





# RayEmb: Arbitrary Landmark Detection in X-Ray Images Using Ray Embedding Subspace

Pragyan Shrestha<sup>1</sup>, Chun Xie<sup>1</sup>, Yuichi Yoshii<sup>2</sup>, and Itaru Kitahara<sup>1</sup>

<sup>1</sup> University of Tsukuba, Tsukuba, Ibaraki, Japan  
shrestha.pragyan@image.iit.tsukuba.ac.jp  
{xiechun,kitahara}@ccs.tsukuba.ac.jp

<sup>2</sup> Tokyo Medical University, Ami, Ibaraki, Japan  
yyoshii@tokyo-med.ac.jp

**Abstract.** Intra-operative 2D-3D registration of X-ray images with pre-operatively acquired CT scans is a crucial procedure in orthopedic surgeries. Anatomical landmarks pre-annotated in the CT volume can be detected in X-ray images to establish 2D-3D correspondences, which are then utilized for registration. However, registration often fails in certain view angles due to poor landmark visibility. We propose a novel method to address this issue by detecting arbitrary landmark points in X-ray images. Our approach represents 3D points as distinct subspaces, formed by feature vectors (referred to as ray embeddings) corresponding to intersecting rays. Establishing 2D-3D correspondences then becomes a task of finding ray embeddings that are close to a given subspace, essentially performing an intersection test. Unlike conventional methods for landmark estimation, our approach eliminates the need for manually annotating fixed landmarks. We trained our model using the synthetic images generated from CTPelvic1K CLINIC dataset, which contains 103 CT volumes, and evaluated it on the DeepFluoro dataset, comprising real X-ray images. Experimental results demonstrate the superiority of our method over conventional methods. The code is available at <https://github.com/Pragyanstha/rayemb>.

**Keywords:** 2D-3D Registration · Landmark Detection · Subspace

## 1 Introduction

Intra-operative 2D-3D registration of X-ray images with pre-operatively obtained CT scans is a widely used technique in orthopedic surgeries. In a typical surgical setting, a fluoroscopy device captures an X-ray image of the target structure within the patient, who is positioned on the operating table. The X-ray image is a projection of the target structure based on the geometric setup of the device, making it challenging for surgeons to visualize the actual 3D anatomy. 2D-3D registration enables overlaying the 3D model, derived from the pre-operative CT scan, onto the X-ray image. Additionally, surgical planning data, such as the placement of implants and pedicle screws, can be visualized similarly through registration. In clinical practice, procedures like total hip arthroplasty, total



sponding points on the surface [14, 15, 42]. Similarly, template matching based approaches [18, 26] work based on similarity of patches across the template and query image. However, these underlying assumptions do not hold in transmissive imaging technique such as X-ray images, making 2D-3D registration a more challenging task.

Recent advances in 2D-3D registration have shown significant progress by adopting learning based methods. One promising approach involves establishing 2D-3D correspondences by localizing predefined landmarks in the X-ray images [2, 6, 11]. These methods train a model to predict a heatmap of landmarks using keypoint-annotated images. During the registration process, the landmark annotations in the CT volume, combined with the predicted 2D heatmap, are used to establish the 2D-3D correspondences, which then serve as input for perspective-n-point (PnP) algorithms with random sample and consensus (RANSAC) to filter outliers. Although this strategy achieves high registration accuracy on successfully registered images, it often fails when landmark visibility is limited due to partial views of the structure. Additionally, annotating landmarks in the CT volume is not trivial, as it requires domain expertise to precisely identify clinically meaningful anatomical landmarks specific to the target structure. The reliance on these annotated landmarks during inference further complicates use in emergency situations, where a quick setup is essential.

To overcome the limitations of current landmark-based registration and enhance applicability across patients, we propose a novel method for localizing the 2D projections of arbitrary 3D points within a specified CT volume. By leveraging the transmissive nature of X-ray images, our approach hypothesize that a patch in an X-ray image contains information about the corresponding ray and its attenuation distribution along that ray. Thus, a 3D point within a CT volume can be uniquely represented by a collection of rays passing through it. In other words, rays intersecting this point can be distinguished by analyzing the information contained in the associated pixels. Specifically, we train an encoder to learn ray embeddings such that embeddings of intersecting rays span a distinctive subspace, which can be used to identify whether a query ray embedding intersects at the same point. This differentiation from other rays not intersecting the point is mathematically realized using subspace and the similarity of a vector with its projection onto the subspace. While some previous works with similar objective have focused on establishing 3D-3D or 2D-2D correspondences [21, 40]. To the best of our knowledge, this is the first attempt at establishing 2D-3D correspondences of arbitrary landmark points.

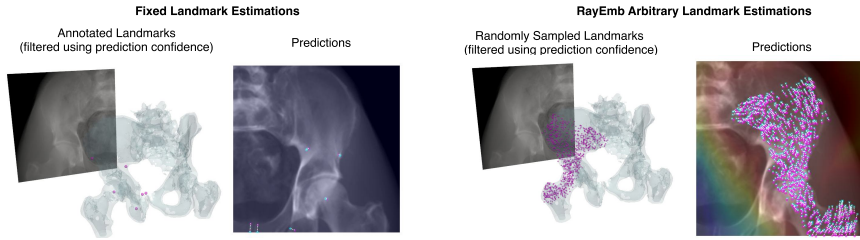
In summary, our main contributions are as follows.

- A novel method for establishing arbitrary landmark correspondences within the field of view, enabling 2D-3D registration even in difficult views with partial visibility.
- A sampling strategy for finding good 2D-3D correspondences through test time augmentation of selected templates.
- A self-supervised training strategy utilizing synthetic data generation, where landmark annotation by domain expert is not required.

