# Deriving Neural Network Architectures using Precision Learning: Parallel-to-fan beam Conversion

Christopher Syben<sup>1,2</sup>, Bernhard Stimpel<sup>1,2</sup>, Jonathan Lommen<sup>1,2</sup>, Tobias Würfl<sup>1</sup>, Arnd Dörfler<sup>2</sup>, and Andreas Maier<sup>1</sup>

<sup>1</sup> Pattern Recognition Lab, Department of Computer Science,

Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany $^2\,$ Department of Neuroradiology, Universitätsklinikum Erlangen,

Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany

Abstract. In this paper, we derive a neural network architecture based on an analytical formulation of the parallel-to-fan beam conversion problem following the concept of precision learning. The network allows to learn the unknown operators in this conversion in a data-driven manner avoiding interpolation and potential loss of resolution. Integration of known operators results in a small number of trainable parameters that can be estimated from synthetic data only. The concept is evaluated in the context of Hybrid MRI/X-ray imaging where transformation of the parallel-beam MRI projections to fan-beam X-ray projections is required. The proposed method is compared to a traditional rebinning method. The results demonstrate that the proposed method is superior to ray-by-ray interpolation and is able to deliver sharper images using the same amount of parallel-beam input projections which is crucial for interventional applications. We believe that this approach forms a basis for further work uniting deep learning, signal processing, physics, and traditional pattern recognition.

Keywords: Machine Learning  $\cdot$  Precision Learning  $\cdot$  Hybrid MRI/X-ray imaging

# 1 Introduction

Deep learning is a game-changer in many perceptual tasks ranging from image classification over segmentation to localization [2]. A major disadvantage of perceptual problems is that no prior knowledge on how the classes and labels are obtained is available. As such a large body of literature exists that investigates different network topologies for different applications. As result, we managed to replace hand-crafted features with hand-crafted networks.

Recently, these techniques also emerge to other fields in signal processing. One of them is medical image reconstruction in which surprising results have been obtained [16,4]. For signal processing, however, we do have prior knowledge available that can be reused in the network design. The use of these prior operators reduces the number of unknowns of the network, therewith the amount of required training samples, and the maximal training error bounds [6]. Up to now, this *precision learning* approach was only used to augment networks with prior knowledge and or to add more flexibility into existing algorithms [14, 15, 10, 3]. In this paper, we want to extend this approach even further: we demonstrate that we can derive a mathematical model to tackle a problem under consideration and use deep learning to formulate different hypothesis on efficient solution schemes that are then found as the point of optimality of a deep learning training process.

In particular, we aim in this paper at an efficient convolution-based solution for parallel-to-fan-beam conversion. Up to now, such an efficient algorithm was unknown and the state-of-the-art to address this problem is rebinning of rays that is inherently connected to interpolation and a loss of resolution.

The problem at hand is not only interesting in terms of algorithmic development, it also has an immediate application. Novel hybrid medical scanners will be able to combine Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) in a single device for interventional applications [8, 13]. While CT offers high spatial and temporal resolution, MRI allows for the visualization of soft-tissue contrast, vessels without the use of contrast agent, and there is no need for harmful ionizing radiation.

However, acquisition on MR devices is slow compared to CT. Flat-panel detectors allow image-guided interventions using fluoroscopic projection images that can be acquired at high frame-rates with up to 30 frames per second. This is a challenging time constraint for MRI. Recent developments indicate that MRI is also able to perform projection imaging at acceptable frame rates [5]. Yet the two modalities are inherently incompatible, as MRI typically operates in a parallel projection geometry and X-rays emerge from a source point that restricts them to fan- and cone-beam geometries.

Recent publications elaborate on the idea of MR/X-ray projection fusion and extend the MR acquisition such that the final MR-projection image shows the same perspective distortion as the X-ray projection [12, 11, 5]. Current approaches, however, rely on rebinning that requires interpolation which inherently reduces the resolution of the generated images. In this paper, we propose to derive an image rebinning method from the classical theory. However, as this would require an expensive inverse of a large matrix, we propose to replace the operation with a highly efficient convolution that is inspired by the classical filtered back-projection solution in CT. Here, we examine two cases for this convolution: a projection-independent and a projection-dependent one.

