the network architecture used in this work. For the encoder–decoder in this paper, we adopt a 6-level U-Net [27] design with 32 features at the top level and 1024 features at the bottom. Our implementation is fully convolutional with learned 2×2 convolutions of stride 2 for downsampling and transposed convolutions for upsampling.

The segmentation path follows directly from the original U-Net design. The differentiable dice score [28] is computed for each class and then averaged. This value is bounded, taking on values in $[0, 1]$, with larger values indicating a higher-quality segmentation.

Segmentation features prior to soft-max are concatenated with the features output from the encoder–decoder and passed through two 1×1 convolutions to obtain a feature map where each channel estimates the heatmap of a landmark.

Ground truth heatmaps for each landmark location are defined by a symmetric 2D normal distribution with mean value equal to the landmark location and standard deviations of $\sigma = 3.88$ mm in each direction.

The value of σ was subjectively chosen to approximate the variance found in manual landmark annotation. Each heatmap is set to be identically zero when the landmark is not visible. Examples of ground truth heatmaps are shown in the third row of Fig. 2. Heatmap loss is computed using the average normalized cross-correlation (NCC) between each ground truth heatmap and the corresponding estimate. This

term is bounded, taking on values in $[-1, 1]$, with larger positive values indicating stronger correlation between ground truth heatmaps and estimated heatmaps.


By scaling and shifting the average NCC value into the range of $[0, 1]$, the heatmap loss may be weighted equally to the dice term without any additional hyper-parameter tuning. Finally, the combined dice and heatmap terms are negated (for minimization).

Networks are trained using stochastic gradient descent, with an initial learning rate of 0.1, Nesterov momentum of 0.9, weight decay of 0.0001, and a batch size of five images. Training and validation data sets are obtained by applying a random 85%/15% split to the data not used for testing. Test data sets are comprised of images collected from a single specimen, and no images derived from this specimen are present in the training and validation data. Extensive online data augmentation is applied to each image with probability 0.5. If an image is to be augmented, the intensities are randomly inverted, random noise is added to the image, the contrast is randomly adjusted, a random 2D affine warp is applied, and a random number of regions are corrupted with very large amounts of noise. Each image is normalized to have zero mean and standard deviation one before input into the network. Training is run for a maximum of 500 epochs, and the learning rate is multiplied by 0.1 after validation loss plateaus. The network expects images of size 192×192 pixels, and fluoroscopy data are downsampled $8\times$ in each dimension, accordingly. PyTorch 1.2 was used to implement, train, and test the networks.

Extracting landmark locations

Both, segmentations and heatmaps, are used to estimate anatomical landmark locations. Candidate locations of the FH landmarks are restricted to pixels labeled as the corresponding femur, and all remaining landmarks are restricted to locations labeled as the corresponding hemipelvis. Restricting candidate locations in this way avoids possible false alarms when the ipsilateral landmark is not in the view and a large heatmap intensity is located about the contralateral landmark. The final proposed location of each landmark is defined as the candidate location with maximal heatmap intensity. In order to distinguish between the cases of landmark detection, no detection, and spurious heatmap values, a 25^2 pixel region of the estimated heatmap, centered around the proposed location, is matched against the 2D symmetric normal distribution template of a detection at the center of the region. A detection is reported when NCC between the two regions is greater than 0.9, and no detection is reported otherwise.

Fig. 2 Example annotations of four specimens. The top row shows the ground truth segmentation labels for each object overlaid onto the fluoroscopic images, along with the landmark locations as yellow circles. The colors of each object correspond to those from Fig. 1. CNN estimates are shown in the second row, with ground truth landmark locations shown as yellow circles and estimated locations shown as yellow crosshairs (+). Missed detections are indicated by a circle without a corresponding cross. Ground truth heatmaps for the R. MOF, L. ASIS, L. GSN, and L. IOF, in **a–d**, respectively, are overlaid and shown in the third row. Estimated heatmaps for these landmarks are shown in the bottom row. The heatmap shown in **b** highlights a successful no detection report for L. ASIS



Intraoperative registration

The intraoperative registration strategies in this paper attempt to solve (1) in a similar fashion as the method used for construction of the training data set: the pelvis is registered first, followed by optimizations of each femur's rotation, followed by a simultaneous optimization over the rigid poses of all objects.

