

OPEN

# Projection-to-Projection Translation for Hybrid X-ray and Magnetic Resonance Imaging

Bernhard Stimpel<sup>1,2\*</sup>, Christopher Syben<sup>1,2</sup>, Tobias Würfl<sup>1</sup>, Katharina Breininger<sup>1</sup>, Philip Hoelter<sup>2</sup>, Arnd Dörfler<sup>2</sup> & Andreas Maier<sup>1</sup>

Hybrid X-ray and magnetic resonance (MR) imaging promises large potential in interventional medical imaging applications due to the broad variety of contrast of MRI combined with fast imaging of X-ray-based modalities. To fully utilize the potential of the vast amount of existing image enhancement techniques, the corresponding information from both modalities must be present in the same domain. For image-guided interventional procedures, X-ray fluoroscopy has proven to be the modality of choice. Synthesizing one modality from another in this case is an ill-posed problem due to ambiguous signal and overlapping structures in projective geometry. To take on these challenges, we present a learning-based solution to MR to X-ray projection-to-projection translation. We propose an image generator network that focuses on high representation capacity in higher resolution layers to allow for accurate synthesis of fine details in the projection images. Additionally, a weighting scheme in the loss computation that favors high-frequency structures is proposed to focus on the important details and contours in projection imaging. The proposed extensions prove valuable in generating X-ray projection images with natural appearance. Our approach achieves a deviation from the ground truth of only 6% and structural similarity measure of  $0.913 \pm 0.005$ . In particular the high frequency weighting assists in generating projection images with sharp appearance and reduces erroneously synthesized fine details.

Medical imaging offers various possibilities to visualize soft and hard tissue, physiological processes, and many more. The range of information that can be acquired is, however, divided between many different modalities. Hybrid imaging has the potential to provide simultaneous access to distributed information for diagnostic and interventional applications<sup>1-4</sup>. For example, future research advances can use the combination of computed tomography (CT) and magnetic resonance imaging (MRI) for clinical applications. This is of particular interest for interventional applications, as a large number of tasks are associated with the manipulation of soft tissue. Despite this fact, X-ray imaging, which is insensitive to soft tissue contrast, is still the workhorse of interventional radiology. The reasons for this are the high spatial and temporal resolution, which make the handling of interventional devices much easier. By complementing these advantages with better soft tissue contrast provided by MRI, the gain from the simultaneous acquisition of soft and dense tissue information through hybrid imaging would offer great opportunities.

Assuming that the information from both modalities is available simultaneously, numerous existing post-processing methods become applicable. Image fusion techniques, such as image overlays, have proven successful in the past. Additionally, techniques for image enhancement, such as noise reduction or super-resolution, can be considered. For many of the previously presented methods it is advantageous to have the data available in the same domain. The generation of CT images from corresponding MRI data was previously presented<sup>5-8</sup>, mainly to create attenuation maps for radiation therapy. However, all are applied to volumetric data, i.e., to tomographic images.

Contrarily, interventional radiology is strongly dependent on line-integral data originating from projection imaging. Images with similar perspective distortion can be acquired directly with an MR device<sup>9,10</sup>. This avoids time-consuming volumetric MRI acquisition and subsequent forward projection. The synthesis of X-ray-like projection images for further processing based on the acquired MRI signal is, however, an inherently ill-posed problem. The X-ray imaging signal is dominated by dense tissue structures, e.g., bone, which provide almost no signal in MRI. As air also does not induce signal in MRI, the materials cannot be discriminated based on intensity

<sup>1</sup>Pattern Recognition Lab, Friedrich-Alexander University Erlangen-Nürnberg, Erlangen, Germany. <sup>2</sup>Department of Neuroradiology, University Hospital Erlangen, Erlangen, Germany. \*email: [bernhard.stimpel@fau.de](mailto:bernhard.stimpel@fau.de)

values alone. Resolving this ambiguity is, therefore, only possible based on the available structural information. In volumetric images, the materials may be unknown but are resolved in distinct regions of the image. While MR projection imaging allows for continuous intraoperative use due to much faster acquisition times compared to volumetric MRI scans, the structural information diminishes in the projection image by integrating the intensity or attenuation values on the detector. This corresponds to a linear combination of multiple slice images with unknown path length which further increases the difficulty of the synthesis task. Enabled by the progress in fast MR projection acquisition, we investigate a solution for generating X-ray projections from corresponding MRI views by projection-to-projection translation.

## Related Work

**MR projection imaging.** Common clinical X-ray systems used for image-guided interventions exhibit a cone-beam geometry with the according perspective distortion in the acquired images. These systems can acquire fluoroscopic sequences with a rate of up to 30 frames per second. In contrast, MR imaging is usually performed in a tomographic fashion. The acquisition of whole volumes and subsequent forward projection is, however, too slow to be applicable in interventional procedures. Fortunately, the direct acquisition of projection images is possible<sup>9,10</sup>. Yet, the resulting images are subject to a parallel-beam geometry and, therefore, incompatible with their cone-beam projected X-ray counterpart. To remedy this contradiction, recent research has addressed the problem of acquiring MR projections with the same perspective distortion as an X-ray system, without the detour of volumetric acquisition<sup>11</sup>. Most approaches rely on rebinning to convert the acquired projection rays from parallel to fan- or cone-beam geometry<sup>9,12</sup>. Unfortunately, this requires interpolation which reduces the resulting image's quality. To avoid loss of resolution, Syben *et al.*<sup>10</sup> proposed a neural network based algorithm to generate this perspective projection.

