Learning the Inverse Weighted Radon Transform

Philipp Roser[1,2], Lina Felsner[1,2], Andreas Maier[1], Christian Riess[1]

[1]Pattern Recognition Lab, FAU Erlangen-Nürnberg
[2]Both authors contributed equally
lina.felsner@fau.de

Abstract. X-ray phase-contrast imaging enhances soft-tissue contrast. The measured differential phase signal strength in a Talbot-Lau interferometer is dependent on the object's position within the setup. For large objects, this affects the tomographic reconstruction and leads to artifacts and perturbed phase values. In this paper, we propose a pipeline to learn a filter and additional weights to invert the weighted forward projection. We train and validate the method with a synthetic dataset. We tested our pipeline on the Shepp-Logan phantom, and found that our method suppresses the artifacts and the reconstructed image slices are close to the actual phase values quantitatively and qualitatively. In an ablation study we showed the superiority of our fully optimized pipeline.

1 Introduction

Medical applications could benefit from the high soft-tissue contrast of X-ray phase contrast imaging techniques [2]. Especially, the Talbot-Lau interferometer (TLI) is a promising setup to acquire phase contrast images in a medical context, due to the comparably low overall system requirements and the high robustness of the setup [1]. The TLI setup contains three gratings that are placed between the source and detector (Fig. 1).

Engelhard et al. reported a correlation between the object magnification and the measured signal strength of the differential phase image in the TLI [3]. Donath et al. confirmed these findings, and clarified that the measured phase value depends on the position of the object relative to the phase grating G1 [4]. More concretely, the interferometer's angular sensitivity is given as

$$S = \frac{1}{2\pi}\frac{\Delta\varphi}{\alpha} \ , \tag{1}$$

where $\Delta\varphi$ is the measured intensity oscillation of the phase scan (normalized to 2π) and α is the refraction angle caused by the object.

The position-dependent sensitivity plays a critical role in the tomographic reconstruction of large objects. The differential phase contrast at tomographic angle θ and detector position t is [5]

$$\varphi(\theta, t) = \frac{\partial}{\partial t}\int_{-\infty}^{\infty} S(r)\delta(t, r)\mathrm{d}r \ , \tag{2}$$

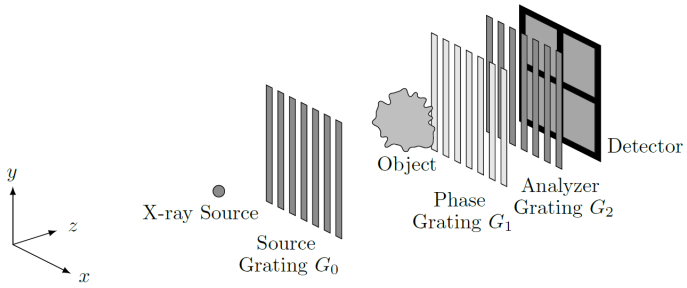

Fig. 1. Schematic setup of a Talbot-Lau interferometer. From [1].

where $\delta(t, r)$ encodes the objects phase values, and integration is along the ray direction r, which depends on the tomographic angle θ. In general, the position-dependent sensitivity changes the forward projection process, such that it is not possible to use conventional filtered backprojection (FBP) algorithms. Chabior et al. [5] first discussed the changing contrast formation in tomographic imaging. One particular effect is that a parallel-beam geometry requires a circular trajectory of 2π instead of π, and Chabior et al. show that a trajectory over π leads to severe reconstruction artifacts. This shows that the reconstruction task changes such that the standard analytic inversion is not possible. Furthermore, Felsner et al. showed that the task can also not be exactly solved if it is split into a standard general-purpose reconstruction and another part that is specific to the differential phase [6]. In view of this challenge to find a direct solution, we propose as an alternative a data-driven approach to differential phase reconstruction. One particularity of our approach is that we propose a specialized neural network architecture that integrates domain knowledge to remain true to the physical model [7]. For standard CT, this has been used to learn redundancy weights [8] or dedicated reconstruction filters [9]. In this work, we learn a tailored reconstruction filter, normalization weights, and voxel weights for differential phase-contrast, and experimentally show that it provides highly accurate reconstructions.

