Shape from projections via differentiable forward projector for computed tomography

Koo, Jakeoung; Dahl, Anders Bjorholm; Bærentzen, J. Andreas; Chen, Qiongyang ; Bals, Sara; Dahl, Vedrana Andersen

Published in:
Ultramicroscopy

Link to article, DOI:
[10.1016/j.ultramic.2021.113239](https://doi.org/10.1016/j.ultramic.2021.113239)

Publication date:
2021

Document Version
Early version, also known as pre-print

[Link back to DTU Orbit](#)

Citation (APA):
Koo, J., Dahl, A. B., Bærentzen, J. A., Chen, Q., Bals, S., & Dahl, V. A. (2021). Shape from projections via differentiable forward projector for computed tomography. *Ultramicroscopy*, 224, [113239]. <https://doi.org/10.1016/j.ultramic.2021.113239>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Shape from Projections via Differentiable Forward Projector for Computed Tomography

Jakeoung Koo^{a,*}, Anders B. Dahl^a, J. Andreas Bærentzen^a, Qiongyang Chen^b, Sara Bals^b, Vedrana A. Dahl^a

^aTechnical University of Denmark, Anker Engelunds Vej 1, Kgs. Lyngby, 2800, Denmark

^bUniversity of Antwerp, Prinsstraat 13, Antwerpen, 2000, Belgium

Abstract

In computed tomography, the reconstruction is typically obtained on a voxel grid. In this work, however, we propose a mesh-based reconstruction method. For tomographic problems, 3D meshes have mostly been studied to simulate data acquisition, but not for reconstruction, for which a 3D mesh means the inverse process of estimating shapes from projections. In this paper, we propose a differentiable forward model for 3D meshes that bridge the gap between the forward model for 3D surfaces and optimization. We view the forward projection as a rendering process, and make it differentiable by extending recent work in differentiable rendering. We use the proposed forward model to reconstruct 3D shapes directly from projections. Experimental results for single-object problems show that the proposed method outperforms traditional voxel-based methods on noisy simulated data. We also apply the proposed method on electron tomography images of nanoparticles to demonstrate the applicability of the method on real data.

Keywords: Computed Tomography, Electron Tomography, Tomographic Reconstruction, Mesh Deformation

1. Introduction

In computed tomography (CT), we aim at solving the inverse problem of computing the 3D structure (shape and attenuation) of an object from a set of projection images (Buzug, 2008) taken from different angles. Here, the geometry and the physics of the imaging system is known, which allows us to model the forward process, i.e. if we have a suggestion for the 3D structure of the imaged object, we can compute the projection images.

We need a data structure to represent the structure of the object that should be reconstructed. The most common data structure is a volumetric image, with voxel intensities representing local attenuation. This approach may be used for reconstructing any type of object. Since each voxel in the volume is a parameter that must be computed, we have a very large number of unknowns. This makes it difficult to reconstruct volumes in situations where we have projections from a limited angular view (e.g., in electron tomography) or noisy data, and it can be difficult to accurately compute the attenuations in all voxels. Therefore, we propose to use a mesh to represent the shape of the object. The mesh separates the object into parts with a constant attenuation.

In tomography, forward projection of 3D meshes has mostly been used for simulating tomographic data acquisition, i.e. modeling the forward projection. This includes modeling X-ray transmission imaging based on Monte-Carlo methods (Bonin et al., 2002; Freud et al., 2006) or ray tracing techniques (Freud et al., 2006; Marinovszki et al., 2018). Furthermore, Vidal et al.

Figure 1: The proposed differentiable forward projector enables optimizing a 3D mesh from projections for tomographic reconstruction. Given a mesh and scanning geometry, the forward projector computes projections for the current mesh, which is updated by deforming vertices in the direction that minimizes our objective function.

(2009); Sujar et al. (2017) took the advantage of the OpenGL library to simulate X-ray images in real time. However, none of the proposed methods are concerned with reconstruction, i.e. solving the inverse problem.

For the mesh-based tomographic reconstruction that we propose, the reconstruction problem is two-fold. The mesh must be deformed to follow the boundaries of the depicted object, and in each part of the object, a single attenuation coefficient must be estimated. To forward project the mesh, we employ rendering techniques, which allow a very efficient projection of the 3D mesh to the detector plane. We extend the differentiable rasterizer recently proposed by Chen et al. (2019) to derive differentiable forward projection. This enables us to compute vertex displacements that deform the mesh based on the difference between the forward projection and projection data. In Fig. 1, we provide an overview of our shape estimation method. In Fig. 2, we illustrate how our work differs from existing image-based reconstruction methods (Andersen and Kak, 1984; Buzug, 2008).