## 2 Related Works

**6DoF Object Pose Estimation.** Various approaches have been proposed for solving 6DoF object pose estimation. Direct methods such as PoseNet [16], PoseCNN [39], Deep-6DPose [5], and EfficientPose [4] regress the rotation and translation vectors of detected objects. In contrast, methods like DeepIM [19], CosyPose [18], and MegaPose [17] estimate the pose updates iteratively by comparing the current rendering with the observed image. Indirect methods solve a proxy problem by finding correspondences, which are then used in PnP-RANSAC framework to derive the final pose. OSOP [35] establishes 2D-3D correspondences indirectly by find a matching template, with known pose, and establishing 2D-2D correspondences with it. Similarly, Gigapose [26] first finds the out-of-plane rotation by matching templates and estimates the remaining 4-DoF using patch correspondences. BB8 [31] locates the 2D projections of the corners of 3D bounding boxes to establish correspondences. To address the challenge of truncated objects where keypoints may lie outside the image, PVNet [30] regresses vectors that point towards the keypoints, determining the final location through a voting mechanism. In dense prediction-based methods, Pix2Pose [29] regresses the 3D coordinates of corresponding pixels to establish dense 2D-3D correspondences. Inspired by DensePose [13], which maps all pixels of a human RGB image to the 3D surface of the human body, DPOD [42] adopts a different approach by estimating UV texture maps of the object. Similarly, EPOS [15] predicts multiple surface fragment coordinates for a pixel, allowing for the modeling of pose ambiguities due to partial symmetries in the object. Building on EPOS, Surfemb [14] estimates a continuous distribution over the object surface and uses this distribution for pose scoring and refinement. GDRNet [38] employs a hybrid approach by regressing dense geometric features for establishing 2D-3D correspondences and uses patch-PnP network for direct pose regression, rather than the conventional PnP-RANSAC framework.

**Learning based 2D-3D Registration.** Traditionally, 2D-3D registration in medical imaging have incorporated two main approaches in clinical practice: image based optimization of digitally reconstructed radiographs (DRRs) from CT volumes against X-ray images [25], and the establishment of 2D-3D correspondences through pre-embedded fiducial landmarks visible in X-ray images [8]. Recent advancements in learning-based methods have aimed to overcome the limitations inherent in these traditional approaches. Some works have focused on extending the capture range of optimization schemes by iteratively updating pose estimates [9, 12]. These learning-based methods, which involve comparing DRRs with actual X-ray images to estimate pose differences, achieve better registration accuracy through further refinement via image similarity-based optimization. However, initial estimates often remain suboptimal compared to those derived from landmark estimation methods, leading to local optima during the image based optimization. While landmark-based registration [2, 11] typically provides more accurate initial estimates, it faces several challenges. These include the requirement for 3D landmark annotation during inference and failures



**Fig. 2:** Comparison of landmark detection results between conventional fixed landmark estimation and our arbitrary landmark estimation method. The 3D landmarks are shown in magenta on the left, while the estimated 2D landmarks are displayed in cyan and the ground truth in magenta on the right. Our method can generate a large number of corresponding pairs of 3D landmarks and 2D projections, whereas the fixed landmark estimation approach is limited to the available 2D projections and does not provide 3D landmark information.

in views with insufficient landmark visibility. One approach to address these issues involves backprojecting 2D landmark estimates from multi-view DRRs and fine-tuning the landmark estimation model with synthetic images [6]. Recent study have shown the effectiveness of training deep models with synthetic images for vision tasks related to X-ray images [7]. Additionally, scene coordinates can be directly estimated from X-ray images [34], eliminating the need for 3D landmark annotation. Although effective, these patient-specific methods are less suitable for urgent scenarios, such as emergency surgeries, due to their preparatory requirements.

### 3 Methods

At a high level, our proposed method involves arbitrary landmark estimation followed by 2D-3D registration. Fig. 1 illustrates an overview of our landmark estimation process. Once the correspondences are established, the perspective-n-point algorithm [24] with marginalizing sample and consensus (MAGSAC) [1] is used to obtain the initial pose estimate. This estimate serves as initialization for DiffDRR [10], a gradient-based optimization refinement module. The core idea is to provide pose estimator with large amount of corresponding pairs of 3D landmarks and its 2D projections. Fig. 2 compares this approach to the conventional fixed landmark estimation method.

#### 3.1 Problem Formulation

The X-ray image formation is typically described using Beer-Lambert’s Law [36].

$$I(\mathbf{d}; \mathbf{T}) = I_0 \exp\left(-\int_0^x \mu(t\mathbf{d}; \mathbf{T}) dt\right) \quad (1)$$

where  $I(\mathbf{d})$  is the intensity of the X-ray after passing through the material.  $I_0$  is the initial intensity of the X-ray before entering the material,  $\mu$  is the linear attenuation coefficient of the volume,  $\mathbf{d}$  is the direction vector towards the detector pixel in camera space, and  $\mathbf{T}$  represents the pose of the volume. Given a target X-ray image  $I_{target}$ , the goal is to find the optimal pose  $\mathbf{T}_{optim}$  that minimizes the discrepancy between the rendered image and the target image.

$$\mathbf{T}_{optim} = \arg \min_{\mathbf{T}} d(I_{target}, I(\mathbf{T})) \quad (2)$$

where  $d$  measures the distance between the two images. In practice, effectively solving Eq. (2) requires an accurate initial pose estimate  $T$  to ensure convergence to the correct solution. We propose to obtain a suitable initialization for this pose by establishing 2D-3D correspondences. However, as evident from Eq. (1), establishing direct correspondence between a pixel and a 3D point is challenging, since the pixel intensity represents the integral of attenuation coefficients along the ray path. In other words, a many-to-one correspondence must be established between the points along the ray and the pixel.