2 Methods

This work investigates the inverse weighted radon transform, more specifically its reconstruction filter and additional pixel and voxel weights. Although the TLI measures the differential phase signal, for simplicity we consider in this work direct refraction angles. We describe the network architecture in Sec. 2.1, introduce our training and test data in Sec. 2.2, and describe the experimental setup in Sec. 2.3.

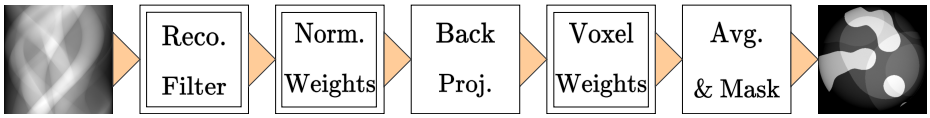

Fig. 2. Proposed Pipeline. The highlighted blocks contain parameters that can be trained in data-driven fashion.

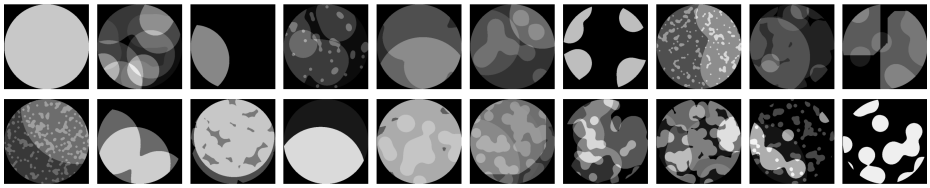

Fig. 3. Exemplary synthetic blobs phantoms used for training our model.

2.1 Inverse Weighted Radon Transform

Our approach is similar to previous works, where either redundancy weights [8] or filter kernels [9] are estimated in a data-driven fashion. To this end, we model the reconstruction task with a neural network (NN), where each step in the pipeline is interpreted as a layer. A schematic sketch of the pipeline can be found in Fig. 2.

The input of the NN is a sinogram that contains line integrals of weighted phase values. The first layer corresponds to the reconstruction filter $\boldsymbol{H}$, which is carried out in the Fourier domain. Then, a layer of pixel-wise multiplicative normalization weighting $\boldsymbol{W}_{\mathrm{N}}$ is added. The next layer represents the (analytical) row-wise parallel backprojection. Afterwards, a voxel-wise multiplicative weighting layer $\boldsymbol{W}_{\mathrm{V}}$ accounts for the position of the phase values in the volume. The final step comprises averaging of the row-wise backprojections and masking of the reconstructed volume. The filter and the weighting layers (as highlighted in Fig. 2) can be optimized.

2.2 Data

For training, we used 1000 synthesized two-dimensional phantoms of 400×400 voxels. Each phantom is a superposition of up to six randomly scaled and sized distributions of binary blobs [10]. Figure 3 shows several examples. Based on the phantoms, we computed an ideal sinogram $\boldsymbol{P} \in \mathbb{R}^{m \times n}$ without voxel weights, and a sinogram $\tilde{\boldsymbol{P}} \in \mathbb{R}^{m \times n}$ with voxel weights $w_{i,j} \in [0.1, 0.9]$, with $m = 400$ projections comprising $n = 400$ pixels acquired over π. We use the weighted sinogram $\tilde{\boldsymbol{P}}$ as input to our method and the FBP of the ideal sinogram $\boldsymbol{P}$ as target for optimizing the model parameters. For validation, we used 20 % of the training data. To evaluate our method's performance on unseen data, we selected the Shepp-Logan phantom [11].