March 12, 2021

*Corresponding author. E-mail: jakoo@dtu.dk

Accepted manuscript, Ultramicroscopy

DOI: <https://doi.org/10.1016/j.ultramicro.2021.113239>

© 2021. This manuscript version is made available under the CC-BY-NC-ND 4.0 license <http://creativecommons.org/licenses/by-nc-nd/4.0/>

Figure 2: Algebraic reconstruction vs. our proposed approach. *Left*: Existing algebraic reconstruction methods use a linear system of equations to model the forward projections typically together with some regularization to constrain the solution. *Right*: We propose a differentiable forward model f which computes projections for an object represented by a triangular mesh and attenuation coefficients μ . This forward model is used to reconstruct the shape from projections data $\hat{p}$.

Our model has been developed for problems like X-ray CT where the damping of the X-ray attenuation coefficient are modeled as linearly dependent on the path length through the sample (Buzug, 2008), which is also assumed in electron tomography up to a certain thickness (Midgley and Weyland, 2003) where coefficient is related to electron scattering.

In summary, our contribution is two-fold. We suggest a differentiable forward projector to generate projections from 3D meshes, and we propose a shape estimation method employing the differentiable forward projector. Our code is available online: <https://github.com/jakeoung/ShapeFromProjections>.

2. Related works

Differentiable rendering based on rasterization. In our work we deform a mesh by changing vertex positions using gradient descent. Therefore, we need the gradient of our objective function with respect to vertex coordinates, which is possible through differentiating the forward projection. As observed in (Vidal et al., 2009), the tomographic forward projection can be implemented by extending a rendering technique in computer graphics based on rasterization. Such rasterization-based rendering technique projects 3D models onto 2D image plane and involves a discrete step to choose the pixels covered by triangles of 3D models. This discrete step is not differentiable in conventional graphics pipelines. Making such rasterization-based rendering differentiable, called *differentiable rasterization*, has been studied by several works to connect rendering and optimization. As our work extends such differentiable rasterization techniques, we first review those works and explain the difference to our work.

The general framework OpenDR, that was proposed by Loper and Black (2014), approximated gradients of pixel values with respect to model parameters. Kato et al. (2018) suggested a heuristic forward and backward pass where blurring is used to avoid zero gradients. This approach has an inconsistency between the forward and backward pass, and to circumvent this inconsistency, Liu et al. (2019) proposed a method called SoftRas by relaxing discrete rasterization process into the aggregation of smooth probability functions. Unlike the conventional rasterization rendering, in SoftRas, each face in the mesh affects many pixels in the image plane, which is computationally costly and memory-demanding.

Chen et al. (2019) suggested an interpolation-based differentiable rasterizer called DIB-R. DIB-R reformulates the barycentric interpolation in the rasterization process to analytically derive the gradients. Our forward projector extends this reformulation when computing the thickness of an object for computed tomography, but differs in some aspects. The rendering techniques such as DIB-R can improve the performance by ignoring invisible faces, but since we model penetrating radiation, our forward projector needs to consider all the faces. DIB-R used the idea of SoftRas (Liu et al., 2019) for background pixels to propagate the gradients on those background, but our forward projector does not use it to reduce the computational cost.

Shape reconstruction from projections. Our proposed method is related to tomographic segmentation, where segments are directly computed from projections. This includes (Elangovan and Whitaker, 2001; Whitaker and Elangovan, 2002; Alvino and Yezzi, 2004) that are based on the Mumford-Shah model (Mumford and Shah, 1989) where boundaries are represented using level-sets (Osher and Fedkiw, 2004). Recently, the parametric level-set method (Aghasi et al., 2011) has been used for tomographic segmentation in (Kadu et al., 2018; Eliasof et al., 2020) where level-sets are represented as an aggregation of radial basis functions. Although the parametric level-set method has fewer unknown variables, its forward projection still depends on a regular grid. Gadelha et al. (2019) used a deep convolutional neural network for 2D tomographic reconstruction, where the forward projection is based on the transformation of a regular grid and resampling. On the other hand, the work (Dahl et al., 2018) based on snakes (Kass et al., 1988) avoids a dense grid – it represents curves explicitly and proposes a direct forward projection of the curves. However, this method is limited to a single 2D curve, while the proposed method supports 3D objects. Another difference is that (Dahl et al., 2018) evolves curves in the normal directions of curve points, while our deformation can displace the vertices in all directions.

In summary, existing shape estimation methods from projections are either based on regular grid or limited to single 2D curve. To our knowledge, our work is the first to propose the differentiable forward projector for 3D triangular mesh and use it for reconstructing shapes from projections.

3. Differentiable forward projector

In this section, we describe our main contribution of the differentiable forward projector. The goal is to forward project triangular meshes and make this process differentiable with respect to 3D vertex positions and attenuation coefficients. This differentiable forward projection will be used for optimizing the mesh shape described in Sec. 4. First we describe the case of a single object and then extend to composite objects.