### 3.2 Arbitrary Landmark Estimation

Estimating correspondences for arbitrary landmark points between 2D and 3D representations poses significant challenges due to the mismatch in spatial dimensions and the transmissive properties of X-rays, which can associate multiple 3D points with a single 2D projection. To address this, we employ a pixel-wise feature extractor and pre-render DRR templates to represent a 3D point as a subspace. The pixel-wise features are referred to as ray embeddings, as we associate the feature vector to its back-projected ray. Since a 3D point can be represented by a collection of intersecting rays, a set of ray embeddings can be associated with a 3D point if their underlying rays intersect. The 2D projection of a 3D point can then be identified by evaluating the closeness of the ray embeddings in the query image to the subspace representing the 3D point. The following sections will provide a detailed explanation of the main components of the proposed method.

**Ray Embeddings.** The input query X-ray image and the pre-rendered templates are both fed into an encoder for obtaining pixel wise ray embedding vectors. Formally, let  $I_q(\mathbf{x})$  be the input query image and  $I'_t(\mathbf{x})$  be the template images, where  $t \in \{1, 2, 3, \dots, N\}$ , with  $N$  being the number of templates. The ray embedding for each image point  $\mathbf{x}$  in query input image is defined as  $\mathbf{e}(\mathbf{x}) = E(I_q(\mathbf{x}); \mathbf{w})$ , where  $E : \mathbb{R} \rightarrow \mathbb{R}^d$ . Similarly, the ray embedding in the template image is given by  $\mathbf{e}'_t(\mathbf{x}) = E(I'_t(\mathbf{x}); \mathbf{w})$ . It should be noted that both the query image and template images are processed using the same encoder to ensure that the resulting embedding vectors reside in the same contextual space.

**Subspace Representation of 3D Landmarks.** Let  $\mathbf{X}$  represent a sampled 3D landmark inside the CT volume (as shown in the left section of Fig. 1). Given the known camera transformation matrices  $\mathbf{T}'_t$  for each template image, we can compute the projection of  $\mathbf{X}$  onto the template image planes as  $\mathbf{x}'_t = \pi_{\mathbf{K}}(\mathbf{X}; \mathbf{T}'_t)$ , where  $\pi_{\mathbf{K}}(\cdot)$  is the camera projection operator. For clarity, the ray embeddings at these projected points are denoted as follows:

$$\mathbf{e}'_t = E(I'_t(\mathbf{x}'_t); \mathbf{w}), \quad t \in \{1, 2, 3, \dots, N\} \quad (3)$$

The embedding vectors are then stacked to form a transformation matrix  $\mathbf{F}$ , whose column space is spanned by the embedding vectors of rays that intersect at  $\mathbf{X}$  in 3D space.

$$\mathbf{F} = (\mathbf{e}'_1, \mathbf{e}'_2, \dots, \mathbf{e}'_N) \quad (4)$$

The orthogonal projection  $\mathbf{P}$  onto the subspace spanned by the embedding vectors can be calculated by right multiplying  $\mathbf{F}$  with its Moore-Penrose inverse  $\mathbf{F}^+ = \mathbf{V}\mathbf{\Sigma}^+\mathbf{U}^T$ , where  $\mathbf{U}$ ,  $\mathbf{\Sigma}$ , and  $\mathbf{V}$  are the matrices containing the left singular vectors, the singular values, and the right singular vectors of  $\mathbf{F}$ , respectively.

$$\mathbf{P} = \mathbf{F}\mathbf{F}^+ = \mathbf{U}\mathbf{\Sigma}\mathbf{\Sigma}^+\mathbf{U}^T \quad (5)$$

The projection matrix  $\mathbf{P}$  now represents an  $N$ -dimensional subspace within a  $D$ -dimensional embedding space. We associate this projection matrix with the 3D point  $\mathbf{X}$ .

**Corresponding 2D Landmark Estimation.** During inference, ray embeddings of the query input image are available for all spatial grid points on the detector plane. For a 3D point  $\mathbf{X}$ , we calculate the cosine similarity between the projection of each ray embedding onto the subspace and the embedding itself:

$$\text{sim}(\mathbf{P}, \mathbf{x}) = \frac{\mathbf{e}^T(\mathbf{x})\mathbf{P}\mathbf{e}(\mathbf{x})}{|\mathbf{e}^T(\mathbf{x})| |\mathbf{P}\mathbf{e}(\mathbf{x})|} \quad (6)$$

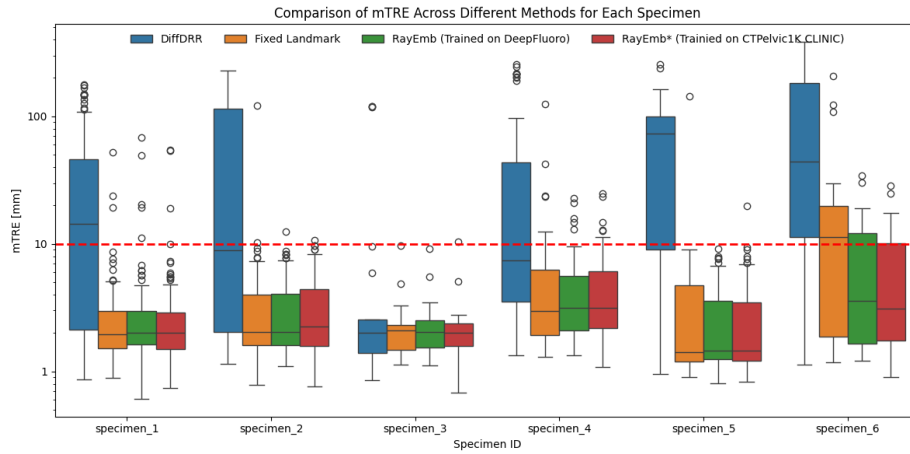
where  $\mathbf{x}$  represents 2D points on the detector plane. The corresponding 2D point  $\hat{\mathbf{x}}$  for the given 3D point  $\mathbf{X}$ , associated with subspace  $\mathbf{P}$ , is identified by finding the location with maximum similarity value across the detector plane.