Method 1: A naive approach for efficient automatic registration uses only intensity information, with uniform patch weightings and no regularization applied. The single landmark initialization described in [7] is used to calculate an initial AP pose of the pelvis, aligning the 3D centroid of the L. ASIS, R. ASIS., L. SPS, and R. SPS with the center of the image.

Method 2: However, a great deal of information about the 2D image is known, courtesy of the segmentation and landmark localizations produced by the CNN. A less naive approach uses detected landmarks to solve the PnP problem [29] and automatically initializes an intensity-based registration. The segmentation is used to apply non-uniform

patch weightings in \mathcal{S} , and soft bounds are applied through a regularization on pose parameters.

Method 3: Instead of treating landmark features and intensity features separately, the detected landmark locations may be incorporated into a robust reprojection regularizer for intensity-based registration. The regularizer is defined in (2), with the l th landmark location in 3D being denoted by $p_{3D}^{(l)}$ and corresponding estimated 2D location $p_{2D}^{(l)}$.

$$\mathcal{R}(\theta_P) = \frac{1}{2\sigma_\ell^2} \sum_{l=1}^{N_L} \left\| \mathcal{P}(p_{3D}^{(l)}; \theta_P) - p_{2D}^{(l)} \right\|_2^2 \quad (2)$$

As with the PnP approach, non-uniform patch weightings are applied using the segmentation. Using one of the estimated 2D landmark locations, the single landmark initialization is used to calculate an initial AP pose of the pelvis.

Pelvis registrations first use a CMA-ES optimization, followed by the BOBYQA strategy at a finer resolution without patch weightings or regularization. When using patch weightings, patches centered at pixels labeled as either hemipelvis are given uniform weight and the remaining patches are weighted zero. Next, femur registration proceeds

Table 1 Landmark detection errors across all trained networks for each landmark

| Landmark | Error | | False negative rate | False positive rate |
|----------|---------------|-----------------|---------------------|---------------------|
| | Pixels | mm | | |
| L. FH | 1.9 ± 0.9 | 3.0 ± 1.5 | 0.02 | 0.02 |
| R. FH | 3.2 ± 1.9 | 5.0 ± 3.0 | 0.04 | 0.01 |
| L. GSN | 4.4 ± 2.0 | 6.8 ± 3.1 | 0.20 | 0.00 |
| R. GSN | 4.7 ± 2.3 | 7.3 ± 3.6 | 0.14 | 0.01 |
| L. IOF | 2.8 ± 4.0 | 4.3 ± 6.2 | 0.23 | 0.01 |
| R. IOF | 2.3 ± 3.3 | 3.5 ± 5.1 | 0.16 | 0.02 |
| L. MOF | 3.7 ± 3.2 | 5.8 ± 5.0 | 0.17 | 0.04 |
| R. MOF | 3.4 ± 1.9 | 5.2 ± 3.0 | 0.17 | 0.02 |
| L. SPS | 3.1 ± 2.4 | 4.7 ± 3.7 | 0.27 | 0.02 |
| R. SPS | 3.7 ± 2.9 | 5.8 ± 4.5 | 0.22 | 0.01 |
| L. IPS | 1.9 ± 1.6 | 3.0 ± 2.4 | 0.17 | 0.02 |
| R. IPS | 1.5 ± 1.0 | 2.3 ± 1.6 | 0.15 | 0.01 |
| L. ASIS | 9.0 ± 9.6 | 14.0 ± 14.9 | 0.29 | 0.01 |
| R. ASIS | 3.9 ± 3.7 | 6.0 ± 5.7 | 0.14 | 0.01 |
| All | 3.2 ± 3.4 | 5.0 ± 5.2 | 0.17 | 0.01 |

identically to that used during construction of the training data set, except for the case of patch weighting. When registering the individual femurs and using patch weightings, patches centered at pixels labeled as either hemipelvis or either femur are given uniform weight and all remaining patches are weighted zero. Multiple resolution levels are used, with either $8\times$ or $4\times$ downsampling applied.