3.1. Single object

Consider an object represented by a watertight triangular mesh. A watertight triangle mesh forms a closed surface that has a well-defined interior and exterior: any path from a point

Figure 3: *Left*: The vertices of the triangle mesh (blue dots) are projected onto the detector. Each detector pixel (red dots) is associated with the projection ray which intersects mesh triangles. *Right*: One triangle j , here given by vertices $k = 0, 1, 2$, and one detector pixel i . Using barycentric coordinates, the distance d_{ij} may be expressed in terms of l_k .

in the interior to a point in the exterior must cross the triangle mesh. We assume that the object is homogeneous, i.e. it has a certain attenuation coefficient μ associated with the volume inside the mesh. For now, we consider μ constant, we will later explain how it is computed in the next subsection. The mesh consists of K vertices, and we write $\mathbf{v}_k$ for the 3D coordinates of the vertex k .

To simplify the explanation, we start with a single projection and later expand to multiple projections from different angles. Let $\mathbf{P}$ and $\mathbf{R}$ be the position of the detector and a matrix that rotates from detector coordinates to the global frame, respectively. If we denote the position of vertex k in global coordinates by $\mathbf{V}_k$, the position in detector coordinates are $\mathbf{v}_k = \mathbf{R}^T(\mathbf{V}_k - \mathbf{P})$. Note that in detector coordinates, the detector itself corresponds to the plane $z = 0$, its center is at the origin, and the positive z -axis points towards the object, see Fig. 3.

Expressed in detector coordinates, the distance of the vertex k from the detector is trivially $l_k = \mathbf{e}_3^T \mathbf{v}_k$ while $\mathbf{s}_k = [\mathbf{e}_1 \ \mathbf{e}_2]^T \mathbf{v}_k$ are the coordinates of the projection of the vertex onto the detector. Here $\mathbf{e}_1$, $\mathbf{e}_2$, and $\mathbf{e}_3$ are unit vectors in x , y and z direction, for example $\mathbf{e}_1 = [1 \ 0 \ 0]^T$.

Projecting the object onto the detector pixel i we consider the projection ray associated with i (slightly sloppy, call it ray i), and its path length in the object. As explained in (Vidal et al., 2009), this can be broken into contribution of all intersections of the ray i with the mesh triangles

$$p_i = \mu \sum_j \text{sign}(\mathbf{e}_3^T \mathbf{n}_j) d_{ij} \quad (1)$$

where $\mathbf{n}_j$ is the normal of the triangle j (needed for determining the sign of the contribution), and d_{ij} is the distance of the intersection point to the detector. Here, we consider the sign value $\text{sign}(\mathbf{e}_3^T \mathbf{n}_j)$ as a constant attribute of each triangle.

Considering now a single triangle j we express d_{ij} using barycentric coordinates

$$d_{ij} = \sum_{k \text{ in } j} w_{ij}^k l_k, \quad (2)$$

where k are the indices of the three vertices of the triangle j and

w_{ij}^k are the corresponding three barycentric coordinates of pixel i with respect to the projection of triangle j onto the detector plane, see Fig. 3, *right*.

To make the forward projection differentiable, we derive

$$\frac{\partial p_i}{\partial \mathbf{v}_k} = \mu \sum_{j \text{ intersects } i} \text{sign}(\mathbf{e}_3^T \mathbf{n}_j) \frac{\partial d_{ij}}{\partial \mathbf{v}_k}, \quad (3)$$

and

$$\frac{\partial d_{ij}}{\partial \mathbf{v}_k} = \sum_{k \text{ in } j} \left(w_{ij}^k \frac{\partial l_k}{\partial \mathbf{v}_k} + \frac{\partial w_{ij}^k}{\partial \mathbf{v}_k} l_k \right) = \sum_{k \text{ in } j} \left([0 \ 0 \ w_{ij}^k] + \frac{\partial w_{ij}^k}{\partial \mathbf{v}_k} l_k \right). \quad (4)$$

For the last step, computation of $\partial w_{ij}^k / \partial \mathbf{v}_k$, we employ the idea from (Chen et al., 2019), which reformulates the barycentric form to express the coefficients w_{ij}^k in terms of 2D projected positions $\mathbf{s}_k$ and the position of detector pixel i .

3.2. Composite objects

The method described above generalizes to composite objects if certain conditions are met. Specifically, we require that we know beforehand the topology of the parts of the composites, and how parts are embedded within one another. Thus, for each interface triangle we will have a suggestion that what class of material is on either side, but we do not know the specific attenuation coefficient of the classes, since we solve for those.

