
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Agrawal, Harshit; Hietanen, Ari; Sarkka, Simo

Deep learning based projection domain metal segmentation for metal artifact reduction in cone beam computed tomography

Published in:
IEEE Access

DOI:
[10.1109/ACCESS.2023.3314700](https://doi.org/10.1109/ACCESS.2023.3314700)

Published: 13/09/2023

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Agrawal, H., Hietanen, A., & Sarkka, S. (2023). Deep learning based projection domain metal segmentation for metal artifact reduction in cone beam computed tomography. *IEEE Access*, 11, 100371-100382.
<https://doi.org/10.1109/ACCESS.2023.3314700>

Received 17 August 2023, accepted 10 September 2023, date of publication 13 September 2023,
date of current version 19 September 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3314700

APPLIED RESEARCH

Deep Learning Based Projection Domain Metal Segmentation for Metal Artifact Reduction in Cone Beam Computed Tomography

HARSHIT AGRAWAL^{1,2}, (Graduate Student Member, IEEE), ARI HIETANEN²,
AND SIMO SÄRKKÄ¹, (Senior Member, IEEE)

¹Department of Electrical Engineering and Automation, Aalto University, 00076 Aalto, Finland

²Department of Research and Technology, Planmeca Oy, 00880 Helsinki, Finland

Corresponding author: Harshit Agrawal (harshit.agrawal@aalto.fi)

This work was supported in part by the Business Finland Grant.

ABSTRACT Metal artifact correction is a challenging problem in cone beam computed tomography (CBCT) scanning. Metal implants inserted into the anatomy cause severe artifacts in reconstructed images. Widely used inpainting-based metal artifact reduction (MAR) methods require segmentation of metal traces in the projections as a first step, which is a challenging task. One approach is to use a deep learning method to segment metals in the projections. However, the success of deep learning methods is limited by the availability of realistic training data. It is laborious and time consuming to get reliable ground truth annotations due to unclear implant boundaries and large numbers of projections. We propose to use X-ray simulations to generate synthetic metal segmentation training dataset from clinical CBCT scans. We compare the effect of simulations with different numbers of photons and also compare several training strategies to augment the available data. We compare our model's performance on real clinical scans with conventional region growing threshold-based MAR, moving metal artifact reduction method, and a recent deep learning method. We show that simulations with relatively small number of photons are suitable for the metal segmentation task and that training the deep learning model with full size and cropped projections together improves the robustness of the model. We show substantial improvement in the image quality affected by severe motion, voxel size under-sampling, and out-of-FOV metals. Our method can be easily integrated into the existing projection-based MAR pipeline to get improved image quality. This method can provide a novel paradigm to accurately segment metals in CBCT projections.

INDEX TERMS CBCT, deep learning, synthetic data, X-ray simulation, metal artifact reduction, metal segmentation.

I. INTRODUCTION

Cone beam computed tomography (CBCT) acquires 2D X-ray projections of an object to reconstruct 3D images. The presence of high-density metallic implant in the object may corrupt the projection data and cause severe artifacts such as blooming, streaking, and shading [1] in the reconstructed images. Several projection-based metal

artifact removal (MAR) methods have been proposed in literature [2], [3], [4], [5], [6], [7]. A crucial step in these methods is the segmentation of corrupt metal traces in the projections. First, the metals are segmented in the 3D image-domain by a threshold. Then the segmented metals are forward-projected to obtain the metal trace in the projections. The data inside the metal trace is discarded and interpolated from the nearby projection data. The corrected projections are back-projected (reconstruction) to generate corrected 3D images.

The associate editor coordinating the review of this manuscript and approving it for publication was Zhaoqing Pan¹.