# 2 Methods

In the first section we shortly describe the link between X-ray and MRI projections using rebinning [11], afterwards we revisit the discrete form of the reconstruction problem, which is then followed by our proposed problem description. Subsequently the network topology will be derived following the precision learning paradigm. This section is concluded by a description of the training process and the used training data.

#### 2.1 Linking MRI and X-ray Acquistion

The link between the X-ray and MRI acquisition is given by the central slice theorem. This has first been demonstrated by Syben et al. [11] for simulation data and was later applied for the construction of X-ray projections from MRI measurement data [5]. Their approach is inspired by the geometric rebinning method which allows the reconstruction of fan-beam data by resampling the fan-beam acquisition to a parallel-beam acquisition.

They follow the central slice theorem which states that the Fourier transform of a 1D projection of a 2D object can be found in the 2D Fourier transform of the object along a radial line with the same orientation as the detector. Because the MRI can sample the Fourier transform of the object, parallel projections can be acquired. This relationship combined with the geometric rebinning method can be used to convert a set of parallel projections to one fan-beam projection as shown in Fig 1.

In their publication they analyze the sub-sampling capability of this method. In this context, full sampling means that the MR device acquires one parallel projection for each fan-beam detector pixel. Thus sub-sampling is related to the case where less parallel projections are acquired with respect to the number of fan-beam detector pixels. They show that only few projections are necessary to create the target fan-beam projection with a small error [11]. Following their geometric rebinning method two steps of interpolation in spatial domain are required: first an interpolation between two projections with different projection angles is carried out followed by an interpolation between the pixels of the parallel-beam projection.

#### 2.2 The Tomographic Reconstruction Problem

The CT imaging procedure from acquiring X-ray projections to the reconstructed object information can be described in discrete linear algebra. The acquisition of the projection images of the object can be described with

$$\boldsymbol{A}\boldsymbol{x} = \boldsymbol{p} \quad , \tag{1}$$

where A is the system matrix describing the geometry of the imaging system. x is the object itself and p are the projections of x under the described geometry A. Correspondingly the reconstruction can be obtained with

$$\boldsymbol{x} = \boldsymbol{A}^{-1} \boldsymbol{p} \quad , \tag{2}$$

where  $A^{-1}$  is the inverse of the system matrix, which can not be inverted since it is a tall matrix. Thus, the reconstruction is conducted using the left-side pseudo





Fig. 1: MRI to X-ray link and geometric rebinning method proposed by Syben et al. [11].

inverse which gives the approximation with minimal distance to the inverse in a  $\ell 2\text{-norm}$  sense.

$$\boldsymbol{x} = \boldsymbol{A}^{\top} (\boldsymbol{A} \boldsymbol{A}^{\top})^{-1} \boldsymbol{p}$$
(3)

where  $\mathbf{A}^{\top}$  is the transposed system matrix, which can be algorithmically described as the back-projection operator. For a full scan with 180° of rotation in parallel geometry, the inverse bracket is a filtering step in the Fourier domain and can be described as

$$\boldsymbol{x} = \boldsymbol{A}^{\top} \boldsymbol{F}^{\boldsymbol{H}} \boldsymbol{K} \boldsymbol{F} \boldsymbol{p} \tag{4}$$

where  $F, F^H$  is the Fourier and inverse Fourier transform, respectively. K is the so called Ramp filter represented in a diagonal matrix. Together the pseudo inverse this describes the filtered back-projection algorithm in a discrete fashion.