In order to extend Eq. (1) to composites, we only have to observe that triangles may now be the interface between two materials and not just air and material. This can be handled simply by letting each triangle contribute twice

$$p_i = \sum_j \mu_j \text{sign}(\mathbf{e}_3^T \mathbf{n}_j) d_{ij} - \sum_j \bar{\mu}_j \text{sign}(\mathbf{e}_3^T \mathbf{n}_j) d_{ij} \quad (5)$$

$$= \sum_j (\mu_j - \bar{\mu}_j) \text{sign}(\mathbf{e}_3^T \mathbf{n}_j) d_{ij}, \quad (6)$$

where μ_j is the attenuation of the interior material and $\bar{\mu}_j$ of the exterior material according to normal orientation. Of course, either attenuation will be zero if the material on the corresponding side of the triangle is air.

The derivative of a pixel value with respect to the contributing attenuation μ_m for material m (by abuse of notation) is

$$\frac{\partial p_i}{\partial \mu_m} = \sum_j \pm \text{sign}(\mathbf{e}_3^T \mathbf{n}_j) d_{ij}, \quad (7)$$

where the $\pm$ is positive if the interior material of face j is labeled as m and negative if m is the exterior material. We also modify Eq. (3) by changing μ to $(\mu_j - \bar{\mu}_j)$.

We have derived the Jacobians in Eq. (3) and (7), which will be used to optimize an objective function E . That is, we can propagate the gradients from E

$$\frac{\partial E}{\partial \mathbf{v}_k} = \sum_i \frac{\partial E}{\partial p_i} \frac{\partial p_i}{\partial \mathbf{v}_k}, \quad \frac{\partial E}{\partial \mu_m} = \sum_i \frac{\partial E}{\partial p_i} \frac{\partial p_i}{\partial \mu_m}, \quad (8)$$

where the summation is over all the detector pixels.

4. Shape from projections

In this section, we use the proposed forward projector to reconstruct shapes from projections. We assume that a template mesh with the correct topology is given. We aim to deform the template mesh and estimate the attenuation coefficients by minimizing the residual between data $\hat{p}$ and our estimation p .

Optimizing only the data fitting term can lead to degenerate meshes. To obtain high-quality meshes, we impose three regularization terms. The first term is the Laplacian term (Wang et al., 2018), which constraints the vertices to move similarly with their neighbors, defined by

$$E_{\text{lap}} = \sum_k \|\mathbf{V}_k - \frac{1}{|\mathcal{N}(k)|} \sum_{n \in \mathcal{N}(k)} \mathbf{V}_n\|^2, \quad (9)$$

where $\mathcal{N}(k)$ is the index set of neighboring vertices to k -th vertex. The second term is the edge length term to penalize long edges

$$E_{\text{edge}} = \sum_{(\mathbf{V}_k, \mathbf{V}_n) \in G} \|\mathbf{V}_k - \mathbf{V}_n\|^2, \quad (10)$$

where G denotes the set of edges. Lastly, we impose the flattening term (Kato et al., 2018; Liu et al., 2019)

$$E_{\text{flat}} = \sum_{e \in G} (1 - \cos \theta_e)^2, \quad (11)$$

where θ_e is the angle between the normal vectors of two faces sharing the edge e . Flattening term is needed to remove near-zero volume spikes. These thin artifacts have negligible contribution to the forward projection, and will be ignored by the data fitting term. As shown in Fig. 5, such artifacts can appear during the deformation, but disappear later.

With the data fidelity and regularization terms, the objective function to minimize is

$$E(\{\mathbf{V}_k\}, \{\mu_m\}) = \|p - \hat{p}\|_2^2 + \alpha E_{\text{lap}} + \beta E_{\text{edge}} + \gamma E_{\text{flat}} \quad (12)$$

where α, β, γ control the relative weights between the terms. Note that the size of projection data $\hat{p}$ is the number of detector pixels times the number of projection angles. We use automatic differentiation to minimize E . For large data, we can use stochastic gradient descent with mini batches in terms of projection angles. In this paper, however, we only consider full batch size of data.

5. Experiments and Results

In this section, we present the experimental results of the proposed method on synthetic data of single objects. We also show the results on real data of some nano particles from electron tomography, which has limited range of angles.

5.1. Shape reconstruction of a set of single objects

Datasets. This experiment is designed to test our shape estimation method on noisy simulated data. We use 5 watertight meshes (closed surfaces without any holes): star, spot, bunny, bob, kitten, shown in the first row of Fig. 6 and the attenuations of the objects are set to 1. Generating projections of those

Figure 4: Qualitative results of estimated projections on noisy data with relative noise level 0.4. The first row shows the ground truth, *i.e.*, noise-free data. 2nd-5th row show the forward projection from the solutions by SIRT, TV and the proposed method, respectively.