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \text{sim}(\mathbf{P}, \mathbf{x}) \quad (7)$$

**Contrastive Learning.** The goal is to train the encoder to learn embeddings such that rays intersecting at the same 3D point are mapped to a distinct subspace. During training, we have ray embedding vectors from the templates that intersect in 3D space, along with a positive sample  $\mathbf{x}^+$  from the query image that also intersects the same 3D point. All other embedding vectors are treated as negative samples  $\mathbf{x}^-$ . The InfoNCE loss [27] is employed to achieve this:

$$\mathcal{L}(\mathbf{e}, \mathbf{e}'_t) = -\log \frac{\exp(\text{sim}(\mathbf{P}, \mathbf{x}^+)/\tau)}{\sum_{\mathbf{x}^- \in \mathcal{N}} \exp(\text{sim}(\mathbf{P}, \mathbf{x}^-)/\tau)} \quad (8)$$

Here,  $\tau$  is the temperature constant, and  $\mathcal{N}$  includes the positive sample as well.



**Fig. 3:** Box plots comparing the mean Target Registration Errors (mTRE) for six specimens using three different registration methods: DiffDRR, Fixed Landmark, and our proposed method, RayEmb (trained on DeepFluoro) and RayEmb\* (trained on CTPelvic1K CLINIC). A red dashed line at 10 mm indicates the critical threshold for acceptable performance limit.

### 3.3 2D-3D Registration

Using the 2D-3D correspondences and similarity scores obtained from Equation 6, we adopt MAGSAC [1], a variant of RANSAC that does not require an explicitly set inlier threshold, to compute the initial pose estimate. Although MAGSAC is capable of rejecting outlier correspondences that arise from points outside the field of view, we pre-filter the correspondences by selecting the top- $k$  points with the highest similarity values. This filtering step ensures that the resulting correspondences are mostly within the field of view, as shown in Fig. 2. To further improve robustness, we apply test-time augmentation during subspace generation by randomly choosing the templates  $n$ -times and selecting the 2D projections with the maximum similarity response. Following this, a gradient-based optimization of the pose is performed using multi-scale normalized cross-correlation with differentiable rendering, in line with the approach of DiffPose [9].

## 4 Experiments

### 4.1 Datasets

Following prior research, we utilized the DeepFluoro dataset [11] to evaluate our proposed method. Although we generated synthetic images using the CT volumes for training, we evaluated the model on real X-rays from this dataset. Additionally, we trained the model on the CTPelvic1K CLINIC [22] dataset to assess its generalization capability when tested on real X-ray image from DeepFluoro dataset.



**DeepFluoro.** This dataset comprises CT scans from six specimens, totaling 366 real X-ray images registered semi-automatically. For each specimen, 14 anatomical landmarks are labeled in the CT volume and their corresponding projections on the real X-ray images. The X-ray images were collected using a Siemens CIOS Fusion mobile C-arm device with  $30 \times 30 \text{ cm}^2$  detector size,  $1536 \times 1536$  pixels with 0.194 mm pixel spacing. The 3D anatomical landmarks were used to train the baseline model; however, these landmarks were not required, and therefore not used, during training or inference with our proposed method. We used DiffDRR [9] to generate 10,000 synthetic images for each specimen by sampling random poses, simulating the C-arm CIOS Fusion device. Synthetic images from specimens 1 to 5 were used for training and validation, while specimen 6 was reserved exclusively for testing.

**CTPelvic1K CLINIC.** To evaluate the generalizability of the proposed method across different patient CT scans, we also trained our model using the CTPelvic1K CLINIC dataset, which is a subset comprising 103 CT scans. This dataset is part of a larger collection originally containing approximately 1184 CT volumes with varying resolutions and a diverse range of subjects, each annotated with pelvic segmentation labels. For this study, we specifically selected the CLINIC subset, which provides a balanced sample size, avoiding extremes of very large or very small numbers of CT volumes, thereby offering a representative sample for effective training. Similar to the DeepFluoro dataset, we used DiffDRR to generate 1000 synthetic images per CT scan, doubling the total number of images than that of the DeepFluoro dataset. The model trained on the CTPelvic1K CLINIC dataset was then evaluated on real X-ray images from the DeepFluoro dataset.

## 4.2 Baselines and Evaluation Metrics

We compare the RayEmb method against a conventional approach that identifies fixed spatial landmarks, referred to here as Fixed Landmark. The Fixed Landmark method estimates heatmaps of 14 anatomical landmarks using a U-Net-based network, following the architecture described in [11] but without the segmentation module, instead regressing heatmaps at the decoder output. Since both methods incorporate pose refinement using DiffDRR, we also compare them against DiffDRR initialized from a standardized anterior-posterior (AP) view to ensure consistency across evaluations. We report two key metrics for registration: mean target registration error (mTRE) and mean projected distance (mPD), defined as follows:

$$\text{mTRE}(\mathbf{T}, \hat{\mathbf{T}}) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{T}\mathbf{X}_i - \hat{\mathbf{T}}\mathbf{X}_i\| \quad (9)$$

$$\text{mPD}(\mathbf{T}, \hat{\mathbf{T}}) = \frac{1}{N} \sum_{i=1}^N \|\pi_{\mathbf{K}}(\mathbf{X}_i; \mathbf{T}) - \pi_{\mathbf{K}}(\mathbf{X}_i; \hat{\mathbf{T}})\| \quad (10)$$

**Table 1:** The mTRE for each 25th, 50th, and 95th percentile of the DeepFluoro dataset, along with GFR at 10mm and 5mm in percentage, are reported. The proposed method trained on DeepFluoro as well as CTPelvic1K achieves the lowest errors in all metrics for specimen 6, which includes challenging views.