Experiments and results

Data collection and training dataset creation

Lower torso CT scans were obtained for three male and three female cadaveric specimens with median age 88 and ranging 57–94 years. Each CT scan was semi-automatically segmented, and 3D landmarks were manually digitized. A total of 399 fluoroscopic images were collected at various C-arm poses. The “offline” *training* data set registration pipeline was run on each image, and a total of 366 images were verified to have registered successfully. For each successfully registered image, manual annotations were made for each femur indicating whether enough of the bone was visible to use for future registration evaluation. These counts are also broken down across each specimen in Supp. Table S-6.

Across all specimens, the intraoperative femur poses differed from their poses during preoperative CT scanning by an average of $15.3^\circ \pm 7.4^\circ$ and 0.7 ± 0.6 mm. For specimens 3 and 6, there were no images which had sufficient views of the left femur needed to evaluate registration. For specimen 3, only two images had sufficient views of the right

femur needed to evaluate registration, and the femur was in the same pose for these two views. In the images which had sufficient views for femoral registration, more than two poses were observed for each femur of specimens 1, 2, 4, and 5.

Poses recovered from successful registrations in this phase were treated as ground truth during intraoperative registrations. Examples of generated 2D ground truth annotations are shown in the top row of Fig. 2. The mean total registration time per image was 4 minutes using a NVIDIA Tesla P100 (PCIe) GPU.

Segmentation and landmark localization

A total of six networks were trained in a leave-one-specimen-out experiment. For each network, the training and validation data consisted of all labeled images from five specimens and all labeled images from the remaining specimen were used as test data. Across all test images, mean dice coefficients of 0.86 ± 0.20 , 0.87 ± 0.18 , 0.90 ± 0.24 , 0.84 ± 0.31 , 0.74 ± 0.19 , and 0.63 ± 0.13 were obtained for the left hemipelvis, right hemipelvis, left femur, right femur, vertebra, and upper sacrum, respectively. A listing of dice coefficients for each object of each specimen is shown in Supp. Table S-7. The average landmark 2D localization error was 5.0 ± 5.2 mm in the detector plane. Table 1 lists the average landmark errors, false negative rates, and false positive rates for each landmark. Figure 3 shows a plot of localization error thresholds and corresponding correct detection rates. The mean time for segmentation and landmark detection per image was 24.0 ± 0.4 ms using a NVIDIA Tesla P100 (PCIe) GPU.




Fig. 3 A plot of 2D landmark detection accuracy given various thresholds in mm. The bilateral cases for each landmark are combined in this plot

Intraoperative registration

Using estimated segmentations and landmark locations of the 366 test images, each intraoperative registration strategy was run and compared to the ground truth pose estimates from the training dataset. Registrations with pelvis rotation error less than 1° were defined as successful. This error was computed in the projective frame with center of rotation at the ground truth midpoint between the two femoral heads. Femur errors were computed using relative poses of the femur with respect to the pelvis in the APP with center of rotation at the ipsilateral femoral head. Table 2 lists the mean registration errors for each left out specimen and the errors over all specimens. Depth estimation inaccuracies accounted for nearly all of the pelvis translation error, with mean errors about the X, Y, and Z axes of 0.1 ± 0.1 mm, 0.1 ± 0.1 mm and 1.4 ± 2.0 mm, respectively, for method 3. For each specimen, the decompositions of method 3's pelvis errors are listed in Supp. Table S-8.

Two-tailed Mann-Whitney U tests were used to compare the magnitudes of the rotation and translation errors

Table 2 Pelvis and femur registration errors from successful pelvis registrations using the three intraoperative approaches and broken down by cadaver specimen. Femur registrations errors are reported for all successful pelvis registrations which have sufficient visibility of a femur