**Image-to-image translation.** Projection-to-projection translation suffers from the problems explained in Section 1, i.e., ambiguous signal and overlapping structures. However, from a machine learning point of view, it is essentially an image-to-image translation task. While we target the application of synthesizing 2D projection images instead of tomographic slice images, the preliminary work on pseudo-CT generation, mostly applied for radiotherapy treatment planning, is still relevant. Two main approaches to the task of estimating 3D CT images from corresponding MR scans have prevailed up to now: atlas-based and learning-based methods. Atlas-based methods as proposed in<sup>13,14</sup> achieve good results, however, their dependence on accurate image registration and often high computational complexity is undesirable in the interventional setting. In contrast, inference is fast in learning-based approaches which makes them suitable for real-time applications. Essentially a regression task, image-to-image translation can be solved using classical machine learning methods like random forests<sup>5</sup>. The advent of convolutional neural networks (CNN) revolutionized the field of image synthesis and image processing in general. Numerous different approaches to image synthesis using deep neural networks have been published since then<sup>15–17</sup>. Lately, generative adversarial networks (GAN)<sup>18</sup> proved valuable for this task. Isola *et al.*<sup>19</sup> presented one of the first approaches to general-purpose image-to-image translation based on GANs and others followed<sup>6,7,20–24</sup>.

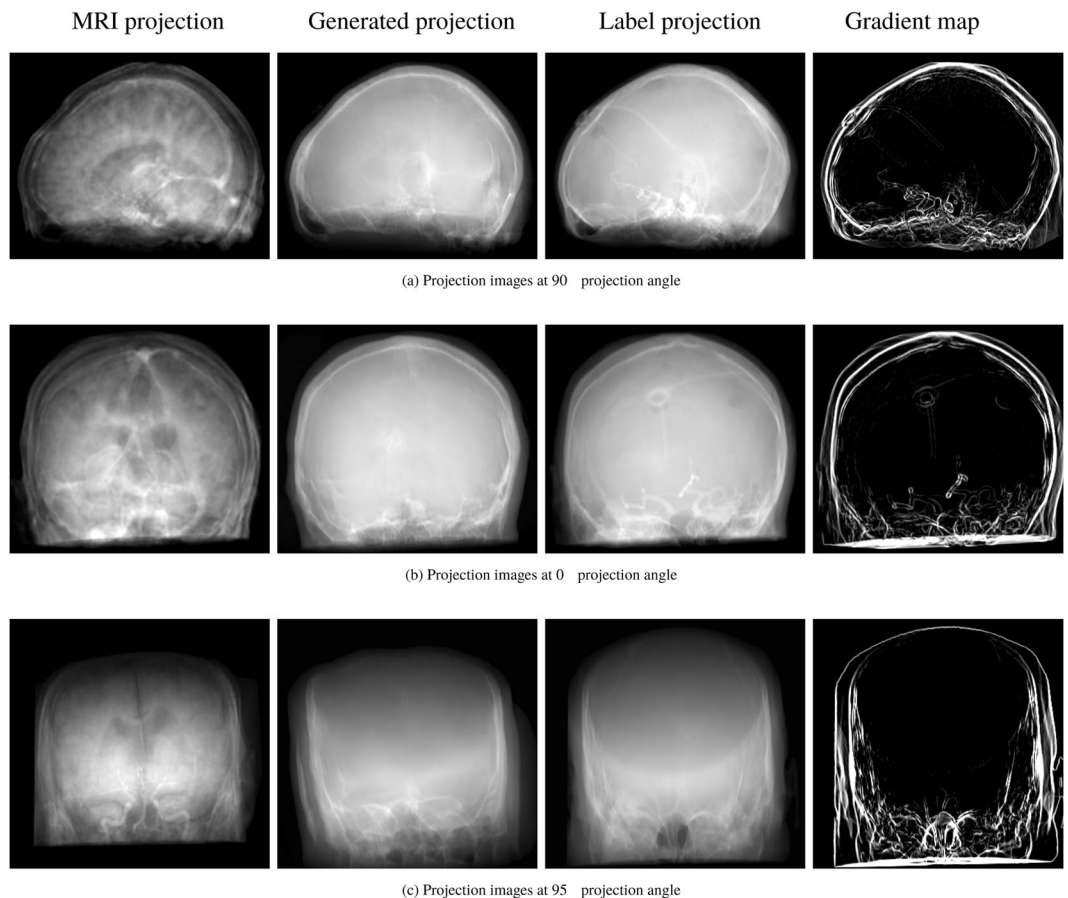
Applying this idea to the field of medical image synthesis, Nie *et al.* enhanced the conditional GAN structure with an auto-context model and achieved good results in the task of tomographic CT image synthesis based on its MR pendant<sup>6</sup>. Furthermore, in<sup>7,25–27</sup>, the successful training of a GAN based on unpaired MR and CT images was shown. As a result of the achieved successes, GANs are now a frequently used tool in medical imaging research that extends beyond the realm of image-to-image translation<sup>28–32</sup>. Inspired by these advances in the field of image synthesis, we seek to find a deep learning-based solution to generate X-ray projections from corresponding MR projections. In previous research<sup>33</sup>, we were able to show a proof of concept, which however was only able to produce mediocre images. We want to remedy this shortcoming in the underlying work and additionally show a profound analysis of the results on a more representative data set.