Figure 5: Deformation examples. The odd rows show the intermediate meshes during deformation and the even rows show the corresponding computed projections for one projection angle. We refine the mesh in finer resolution at iteration 60 and fix the mesh at iterations 120 and 180.

meshes using our forward model may resemble to the so-called inverse crime (Mueller and Siltanen, 2012). To avoid it, we employ the Blender software to make projection data based on

Figure 6: Qualitative results of extracted meshes on noisy data with relative noise level 0.4. The top row shows the ground truth meshes. Rows 2 and 3 show the extracted isosurface from the results of TV reconstruction with the optimal regularization parameter (best) and a high regularization parameter (high), respectively. The last row shows our results.

ray casting methods similar to (Marinovszki et al., 2018). We use 3D parallel projection geometry with 30 projection angles and the detector of size 192×192 pixels. Some projection images without noise are shown in the first row of Fig. 4 for one projection angle.

Evaluation metrics. We compare our result to two standard reconstruction methods: simultaneous iterative reconstruction technique (SIRT) (Andersen and Kak, 1984) and total variation (TV) based reconstruction (Chambolle and Pock, 2011; Sidky et al., 2012). These methods yield 3D images, whereas the proposed method produces surface mesh, making direct comparison of the main output challenging. For consistent comparison, we employ a residual-based metric: *residual projection error* (Roelandts et al., 2014), which measures the L^2 norm difference of data and computed projections. We impose the relative Gaussian noise on the original data and calculate the residual projection error between noise-free projections and the estimations of other methods and the proposed method. For SIRT and TV, the voxel size is set as $192 \times 192 \times 192$ and the algorithm parameters are chosen carefully.

Experimental details. Our implementation relies on PyTorch (Paszke et al., 2019) and uses Adam (Kingma and Ba, 2015) as an optimizer. The proposed forward projector is implemented as a module in PyTorch. As for the regularization parameters, we fix $\alpha = 10$, $\gamma = 0.01$ and iterate 500 times. The step size τ (learning rate in PyTorch) is set to 0.01 and reduced by half at 400 iteration. This reduction step is not really necessary but can yield a more stable result. We observe that 500 iterations are needed for capturing fine details of the complicated objects such as the *bunny* data. During the experiment, we only vary the edge length parameter β among the values of

1, 2, 4, 8, 16 and 32 and the optimal parameter would depend on the data. Finding the optimal regularization parameter is itself a research topic and not straightforward also in regularization-based image reconstruction methods such as TV. As for the initialization, the proposed method begins from an icosphere for genus-0 objects (*star*, *spot*, *bunny*), and from a torus for genus-1 objects (*bob*, *kitten*). Except for *star* data, we refine the mesh by (Huang et al., 2018) at iteration 60 and improve the mesh quality 3 times by a lightweight repair method (Attene, 2010) at iteration 60, 120, 180. These refine and repair steps help remove some artifacts and lead to fast convergence. As also observed in (Vidal et al., 2009), some artifact pixels can appear (e.g., when the ray hits an odd number of times). When a large deformation happens, we may observe some artifact pixels, which we exclude in the objective function. However, in the end, we only observe around 2 artifact pixels. In Fig. 5, we show the deformation of the estimated meshes with the corresponding computed projections.

Robustness to noise. The experiments show that the proposed method is robust to noise. Fig. 4 shows some computed projection images of the reconstruction results from noisy data with relative noise level 0.4 achieved using SIRT, TV reconstruction and the proposed method. Since SIRT has no regularization, it fits closely to the highly noisy data. The results of TV are relatively smooth, but sharp transitions appear blurred. On the other hand, the proposed method yields projections similar to noise-free data. Fig. 6 shows the final mesh results, where the proposed method yields qualitatively better results than TV – we omit highly noisy SIRT results. In Fig. 7, we provide the quantitative results of residual projection error with respect to relative noise levels. Without noise, SIRT gives the superior result as it fits to data without regularization. However, as noise

Figure 7: Quantitative results with varying relative noise level over 5 datasets. The *residual projection error* represents the error between the noise-free projections and the estimated projections. The error bars denote the average values with the maximum and minimum value and the y-axis is in logarithmic scale.

increases, the results of SIRT and TV are shown to be poorer.

Effect of parameters. Here, we investigate the effects of the parameters. In this experiment, we use data with a fixed relative noise level of 0.4 and use *residual projection error* to measure performance. Fig. 8 shows the effect of initializing the mesh by varying the radius of the icosphere used for initialization and then iterating 500 times. The final result is mostly not affected by the initial radius size thanks to the refinement steps during the optimization. We choose the *bunny* data model to generate the results shown in Fig. 8, because of its complicated shape.

In Fig. 9 we show the effects of the regularization parameters α , β and γ , while keeping other parameters fixed. Here we use all data models and the residual projection error for the *star* is large due to the coarse resolution of the mesh. We observe that if $\beta < 1.0$ or $\gamma > 0.1$, the final meshes might have some artifacts. As mentioned before, we use the default settings ($\alpha = 10$ and $\gamma = 0.01$), which in most experiments give a stable result. With these fixed values, we observe that it is enough to only vary the edge length parameter, depending on the desired degree of smoothness.