Specimen	Method	25 <sup>th</sup>	50 <sup>th</sup>	95 <sup>th</sup>	GFR@10↓	GFR@5↓
1	DiffDRR	2.125	14.259	147.349	54.054	63.964
	Fixed Landmark	1.520	<b>1.960</b>	6.701	<b>2.703</b>	<b>9.009</b>
	RayEmb	1.625	2.010	<b>6.520</b>	4.505	9.910
	RayEmb*	<b>1.502</b>	2.011	6.535	<b>2.703</b>	10.811
2	DiffDRR	2.049	8.879	203.835	48.077	59.615
	Fixed Landmark	<b>1.602</b>	<b>2.027</b>	<b>7.660</b>	1.923	19.231
	RayEmb	1.604	2.049	7.772	<b>0.962</b>	<b>17.308</b>
	RayEmb*	1.578	2.264	8.281	<b>0.962</b>	20.192
3	DiffDRR	<b>1.402</b>	2.007	119.635	12.500	20.833
	Fixed Landmark	1.474	2.100	<b>4.628</b>	<b>0.000</b>	<b>4.167</b>
	RayEmb	1.549	2.053	5.239	<b>0.000</b>	8.333
	RayEmb*	1.583	<b>2.019</b>	4.737	4.167	8.333
4	DiffDRR	3.529	7.402	216.490	41.667	66.667
	Fixed Landmark	<b>1.937</b>	<b>2.965</b>	23.770	14.583	35.417
	RayEmb	2.094	3.157	15.375	<b>10.417</b>	<b>29.167</b>
	RayEmb*	2.186	3.167	<b>14.103</b>	12.500	33.333
5	DiffDRR	9.062	73.597	140.134	74.545	76.364
	Fixed Landmark	<b>1.190</b>	<b>1.414</b>	8.172	1.818	25.455
	RayEmb	1.248	1.447	<b>7.633</b>	<b>0.000</b>	<b>18.182</b>
	RayEmb*	1.211	1.452	8.313	1.818	21.818
6	DiffDRR	11.263	43.853	375.553	75.000	79.167
	Fixed Landmark	1.875	11.250	121.944	50.000	62.500
	RayEmb	<b>1.647</b>	3.580	28.496	33.333	<b>45.833</b>
	RayEmb*	1.754	<b>3.090</b>	<b>23.832</b>	<b>25.000</b>	<b>45.833</b>

where  $\mathbf{X}_i$  is the 3D landmark,  $\mathbf{TX}_i$  represents landmark in camera space, and  $\pi_{\mathbf{K}}(\mathbf{X}_i; \mathbf{T})$  denotes its projection onto the detector plane. Both mTRE and mPD are measured in millimeters. Additionally, we report the gross failure rate (GFR), defined as the proportion of images that could not be registered within a specified tolerance threshold (10mm or 5mm), for each specimen.

### 4.3 Implementation Details

**Training.** We use a DINOv2 [28] pretrained Vision Transformer encoder with a fully connected layer to convert 768-dimensional descriptors into 32-dimensional embedding vectors. Each input image is resized to  $224 \times 224$  pixels. A total of 324 template DRRs are generated from equally spaced pose samples, covering 45 degrees in the LAO/RAO directions and 22.5 degrees in the CRA/CAU directions, with 18 steps taken in each direction. During training, 4 templates are randomly selected to generate subspaces for the given sampling points. PyTorch is used to compute the backpropagatable pseudo-inverse, and gradient clipping with a maximum value of 10.0 is applied. The model is trained using the Adam optimizer with a learning rate of  $1e-4$  for 50 epochs on an RTX 3090 GPU Ti, with a batch size of 8, 40 3D sampling points, and a temperature parameter of

**Table 2:** The mPD for each 25th, 50th, and 95th percentile of the DeepFluoro dataset, along with GFR at 10mm and 5mm threshold in percentage, are reported. Best values highlighted for each metric across specimens.


Specimen	Method	25 <sup>th</sup>	50 <sup>th</sup>	95 <sup>th</sup>	GFR@10↓	GFR@5↓
1	DiffDRR	1.681	5.476	172.835	45.045	52.252
	Fixed Landmark	1.417	1.696	<b>2.598</b>	<b>1.802</b>	<b>1.802</b>
	RayEmb	1.420	1.722	2.676	<b>1.802</b>	3.604
	RayEmb*	<b>1.340</b>	<b>1.676</b>	<b>2.636</b>	<b>1.802</b>	2.703
2	DiffDRR	1.868	2.877	240.662	45.192	47.115
	Fixed Landmark	1.589	1.808	<b>2.692</b>	0.962	0.962
	RayEmb	1.610	1.817	2.741	<b>0.000</b>	<b>0.000</b>
	RayEmb*	<b>1.515</b>	<b>1.793</b>	2.852	<b>0.000</b>	<b>0.000</b>
3	DiffDRR	1.311	1.722	166.758	12.500	12.500
	Fixed Landmark	1.653	1.739	2.962	<b>0.000</b>	<b>0.000</b>
	RayEmb	1.523	<b>1.824</b>	<b>2.950</b>	<b>0.000</b>	<b>0.000</b>
	RayEmb*	<b>1.393</b>	1.846	3.313	<b>0.000</b>	<b>0.000</b>
4	DiffDRR	2.984	5.100	290.629	35.417	50.000
	Fixed Landmark	2.146	<b>2.519</b>	7.091	2.083	10.417
	RayEmb	<b>2.046</b>	2.572	<b>4.755</b>	<b>0.000</b>	<b>4.167</b>
	RayEmb*	2.314	2.783	6.662	<b>0.000</b>	14.583
5	DiffDRR	2.405	56.259	171.329	70.909	70.909
	Fixed Landmark	1.295	<b>1.536</b>	2.654	1.818	1.818
	RayEmb	<b>1.288</b>	1.561	<b>2.152</b>	<b>0.000</b>	<b>0.000</b>
	RayEmb*	1.355	1.630	2.705	<b>0.000</b>	<b>0.000</b>
6	DiffDRR	3.232	22.468	257.309	62.500	66.667
	Fixed Landmark	2.251	3.433	152.635	20.833	20.833
	RayEmb	2.006	<b>2.567</b>	6.096	<b>0.000</b>	20.833
	RayEmb*	<b>1.838</b>	2.605	<b>5.093</b>	<b>0.000</b>	<b>8.333</b>