## Methods

In contrast to most natural image synthesis problems we do not desire to learn a one-to-many mapping. While there are different “correct” solutions for synthesized images from, e.g., a semantic layout, only one solution is correct for projection-to-projection translation tasks. The underlying network can, therefore, be trained in a supervised manner based on corresponding pairs of MR and X-ray projection images. Please note that the mapping itself remains an ill-posed problem due to the previously discussed signal ambiguities and structural overlaps. Formally, we seek to learn a mapping that generates an image  $G$  based on an input image  $I$ . The mapping is trained based on corresponding input and label image pairs,  $I$  and  $L$ , such that the generated image  $G$  is as close as possible to  $L$ . In the underlying task,  $I$  and  $L$  are the MR and X-ray projection images, respectively, and  $G$  is the generated X-ray projection.

**Network architecture.** We employ a fully convolutional neural network as our image generator. The general considerations regarding the network's architecture are based on the popular approaches proposed by<sup>7,21,22,34</sup> which have shown promising results also in medical imaging applications<sup>7</sup>. The network's architecture is designed in an encoder-decoder fashion. In the originally proposed approaches, the lowest resolution level consists of a series of subsequent residual blocks<sup>35</sup>. These blocks allocate the majority of network capacity that is utilized for the projection-to-projection translation. At the lowest, most coarse resolution layer the general outline of structures and their alignment in the image are determined. Considering the overwhelming variety of possible solutions that exist in the synthesis of natural images, the use of a large part of the available resources for the basic determination of the images seems reasonable. In contrast, the underlying variance observed in medical projection imaging is limited. In addition, the most insightful information during interventional procedures is





**Figure 2.** Representative examples of the projection-to-projection translation for different projection angles and patients. The top and middle row are projections originating from the first test patient and the bottom row from the second test patient.

	CT data	MR data
Type	unenhanced spiral h70 (“sharp”) reconstruction kernel	time-of-flight angiography
Dimensions	$512 \times 512 \times (120-368)$	$(256-412) \times (256-512) \times (176-188)$
Spacing	$(0.39-0.45) \times (0.39-0.45) \times (0.6-1.2) \text{ mm}^3$	$0.39 \times 0.39 \times 0.5 \text{ mm}^3$

**Table 1.** A selection of the important properties of the data used for the underlying work.

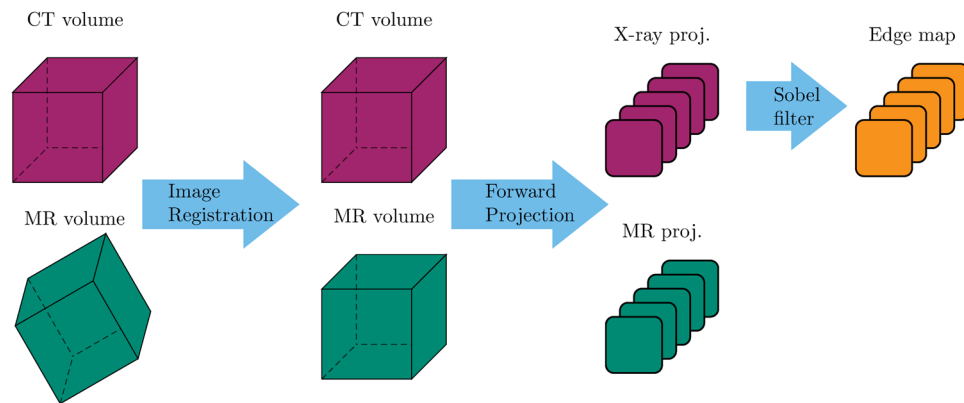
where  $V_s(\mathbf{L})$  and  $V_s(\mathbf{G})$  are the feature activation maps of the VGG-19 network<sup>39</sup> at the layer  $s \in S$ . This leads to the final objective functions

$$\ell(\mathbf{L}, \mathbf{G}, \mathcal{D}) = (\ell_{\text{GAN}}(\mathbf{L}, \mathbf{G}, \mathcal{D}) + \ell_{\text{FM}}(\mathbf{L}, \mathbf{G})) \cdot E_L, \quad (3)$$

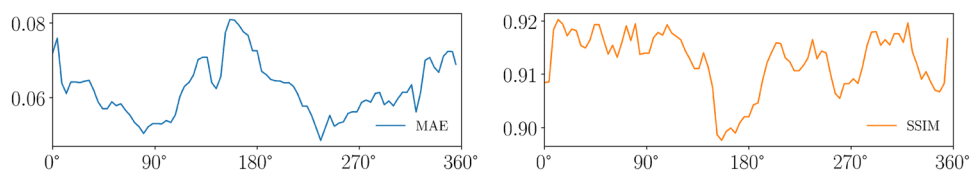
which is the combination of the GAN and feature matching loss weighted by the gradient map  $E_L$  of the label image. For a simplified representation, we assumed here that all elements have the same dimensions. In practice, the outputs of the feature matching and GAN loss are multi- and low-dimensional, respectively, which is why the precalculated edge map has to be adjusted. We use bilinear downsampling for this.