Figure 8: Effect of initializing the mesh with spheres of varying diameter. The x-axis represents the radius of the initialized sphere. The radius is determined relatively to an object space that is normalized to $(-1, 1)^3$. The experiments are based on the *bunny* data model.

(a) Effect of α

(b) Effect of β

(c) Effect of γ

Figure 9: Effect of the regularization parameters: (a) the Laplacian regularization parameter α given $\beta = 4$, $\gamma = 0.01$, (b) the edge length parameter β given $\alpha = 10$, $\gamma = 0.01$, and (c) the flattening parameter γ given $\alpha = 10$, $\beta = 4$. The vertical bars in (a) and (c) correspond to the values we keep fixed, when our method is compared with other methods.

Computational cost. Table 1 shows the size of the initial and final mesh and the corresponding running times. The mesh size is one of the major factors that contribute to computational time. For example, as we do not refine the mesh for *star* object, its final mesh size and its computational time is lower than others. We do the experiment on a Ubuntu server with 256GB RAM and Titan X GPU.

For a reference, the running times for the image-based methods are 38 seconds for SIRT and 152 seconds for TV which are implemented based on multithreading on 8 CPU cores, not GPU. For image-based methods, we use the fixed grid of size $192 \times 192 \times 192$ for all the data.

Table 1: Size of meshes and running times

Data	Star	Spot	Bunny	Bob	Kitten
Initial mesh	Icosphere (1280 faces)			Torus (3200 faces)	
Final mesh	no refine.	18680	17908	19308	20904
Run time	244.1 (sec.)	730.7	719.6	816.3	816.8

5.2. Application to electron tomography

The goal of this experiment is to estimate the shape of a bimetallic nanoparticle having a Au-core and a Ag-shell nano particles. We obtained 2 tilt series datasets of a nano triangular bipyramid and a nanocube using high-angle annular dark-field scanning transmission electron microscopy (HAADF-STEM). We hereby used a Thermo Fisher Tecnai Osiris electron microscope generated at 200kV. Each dataset contains 49 projection images acquired over $\pm 72^\circ$ with a tilt increment of 3° and a frame time of 4 seconds. This small range of projection angles is typical for electron tomography and makes the reconstruction challenging. Fig. 10 shows three images for each of the two datasets.

Figure 10: Projection images along 3 different tilt angles for a nano triangular bipyramid (top) and a nano cube (bottom).

To test the proposed method on the data, we initialize the meshes as two icospheres having total 2,560 faces. We use the collision detection method in (Lauterbach et al., 2009), to make sure that the core and the shell part do not collide. In this experiment, we set the number of iterations as 300, the step size τ as 0.005 and observe no collisions when a proper regularization weight is used. As the unknown shapes are relatively simple, we impose only the Laplacian regularization term, by setting $\alpha = 15$ for the nano bipyramid and $\alpha = 5$ for the nanocube. The nanocube data is not much affected by this parameter α , but for the nano bipyramid data, the parameter α should be greater than or equal to 10 to avoid the collision and obtain a high quality mesh when a small range of tilt angles is used. As for the initialization of two icospheres, the inner icosphere’s size is set to half the size of the outer icosphere. This outer icosphere is initialized with a relative radius 0.4 (when the object space is normalized to $(-1, 1)^3$), but we obtain a similar result with the radius range of 0.2 ~ 0.8 for the nano cube and 0.2 ~ 0.5 for the nano triangular data.

Fig. 11 and Fig. 12 show the 3D reconstruction for the nano triangular bipyramid and the nano cube data, respectively, by SIRT, TV, and the proposed method. To evaluate the effect of tilt angle range, we test two cases: an angular range of $\pm 72^\circ$ and a subsampled tilt angles of $\pm 18^\circ$. To compare our mesh result to the voxel-based methods, we extract meshes from the reconstruction images of SIRT and TV, by median filtering, thresholding, filling holes, and extracting the isosurfaces. For the highly-limited angle case with a tilt range of $\pm 18^\circ$, we additionally apply Gaussian smoothing on the results of SIRT and TV before the thresholding step, to obtain the visually-appealing meshes.

For the tilt angle ranges of $\pm 72^\circ$, the results of SIRT, TV, and the proposed method are in good comparison. However, when the range of tilt angles is reduced to $\pm 18^\circ$, SIRT and TV yield a degenerated reconstruction, which also affects the quality of the final extracted surfaces. On the other hand, the proposed method is shown to be less affected by a tilt range. For the bipyramid in Fig. 11, the volume of the Au-core decreased for the limited angle case, due to the high regularization effect. In Table 2, we provide the estimated volumes of the Au-core and Ag-shell particles by the proposed method and the extracted isosurfaces from the images by SIRT and TV.