## 2.3 Rebinning using Tomographic Reconstruction

As shown in [10], the discrete description of the reconstruction problem can be used to derive a network topology and to learn the reconstruction filter. In the following, we use this idea to derive an optimization problem to find a filter which can be used to transform several parallel projections to one fan-beam projection. A fan-beam projection can be created by

$$\boldsymbol{A}_f \boldsymbol{x} = \boldsymbol{p}_f \quad , \tag{5}$$

where  $A_f$  describes the system matrix for a fan-beam projection and  $p_f$  is the respective projection. The necessary parallel projections which contain the information for the fan-beam projection can be found in the Fourier domain (or K-space of the MRI system) in a wedge region [11] which is defined by the fan angle of the fan-beam geometry. These parallel projections can be described with

$$\boldsymbol{A}_{p}\boldsymbol{x} = \boldsymbol{p}_{p} \quad , \tag{6}$$

where  $A_p$  is the system matrix generating the projections  $p_p$  from object x under the parallel-beam geometry. The object x in Eq. 5 can be substituted by the reconstruction using the inverse of the system matrix and the projections from Eq. 3 in Section 2.2:

$$\boldsymbol{A}_{f} \underbrace{\boldsymbol{A}_{p}^{\top} (\boldsymbol{A}_{p} \boldsymbol{A}_{p}^{\top})^{-1} \boldsymbol{p}_{p}}_{\text{Parallel reconstruction } \boldsymbol{x}} = \boldsymbol{p}_{f} \quad . \tag{7}$$

In principle, the above equation is hard to solve, as the reconstruction task from this very small set of projections is ill-posed and there is no analytical closed-from solution known. However, we now simply postulate that there exists a projection-independent filter which is a close approximation the above inverse bracket. As in Section 2.2, this allows us to express the solution as a multiplication with an diagonal filter matrix  $\boldsymbol{K}$  in Fourier domain:

$$\boldsymbol{A}_{f}\boldsymbol{A}_{p}^{\dagger}\boldsymbol{F}^{H}\boldsymbol{K}\boldsymbol{F}\boldsymbol{p}_{p}=\hat{\boldsymbol{p}}_{f} \quad , \tag{8}$$

where  $\hat{p}_f$  is the approximated fan-beam projection under the above stated assumption. Now the only unknown operation in above equation is K that can be determined using an objective function:

$$f(\boldsymbol{K}) = \frac{1}{2} \|\boldsymbol{A}_{f} \boldsymbol{A}_{p}^{\top} \boldsymbol{F}^{\boldsymbol{H}} \boldsymbol{K} \boldsymbol{F} \boldsymbol{p}_{p} - \boldsymbol{p}_{f}\|_{2}^{2} .$$

$$(9)$$

The gradient of function f is with respect to K is

$$\frac{\partial f(\boldsymbol{K})}{\partial \boldsymbol{K}} = \boldsymbol{F} \boldsymbol{A}_{p} \boldsymbol{A}_{f}^{\top} (\boldsymbol{A}_{f} \boldsymbol{A}_{p}^{\top} \boldsymbol{F}^{\boldsymbol{H}} \boldsymbol{K} \boldsymbol{F} \boldsymbol{p}_{p} - \boldsymbol{p}_{f}) (\boldsymbol{F} \boldsymbol{p}_{p})^{\top} \quad . \tag{10}$$

Note that this gradient is determined automatically by back-propagation to update the weights of layer K, if Eq.8 is implemented by means of a neural network as already observed for a different application in [10]. Thus, the network topology for a network which learns the transformation from several parallel projections to one fan-beam projection could be derived by the presented approach.

#### 2.4 Network

The network topology can be directly derived from the description of the objective function in Eq. 9 and is shown in Fig. 2. Projectors and back-projectors are scaled to each other in terms of sampling density and number of projections. Since we mix a parallel back-projector with a fan-beam forward projector and aim at different sampling densities we added an additional scaling layer S to the network to compensate accordingly.



Fig. 2: Network based on objective function (Eq. 9).

## **Implementation Details**

We have implemented the network using Tensorflow [1]. Thus, the Fourier and inverse Fourier transform are layers provided by the Tensorflow framework. The parallel projector and back-projector as well as the fan-beam projector and backprojector are unmatched pairs and are implemented as custom ops in Tensorflow using Cuda kernels. For the back-propagation the respective operation is assigned to the layers for the gradient calculation.

