

Deep Learning compatible Differentiable X-ray Projections for Inverse Rendering

Karthik Shetty¹, Annette Birkhold², Norbert Strobel³, Bernhard Egger⁴,
Srikrishna Jaganathan¹, Markus Kowarschik², Andreas Maier¹

¹Pattern Recognition Lab, FAU Erlangen-Nürnberg

²Siemens Healthcare GmbH, Forchheim

³Fakultät Elektrotechnik, HS für angewandte Wissenschaften Würzburg-Schweinfurt

⁴MIT - BCS, CSAIL & CBMM, USA

karthik.shetty@fau.de

Abstract. Many minimally invasive interventional procedures still rely on 2D fluoroscopic imaging. Generating a patient-specific 3D model from these X-ray projection data would allow to improve the procedural workflow, e.g. by providing assistance functions such as automatic positioning. To accomplish this, two things are required. First, a statistical human shape model of the human anatomy and second, a differentiable X-ray renderer. In this work, we propose a differentiable renderer by deriving the distance travelled by a ray inside mesh structures to generate a distance map. To demonstrate its functioning, we use it for simulating X-ray images from human shape models. Then we show its application by solving the inverse problem, namely reconstructing 3D models from real 2D fluoroscopy images of the pelvis, which is an ideal anatomical structure for patient registration. This is accomplished by an iterative optimization strategy using gradient descent. With the majority of the pelvis being in the fluoroscopic field of view, we achieve a mean Hausdorff distance of 30 mm between the reconstructed model and the ground truth segmentation.

1 Introduction

Over the last decades, the amount of X-ray guided interventional procedures has increased steadily, raising the awareness for optimized procedural workflows and radiation-dose induced adverse effects. Assistance systems based on a 3D digital twin of the patient have the potential to speed up procedures and minimize the required radiation. However, precisely representing the individual human anatomy based on the commonly acquired 2D projection images is an ill-posed problem. Ehlke et al. showed that an accurate and fast 3D reconstruction of the human pelvis from 2D projection images is possible by learning its shape space in the form of tetrahedral mesh structures from a large set of pelvis scans including bone density information [1]. In theory, this approach has the potential to generate digital twins of patients from live fluoroscopic imaging data and may allow to develop advanced assistance applications in the interventional environment.

The process of generating a 3D model from 2D projection images is known as inverse-rendering and has received considerable attention in computer vision [2,3]. However, achieving this usually means solving an ill-posed problem, and rendering pipelines are not necessarily differentiable. This is usually overcome by softening the non-differentiabilities by soft-rasterization [2]. Usually, a 3D scene is represented by a polygon mesh, which can be expressed by a set of vertices and faces. Such a representation is usually chosen for statistical shape and pose modeling. Skinned Multi-Person Linear Model (SMPL) is an example of such a parameterized model, which is deformable and accurately represents a large variation of human shapes [4]. Employing such a model enables more precise and faster 3D reconstruction from a 2D projection image space due to the incorporation of prior knowledge [3]. To carry out the 3D reconstruction from X-ray projections, we need to calculate artificial X-ray projections from a 3D model, e.g., a Statistical Shape Model (SSM). Generation of transmission images from mesh structures has been previously proposed [1,5]. For a viable reconstruction, we seek to optimize the shape β and pose θ parameters by minimizing a similarity function between the simulated image and the acquired image. To this end, we propose a framework to combine existing X-ray rasterization approaches with a differentiable rasterizer, to create a setup that can be iteratively optimized by gradient descent techniques. The proposed renderer enables future end-to-end frameworks.

2 Methods

2.1 Differentiable Rasterizer

A watertight manifold surface can be represented as a triangular mesh by a set of faces $\mathbf{f} \in \mathbb{N}^{N_f \times 3}$ and vertices $\mathbf{v} \in \mathbb{R}^{N_v \times 3}$, where the object has N_f faces and N_v vertices. Given a projection matrix, it is possible to transform the mesh from world coordinates to detector coordinates while maintaining the actual depth of the vertex. In traditional rasterization of 3D scenes, each pixel P_{xy} on the screen is affected by only the face $\mathbf{f}_j$ nearest to the ray from the camera. This is usually achieved by storing the depth information for all the faces influenced by a given pixel in the form of *z-buffer*. However, in the case of transmission imaging we need to know the distance traveled by a ray through an object in the form of *l-buffer* [5]. For each face $\mathbf{f}_j$, the normals $\vec{N}_{\mathbf{f}_j}$ in the projective space are computed. The sign of the dot product $D_{\mathbf{f}_j R_{xy}} = \text{sign}(\vec{N}_{\mathbf{f}_j} \cdot \vec{N}_{view})$ between the face normal $\vec{N}_{\mathbf{f}_j}$ and detector plane normal $\vec{N}_{view} = (0, 0, 1)^T$ determines if the ray R_{xy} is either entering or exiting the object for a given face. Hence, the total path length L_{xy} traveled by a ray for a given object p from the source to the detector pixel position P_{xy} is visualized in Fig. 1 and can be formulated as