| Regi. method | Spec. | Pelvis errors | | | Femur errors | | |
|--------------|-------|---------------|-----------|-------------|--------------|-----------|-------------|
| | | # Success | Rot. (°) | Trans. (mm) | # | Rot. (°) | Trans. (mm) |
| 1: Naive | | | | | | | |
| | 1 | 32 (29%) | 0.1 ± 0.1 | 0.3 ± 0.2 | 13 | 0.4 ± 0.2 | 0.3 ± 0.3 |
| | 2 | 15 (14%) | 0.1 ± 0.2 | 0.8 ± 1.9 | 5 | 0.7 ± 0.4 | 0.4 ± 0.5 |
| | 3 | 1 (4%) | < 0.1 | 0.2 | 0 | — | — |
| | 4 | 4 (8%) | 0.1 ± 0.1 | 0.4 ± 0.4 | 0 | — | — |
| | 5 | 13 (24%) | 0.1 ± 0.1 | 0.4 ± 0.3 | 3 | 0.4 ± 0.2 | 0.7 ± 0.4 |
| | 6 | 1 (4%) | 0.1 | 0.3 | 1 | 0.6 | 0.2 |
| | All | 66 (18%) | 0.1 ± 0.1 | 0.4 ± 0.9 | 22 | 0.4 ± 0.3 | 0.4 ± 0.3 |
| 2: PnP Init. | | | | | | | |
| | 1 | 99 (89%) | 0.1 ± 0.1 | 0.8 ± 1.1 | 73 | 1.7 ± 5.2 | 0.6 ± 0.5 |
| | 2 | 96 (92%) | 0.1 ± 0.2 | 1.0 ± 1.4 | 59 | 1.2 ± 1.0 | 0.5 ± 0.4 |
| | 3 | 19 (79%) | 0.2 ± 0.2 | 1.6 ± 3.1 | 2 | 0.9, 0.8 | 0.4, 1.0 |
| | 4 | 38 (79%) | 0.2 ± 0.2 | 1.4 ± 1.8 | 27 | 1.3 ± 1.2 | 0.4 ± 0.4 |
| | 5 | 40 (73%) | 0.1 ± 0.1 | 0.7 ± 0.9 | 20 | 0.8 ± 0.8 | 0.6 ± 0.7 |
| | 6 | 7 (29%) | 0.1 ± 0.1 | 0.8 ± 1.0 | 2 | 1.3, 0.6 | 0.4, 0.1 |
| | All | 299 (82%) | 0.1 ± 0.2 | 1.0 ± 1.5 | 183 | 1.4 ± 3.4 | 0.5 ± 0.5 |
| 3: Combined | | | | | | | |
| | 1 | 101 (91%) | 0.1 ± 0.1 | 1.0 ± 1.5 | 73 | 1.8 ± 5.2 | 0.6 ± 0.5 |
| | 2 | 99 (95%) | 0.2 ± 0.2 | 1.4 ± 1.7 | 61 | 1.3 ± 1.0 | 0.7 ± 0.8 |
| | 3 | 18 (75%) | 0.2 ± 0.2 | 2.8 ± 3.4 | 2 | 1.1, 1.1 | 1.0, 1.3 |
| | 4 | 41 (85%) | 0.2 ± 0.2 | 2.1 ± 2.9 | 29 | 1.6 ± 1.3 | 0.6 ± 1.0 |
| | 5 | 47 (85%) | 0.1 ± 0.1 | 0.9 ± 1.2 | 24 | 0.8 ± 0.8 | 0.5 ± 0.6 |
| | 6 | 7 (29%) | 0.3 ± 0.3 | 3.0 ± 3.2 | 3 | 1.0 ± 0.7 | 0.3 ± 0.2 |
| | All | 313 (86%) | 0.2 ± 0.2 | 1.4 ± 2.0 | 192 | 1.5 ± 3.3 | 0.6 ± 0.7 |




Fig. 4 Abnormal cases with C-arm poses different from the training dataset. The lower hip is not visible in **a**; however, 2 landmarks were accurately detected allowing a successful pelvis registration. Excessive pelvic tilt is shown in **b**, **c**, while **d** shows large magnification. Detections with large errors are highlighted by yellow boxes. In **b**, a single

landmark, out of five detections, had large error, which allowed the registration strategy to succeed. The C-arm pose of **c** causes the boundary along the left femoral neck to appear similar to that adjacent to the IOF in an AP view, resulting in a detection with large error. Pelvis registration in **d** fails due to the large localization errors in each landmark

between methods 2 and 3. Using a 0.005 threshold, a significant difference was found between the pelvis rotation errors ($p < 0.001$), while no significant differences were found between pelvis translation errors ($p = 0.045$), femur rotation errors ($p = 0.089$), and femur translation errors ($p = 0.268$).