## 2.5 Training Process

## **Training Data**

For the training we use numerical phantoms which bring their different characteristics into the training process (Fig.3). The first type of phantoms are homogeneous objects that fill the field of view like ellipses and circles. The second type contain a homogeneous field of view filling ellipse with contains varying number of elongated ellipses (in the following called bars). The third type of phantoms uses only bars without the surrounding ellipsoid. As a last type we use phantoms which contain normal distributed noise.

In the following list, the number test phantoms are listed:

- 1 Ellipse phantom
- 1 Circle phantom
- -8 Ellipse-bar phantoms (with increasing number of bars from 1 up to 8)
- -5 Bar phantoms (with increasing number of bars from 1 to 5
- 50 Noise phantoms (Normal distributed noise)

The parallel projections and respective label projections (fan-beam) are based on the following geometry:

- Trajectory:  $[0^{\circ}, 25^{\circ}, 45^{\circ}, 65^{\circ}, 90^{\circ}]$
- Source detector distance (SDD): 1200 mm
- Source isocenter distance (SID): 900 mm
- Parallel and fan-beam detector size: 512 pixel
- Reconstruction size:  $256 \times 256$



Fig. 3: Phantom types in training set.

Thus, the training data set consists of partial parallel projections according to the method described in [10] using the given angles in the trajectory for  $\beta$ . The respective label fan-beam projection is generated for each angle of the trajectory for each phantom. All projections are generated using the implemented projection layers. The performance of the network is validated using the Shepp-Logan phantom [9].

## **Training Setup**

The training process is divided into two steps. In the first step, the scaling layer S is trained while K remains fixed. After the training of the scaling layer Sconverges, the scaling factor is fixed and the training of the filter K is started. This separation is based on two thoughts. First, the scaling layer fixes an occurring problem due to the mix up of the different forward- and backprojection geometries and is not part of the unknown operator. The second point is that by dividing the learning process into two parts the learning rate for the scaling layer can be much higher and therefore speed up the whole training process. Furthermore the separation ensures that the calculated loss w.r.t. the label projection can express the deviation from the real fan-beam projection and is not distracted by a scaling factor due to the mixed forward- and backprojection. The filter K is initialized with the Ram-Lak filter [7], which is an optimal discrete reconstruction filter for a complete data acquisition and therefore can be interpret as a strong pre-training of the network. We train on different sub-sampling factors, starting with full sampling and continuing by successively sub-sampling to 15, 7, 5 and 3 projections. This allows us to compare with the geometrical rebinning approach [11].

#### Projection-dependent Vs. Projection-independent

To determine which type of filter performs the best, we performed all experiments on the different sub-sampling levels using both a projection-dependent and a projection-independent of version of K.

#### Regularization

To achieve smooth filter weights we use a Gaussian smoothing after each training epoch.

#### 8 Syben et al.

# 3 Results

The performance of the network is evaluated in three steps. First we analyze the performance for the different sub-sampling stages using the Shepp-Logan phantom and the fan-beam forward projection of the phantom as ground truth (GT). Afterwards, the results are compared with the geometrical approach with certain sub-sampling factors. To provide a better qualitative impression of the performance we subsequently present a comparison based on a 3D phantom using a stacked fan-beam approach. The network performance analysis is followed by a presentation of the learned filter types.

## **Network Performance**

In Fig. 4 the rebinning performance of the learned network for the projectiondependent filter using the Shepp-Logan phantom with different sub-sampling factors is shown. All results show a similar shape as the line profile of the GT projection. The full sampling as well as the sub-sampling case using 15 projections show a noisy behavior. The noise is less for the sub-sampling cases using 7, 5 and 3 projections, respectively. For all versions, except for the case with 7 projections, the rebinned signal overshoots the GT signal at the edges of the object. For the projection-independent version of the filter (Fig. 5) simi-



Fig. 4: Sub-sampling comparison of the projection-dependent filter with the GT projection of the Shepp-Logan phantom. The plot colors are red for the reference, blue for the respective line profile and green for the difference.

lar but strengthened behavior can be observed. For all four rebinning types the projection-independent counterpart is more noisy and overshoots or undershoots more extensively, especially for the rebinning with 5 projections.