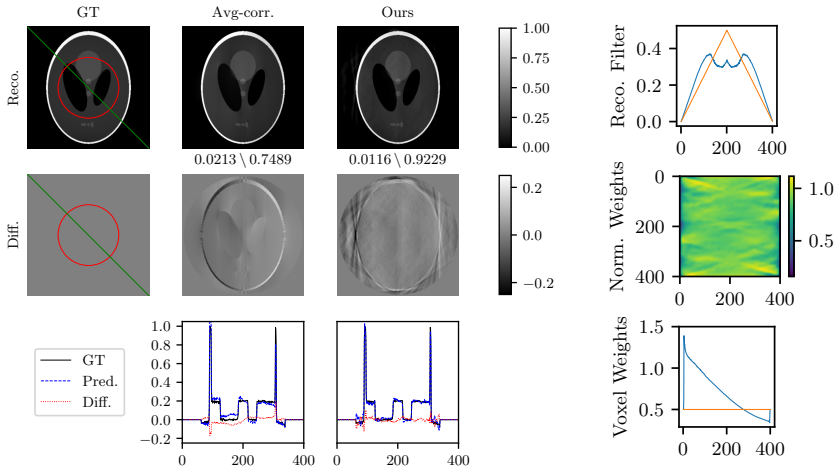

(a) Reconstruction results (b) Filter and weights

Fig. 4. Results of the baseline method (Avg. corr.) and our trained model (Ours): (a) Reconstructions, differences to the ground truth, and corresponding diagonal line profiles. MAE and SSIM (<MAE> \ <SSIM>) refer to the circular ROI. (b) Averaged reconstruction filter (top), the learned normalization weights (mid), and voxel weights (bottom) of our trained model (blue) and the baseline/initialization (orange).

2.3 Experiments

We initialized the reconstruction filter with the Ram-Lak filter [12], which is the optimal filter for the conventional reconstruction problem. The normalizing weights were initialized with ones and the voxel weights were initialized homogeneously to the average of the weight interval, i.e., $\bar{w} = 0.5$ on our data. Overall, we performed three experiments. First, as a baseline, we used the the pipeline with the initialized weights for reconstruction. This correspond to an offset corrected FBP reconstruction (conf. [5]). Second, in our method, the filter and weights of the baseline were optimized on the training data (Sec. 2.2). Third, to provide an ablation, we also separately trained for (i) a projection-wise reconstruction filter, (ii) a global reconstruction filter, (iii) the normalizing weights, and (iv) the voxel weights. Where applicable, we optimized the free parameters with respect to the mean absolute error (MAE) using stochastic gradient descent with adaptive moments [13] with 10^{-4} learning rate. In addition to the MAE, we investigate the structural similarity index (SSIM) concerning our test dataset within a centered ciruclar region of interest (ROI).

3 Results

The reconstruction of the Shepp-Logan phantom is shown for the average-corrected baseline and our proposed method in top row of Fig. 4(a). The difference images to the ideal reconstruction are shown below. We found that our proposed method

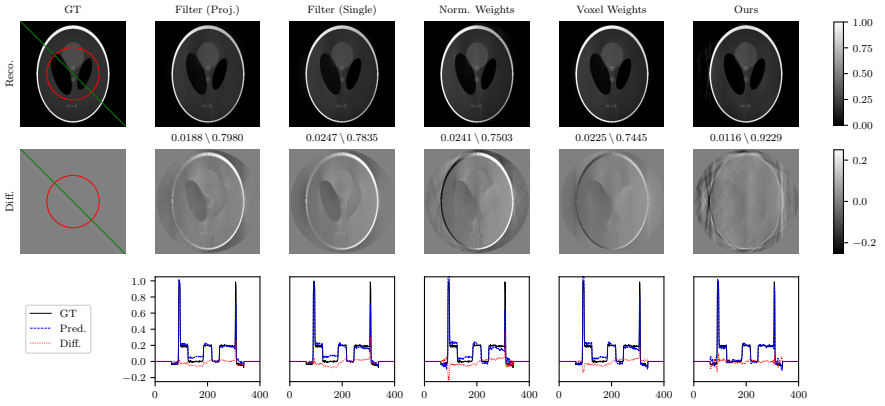

Fig. 5. Ablation study for the proposed pipeline. Reconstructions, differences to the ground truth, and corresponding diagonal line profiles. The MAE and SSIM (<MAE> \ <SSIM>) are given for the circular ROI.

remarkably reduces the error in the ROI compared to the average-corrected reconstruction. This can be seen from the line plots in Fig. 4(a), as well as from the quantitative values. The MAE is almost reduced by half, and the SSIM is increased by over 20 %. Figure 4(b) (top) shows the learned reconstruction filter averaged over all projection angles in blue. The comparison to the Ram-Lak (orange) shows that especially the high-frequency components get reduced in favor of low to medium frequencies. Figure 4(b) (mid) shows the normalization weights that were applied in the projection domain. Although we found that the sinogram is nearly point-symmetric, it is interesting to note that the weights along one projection are *not* symmetric. Figure 4(b) (bottom) shows the voxel weights along (each) ray, plotted from the source to the detector. As expected, we observed a strong dependence of voxel weights on the position along the ray.