$$L_{xy} = \sum_{j=1}^K D_{\mathbf{f}_j R_{xy}} Z_{xy} \mathbf{f}_j \quad (1)$$

Here $Z_{xy}f_j$ represents the distance between the face f_j and the detector pixel position P_{xy} , while j and K represents the current valid intersection and the fixed number of mesh overlaps stored in the z-buffer, respectively. As a result, there is no substantial requirement for the elements of the z-buffer to be stored in a sorted manner. The only constraint is the number of faces K influencing a given pixel. To determine the gradients for the backward pass we follow the same procedure as implemented by Ravi et al [2], with an exception for the gradients for the z-buffer:

$$\frac{\partial C}{\partial z_{xyj}} = \frac{\partial C}{\partial L_{xy}} D_{f_j R_{xy}} \quad (2)$$

Here, z_{xyj} represents the z-buffer distance from the pixel P_{xy} corresponding to face f_j with C as some cost function, e.g., the normalized gradient correlation (NGC) explained below.

It is generally assumed that the number of incoming rays for a given object is the same as the outgoing rays. However, artifacts occur when the face normal is perpendicular to the detector plane or in the presence of a non-manifold surface. The occurrence of such a situation can be determined when the summation over the pixel sign function $\sum_j^K R_{xy}f_j$ is non-zero. To overcome this, we set the image pixel value to a neighbouring pixel with zero sum and set the gradient for that particular pixel to zero.




Fig. 1. Visualization for the determination of the ray distance for a simplified structure. Here, total distance covered by the ray is $L_{xy} = d_1 - d_2 + d_3 - d_4$.

2.2 X-ray rendering

The pipeline for generating the X-ray projection image for a human patient model is depicted in Fig. 2. Distance projection maps L_p are generated individually for air (L_{air}), body (L_{body}), bones (L_{bones}) and organs (L_{organs}). L_{air} represents the Euclidean distance from a detector pixel P_{xy} in world coordinates to the X-ray source. L_{body} , L_{bones} and L_{organs} are generated as described in Sec. 2.1 from individual meshes respectively. For ease of explanation, we describe L_{organs} as a single organ distance map, however it is a subset of lungs, heart, liver, kidney and spleen. As L_{air} is inclusive of L_{body} , we define the actual air distance map as $L_{air} - L_{body}$. For the body we proceed similarly as shown in Fig. 2.




Fig. 2. Left: Projection performed around the thorax section from a mesh representation of a patient model. Right: Pipeline to generate a transmission image from the given mesh structures and projection matrix.

The X-ray projection is generated using a primary signal I_{xy} for a pixel P_{xy} described using the Beer-Lambert law as shown in Eq. 3 with a photon energy E , X-ray energy $I_o(E)$, and linear attenuation coefficients $\mu(p, E)$ for objects p

$$I_{xy} = \sum_E I_o(E) e^{-\sum_p (\mu(p, E) L_{p_{xy}})} \quad (3)$$

Making use of the aforementioned distance maps L_p , the simulation can be extended for polychromatic X-ray beam spectra.

2.3 Application: Inverse Rendering

We build a statistical parametric model $M(\beta, \theta; \Phi)$, similar to SMPL [4] extended for internal anatomy using principal component analysis to characterize the human anatomy shape space from a large set of segmented whole-body CT scans. This makes it possible to generate realistic X-ray projections with the entire pipeline being differentiable in nature. Here, Φ represents the learned parameters of the model. To reconstruct a 3D model from a projection image we fix the camera projection matrix of the simulation system as defined by the target along with an initial mean shape β_m and rest pose θ_r . A gradient descent based optimization with backpropagation is applied to the simulation model. NGC is used as a similarity measure [6], which is defined by the normalized cross-correlation of the image gradient between two images. The model is optimized until convergence to obtain the optimal pose θ and shape β parameters.

For evaluation, we used the dataset provided by Grupp et al. made up of the hip anatomy from 6 patients [7]. For each patient, a 3D segmentation of the pelvis along with 14 landmarks and a total of 366 fluoroscopy images in varying orientations with their respective projection matrices are available.

3 Evaluation and Results

The SSM is registered to each patient from the dataset. This resulted in a base score with an average Hausdorff distance of 17.79 mm and a mean landmark error of 11.81 mm. The rotation and translation parameters obtained are assumed as

Table 1. Overview of the reconstruction and landmark error. All distances are measured in mm. The lower bound error is dependent on the shape model, with a mean Hausdorff distance of 17.79 mm and a mean landmark error of 11.81 mm.