Correlation coefficients between dice scores and the pelvis rotation and translation errors were calculated using Spearman's rank correlation coefficient. For method 2, the correlation coefficients for the pelvis rotation and translation errors were -0.32 and -0.29 , respectively, and -0.31 and -0.33 , respectively, for method 3. The average of dice scores from the segmentations of hemipelves and femurs was used for this calculation.

The mean runtime for method 3 was 7.2 ± 0.7 s using a NVIDIA Tesla P100 (PCIe) GPU. Examples of automatic annotation and registration with method 3 are shown in the supplementary video.¹

Discussion and conclusion

The naive intraoperative registration performed poorly, succeeding in only 18% of trials, while the methods leveraging CNN annotations succeeded over 4 times as often. Despite the fairly large false negative detection rate of 17%, an average of 7 landmarks per image were detected, allowing methods 2 and 3 to perform well. Method 3's performance was robust when only 2, 3, and 4 landmarks were detected, reporting success in 2, 15, and 30 cases, respectively. Figure 4a shows an image with 2 detected landmarks and was registered successfully. Highlighting the robustness gained from mixing intensity features with landmark

features, method 2 was only successful with these number of detections in 0, 7, and 26 cases, respectively. The low false positive detection rate ensured that inconsistent features would not confound the registration.

Although the naive registration only succeeded in 66 cases, the mean femur rotation errors were about 1° smaller than those of methods 2 and 3. However, methods 2 and 3 were also successful in 62 and 64 of method 1's successes, each with a mean femur rotation error of $0.8^\circ \pm 0.5^\circ$. This indicates that the three methods perform comparably on images for which the naive approach succeeds. Moreover, the larger errors of methods 2 and 3 in the remaining cases are in part caused by the more challenging pelvis registration problems presented in these images, for which the naive method was unable to solve.

Method 3 was robust to poor initializations, which most likely resulted in the larger number of successful pelvis registrations compared to those resulting from method 2. In contrast to method 2, the objective function of method 3 never places penalties on the offsets of poses from their initial estimates. The registration is free to minimize the image similarity term on the condition that known 3D landmarks project to approximately the correct location in 2D. Conversely, the registration is free to minimize landmark reprojection error, so long as the candidate poses produce images that approximately match the observed image. This is contrary to the standard approach for regularization used by method 2, which imposes limits on the amount of rigid movement, even when the initial estimates are far from the true poses.

Despite the significant difference in pelvis rotation errors between methods 2 and 3, we believe that the small error magnitudes resulting from both methods should not negatively impact the clinical application of either approach.

¹ <https://youtu.be/5AwGIKncp9o>.

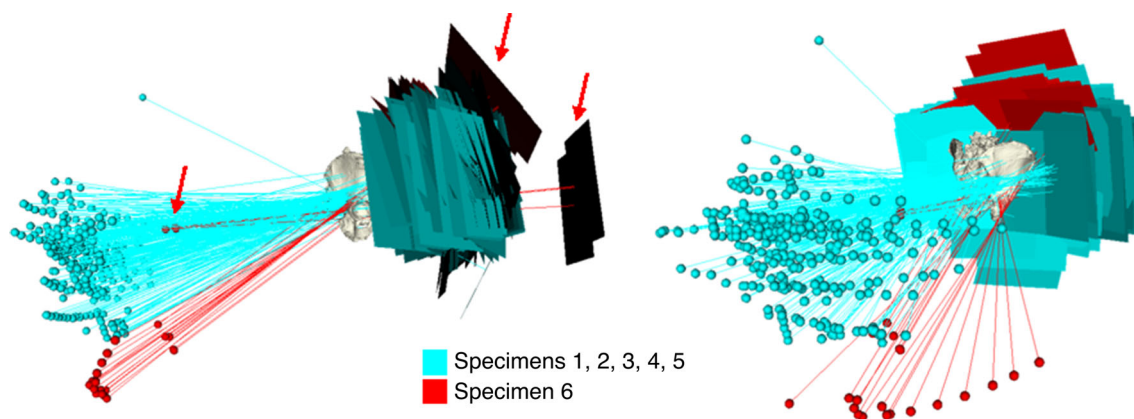


Fig. 5 A visualization of all ground truth projection geometries using the APP as the world coordinate frame. Each sphere represents a position of the X-ray source, each square represents the position of the X-ray detector, and each line connects the X-ray source to the principal point

on the detector. Red arrows highlight difficult to see geometries of specimen 6. Most of the poses are contained within a 60° range of C-arm orbital rotations and a 30° range of pelvic tilts