## Experiments

**Data generation.** Evaluating the performance based on real-world data is crucial. However, since the acquisition of corresponding projection images is not a part of the clinical routine up to now, these have to be artificially created. To this end, MRI and CT head datasets of 13 individuals with different pathologies are provided by the Department of Neuroradiology, University Hospital Erlangen, Germany. (MR: 1.5T MAGNETOM Aera/CT: SOMATON Definition, Siemens Healthineers, Erlangen / Forchheim, Germany). Table 1 shows important details about the type and properties of the datasets. However, please note that no special imaging or selection of data is performed for this work. All datasets were acquired as a part of routine clinical imaging procedures.



**Figure 3.** A schematic overview of the data generation process. Note that ultimately the goal is to acquire simultaneous projection images in hybrid X-ray and MR imaging (cf. Section 2.1). To generate training data for the underlying work, simulation was necessary.



**Figure 4.** Evaluation metrics with respect to the projection angle.

	MAE in [%]	SSIM	PSNR
Reference w/o edge-weighting	$7.7 \pm 1.7$	$0.884 \pm 0.011$	$20.132 \pm 1.644$
Reference w/ edge-weighting	$6.6 \pm 1.5$	$0.902 \pm 0.013$	$21.558 \pm 1.802$
Ours w/o edge-weighting	$7.4 \pm 3.4$	$0.909 \pm 0.011$	$21.480 \pm 3.187$
Ours w/ edge-weighting	$5.7 \pm 0.9$	$0.913 \pm 0.005$	$21.859 \pm 0.676$

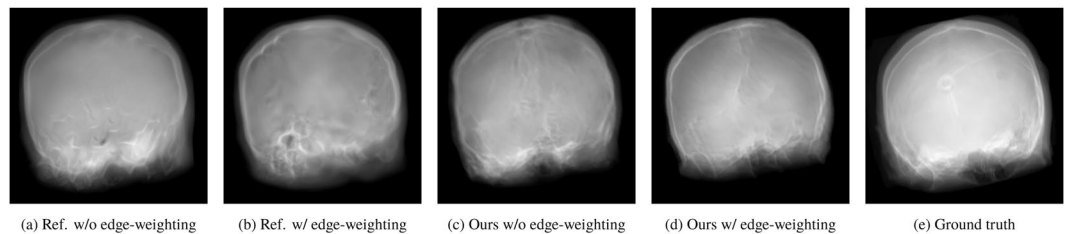
**Table 2.** MAE, SSIM, and PSNR of the different network architectures.

Rigid registration of the tomographic 3D data is performed using 3D Slicer<sup>40</sup>. Subsequently, the registered datasets are forward projected using the CONRAD framework<sup>41</sup> to create matching cone-beam MR and X-ray projection images. The projection geometry and according parameters are set to closely match clinical C-arm CT systems<sup>42</sup>. To ensure substantial variation while training, 450 projections per patient are created that are distributed equiangularly along the azimuthal and inclination angle. During testing, projections from a 360° rotation in the transversal plane are used which corresponds to common clinical acquisition procedures. In Fig. 3, a visual representation of the data generation process is given.

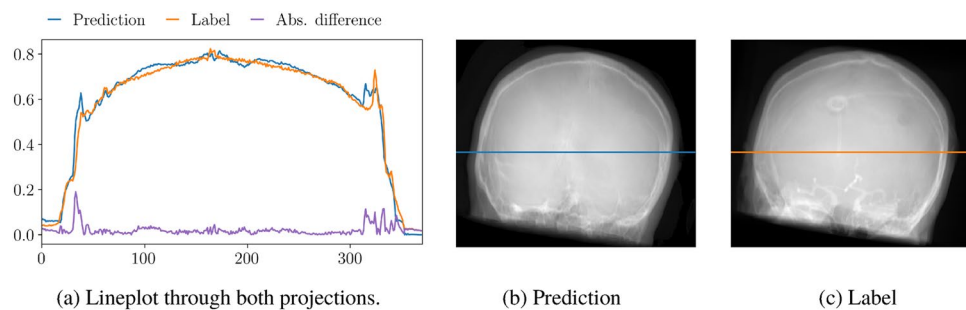
All projections are normalized prior to training by zero-centering and division by standard deviation. For the input data, i.e., the MRI projections, this preprocessing is applied to each subject's data individually and not on the whole dataset to account for differences in the MR protocols. Training is performed for a fixed number of 300 epochs using the ADAM optimizer<sup>43</sup> with a learning rate of  $1e^{-4}$  and a batch size of one. The edge map generated by the Sobel filter is thresholded with a value of 0.4 (data range 0 to 1) in order to abandon pixels with low likelihood of belonging to an edge.