Table 2: Estimated volumes in the unit of 10^3 nm by SIRT reconstruction image followed by isosurface extraction, and the proposed method.

Tilt angles	Method	bipyramid		nano cube	
		core	shell	core	shell
$-72^\circ \sim 72^\circ$	SIRT+Iso.	3.6	33.2	2.4	29.8
	TV+Iso.	3.7	33.5	2.5	29.8
	Proposed	4.2	36.7	2.3	31.2
$-18^\circ \sim 18^\circ$	SIRT+Iso.	1.3	31.5	1.2	14.9
	TV+Iso.	3.5	25.0	1.3	11.2
	Proposed	2.7	37.7	2.5	29.8

6. Conclusion and Discussion

In this paper, we have developed a forward projector to map a triangular mesh onto the projections domain which is differentiable with respect to the mesh vertices and the attenuation values. Based on our differentiable forward model, we have suggested an optimization-based shape estimation method from projections. Our model should be chosen for reconstructing homogeneous objects with relatively simple geometry, but in situations where the data foundation is so noisy or limited that other methods will not allow a reconstruction. While a conventional approach most often consists of the two steps of image reconstruction followed by a segmentation to obtain a surface, the proposed method directly yields the surface, which among others allows imposing a shape prior directly on the reconstructed shapes. Our experiments on synthetic data show that the proposed method is robust to noise when reconstructing single objects. The experiments on the electron tomography data show how our method is robust even when the range of tilt angles is highly limited, compared with SIRT and TV. In this challenging case, the conventional image reconstructions by SIRT and

TV are highly degraded and these degenerate results propagate errors to the surface estimation from the reconstructed images. Our method can, however, capture the overall shapes well with the regularization term which can impose a smoothness prior directly on the object shapes.

To discuss the limitation of the proposed method, we divide our contributions into two parts: the differentiable forward model and the shape estimation method. The proposed forward model estimates the displacement direction of each vertex for minimizing the objective function. These displacements will only allow change in shape and not collisions of splitting of the mesh that would require topological changes. Handling such topological changes and collisions is itself an active research topic in mesh deformation in computer graphics or computer vision. Investigating how to handle topology changes for a more general mesh adaption and deformation is for future research.

The current differentiable forward model has only been tested on parallel beam geometry where rays are perpendicular to the detector plane. This can be extended to a more general setting such as the cone beam geometry (Buzug, 2008), but we have not investigated the effect of that, as electron tomography data are based on parallel beam geometry.

As a potential application, the proposed method can be useful for estimating the geometry precisely, e.g. for registering the reconstructed object to a 3D CAD model. Another future application can be dynamic tomography where we aim to analyze the objects which may change during scanning.

Acknowledgments

This work is funded by EU Horizon 2020 MSCA Innovative Training Network MUMMERING Grant Number 765604.

References

- Aghasi, A., Kilmer, M., Miller, E.L.. Parametric Level Set Methods for Inverse Problems. *SIAM J Imaging Sci* 2011;4(2).
- Alvino, C.V., Yezzi, A.J.. Tomographic reconstruction of piecewise smooth images. In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2004. .
- Andersen, A.H., Kak, A.C.. Simultaneous Algebraic Reconstruction Technique (SART): A superior implementation of the ART algorithm. *Ultrason Imaging* 1984;6(1).
- Attene, M.. A lightweight approach to repairing digitized polygon meshes. *Vis Comput* 2010;26(11).
- Bonin, A., Chalmond, B., Lavayssière, B.. Monte-Carlo simulation of industrial radiography images and experimental designs. *NDT & E International* 2002;35(8).
- Buzug, T.M.. *Computed Tomography from Photon Statistics to Modern Cone Beam CT*. Springer, 2008.
- Chambolle, A., Pock, T.. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *J Math Imaging Vis* 2011;40(1).
- Chen, W., Gao, J., Ling, H., Smith, E.J., Lehtinen, J., Jacobson, A., Fidler, S.. Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer. In: *Neural Information Processing Systems*. 2019. .
- Dahl, V.A., Dahl, A.B., Hansen, P.C.. Computing segmentations directly from x-ray projection data via parametric deformable curves. *Meas Sci Technol* 2018;29(1).
- Elangovan, V., Whitaker, R.T.. From Sinograms to Surfaces: A Direct Approach to the Segmentation of Tomographic Data. In: *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. 2001. .
- Eliasof, M., Sharf, A., Treister, E.. Multi-modal 3D shape reconstruction under calibration uncertainty using parametric level set methods. *SIAM J Imaging Sci* 2020;13(1).
- Freud, N., Duvauchelle, P., Létang, J.M., Babot, D.. Fast and robust ray casting algorithms for virtual X-ray imaging. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 2006;248(1).
- Gadella, M., Wang, R., Maji, S.. Shape Reconstruction Using Differentiable Projections and Deep Priors. In: *International Conference on Computer Vision*. 2019. .
- Huang, J., Su, H., Guibas, L.. Robust Watertight Manifold Surface Generation Method for ShapeNet Models. *ArXiv180201698 Cs* 2018;.
- Kadu, A., van Leeuwen, T., Batenburg, K.J.. A parametric level-set method for partially discrete tomography. In: *International Conference on Discrete Geometry for Computer Imagery*. 2018. .
- Kass, M., Witkin, A., Terzopoulos, D.. Snakes: Active contour models. *Int J Comput Vis* 1988;1(4).
- Kato, H., Ushiku, Y., Harada, T.. Neural 3D Mesh Renderer. In: *IEEE Conference on Computer Vision and Pattern Recognition*. 2018. .
- Kingma, D., Ba, J.. Adam: A method for stochastic optimization. In: *International Conference on Learning Representations*. 2015. .
- Lauterbach, C., Garland, M., Sengupta, S., Luebke, D.P., Manocha, D.. Fast BVH Construction on GPUs. *Comput Graph Forum* 2009;28(2).
- Liu, S., Chen, W., Li, T., Li, H.. Soft Rasterizer: Differentiable Rendering for Unsupervised Single-View Mesh Reconstruction. In: *International Conference on Computer Vision*. 2019. .
- Loper, M.M., Black, M.J.. OpenDR: An Approximate Differentiable Renderer. In: *European Conference on Computer Vision*. 2014. .
- Marinovszki, Á., Beenhouwer, J.D., Sijbers, J.. An efficient CAD projector for X-ray projection based 3D inspection with the AS-TRA Toolbox. In: *Conference on Industrial Computed Tomography*. 2018. .
- Midgley, P., Weyland, M.. 3d electron microscopy in the physical sciences: the development of z-contrast and efem tomography. *Ultramicroscopy* 2003;96(3):413–431.
- Mueller, J.L., Siltanen, S.. *Linear and Nonlinear Inverse Problems with Practical Applications*. volume 10. Society for Industrial and Applied Mathematics, 2012.