Table 2 shows that the mean and standard deviation of method 3's pelvis translation errors were both larger than those of method 2 by approximately 0.5 mm. It is possible that some cases of inaccurate landmark point estimates may have limited the influence of image similarities, resulting in the larger errors for method 3. We believe that this issue may be overcome by replacing the landmark reprojection distances of (2) with the heatmap values at each projected 3D landmark. Since the heatmaps encode landmark localization uncertainties, this modification should reduce the penalty of reprojection distances for inaccurately estimated landmarks.

In comparison with the femur, the pelvis is a larger, more complex shape, which occupies larger regions of the fluoroscopic views. The pelvis' fluoroscopic appearance is thus more sensitive to rotational changes than that of the femur, which is consistent with the smaller rotation errors observed with registrations of the pelvis and shown in Table 2.

For the intact hip anatomy that is considered in this paper, the connective tissues joining each femoral head to the acetabular regions of the pelvis cause each femur and the pelvis to mostly translate together. As a result, the relative pose of a femur with respect to the pelvis contains very little translation. This prior knowledge is incorporated into the registration strategies and causes mostly small femoral translations to be reported, which results in the small translation errors for femur registrations listed in Table 2. The influence of pelvis translation errors on these estimates was minimal, since nearly all of the pelvis translation errors were found in the depth direction, which has a minor impact on single-view appearance.

The performance of our approach degrades as images collected with C-arm poses not found in the training data set are processed. This is highlighted in Table 2, showing the poor performance of specimen 6 compared to all other speci-

mens. When testing on spec. 6, average landmark localization error was 10.0 mm, with a false negative rate of 30% and successful pelvis registration rate of 29% for both methods 2 and 3. Figure 5 shows a visualization of all 366 ground truth projection geometries. The geometries associated with spec. 6 are clearly collected at different C-arm poses than those used for training the network tested on spec. 6. Three examples of spec. 6 are shown in Fig. 4b–d. This limitation may be overcome by collecting more fluoroscopy data for training. However, it is conceivable that some C-arm poses encountered during testing will still be absent during training. By augmenting actual fluoroscopy with realistic synthetic fluoroscopy [30] during training, we believe that quality performance at these “missing” poses may be achieved.

Large variations in dice scores may result from the slight mislabeling of a narrow structure, such as the ilium in some views. These small segmentation errors are not expected to negatively impact registration performance, as the estimated labels are used to weight the contribution of local, overlapping, patches to the image similarity term. Therefore, it is not surprising that a weak correlation between dice scores and registration errors was indicated by the Spearman rank coefficient values.

The automated annotation and registration techniques proposed in this paper could streamline intraoperative workflows related to intact hip anatomy, such as osteotomy planning [8], robotic drilling [31], and 3D reconstructions of bone [32] or implanted tracking fiducials [33]. Extending the annotation component to label additional objects, such as surgical instruments, artificial implants, and bone fragments could enable automatic registration of surgically modified hip anatomy and is the subject of future research. Although all possible patches are currently evaluated when computing image similarities, the semantic labeling of fluoroscopy should enable

much smaller subsets of anatomically relevant patches to be used during the registration. By only rendering the DRR pixels which intersect this subset of patches, significant reductions in registration runtimes may be possible.

In conclusion, this paper has demonstrated that small annotated datasets of actual hip fluoroscopy may be used to train CNN models capable of state-of-the-art segmentation and landmark localization results. Furthermore, we have shown that by coupling the automatic annotations produced by the CNN models with the image intensities used during 2D/3D registrations, robustness against poor initializations is possible. This is a clinically relevant result, as this robustness removes the need for manual initialization and allows navigation to be more naturally integrated into existing surgical workflows. To our knowledge, the dataset presented in this paper is the first annotated dataset of *actual* hip fluoroscopy, consisting of individual bone segmentations and anatomical landmark locations. We have also made this dataset publicly available.² Creation of the precise labels found in the training dataset was made possible by extending existing 2D/3D registration technology into a new *offline* and semi-automatic annotation pipeline. We believe this ground truth labeling method will translate to fluoroscopy of other anatomy and enable machine learning applications in other specialties.