However, the noisiness of the 1D plots is misleading as the visual impression of the rebinned MR-projections from the head phantom show in Fig. 6. Even though the noisy behavior of the previous evaluation can be observed in the line profiles of the different sub-sampling methods, the noise level is not the main factor of the observed image impression. The experiment with 15 projections gives a sharp visual impression of the object, although it suffers from the strongest noise. The line profiles of the network trained with 5 and 3 projections show a reduced noise



Fig. 5: Sub-sampling comparison of the projection-independent filter applied on the Shepp-Logan phantom. The plot colors are red for the reference, blue for the respective line profile and green for the difference.

level compared to the network using 15 projections, but high-frequency artifacts and blurriness towards the edges of the image can be observed in the image.



Fig. 6: Sub-sampling comparison of the projection-dependent filter applied on the MR head phantom. The plot colors are red for the reference, blue for the respective line profile and green for the difference.

For the projection-independent filter results, a similar but strengthened behavior can be observed in Fig. 7. The filter for 15 projections provides a similar visual impression as the projection-dependent counterpart. The strength of the noise is stronger for the filter with 7, 5, and 3 projections than their respective projection-dependent counterpart. The high frequency artifacts are much stronger for the case with 5 and 3 projections.

In Fig. 8 both filter types, projection-independent and projection-dependent are compared to the performance of the geometrical rebinning [11]. For the experiment 15 out of the acquired 121 projections of the head phantom are used. Both filters provide a sharper image impression compared to the reference method. In comparison with the geometrical rebinning method the results of



Fig. 7: Sub-sampling comparison projection-independent filter, real data. The plot colors are red for the reference, blue for the respective line profile and green for the difference.

both filters show high frequency artifacts at the edges of the phantom, which can be also seen in the line profiles.

#### **Filter Appearance**

In Fig. 9 the different learned projection-independent filters are shown. The filter using 512 projections is very smooth, while the filters with 15, 7, and 5 projections show high frequency components with a large amplitude. The filter for 3 projections has a high frequency component too, but with a much smaller amplitude. Furthermore, the amplitude of the filter is decreased compared to the initialization and the other filters. The learned projection-dependent filters are shown in Fig. 10. The filter for 512 projections shows in the middle a shape like the projection-independent counterpart, but drops off at the edges. While this is also true for the filter for 15 projections, the filter for 7, 5 and 3 projections are converging towards a U-shape.

## 4 Discussion

The results of the 1D fan-beam projections prove that our proposed analytical description of the rebinning process can be carried out learning the unknown operators in the problem description. The results of the MR head phantom provide a sharper visual impression than the rebinning method proposed Syben et al., although the noise level in the line profiles is much higher. The blurry visual impression of their approach is linked to the necessary interpolation in their method. Especially for image-guided interventions the sharpness is important to provide a clear impression of the vessels and interventional devices. Although the line profile for geometric rebinning overlaps very well with the projection



Fig. 8: Comparison of the geometrical rebinning methods and filter type using 15 projections out of the acquired 121 MR. The reference image with 121 projections is created by the geometrical rebinning method. The plot colors are red for the reference, blue for the respective line profile and green for the difference.

11



Fig. 9: Learned projection-independent filter in frequency domain for different sub-sampling factors.



Fig. 10: Learned projection-dependent filter in frequency domain for different sub-sampling factors.

reference, it must be taken into account that the reference case was already rebinned with this method itself based on all existing MR projections and is, therefore, already smoothed. Overall, using the learned filter based on 15 projections provides the best visual impression, while the amount of necessary MR projections is small. These observation are confirmed by the analysis published in [11]. In general, additional reduction of the number of projections is desirable, which could be achieved by further improving the filter learning process, e.g. by linking it directly to the k-space acquisition scheme.