Initial	Initial	Reconstructed	Reconstructed
Hausdorff Distance	Landmark Error	Hausdorff Distance	Landmark Error
46.59 ± 4.64	30.00 ± 8.30	29.11 ± 4.56	22.06 ± 6.13

the ground truth irrespective of the shape parameters. The 2D-3D reconstruction is tested by applying a random translation in the range of $[-10,10]$ mm and a rotation in the range of $[-5,5]^\circ$ on all three axes for the SSM. The randomization is small due to the nature of the loss function being sensitive to changes only around a small region, such precision of initialization could be either provided manually or achieved with a network learning those initialization parameters. Fig. 3 shows sample outputs of our proposed method along with the quantitative results in Table 1. For evaluation, we consider only the images where the region of pelvis visible is greater than 50%.




Fig. 3. Samples for 3D reconstruction from 2D projection images. (a) Target image, (b) Projection image with a random translation and rotation, (c) Projection image after registration, (d) Initial 3D overlay of template mesh, (e) 3D overlay of the meshes after registration, (f) NGC map after registration.

4 Discussion and Conclusion

We presented an approach for differential rendering and reconstructing a patient-specific 3D mesh model from a single 2D X-ray projection image. We demonstrated that using a deformable human model parameterized by shape and pose, we were able to approximately register patients to an X-ray projection. We found that given a visibility above 50% and an initialization of a generic model shape positioned close to the actual image, the 3D reconstruction error is significantly reduced with our optimization approach. The remaining error may be due to the simultaneous unconstrained optimization of shape and pose. This is because changes in shape and translation across depth, i.e., in z-direction, results in a

large range of solutions causing the simulated image to appear similar to the target image. For now, we ignore the case when less than 50% of the pelvis is visible due to the intrinsic nature of the loss function being sensitive to the initial position. In general, the error could be reduced with an additional prior to determine a good initial pose [8].

As the registration is independent of the patient pose we can overcome possible pose mismatches resulting from employing preoperative CT scan for registration. Still, the pipeline has limitations. First, the X-ray simulation is an approximation of a real X-ray image, as it does not take into account any scattering mechanisms and other effects causing noise. This can be possibly improved by employing a deep learning model for predicting patient-specific scatter [9]. Second, the renderer assumes a constant density for any given material, resulting in a pseudo-realistic rendering. One way to overcome this problem would be to use a tetrahedral mesh with learned bone density distributions. Third, due to ambiguous depth information in fluoroscopy images, accurate registration becomes a challenging task. This error could be minimized with multi-view images or by using the table position as a constraint. To summarize, we proposed a differentiable renderer by deriving the distance travelled by a ray inside mesh structures. This was further extended to generate X-ray images from human models. Finally, we showed that it could be successfully used to reconstruct 3D models from 2D fluoroscopy images.

Disclaimer: The concepts and information presented in this paper are based on research and are not commercially available.

References

1. Ehlke M, Ramm H, Lamecker H, et al. Fast generation of virtual X-ray images for reconstruction of 3D anatomy. *IEEE Trans Vis Comput Graph*. 2013; p. 2673–2682.
2. Ravi N, Reizenstein J, Novotny D, et al. Accelerating 3d deep learning with pytorch3d. *arXiv preprint arXiv:200708501*. 2020;.
3. Liu S, Li T, Chen W, et al. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In: *Proc IEEE Int Conf Comput Vis*; 2019. p. 7708–7717.
4. Loper M, Mahmood N, Romero J, et al. SMPL: A Skinned Multi-Person Linear Model. *ACM Trans Graph (Proc SIGGRAPH Asia)*. 2015 Oct;34(6):248:1–248:16.
5. Vidal FP, Garnier M, Freud N, et al. Simulation of X-ray Attenuation on the GPU. In: Tang W, Collomosse J, editors. *Theory and Practice of Computer Graphics*. The Eurographics Association; 2009. .
6. Penney GP, Weese J, Little JA, et al. A comparison of similarity measures for use in 2-D-3-D medical image registration. *IEEE Trans Med Imaging*. 1998;17(4):586–595.
7. Grupp RB, Unberath M, Gao C, et al. Automatic annotation of hip anatomy in fluoroscopy for robust and efficient 2D/3D registration. *Int J Comput Assist Radiol Surg*. 2020 May;15(5):759–769.
8. Bier B, Unberath M, Zaech JN, et al. X-ray-transform invariant anatomical landmark detection for pelvic trauma surgery. In: *Med Image Comput Comput Assist Interv*. Springer; 2018. p. 55–63.
9. Roser P, Zhong X, Birkhold A, et al. Physics-driven learning of x-ray skin dose distribution in interventional procedures. *Med Phys*. 2019;46:4654–4665.