1e-4. The Fixed Landmark method is also trained using the same batch size and optimizer settings.

**Inference.** We use the PnP and RANSAC implementation from OpenCV for initial registration. Specifically, we sample 3000 points inside the bone segmentation mask of the CT volume, run our pipeline, and select top- $k$  points based on similarity scores for initial filtering, with  $k = 600$  in our experiments. For these top- $k$  points, we run the pipeline  $n$ -times ( $n = 10$ ), using randomly sampled templates in each iteration, and choose the 2D projection with the highest similarity values. It should be noted that these hyperparameters were manually chosen based on testing with DeepFluoro synthetic validation cases. For the refinement procedure, we follow the same protocol as in DiffPose [9], running the optimization for 100 iterations for both the proposed method and the Fixed Landmark method. To compare with DiffDRR-only method, we initialize the pose with AP view for each specimen and run optimization for 150 iterations.

#### 4.4 Registration Results

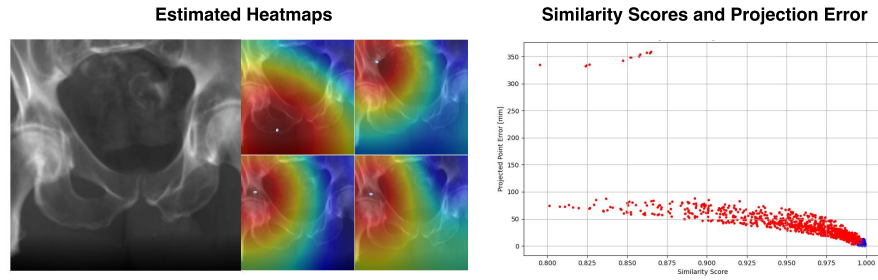
**Quantitative.** We evaluated the mean Target Registration Errors (mTRE) across three registration methods – DiffDRR, Fixed Landmark, RayEmb – for



**Fig. 4:** Two example cases (Top and Bottom) where RayEmb provides better initial pose estimates compared to other methods. The ground truth landmarks are shown in magenta, and the estimated landmarks are in cyan (left). A 3D visualization of the initial pose estimate is displayed in green, the optimized pose in blue, and the ground truth pose in red (right).

six different specimens. RayEmb refers to the proposed method trained on DeepFluoro dataset, while RayEmb\* indicates the method trained on the CTPelvic1K CLINIC dataset. The results are presented as box plots in Fig. 3. DiffDRR exhibited the highest median errors and variability across all specimens. In particular, specimens 2, 4, and 6 showed significantly higher mTREs with this method, far exceeding the 10 mm threshold, suggesting limitation in initializing an image-based similarity optimization scheme with a constant AP pose for achieving precise registration. Table. 1 and Table. 2 show that both RayEmb and Fixed Landmark maintain median mTREs below the 10 mm and mPD below 5mm threshold across all specimens. The mPD values are generally lower than mTRE because mPD is less sensitive to error in the depth direction. Notably, for specimen 6, which presents a challenging test case due to patient and pose variability, the proposed method achieves lower mTRE, mPD, and GFR, suggesting better generalizability to unseen patient CTs.

**Qualitative.** Fig. 4 shows two example test cases from specimen 6 where initial pose estimation is crucial for successful pose refinement. In the left section, the ground truth landmarks are shown in magenta, while the estimated landmarks from each method are indicated in cyan. The top case demonstrates the proposed method providing a pose estimate closer to the ground truth than the Fixed Landmark method, while the bottom case illustrates the proposed method’s ability to generate pose estimates even in situations where both DiffDRR and Fixed Landmark fail. RayEmb achieves significantly closer alignment to the ground truth, with mPD values of 3.77 mm and 4.06 mm, which are sub-



**Fig. 5:** Examples of the estimated heatmaps (left) and the similarity score versus projection error plot (right). In the heatmaps, the cyan point indicates the predicted point, while the magenta point represents the corresponding ground truth. In the plot, red points correspond to sampling points outside the field of view, and blue points correspond to visible sampling points.


stantially lower than those of DiffDRR and Fixed Landmark. In the right section, the initial pose estimates are shown in green, while the ground truth poses are displayed in red. The trajectories between the estimated and true poses illustrate the deviation of each method from the ideal alignment. RayEmb shows smaller deviations compared to the other methods, highlighting its superior precision in initial pose estimation. Additional examples are provided in the supplemental material.