Acknowledgements We thank Mr. Demetrios Boston for assisting with the cadaveric data acquisition. This research was supported by NIH/NIBIB grants R01EB006839, R21EB020113, Johns Hopkins University Internal Funds, and a Johns Hopkins University Applied Physics Laboratory (Grant No. FNACCX24) Graduate Student Fellowship. Part of this research project was conducted using computational resources at the Maryland Advanced Research Computing Center (MARCC).

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Ethical approval This article does not contain any studies with human participants performed by any of the authors.

Informed consent This article does not contain patient data.

References

1. Woerner M, Sendtner E, Springorum R, Craiovan B, Worlicek M, Renkawitz T, Grifka J, Weber M (2016) Visual intraoperative estimation of cup and stem position is not reliable in minimally invasive hip arthroplasty. *Acta Orthop* 87(3):225–230
2. Slotkin EM, Patel PD, Suarez JC (2015) Accuracy of fluoroscopic guided acetabular component positioning during direct anterior total hip arthroplasty. *J Arthroplasty* 30(9):102–106
3. Troelsen A (2009) Surgical advances in periacetabular osteotomy for treatment of hip dysplasia in adults. *Acta Orthop* 80(sup332):1–33
4. Kelley TC, Swank ML (2009) Role of navigation in total hip arthroplasty. *J Bone Joint Surg Am* 91(Supplement):153–158
5. Belei P, Skwara A, Fuente MDL, Schkommodau E, Fuchs S, Wirtz DC, Kämper C, Radermacher K (2007) Fluoroscopic navigation system for hip surface replacement. *Comput Aided Surg* 12(3):160–167
6. Malan DF, van der Walt SJ, Raidou RG, van den Berg B, Stoel BC, Botha CP, Nelissen RG, Valstar ER (2016) A fluoroscopy-based planning and guidance software tool for minimally invasive hip refixation by cement injection. *Int J Comput Assist Radiol Surg* 11(2):281–296
7. Grupp RB, Hegeman R, Murphy R, Alexander C, Otake Y, McArthur B, Armand M, Taylor RH (2020) Pose estimation of periacetabular osteotomy fragments with intraoperative X-ray navigation. *IEEE Trans Biomed Eng* 67(2):441–452
8. Gottschling H, Roth M, Schweikard A, Burgkart R (2005) Intraoperative, fluoroscopy-based planning for complex osteotomies of the proximal femur. *Int J Med Robot Comput Assist Surg* 1(3):67–73
9. Markelj P, Tomaževič D, Likar B, Pernuš F (2012) A review of 3D/2D registration methods for image-guided interventions. *Med Image Anal* 16(3):642–661
10. Newell A, Yang K, Deng J (2016) Stacked hourglass networks for human pose estimation. In: *Proceedings of European conference on computer vision*. Springer, pp 483–499
11. Shelhamer E, Long J, Darrell T (2017) Fully convolutional networks for semantic segmentation. *IEEE Trans Pattern Anal Mach Intell* 4:640–651
12. Otake Y, Takao M, Yokota F, Fukuda N, Uemura K, Sugano N, Sato Y (2018) Construction and application of large-scale image database in orthopedic surgery. In: *Computer assisted orthopaedic surgery for hip and knee*. Springer, pp 191–197
13. Miao S, Wang ZJ, Liao R (2016) A CNN regression approach for real-time 2D/3D registration. *IEEE Trans Med Imag* 35(5):1352–1363
14. Ambrosini P, Ruijters D, Niessen WJ, Moelker A, van Walsum T (2017) Fully automatic and real-time catheter segmentation in X-ray fluoroscopy. In: *Proceedings of medical image computing and computer-assisted intervention*. Springer, pp 577–585
15. Breininger K, Albarqouni S, Kurzendorfer T, Pfister M, Kowarschik M, Maier A (2018) Intraoperative stent segmentation in X-ray fluoroscopy for endovascular aortic repair. *Int J Comput Assist Radiol Surg* 13(8):1221–1231
16. Gao C, Unberath M, Taylor R, Armand M (2019) Localizing dexterous surgical tools in X-ray for image-based navigation. *arXiv preprint arXiv:1901.06672*
17. Laina I, Rieke N, Rupprecht C, Vizcaíno JP, Eslami A, Tombari F, Navab N (2017) Concurrent segmentation and localization for tracking of surgical instruments. In: *Proceedings of medical image computing and computer-assisted intervention*. Springer, pp 664–672
18. Kordon F, Fischer P, Privalov M, Swartman B, Schnetzke M, Franke J, Lasowski R, Maier A, Kunze H (2019) Multi-task localization and segmentation for X-ray guided planning in knee surgery. In: *Proceedings of medical image computing and computer-assisted intervention*. Springer, pp 622–630
19. Bier B, Goldmann F, Zaech JN, Fotouhi J, Hegeman R, Grupp R, Armand M, Osgood G, Navab N, Maier A, Unberath M (2019) Learning to detect anatomical landmarks of the pelvis in X-rays from arbitrary views. *Int J Comput Assist Radiol Surg* 14(9):1463–1473
20. Esteban J, Grimm M, Unberath M, Zahnd G, Navab N (2019) Towards fully automatic X-ray to CT registration. In: *Proceedings*