**Evaluation.** Evaluation of the proposed approach is performed quantitatively as well as qualitatively based on two patient datasets that were reserved for testing. Quantitative evaluation is done based on the mean absolute error (MAE), structural similarity index (SSIM), and peak signal-to-noise ratio (PSNR). In addition, the performance of these metrics with respect to the projection angle is investigated. All projections are normalized beforehand. For MAE and PSNR, only pixels are considered that are nonzero in the label images to limit the optimistic bias caused by the large homogeneous air regions. For treatment planning in radiotherapy, a precise estimation of pixel-wise attenuation values is crucial. In interventional projection imaging, however, also the visual appearance of the generated images is of interest for the surgeon. Pixel-wise error metrics prefer smooth images, as explained previously in Section 3.2. Therefore, qualitative evaluation of the generated X-ray projection images is performed to make sure that the generated and reference images also correspond visually for the human observer.





**Figure 5.** Influence of reference (ref.) and proposed architecture and loss functions on the generated results.



**Figure 6.** Exemplary lineplots of a pair of generated and label projection images.

The proposed modifications of our approach are, first, a relocation of the residual blocks to higher resolution layers to increase the network capacity at the relevant stages and, second, a high-frequency weighting of the computed loss term.

We compare our approach to the conditional GAN architecture as used in<sup>7,21,22,34</sup>, which we will refer to as reference architecture. The reference is similar to our proposed architecture (see Fig. 1) with the main difference that all residual blocks are placed at the lowest resolution layers. To ensure comparability, the network parameters of both approaches are chosen such that all available graphics card memory is utilized.

## Results

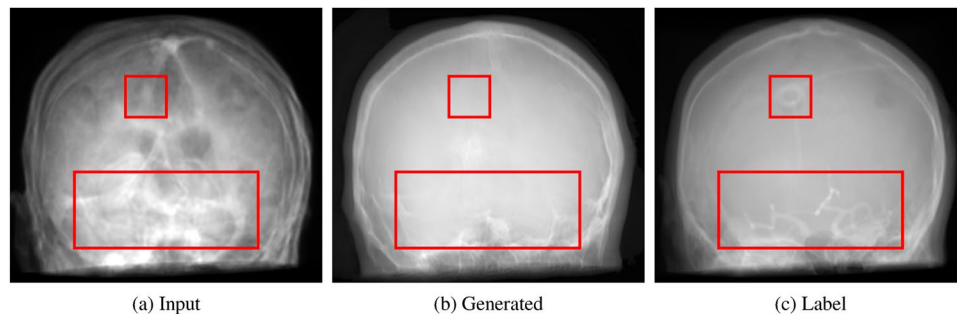
All quantitative metrics and observations reported in the following are computed or made based on the two representative test datasets. In Table 2, the quantitative results of the projection-to-projection translation are presented. We evaluated the proposed changes in architecture and the edge-weighting against the reference architecture. Shifting the residual blocks shows a small improvement in MAE as well as SSIM. The combination of the proposed architecture and the edge-weighted loss function yields an improvement of 25% compared to the reference architecture without the weighted loss term. The gain in SSIM is not as large but still clear. In Fig. 4, the error metrics are plotted w.r.t. the projection angle. In this figure,  $0^\circ$  denotes an RAO angle of  $90^\circ$  and, consequently,  $180^\circ$  represents  $90^\circ$  LAO. Note that opposing projections are not equal due to tilt and the projective geometry. The qualitative examples give a good intuition of the effects of the proposed changes. Representative examples of the generated images are given in Fig. 2 for qualitative examination. The influence of the proposed changes regarding the network architecture and the edge-weighting can be observed in Fig. 5. Fig. 6 shows an example lineplot through the generated and label image and their absolute difference.

Computation time is a limiting factor in image-guided interventions. Our unoptimized project-to-projection translation is capable of processing of  $\sim 24$  frames per second on a Nvidia Tesla V100 gpu. Naturally, this processing rate would be subject to additional delay or latency caused by the acquisition and preprocessing steps on the scanner. Nevertheless, the achievable processing time is fast enough to cope with common clinical X-ray fluoroscopy frame rates and will likely exceed the expected acquisition time for cone-beam MR projection images on a real scanner by far.

## Discussion

Quantitative comparison with other methods is difficult to conduct at this point. Numerous methods for image synthesis in medical imaging were proposed up to now, however, these exclusively target tomographic images. In contrast, projection images are significantly different in their general look regarding contrast, edges and superimpositions. A direct comparison between corresponding evaluation metrics would, therefore, be not meaningful for either side.

**Data generation.** While the amount of patient scans that were available for this task is limited, our method of augmenting the data set by using multiple angles for forward projection is powerful. Data augmentation itself was proven remarkably beneficial in prior work<sup>44</sup>. Thereby, translation, rotation or simple warping are only the baseline and represent modified versions of the original image. In contrast, a forward projection is an X-ray



**Figure 7.** An example of missing information in the generated X-ray projection images. Details that are unobservable in the input MRI projections can, naturally, not be translated in the resulting generated projection images. The small rectangle in the label image (c) outlines the entry point of a ventricular shunt while the larger rectangle frames contrasted cerebral arteries. While the devices were not present at the time of the acquisition of the underlying images, it is likely that none of these details will be captured by the MR imaging protocol.