#### 4.5 Analyzing the Heatmap and Similarity Scores

Fig. 5 presents estimated heatmaps alongside a plot of similarity scores versus projection errors, derived from the DeepFluoro synthetic dataset. The heatmaps demonstrate RayEmb’s ability to accurately estimate landmark positions. In the plot, red points represent sampling points outside the field of view, while blue points indicate visible sampling points within the projected image field. This distinction is crucial as it provides insights into the spatial distribution and accuracy of the landmark estimations relative to the actual imaging area. The plot shows a strong correlation between high similarity scores and low projection errors, validating the effectiveness of selecting top- $k$  correspondences based on high similarity metrics. By prioritizing correspondences with the highest similarity scores, RayEmb effectively minimizes errors in subsequent steps, thereby enhancing both the accuracy and reliability of the 2D-3D registration process.

## 5 Limitations

While the proposed method achieves precise registration with enhanced generalizability, the runtime of the landmark estimation pipeline depends on the number of sampled points; a higher number of samples leads to increased computation



**Fig. 6:** Inference runtime of the proposed method, excluding the DiffDRR optimization.

time. As shown in Fig. 6, current inference setup with 3000 sampling points require approximately 1000ms for initial pass and another 2000ms ( $200ms \times 10$ ) for test time augmentation, excluding the DiffDRR optimization, on an RTX 3090 GPU Ti. Additionally, all three methods employ DiffDRR with a multi-scale normalized cross-correlation loss, which requires few seconds to complete 100 iterations. The efficiency related to sampling points could be improved by implementing hierarchical sampling and filtering, which would use similarity scores to exclude points outside the specific region of interest for a given view. Furthermore, the refinement process could be significantly accelerated by adopting a sparse implementation of normalized cross-correlation (sparse mNCC), as used in DiffPose.

## 6 Conclusion

We demonstrated the effectiveness of arbitrary landmark detection using the ray embedding subspace method to achieve high-precision registration across various specimens, with notable success in the challenging test case of specimen 6. Additionally, we validated the generalizability of the proposed method by training on the CTPelvic1K dataset and evaluating on the DeepFluoro dataset, resulting in a patient-wide model that performs well without any patient-specific fine-tuning steps required by previous approaches. This approach addresses critical needs in emergency treatments and diagnostics, where both accuracy and broad generalizability across patients are essential. Future work will focus on further refining the method and exploring its applicability to other anatomical regions and clinical scenarios.

**Acknowledgments.** This work was partially supported by a grant from JSPS KAKENHI Grant Number JP23K08618. This work used computational resources of Pegasus provided by Multidisciplinary Cooperative Research Program in Center for Computational Sciences, University of Tsukuba.

## References

1. Barath, D., Matas, J., Noskova, J.: MAGSAC: Marginalizing sample consensus. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10197–10205 (2019)
2. Bier, B., Unberath, M., Zaech, J.N., Fotouhi, J., Armand, M., Osgood, G., Navab, N., Maier, A.: X-ray-transform invariant anatomical landmark detection for pelvic trauma surgery. arXiv [cs.CV] (Mar 2018)
3. Bradley, M.P., Benson, J.R., Muir, J.M.: Accuracy of acetabular component positioning using computer-assisted navigation in direct anterior total hip arthroplasty. *Cureus* **11**(4), e4478 (Apr 2019)
4. Bukschat, Y., Vetter, M.: EfficientPose – an efficient, accurate and scalable end-to-end 6D multi object pose estimation approach. arXiv [cs.CV] (Nov 2020)
5. Do, T.T., Cai, M., Pham, T., Reid, I.: Deep-6DPose: Recovering 6D object pose from a single RGB image. arXiv [cs.CV] (Feb 2018)
6. Esteban, J., Grimm, M., Unberath, M., Zahnd, G., Navab, N.: Towards fully automatic X-ray to CT registration. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2019. pp. 631–639. Springer International Publishing (2019)
7. Gao, C., Killeen, B.D., Hu, Y., Grupp, R.B., Taylor, R.H., Armand, M., Unberath, M.: Synthetic data accelerates the development of generalizable learning-based algorithms for X-ray image analysis. *Nat. Mach. Intell.* **5**(3), 294–308 (Mar 2023)
8. George, A.K., Sonmez, M., Lederman, R.J., Faranesh, A.Z.: Robust automatic rigid registration of MRI and X-ray using external fiducial markers for XFM-guided interventional procedures. *Med. Phys.* **38**(1), 125–141 (Jan 2011)
9. Gopalakrishnan, V., Dey, N., Golland, P.: Intraoperative 2D/3D image registration via differentiable X-ray rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11662–11672 (2024)
10. Gopalakrishnan, V., Golland, P.: Fast auto-differentiable digitally reconstructed radiographs for solving inverse problems in intraoperative imaging. In: Clinical Image-Based Procedures, pp. 1–11. Lecture notes in computer science, Springer Nature Switzerland, Cham (2023)
11. Grupp, R.B., Unberath, M., Gao, C., Hegeman, R.A., Murphy, R.J., Alexander, C.P., Otake, Y., McArthur, B.A., Armand, M., Taylor, R.H.: Automatic annotation of hip anatomy in fluoroscopy for robust and efficient 2D/3D registration. *Int. J. Comput. Assist. Radiol. Surg.* **15**(5), 759–769 (May 2020)
12. Gu, W., Gao, C., Grupp, R., Fotouhi, J., Unberath, M.: Extended capture range of rigid 2D/3D registration by estimating riemannian pose gradients. *Mach Learn Med Imaging* **12436**, 281–291 (Oct 2020)
13. Güler, R.A., Neverova, N., Kokkinos, I.: DensePose: Dense human pose estimation in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7297–7306 (2018)
14. Haugaard, R.L., Buch, A.G.: SurfEmb: Dense and continuous correspondence distributions for object pose estimation with learnt surface embeddings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6749–6758 (2022)
15. Hodan, T., Barath, D., Matas, J.: EPOS: Estimating 6D pose of objects with symmetries. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11703–11712 (2020)