The results of the 1D projections as well as the stacked fan-beam experiment encourage a detailed discussion of the filter, its shape and the applied regularization. The smoothing after each epoch leads to smooth filter weights for the projection-dependent case and also for the projection-independent filter for the full sampling. However, the smoothing does not enforce a smooth filter function for the projection-independent sub-sampling filters. Especially the 7 and 5 projection case show strong changes in amplitude. In the course of the experiments, we investigated different regularization terms, like the  $\ell$ 2-norm on the filter weights or the  $\ell$ 1-norm of the first derivative of the filter. However, regularization with the aformentioned methods performed not as well as expected. Despite thorough analysis of other regularization terms and corresponding weighting factors, the Gaussian smoothing lead to a more stable learning process and better results. However, a more profound method to achieve smooth filter weights is desirable. For this we started to look closer into regularizing the filter using the Lipschitz continuity. Certainly a more consistent regularization especially for the projection-independent filter has to be found. Such a regularization could open the opportunity to reduce the number of used projections in the rebinning process while preserving the sharp visual impression. Furthermore, introducing a symmetry constraint for the filter could improve the learning behavior and the outcome filter shape, while at the same time the number of parameters which have to be learned are reduced by a factor of 2.

The results lead to several interesting questions which should be considered in further research. The impact of the number of used projections on the rebinning process as well as the covered frequency spectrum of the used phantoms on the filter shape are a promising line for subsequent work. The observed artifacts and the high frequency component in the projection-independent filter could be caused by insufficient coverage of the frequency space in the training process. Also the selection of the projections means a certain coverage of the wedge in the Fourier space as proposed by Syben et al. Furthermore the shape of the projection-dependent filter compared the 512 projection with the 3 projection filter version invites for further experiments. The U-like shape of the projectiondependent filter in the 5 and 3 projection cases removes large amounts of low frequencies. With regard to MR acquisition, this could lead to a higher acquisition speed, as fewer frequencies have to be acquired in the K-space. Similar thoughts can be made with respect to the projection-independent filter with 7 projections. While it is more likely that the high change in amplitude is linked to the above discussion of the frequency spectrum and selected projections the question arises if a introduction of sparsity could lead to a sparse selection of frequencies.

Note that this is not the only approach for fan-beam MR imaging. Wachowicz et al. [12] propose a method using additional non-linear gradient coils to directly acquire distorted images. Their approach is based on additional hardware, while we are demonstrate an acquisition approach which can be achieved without additional hardware.

An overall interesting observation is the performance of the derived network topology. The results show that we can substitute the inverse bracket of right inverse of the system matrix by a filter in the frequency domain. The network topology to learn such a filter could be derived used the precision learning approach introduced in [6].

# 5 Conclusion

We presented an alternative description of the rebinning process in terms of a projection-dependent or -independent filter. Based on the reconstruction problem and the problem description, we derived a network topology which allows to learn the unknown operators. Our proposed method provides a sharper image impression than the state-of-the-art method, since the necessary interpolation and thus smoothing steps can be avoided. Furthermore, the filter design is done entirely data-driven. The presented results encourage further investigation of the method. With deeper insight in the learning process, we assume that a further 14 Syben et al.

reduction of the necessary number of projections without losing the sharp image impression is possible. Additionally, as a next step the filter learning process may be extended to cone-beam projections. We hope that a better understanding of the filter will enable us to further reduce the number of data points to be recorded in k-space and, in the best case, to reduce them to points analytically determined by the filter. In the future, we want to combine our approach with MR acquisition trajectories specially adapted to our case.

Overall the results encourage to apply the proposed concept of learning unknown operators in domains where prior knowledge is available.

# Acknowledgement

This work has been supported by the project P3-Stroke, an EIT Health innovation project. EIT Health is supported by EIT, a body of the European Union.