² <https://github.com/rg2/DeepFluoroLabeling-IPCAI2020>

- of medical image computing and computer-assisted intervention. Springer, pp 631–639
21. Nikou C, Jaramaz B, DiGioia AM, Levison TJ (2000) Description of anatomic coordinate systems and rationale for use in an image-guided total hip replacement system. In: Proceedings of medical image computing and computer-assisted intervention, pp 1188–1194
 22. Grupp RB, Armand M, Taylor RH (2018) Patch-based image similarity for intraoperative 2D/3D pelvis registration during periacetabular osteotomy. In: Proceedings of international workshop on clinical image-based procedures. Springer, pp 153–163
 23. Storn R, Price K (1997) Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *J Glob Optim* 11(4):341–359
 24. Shi Y, Eberhart R (1998) A modified particle swarm optimizer. In: Proceedings of IEEE international conference on evolutionary computation. IEEE, pp 69–73
 25. Hansen N, Ostermeier A (2001) Completely derandomized self-adaptation in evolution strategies. *Evolut Comput* 9(2):159–195
 26. Powell MJ (2009) The BOBYQA algorithm for bound constrained optimization without derivatives. Cambridge NA Report NA2009/06. University of Cambridge, Cambridge
 27. Ronneberger O, Fischer P, Brox T (2015) U-net: convolutional networks for biomedical image segmentation. In: Proceedings of medical image computing and computer-assisted intervention. Springer, pp 234–241
 28. Milletari F, Navab N, Ahmadi SA (2016) V-net: fully convolutional neural networks for volumetric medical image segmentation. In: Proceedings of 2016 fourth international conference on 3D vision (3DV). IEEE, pp 565–571
 29. Hartley R, Zisserman A (2003) Multiple view geometry in computer vision. Cambridge University Press, Cambridge
 30. Unberath M, Zaech JN, Gao C, Bier B, Goldmann F, Lee SC, Fotouhi J, Taylor R, Armand M, Navab N (2019) Enabling machine learning in X-ray-based procedures via realistic simulation of image formation. *Int J Comput Assist Radiol Surg* 14(9):1517–1528
 31. Gao C, Grupp RB, Unberath M, Taylor RH, Armand M (2020) Fiducial-free 2D/3D registration of the proximal femur for robot-assisted femoroplasty. In: Proceedings of SPIE, vol 11315, p 113151C
 32. Reyneke CJF, Lüthi M, Burdin V, Douglas TS, Vetter T, Mutsvangwa TE (2018) Review of 2-D/3-D reconstruction using statistical shape and intensity models and X-ray image synthesis: toward a unified framework. *IEEE Rev Biomed Eng* 12:269–286
 33. Grupp R, Murphy R, Hegeman R, Alexander C, Unberath M, Otake Y, McArthur B, Armand M, Taylor R (2019) Fast and automatic periacetabular osteotomy fragment pose estimation using intraoperatively implanted fiducials and single-view fluoroscopy. arXiv preprint [arXiv:1910.10187](https://arxiv.org/abs/1910.10187)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.