16. Kendall, A., Grimes, M., Cipolla, R.: PoseNet: A convolutional network for real-time 6-DOF camera relocalization. In: 2015 IEEE International Conference on Computer Vision (ICCV). pp. 2938–2946 (Dec 2015)
17. Labb'e, Y., Manuelli, L., Mousavian, A., Tyree, S., Birchfield, S., Tremblay, J., Carpentier, J., Aubry, M., Fox, D., Sivic, J.: MegaPose: 6D pose estimation of novel objects via render & compare. *CoRL* **205**, 715–725 (Dec 2022)
18. Labb e, Y., Carpentier, J., Aubry, M., Sivic, J.: CosyPose: Consistent multi-view multi-object 6D pose estimation. In: *Computer Vision – ECCV 2020*, pp. 574–591. *Lecture notes in computer science*, Springer International Publishing, Cham (2020)
19. Li, Y., Wang, G., Ji, X., Xiang, Y., Fox, D.: DeepIM: Deep iterative matching for 6D pose estimation. In: *Computer Vision – ECCV 2018*, pp. 695–711. *Lecture notes in computer science*, Springer International Publishing, Cham (2018)
20. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2980–2988 (2017)
21. Liu, F., Yan, K., Harrison, A., Guo, D., Lu, L., Yuille, A., Huang, L., Xie, G., Xiao, J., Ye, X., Jin, D.: SAME: Deformable image registration based on self-supervised anatomical embeddings. *arXiv [eess.IV]* (Sep 2021)
22. Liu, P., Han, H., Du, Y., Zhu, H., Li, Y., Gu, F., Xiao, H., Li, J., Zhao, C., Xiao, L., Wu, X., Zhou, S.K.: Deep learning to segment pelvic bones: large-scale CT datasets and baseline models. *Int. J. Comput. Assist. Radiol. Surg.* **16**(5), 749–756 (May 2021)
23. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: SSD: Single shot MultiBox detector. In: *Computer Vision – ECCV 2016*, pp. 21–37. *Lecture notes in computer science*, Springer International Publishing, Cham (2016)
24. Lu, X.X.: A review of solutions for perspective-n-point problem in camera pose estimation. *J. Phys. Conf. Ser.* **1087**(5), 052009 (Sep 2018)
25. Markelj, P., Tomaževič, D., Likar, B., Pernuš, F.: A review of 3D/2D registration methods for image-guided interventions. *Med. Image Anal.* **16**(3), 642–661 (Apr 2012)
26. Nguyen, V.N., Groueix, T., Salzmann, M., Lepetit, V.: GigaPose: Fast and robust novel object pose estimation via one correspondence. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9903–9913 (2024)
27. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv [cs.LG]* (Jul 2018)
28. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Assran, M., Ballas, N., Galuba, W., Howes, R., Huang, P.Y., Li, S.W., Misra, I., Rabbat, M., Sharma, V., Synnaeve, G., Xu, H., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: DINOv2: Learning robust visual features without supervision. *arXiv [cs.CV]* (Apr 2023)
29. Park, K., Patten, T., Vincze, M.: Pix2Pose: Pixel-wise coordinate regression of objects for 6D pose estimation. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). IEEE (Oct 2019)
30. Peng, S., Liu, Y., Huang, Q., Zhou, X., Bao, H.: PVNet: Pixel-wise voting network for 6DoF pose estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4561–4570 (2019)
31. Rad, M., Lepetit, V.: BB8: A scalable, accurate, robust to partial occlusion method for predicting the 3D poses of challenging objects without using depth. In: Pro-



- ceedings of the IEEE International Conference on Computer Vision. pp. 3828–3836 (2017)
32. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 779–788 (2016)
  33. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems* **28** (2015)
  34. Shrestha, P., Xie, C., Shishido, H., Yoshii, Y., Kitahara, I.: X-ray to CT rigid registration using scene coordinate regression. In: *Lecture Notes in Computer Science*, pp. 781–790. Lecture notes in computer science, Springer Nature Switzerland, Cham (2023)
  35. Shugurov, I., Li, F., Busam, B., Ilic, S.: OSOP: A multi-stage one shot object pose estimation framework. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6835–6844 (2022)
  36. Swinehart, D.F.: The beer-lambert law. *J. Chem. Educ.* **39**(7), 333 (Jul 1962)
  37. Wada, K., Mikami, H., Hamada, D., Yamazaki, T., Tomita, T., Sairyo, K.: Can intraoperative kinematic analysis predict postoperative kinematics following total knee arthroplasty? a preliminary. *J. Med. Invest.* **65**(1.2), 21–26 (2017)
  38. Wang, G., Manhardt, F., Tombari, F., Ji, X.: GDR-net: Geometry-guided direct regression network for monocular 6D object pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16611–16621 (2021)
  39. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: PoseCNN: A convolutional neural network for 6D object pose estimation in cluttered scenes. *arXiv [cs.CV]* (Nov 2017)
  40. Yan, K., Cai, J., Jin, D., Miao, S., Guo, D., Harrison, A.P., Tang, Y., Xiao, J., Lu, J., Lu, L.: SAM: Self-supervised learning of pixel-wise anatomical embeddings in radiological images. *IEEE Trans. Med. Imaging* **41**(10), 2658–2669 (Oct 2022)
  41. Yoshii, Y., Kusakabe, T., Akita, K., Tung, W.L., Ishii, T.: Reproducibility of three dimensional digital preoperative planning for the osteosynthesis of distal radius fractures. *J. Orthop. Res.* **35**(12), 2646–2651 (Dec 2017)
  42. Zakharov, S., Shugurov, I., Ilic, S.: DPOD: 6D pose object detector and refiner. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1941–1950 (2019)