# References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al.: Tensorflow: a system for large-scale machine learning. In: OSDI. vol. 16, pp. 265–283 (2016)
- Christlein, V., Ghesu, F.C., Würfl, T., Maier, A., Isensee, F., Neher, P., Maier-Hein, K.: Tutorial: Deep learning advancing the state-of-the-art in medical image analysis. In: Bildverarbeitung für die Medizin 2017, pp. 6–7. Springer (2017)
- Fu, W., Breininger, K., Schaffert, R., Ravikumar, N., Würfl, T., Fujimoto, J., Moult, E., Maier, A.: Frangi-Net: A Neural Network Approach to Vessel Segmentation. In: Maier, A., T., D., H., H., K., M.H., C., P., T., T. (eds.) Bildverarbeitung für die Medizin 2018. pp. 341–346 (2018)
- 4. Huang, Y., Würfl, T., Breininger, K., Liu, L., Lauritsch, G., Maier, A.: Some investigations on robustness of deep learning in limited angle tomography. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer (2018), to appear
- Lommen, J., Syben, C., Stimpel, B., Bayer, S., Nagel, A., R., F., Dörfler, A., Maier, A.: Mr-projection imaging for interventional x/mr-hybrid applications. In: Proceedings of the 49th Annual Meeting of the German Society for Medical Physics (2018)
- Maier, A.K., Schebesch, F., Syben, C., Würfl, T., Steidl, S., Choi, J.H., Fahrig, R.: Precision learning: Towards use of known operators in neural networks. CoRR abs/1712.00374 (2017), http://arxiv.org/abs/1712.00374
- Ramachandran, G., Lakshminarayanan, A.: Three-dimensional reconstruction from radiographs and electron micrographs: application of convolutions instead of fourier transforms. Proceedings of the National Academy of Sciences 68(9), 2236–2240 (1971)
- Rebecca, F., Kim, B., A., R.J., Rowland, S., John, S., M., S.G., L., D.B., Zhifei, W., L., E.D., J., P.N.: A truly hybrid interventional mr/x-ray system: Feasibility demonstration. Journal of Magnetic Resonance Imaging 13(2), 294–300 (2001). https://doi.org/10.1002/1522-2586(200102)13:2j294::AID-JMRI1042;3.0.CO;2-X
- Shepp, L.A., Logan, B.F.: The fourier reconstruction of a head section. IEEE Transactions on Nuclear Science 21(3), 21–43 (1974)

- Syben, C., Stimpel, B., Breininger, K., Würfl, T., Fahrig, R., Dörfler, A., Maier, A.: Precision Learning: Reconstruction Filter Kernel Discretization. In: Noo, F. (ed.) Proceedings of the Fifth International Conference on Image Formation in X-Ray Computed Tomography. pp. 386–390 (2018)
- Syben, C., Stimpel, B., Leghissa, M., Dörfler, A., Maier, A.: Fan-beam Projection Image Acquisition using MRI. In: Skalej, M., Hoeschen, C. (eds.) 3rd Conference on Image-Guided Interventions & Fokus Neuroradiologie. pp. 14–15 (2017)
- Wachowicz, K., Murray, B., Fallone, B.: On the direct acquisition of beams-eyeview images in mri for integration with external beam radiotherapy. Physics in Medicine & Biology 63(12), 125002 (2018)
- Wang, G., Kalra, M., Murugan, V., Xi, Y., Gjesteby, L., Getzin, M., Yang, Q., Cong, W., Vannier, M.: Vision 20/20: Simultaneous CT-MRI – Next chapter of multimodality imaging. Medical Physics 42, 5879–5889 (Oct 2015). https://doi.org/10.1118/1.4929559
- Würfl, T., Ghesu, F.C., Christlein, V., Maier, A.: Deep Learning Computed Tomography. In: Springer (ed.) Medical Image Computing and Computer-Assisted Intervention MICCAI 2016. vol. 3, pp. 432–440 (2016)
- 15. Würfl, T., Hoffmann, M., Christlein, V., Breininger, K., Huang, Y., Unberath, M., Maier, A.: Deep Learning Computed Tomography: Learning Projection-Domain Weights from Image Domain in Limited Angle Problems. IEEE Transactions on Medical Imaging 37(6), 1454–1463 (2018). https://doi.org/10.1109/TMI.2018.2833499
- Zhu, B., Liu, J.Z., Cauley, S.F., Rosen, B.R., Rosen, M.S.: Image reconstruction by domain-transform manifold learning. Nature 555(7697), 487 (2018)