Several deep learning-based MAR methods have been proposed recently. They either work in projection-domain, image-domain, or both. Projection-domain methods [8], [9], [10], [11], [12] inpaint the metal trace in the projection (CBCT) or sinogram (CT) using deep learning architectures. All of these methods rely on the availability of segmented metal traces. Image-domain methods, such as [13] and [14] directly translate from artifact to artifact-free images. Image-domain methods can effectively reduce metal artifacts in some cases, however, their effect can be limited by the presence of scatter, large and varying shapes of metals, multiple metal implants, and motion artifacts [9]. While inpainting-based methods deal better with these problems, they still require accurate metal segmentation in the projection [15]. In [16], [17], and [18], it was proposed to incorporate both image and sinogram-domain deep learning. These methods use filtered back projection (FBP) and forward projection (FP) layers to propagate the loss from the image-domain to the sinogram-domain. These methods were implemented for CT images (slice-wise reconstruction) though full volume CBCT reconstruction is not feasible due to its large memory requirements [19].

It was demonstrated in [20] that precise segmentation of metal traces can enhance the image quality, and conversely, less accurate segmentation may cause additional artifacts or even remove anatomical details. As forward projection of the image-domain metal is used to obtain the metal traces in the projections, it has several problems. When metal lies outside of the field-of-view (FOV), image-domain segmentation methods cannot segment metals in the projections [21]. This may cause image inhomogeneity in the center and among the different zones of the FOV [22], [23]. Furthermore, in the presence of motion artifacts [24], [25], [26], the metals are blurred and it is difficult to obtain reliable forward projections of the image-domain segmented metals. Pre- and post-processing steps were proposed in [27] and [28] to mitigate motion affected metal artifacts. Moreover, the forward projection step of the metals need to account for the partially covered pixels. This can be done by reconstructing the voxels with smaller size before forward projection or by taking the voxel diameter into account [29] at the cost of increased complexity and processing time.

While several deep learning methods have been investigated for MAR, few have been proposed to directly segment the projection-domain metals in CBCT scans, owing to the unavailability of accurate metal labels [30], [31]. A U-Net [32] architecture was used in [30] to segment metals in dental CT projections. The network was trained and tested on a very small dental CT dataset (five patients for training and four patients for testing) due to the time consuming process of creating ground truth labels. Pairs of metal and metal free projections from cadaver CBCT scans were used in [31] to train a U-Net architecture. Acquiring data from

cadaver images is a rare process and the data might still lack the quantity and variety of images. They used a consistency check condition to reduce false positives which provides more consistent segmentation of metals in the projections. The consistency check involved extra steps to reconstruct larger volumes and needed to calculate accurate thresholds for the metals in the reconstruction. Furthermore, we found that consistency check did not work well for motion-affected scans.

Monte Carlo simulations have been used to generate training data for deep learning-based metal in-painting and scatter correction methods [10], [17], [33], but the use of simulated dataset in metal segmentation training has not been investigated yet. Motivated by the importance of accurate training data for metal segmentation, we used Monte Carlo simulations to generate metal-corrupted CBCT projections and corresponding metal labels for network training. For the tasks of in-painting and scatter correction, simulations with low noise levels are needed [10], [17], [33]. The noise in the simulated projections decreases with the increase in numbers of photons per detector pixel [33], however, running the simulations with large number of photons is computationally expensive. In contrast, for metal segmentation training task, we used simulations with only 300 to 1400 photons per detector pixel (noisy) and compared the results with the simulations of 5000 photons per detector pixel (clean). Noisy simulations are faster to compute and add realistic noise to the dataset.

We trained a modified version of U-Net only on the simulated data to segment metal traces in the projections. In addition, we used a simple strategy to augment our simulated dataset further by training the network on full size and cropped projections together. Our simulated dataset included multiple anatomies such as knee, wrists, ankle, palm, and foot. We demonstrate the robustness of our metal segmentation model on 10 metal-affected and 6 non-metal clinical scans.