transform, i.e., an integral transformation. For a function  $f(L)$  and the line  $L$ , the X-ray transform  $Xf(L)$  is the mapping

$$Xf(L) = \int_{\mathbb{R}} f(\mathbf{x}_0 + t\boldsymbol{\theta}) dt, \quad (4)$$

where the unit vector  $\boldsymbol{\theta}$  denotes the direction of the line  $L$  and  $\mathbf{x}_0$  is the starting point on the line. If  $f(L)$  is a constant function, i.e., one image, the resulting X-ray-transformed function is different for every single  $\boldsymbol{\theta}$ . Thereby, the previously mentioned popular augmentation methods are intrinsically part of the X-ray transform due to effects like perspective distortion, the projection trajectory, and more. What is, unfortunately, not solved by this data augmentation process are different medical indications of the patients. These come in various different appearances and shapes, e.g., only considering the cranial region, aneurysms, lesions, fractures, and many more. As the underlying dataset is composed of scans from different points in time, structural differences occur (see Fig. 7). Like with other learning-based algorithms, handling of these examples can only be covered by including them in the training database. It is, therefore, important to keep this fact in mind when dealing with data-based systems. Furthermore, multiple patient datasets are truncated in axial direction (see Fig. 2(c)) due to incomplete acquisition of the head. Nevertheless, our approach was able to deal with this challenge and produce the according truncated or untruncated synthesized projections.

**Architecture and objective function.** As seen in Fig. 5(a,c), the redistribution of the residual blocks improves the visual impression of the generated images. The projections generated with the residual blocks at the high resolution levels resemble real head projections more closely. According to the motivation of this change described in Section 3.1, we conclude that this improvement is caused by the increased network capacity at the high-resolution layers. The consequent reduction in capacity in the coarser layers does not lead to a deterioration of the results. This may be due to the fact that the rough outline of the structures to be generated, i.e., the head, is the same in the MR and X-ray projections. Observation of the training process also showed that the general arrangement was already fixed after the first few epochs.

High-frequency details are most of the time the most important part in X-ray projection imaging. Yet, most networks used for image synthesis are built in an encoder-decoder fashion which accumulates large portions of the available network capacity at the lower resolution layers to allow for an increase in filter dimension. Considering recent trends in super-resolution<sup>45</sup>, the usage of sub-pixel convolutions<sup>46</sup> might be of interest for future work. This technique performs the image generation almost completely on lower resolution levels without sacrificing on high-frequency details by reordering multiple low-dimensional feature maps into one high-dimensional image. Given the high memory requirements of image synthesis, this could be a valuable tool.

Yet, only redistributing the network's capacity to higher resolution levels is not enough to guide the generation process in the desired direction. A more pronounced improvement can be observed by addition of the edge-weighting to the objective function. In Fig. 5(d), the fine details, i.e., edges, are much sharper compared to its counterpart (Fig. 5(c)) which was not trained using this edge-weighting. Additionally, the amount of erroneously synthesized edges and borders is decreased when applying this weighting scheme. This improvement is not surprising when considering that only  $\sim 9.8\%$  of the total image points belong to an edge according to our computed edge map. If the loss was not weighted, the error resulting from this small portion of image points would simply be overshadowed by the vast majority of low-frequency image points in the images.

Another potentially interesting complement to our approach could be the inclusion of the temporal dimension. Fluoroscopic images consist usually of a sequence of images obtained from the same view. Similarly, while acquiring 3D tomographic scans, successive projections from slightly changed positions are recorded. In both cases, there is a certain degree of consistency between the successive projections. This can potentially be exploited by simultaneous processing of several projections in order to transfer this consistency to the results<sup>47–49</sup>.

Despite the proposed additions and the resulting improvements, the projection-to-projection translation is not yet perfect. The high number of fine structures that are integrated on to the detector during the projection and thus overlap represent a great challenge for the synthesis. Combined with the initially discussed fact that the

acquired signal is partly ambiguous, e.g., air and bone map to similar regions in most MR imaging sequences, providing a conclusive solution is a hard task. However, many interventional tasks do not require a perfect one-to-one correspondence. For example, smooth changes in the intensities may not be of great interest, assuming they are even perceivable for the human observer in the first place. In fact, an exact match of the images is not always the desired result. To apply many post-processing methods, subtle differences are needed to make these techniques possible. Tasks like one-shot digital subtraction angiography become feasible with the underlying synthesized and real contrasted X-ray projection images if slight differences between the projections are present.

For future work, we hope to gain access to further corresponding data sets, especially those of body parts with more diverse information contained in the respective modalities. For example the chest features a greater diversity between the modalities with the X-ray based modality containing the ribs and spine, while all soft-tissue structures like the heart are better observable in the MR images. Additionally, we would like to evaluate the high-frequency component weighting also for other tasks that rely on accurate synthesis of structural details besides projection-to-projection translation.