The results of the ablation study are shown in Fig. 5. For all the configurations, we show from top to bottom reconstruction, difference image, and line plot. Overall, we found that the pipeline with different single layers does not improve the reconstruction. The MAE and SSIM are close to MAE and SSIM for the average-corrected reconstruction (Fig. 4(a)). Especially the offset of the left air cavity is prominent in the line plots, and only reduced with our proposed pipeline.

4 Discussion and Outlook

In a proof-of-concept study, we showed that optimizing reconstruction parameters in a data-driven manner has the potential to improve generally ill-posed reconstruction problems. We improved average error rates and SSIM values by large margins without utilizing costly iterative approaches. Still, our approach preserves the interpretability of the conventional reconstruction pipeline and somewhat ensuring data integrity. However, there are certain limitations that

need to be addressed. First, this study is performed on simulated data. Second, especially close to the border of the reconstructed volume, our method introduced considerable artifacts, whose origin and severity need to be further investigated. Furthermore, we investigated direct phase measurements instead of differential phase, which shall be investigated in the future. In addition, we want to point out promising directions for future studies. We believe the combination of a bilateral or guided filter, or additional regularization with our method is desirable. Using a backprojection filtering approach instead of the FBP-like pipeline aims in the same direction. Also, the learned filter and weights can give insight of a theoretically/physically sound analytical solution.

References

1. Maier A, Steidl S, Christlein V, et al. Medical imaging systems: An introductory guide. vol. 11111. Springer; 2018.
2. Pfeiffer F, Weitkamp T, Bunk O, et al. Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources. Nat Phys. 2006;2(4):258–261.
3. Engelhardt M, Baumann J, Schuster M, et al. High-resolution differential phase contrast imaging using a magnifying projection geometry with a microfocus X-ray source. Appl Phys Lett. 2007;90(22):224101.
4. Donath T, Chabior M, Pfeiffer F, et al. Inverse geometry for grating-based X-ray phase-contrast imaging. J Appl Phys. 2009;106(5):054703.
5. Chabior M, Schuster M, Schroer C, et al. Grating-based phase-contrast computed tomography of thick samples. Nucl Instrum Methods Phys Res A. 2012;693:138–142.
6. Felsner L, Würfl T, Syben C, et al. Reconstruction of voxels with position- and angle-dependent weightings. In: The 6th Int. Conf. on Image Formation in X-Ray Computed Tomography; 2020. p. 502 – 505.
7. Maier AK, Syben C, Stimpel B, et al. Learning with known operators reduces maximum error bounds. Nat Mach Intell. 2019;1(8):373–380.
8. Würfl T, Hoffmann M, Christlein V, et al. Deep learning computed tomography: Learning projection-domain weights from image domain in limited angle problems. IEEE Trans Med Imaging. 2018;37(6):1454–1463.
9. Syben C, Stimpel B, Roser P, et al. Known operator learning enables constrained projection geometry conversion: Parallel to cone-beam for hybrid MR/X-ray imaging. IEEE Trans Med Imaging. 2020;.
10. Van der Walt S, Schönberger JL, Nunez-Iglesias J, et al. Scikit-image: Image processing in Python. PeerJ. 2014;2:e453.
11. Shepp LA, Logan BF. The Fourier reconstruction of a head section. IEEE Trans Nucl Sci. 1974;21(3):21–43.
12. Ramachandran GN, Lakshminarayanan AV. Three-dimensional reconstruction from radiographs and electron micrographs: Application of convolutions instead of Fourier transforms. Proc Natl Acad Sci. 1971;68(9):2236–2240.
13. Kingma DP, Ba J. Adam: A method for stochastic optimization. In: Bengio Y, LeCun Y, editors. 3rd Int. Conf. on Learning Representations; 2015. p. 1–15.