As U-Net architectures have been used in [30] and [31] for the metal segmentation in projection data, using U-Net is not a contribution of this work. However, our main contributions are as following:

- 1) We propose a new training approach for metal segmentation in CBCT projection that uses simulated dataset (with different noise levels) obtained from real metal-affected cases.
- 2) We propose a simple data augmentation strategy by training the network with crops and full size images. This strategy effectively reduces false positives in 6 unseen clinical CBCT scans without metals.
- 3) We demonstrate the metal artifact reduction on 10 unseen clinical CBCT scans affected by metals. The results show good performance in challenging cases such as motion-affected and out-of-FOV metals, and they are robust to voxel size changes.

TABLE 1. Description of scans used for simulated training data generation and testing of the trained models. kV is the tube voltage in kilovolt, voxel is the size of reconstructed voxels in mm, views are the number of projections per scan and anatomy is location of the body part irradiated with X-rays. kV values were not available for set 1 and set 2 in train set.

	Train				Test			
	kV	voxel	views	anatomy	kV	voxel	views	anatomy
set1	96	0.4	400	knee	92	0.2	300	foot
set2	-	0.2	300	wrist	96	0.2	400	ankle
set3	-	0.2	300	ankle	96	0.4	400	ankle
set4	96	0.2	300	knee	96	0.2	300	leg
set5	96	0.2	450	foot	90	0.2	300	forearm
set6	90	0.2	300	wrist	90	0.4	400	knee
set7	90	0.2	300	wrist	96	0.2	400	foot
set8	90	0.2	300	wrist	96	0.6	400	leg
set9	96	0.2	300	ankle	92	0.2	300	knee
set10	90	0.2	300	palm	90	0.2	450	wrist

II. METHODS

In this section, we explain the dataset used for the experiments, our pipeline to generate the training data, the simulation process, the network architecture used for segmentation of metals, implementation of the network, and the evaluation methods in detail.

A. DATASET

We used an unidentifiable and pseudonymized clinical dataset acquired by Plannet Verity® scanner (Plannet Oy, Helsinki, Finland) with patient permission for use in research and development studies. This study was performed in line with the principles of the Declaration of Helsinki. The dataset had 26 clinical scans taken at different extremity anatomies and was collected over multiple years, which makes the dataset diverse and realistic. Out of 26 scans, 20 were metal-affected and 6 were without any metal. Out of 20 metal-affected scans, we included 10 scans for simulated training data generation. The remaining 10 metal-affected and 6 non-metal scans were used for evaluation of the trained models. Source-to-detector distance was 580 mm and source-to-isocenter distance was 392 mm. The description of the training and testing dataset is given in Table 1. Each scan had multiple projection views, ranging from 300 to 450. The tube voltage varied from 90 kV to 96 kV. All scans had isotropic spatial resolution, with voxel spacings of 0.6mm, 0.4 mm or 0.2 mm. Different anatomy locations were scanned, including, knee, wrist, foot, ankle, palm, and forearm. The detector size was 238.76 mm × 189.99 mm. The size of each detector pixel was 0.254 mm × 0.254 mm. The dimension of each projection was 948 × 740 pixels.

From the training set of 10 scans in Table 1, we created 10 reconstructions. Two more reconstruction were created by combining metals from multiple reconstructions. These reconstructions were used to generate 3450 projections using simulation procedure described in Section II-C. As we did not want to compromise the quality of segmentation near metal boundaries, we did not downsample projections for training as in [30]. To increase the data variability, we created

four crops of size 474 × 370 from each projection. The crop was included in the training only if the sum of metal pixels was at least 100 pixels in the cropped projection. Thus we had 9305 crops for training. From the test set of 10 metal-affected scans in Table 1, we created metal labels manually. Since it was time consuming to segment metals in all of the projections, we manually segmented metals in 10 projections from each of the 10 scans. So, in total we gathered 100 pairs of metal-corrupted projections and corresponding ground truth metal labels from the clinical dataset. For the remaining test set of 6 metal-free scans, we did not require ground truth labels. Those scans had a total of 2400 metal-free projections.