- Mumford, D., Shah, J.. Optimal approximations by piecewise smooth functions and associated variational problems. *Commun Pure Appl Math* 1989;42(5).
- Osher, S., Fedkiw, R.. Level set methods and dynamic implicit surfaces. *Appl Mech Rev* 2004;57(3).
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.. PyTorch: An imperative style, high-performance deep learning library. In: *Neural Information Processing Systems*. 2019. .
- Roelandts, T., Batenburg, K.J., den Dekker, A.J., Sijbers, J.. The reconstructed residual error: A novel segmentation evaluation measure for reconstructed images in tomography. *Comput Vis Image Underst* 2014;126.
- Sidky, E.Y., Jørgensen, J.H., Pan, X.. Convex optimization problem prototyping for image reconstruction in computed tomography with the Chambolle–Pock algorithm. *Phys Med Biol* 2012;57(10).
- Sujar, A., Meuleman, A., Villard, P.F., Garcia, M., Vidal, F.. gVirtualXRy: Virtual x-ray imaging library on GPU. In: *Computer Graphics and Visual Computing (CGVC)*. 2017. .
- Vidal, F.P., Garnier, M., Freud, N., Létang, J.M., John, N.W.. Simulation of X-ray attenuation on the GPU. In: *Proceedings of Theory and Practice of Computer Graphics*. Cardiff, UK: Eurographics Association; 2009. .
- Wang, N., Zhang, Y., Li, Z., Fu, Y., Liu, W., Jiang, Y.G.. Pixel2Mesh: Generating 3D Mesh Models from Single RGB Images. In: *European Conference on Computer Vision*. 2018. .
- Whitaker, R.T., Elangovan, V.. A direct approach to estimating surfaces in tomographic data. *Med Image Anal* 2002;6(3).

Figure 11: Reconstruction results on nano triangular bypyramid data. The first and the 4th rows show the reconstruction results by SIRT and from projections with tilt angles from -72° to 72° and -18° to 18° , respectively, where two central slices are visualized in the first and second column. The 2nd and the 5th rows show the reconstruction results by TV. The last column shows the extracted mesh from the reconstructed image by SIRT or TV with some post-processing procedures. The third and the last row show our direct shape estimation results, where two central slices are visualized in blue and yellow.

Figure 12: Reconstruction results on nano cube data. The first and the 4th rows show the reconstruction results by SIRT and from projections with tilt angles from -72° to 72° and -18° to 18° , respectively, where two central slices are visualized in the first and second column. The 2nd and the 5th rows show the reconstruction results by TV. The last column shows the extracted mesh from the reconstructed image by SIRT or TV with some post-processing procedures. The third and the last row show our direct shape estimation results, where two central slices are visualized in blue and yellow. The results on the second and the fourth row appear to be the same, but they are slightly different.