## Conclusion

The underlying work presents an approach to synthesize X-ray projection images based on corresponding MR projections, which is potentially advantageous for hybrid MR and CT imaging applications. In doing so, we have identified distinct differences between the synthesis of projection images and volumetric images. In projection images, usually only edges and similar high-frequency structures contain important information. In order to emphasize this in the translation, we proposed an increase of the network capacity at higher resolution layers. In addition, a weighting of the high-frequency components in the image was introduced when calculating the error in the training process. The proposed modifications and extensions proved to be valuable and showed clear improvements in the quantitative metrics with a decrease in the deviation from the ground truth by  $\sim 25\%$ . In addition, particularly using the proposed weighting of high-frequency components in the training process resulted in a superior qualitative image impression. These results proved to yield sharper edges and fewer erroneously generated fine details. The advancing development in the fields of interventional MR and hybrid imaging motivates increased effort to solve the remaining problems of ambiguous signal and overlapping structures which still represent a challenging problem.

## Data availability

For legal reasons, the authors are not able to share the underlying clinical data sets.

Received: 7 June 2019; Accepted: 25 November 2019;

Published online: 11 December 2019

## References

- Fahrig, R. *et al.* A truly hybrid interventional MR/x-ray system: Feasibility demonstration. *Journal of Magnetic Resonance Imaging* **13**, 294–300 (2001).
- Wang, G. *et al.* Top-Level Design of the first CT-MRI scanner. In *Proceedings of the 12th Fully 3D Meeting*, 5–8 (2013).
- Wang, G. *et al.* Vision 20/20: Simultaneous CT-MRI - Next chapter of multimodality imaging. *Medical Physics* **42**, 5879–5889 (2015).
- Gjesteby, L., Xi, Y., Kalra, M., Yang, Q. & Wang, G. Hybrid Imaging System for Simultaneous Spiral MR and X-ray (MRX) Scans. *IEEE Access* **5**, 1050–1061 (2016).
- Navalpakkam, B. K., Braun, H., Kuwert, T. & Quick, H. H. Magnetic resonance-based attenuation correction for PET/MR hybrid imaging using continuous valued attenuation maps. *Investigative Radiology* **48**, 323–332 (2013).
- Nie, D. *et al.* Medical Image Synthesis with Context-Aware Generative Adversarial Networks. *Medical Image Computing and Computer-Assisted Intervention - MICCAI* 417–425 (2017).
- Wolterink, J. M. *et al.* Deep MR to CT Synthesis Using Unpaired Data. In *International Workshop on Simulation and Synthesis in Medical Imaging*, 14–23 (2017).
- Xiang, L. *et al.* Deep embedding convolutional neural network for synthesizing CT image from T1-Weighted MR image. *Medical Image Analysis* **47**, 31–44 (2018).
- Lommen, J. M. *et al.* MR-projection Imaging with Perspective Distortion as in X-ray Fluoroscopy for Interventional X/MR-hybrid Applications. In *Proceedings of the 12th Interventional MRI Symposium*, **54** (2018).
- Syben, C. *et al.* Deriving Neural Network Architectures using Precision Learning: Parallel-to-fan beam Conversion. In *Proceedings Pattern Recognition, 40th German Conference*, 503–517 (Stuttgart, (2018)).
- Wachowicz, K., Murray, B. & Fallone, B. G. On the direct acquisition of beam's-eye-view images in MRI for integration with external beam radiotherapy. *Physics in Medicine and Biology* **63**, 125002 (2018).
- Syben, C., Stimpel, B., Leghissa, M., Dörfler, A. & Maier, A. Fan-beam Projection Image Acquisition using MRI. In *3rd Conference on Image-Guided Interventions & Fokus Neuroradiologie*, 14–15 (2017).
- Uh, J., Merchant, T. E., Li, Y., Li, X. & Hua, C. MRI-based treatment planning with pseudo CT generated through atlas registration. *Medical Physics* **41**, 051711 (2014).
- Degen, J. & Heinrich, M. P. Multi-Atlas Based Pseudo-CT Synthesis Using Multimodal Image Registration and Local Atlas Fusion Strategies. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 160–168 (2016).
- Yan, X., Yang, J., Sohn, K. & Lee, H. Attribute2Image: Conditional Image Generation from Visual Attributes. In *European Conference on Computer Vision*, 776–7 (2015).
- Zhang, R., Isola, P. & Efros, A. A. Colorful image colorization. In *European Conference on Computer Vision*, 649–666 (2016).
- Chen, Q. & Koltun, V. Photographic Image Synthesis with Cascaded Refinement Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, vol. 2017–October, 1520–1529 (2017).
- Goodfellow, I. J. *et al.* Generative Adversarial Networks. In *Advances in Neural Information Processing Systems*, 2672–2680 (2014).
- Isola, P., Zhu, J. Y., Zhou, T. & Efros, A. A. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, vol. 2017–January, 1125–1134 (2017).
- Liu, M.-Y., Breuel, T. & Kautz, J. Unsupervised Image-to-Image Translation Networks. In Guyon, I. *et al.* (eds) *Advances in Neural Information Processing Systems* **30**, 700–708 (Curran Associates, Inc., 2017).
- Zhu, J. Y., Park, T., Isola, P. & Efros, A. A. Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2242–2251 (2017).