B. TRAINING-DATA GENERATION PIPELINE

The pipeline to generate pairs of metal corrupted projections and corresponding ground truth metal traces is shown in Fig. 1. The pipeline had two main parts. In the first part (shown in the colored box), a segmentation of metals in 3D and an initially corrected volume were obtained. In the second part, a simulation process was used to generate pairs of training data which is described in detail in Section II-C.

First, the 3D image volume was reconstructed from real CBCT projections. Then we segmented the image-domain metals in the 3D reconstructed images using a global threshold. The segmented metals were not smooth and complete due to the artifacts remaining in the reconstruction. To smoothen and complete the metal boundaries, we applied median filtering and binary dilation to the segmented metals in 3D. The segmented 3D metals were further clipped to a minimum value of 3400 Hounsfield Units (HU). The segmented 3D metals were simulated to get only-metal projections. Binary metal traces for the ground truth were obtained by applying a threshold to the simulated metal projections. This thresholding step was straightforward and was done to assign a value of 1 to the metal pixels and 0 to the background. The binarized metal traces constituted ground truths for the model training.

To obtain the metal corrupted projections, we first created metal free reconstruction. The 3D metals obtained from the first step were put back into the metal free reconstruction and the metal-inserted volume was used to simulate metal-corrupted projections. An example of original projection, a simulated projection and the corresponding ground truth metal trace is shown in Fig. 1. The simulated ground truth metals do not always align with the original projection due to morphological operations and inaccuracies in the 3D metal segmentation step. However, the metal ground truth aligns accurately with the metals in the simulated metal corrupted projection. An example of such simulations is shown in Fig. 2 where third small metal is not present in the target metal mask and the two metals are dilated in comparison to the metals in the original projection. However, the metals in the target metal trace aligns well with the simulated projections. This is the reason why we used simulated projections for the training instead of original projections.




FIGURE 1. Pipeline for generating training pairs with examples of intermediate results. The yellow colored box illustrates conventional method for metal artifact reduction. The metal corrupted original 2D projections were first reconstructed using FDK [34]. The metal was segmented and refined. The metals are put back in the initial corrected volume. Then metal and metal inserted volumes were simulated using Monte Carlo simulations. The ground truth binary traces were obtained by applying a threshold to the simulated metal masks.




FIGURE 2. An example of original projection, simulated projections, and the corresponding target metal trace. The noisy projection was generated with 300 photons per detector pixel and the clean projection was generated with 5000 photons per detector pixel. The area around metals is enlarged to show the differences. The simulated projections and the corresponding mask were used for training.

C. SIMULATION PROCESS

We used Monte Carlo simulations to generate accurate metal segmentation training dataset. The implementation was similar to the one described in [35] and [36]. We simulated path of X-ray photons through the voxelized reconstructions obtained from 10 metal-affected clinical CBCT scans as described in Section II-B. Such clinical scans represent realistic variations in the metal shapes and sizes. The reconstruction volumes were first segmented into four different materials by applying thresholds on HU values. The volumes were segmented into air, soft tissue, bone, and metal as follows:

$$\mathbf{M}_{x,y,z} = \begin{cases} \text{air,} & \text{if } \mathbf{V}_{x,y,z} < -500, \\ \text{soft tissue,} & \text{if } -500 \leq \mathbf{V}_{x,y,z} < 500, \\ \text{bone,} & \text{if } 500 \leq \mathbf{V}_{x,y,z} < 3400, \\ \text{metal,} & \text{if } \mathbf{V}_{x,y,z} \geq 3400, \end{cases} \quad (1)$$

where $\mathbf{V}_{x,y,z}$ is the voxel value in Hounsfield Units (HU) at a location x, y, z in the reconstructed volume \mathbf{V} and $\mathbf{M}_{x,y,z}$ is the corresponding material obtained from thresholds.

We simulated the voxelized geometries of CBCT scanner used to acquire the clinical scans to keep the dataset realistic.