22. Wang, T.-C. *et al.* High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8798–8807 (2018).
23. Huang, X., Liu, M.-Y., Belongie, S. & Kautz, J. Multimodal Unsupervised Image-to-image Translation. In *The European Conference on Computer Vision (ECCV)* (2018).
24. Choi, Y. *et al.* StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018).
25. Yang, H. *et al.* Unpaired Brain MR-to-CT Synthesis using a Structure-Constrained CycleGAN. In *International Workshop on Multimodal Learning for Clinical Decision Support*, 174–182 (Springer, Cham, (2018).
26. Hiasa, Y. *et al.* Cross-Modality Image Synthesis from Unpaired Data Using CycleGAN. In Gooya, A., Goksel, O., Oguz, I. & Burgos, N. (eds.) *Simulation and Synthesis in Medical Imaging*, 31–41 (Springer International Publishing, Cham, (2018).
27. Chartsias, A., Joyce, T., Dharmakumar, R. & Tsiftaris, S. A. Adversarial image synthesis for unpaired multi-modal cardiac data. In *International Workshop on Simulation and Synthesis in Medical Imaging*, 3–13 (Springer, (2017).
28. Shin, H.-C. *et al.* Medical Image Synthesis for Data Augmentation and Anonymization Using Generative Adversarial Networks. In Gooya, A., Goksel, O., Oguz, I. & Burgos, N. (eds) *Simulation and Synthesis in Medical Imaging*, 1–11 (Springer International Publishing, Cham, (2018).
29. Armanious, K. *et al.* MedGAN: Medical Image Translation using GANs. *arXiv:1806.06397v2* (2018).
30. Yang, Q. *et al.* Low-Dose CT Image Denoising Using a Generative Adversarial Network With Wasserstein Distance and Perceptual Loss. *IEEE Transactions on Medical Imaging* **37**, 1348–1357 (2018).
31. Yu, B. *et al.* Ea-GANs: Edge-aware Generative Adversarial Networks for Cross-modality MR Image Synthesis. *IEEE Transactions on Medical Imaging* **1–1** (2019).
32. Yi, X., Wallia, E. & Babyn, P. Generative adversarial network in medical imaging: A review. *Medical image analysis* **58** (2019).
33. Stimpel, B. *et al.* Projection image-to-image translation in hybrid X-ray/MR imaging. *Medical Imaging 2019: Image Processing* **10949**, 90 (2018).
34. Johnson, J., Alahi, A. & Fei-Fei, L. Perceptual Losses for Real-Time Style Transfer and Super-Resolution. In *European Conference on Computer Vision*, 694–711 (2016).
35. He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, 770–778 (2016).
36. Dosovitskiy, A. & Brox, T. Generating Images with Perceptual Similarity Metrics based on Deep Networks. *Advances in Neural Information Processing Systems* 658–666 (2016).
37. Gatys, L. A., Ecker, A. S. & Bethge, M. Image style transfer using convolutional neural networks. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 2414–2423 (2016).
38. Sobel, I. & Feldman, G. A  $3 \times 3$  Isotropic Gradient Operator for Image Processing. In *Stanford Artificial Intelligence Project (SAIL)* (1968).
39. Simonyan, K. & Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556* (2015).
40. Fedorov, A. *et al.* 3D Slicer as an image computing platform for the Quantitative Imaging Network. *Magnetic Resonance Imaging* **30**, 1323–1341 (2012).
41. Maier, A. *et al.* CONRAD - A software framework for cone-beam imaging in radiology. *Medical Physics* **40** (2013).
42. Artis zeego Multi-axis system for interventional imaging, URL <https://www.deltamedicalsystems.com/DeltaMedicalSystems/media/Product-Details/Artis-zeego-Data-Sheet.pdf>. Date accessed - Aug-09. (2019)
43. Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations* (2015).
44. Perez, L. & Wang, J. The Effectiveness of Data Augmentation in Image Classification using Deep Learning. *arXiv preprint arXiv:1712.04621* (2017).
45. Fan, Y., Yu, J. & Huang, T. S. Wide-activated Deep Residual Networks based Restoration for BPG-compressed Images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2621–2624 (2018).
46. Shi, W. *et al.* Real-Time Single Image and Video Super-Resolution Using an Efficient Sub-Pixel Convolutional Neural Network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1874–1883 (2016).
47. Engelhardt, S., De Simone, R., Full, P. M., Karck, M. & Wolf, I. Improving Surgical Training Phantoms by Hyperrealism: Deep Unpaired Image-to-Image Translation from Real Surgeries. In *Medical Image Computing and Computer Assisted Intervention – MICCAI*, 747–755 (Cham, (2018).
48. Lai, W.-S. *et al.* Learning Blind Video Temporal Consistency. In *The European Conference on Computer Vision (ECCV)*, 170–185 (2018).
49. Aichert, A. *et al.* Epipolar Consistency in Transmission Imaging. *IEEE Transactions on Medical Imaging* **34**, 2205–2219 (2015).

## Acknowledgements

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. Additional financial support for this project was granted by the Emerging Fields Initiative (EFI) of the Friedrich-Alexander University Erlangen-Nürnberg (FAU) as well as Siemens Healthineers. Furthermore, we thank the NVIDIA Corporation for their hardware donation.

## Author contributions

Conceptualization, B.S., C.S., T.W., K.B.; Data curation, P.H. and A.D.; Funding acquisition, A.D. and A.M.; Methodology, B.S.; Validation, B.S.; Writing—original draft, B.S.; Writing—review and editing, B.S., C.S., T.W., K.B., P.H., A.D., A.M.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to B.S.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019