Each X-ray photon was started from a point source located at the position given by the scanner geometry. The path of the photon was simulated along a straight line from the X-ray source location to the rectangular detector grid as shown in Fig. 3. The initial energy of each photon was randomly sampled from a polychromatic energy distribution according to the kV used during the scan. When the X-ray photon passes through the object volume, the photon interacts with the object's material according to object's density and photon energy. Three types of X-ray interactions were simulated, namely, photoelectric absorption, Compton scattering, and Rayleigh scattering based on [37]. Energy-dependent attenuation curves provided in [37] were used for each of the material obtained from the segmentation in (1) to model interactions of the polychromatic X-ray photons during transport. Projections were formed by integrating the energy entering detector pixels. More detailed explanation about the Monte Carlo simulations can be found in [35] and [36]. We simulated multiple 2D projection views for each static reconstruction volume by rotating the source and detector according to the original geometry used by CBCT scanner. Furthermore, a simulation was done for each projection view without the object in the field-of-view to obtain 2D projections, called flat field. Projections were normalized and linearized, according to

$$p_{\text{sim}} = -\log\left(\frac{S}{F}\right), \quad (2)$$

where S is the simulated projection and F is the simulated flat field projection.

D. NETWORK ARCHITECTURE

The modified U-Net architecture used for metal segmentation training is shown in Fig. 4. It consisted of 4 downsampling




FIGURE 3. Illustration of CBCT geometry for Monte Carlo simulations. Incident rays leaving the object along a straight line primary (blue) and scattered rays (red) are shown.

(encoder) and 4 upsampling (decoder) blocks. Each down-sampling block had two convolutional layers (kernel size = 3×3 , stride = 1×1 , zero padding), each followed by a rectified linear unit (ReLU) and an instance normalization layer [38]. In the first block, the network had 16 channels. The number of channels was doubled in each block. After every two convolutional layers, a 2×2 max-pooling layer was applied to downsample features. Each upsampling block had a bilinear upsampling layer followed by a convolutional layer (kernel size = 3×3 , stride = 1×1 , zero padding). The features from encoder blocks were concatenated to the features of decoder blocks to allow flow of high resolution information from earlier layers. Output of each upsampling block was padded to match the dimensions of the skip connection from the encoder side. This padding scheme ensured the size of input and output remain same and varying sizes of input images could be given to the network during training. The final layer in the network was a 1×1 convolutional layer without any activation. A sigmoid function was applied after the final layer of the network to convert the network output to a probability map. To obtain metal segmentation during the inference, pixels with probability equal to or above 0.5 were classified as metal and pixels below the probability of 0.5 were classified as background.

E. IMPLEMENTATION DETAILS

Our network was implemented using PyTorch [39] deep learning library and was trained on a single GeForce RTX 2080 Ti Rev. A (11GB), with a batch size of 4. We used Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$) [40] to optimize the network parameters to minimize loss between the network output and ground truth. The initial learning rate was set to e^{-4} and reduced after each epoch in logarithmic steps down to e^{-6} . The learning rate was reduced during the first 25 epochs and was fixed to e^{-6} after that. As [41] showed that stacked data augmentations are effective for the model's generalization in medical segmentation tasks, we applied a number of stacked data augmentations during training. Each input projection was normalized by its maximum value before

augmentation. We used horizontal and vertical flips, image shift, rescale, rotation, Gaussian noise, multiplicative noise, elastic transformation, and mask dropout augmentations. The augmentations were applied on-the-fly during training with a probability of 0.2 for each augmentation. Therefore, each epoch had a slightly different set of inputs. To avoid over-fitting, the training was stopped if the loss did not decrease for five epochs (early stopping). The model parameters and data augmentations were initialized with random seed 2060. To account for stochastic training process, we trained each model two more times with random seeds 12060 and 22060.

We used binary cross entropy loss for training the neural network, defined as

$$l = \frac{1}{N \times w \times h} \sum_{i=0}^{N \times w \times h} y_i \log(\sigma(x_i)) + (1 - y_i) \log(1 - \sigma(x_i)), \quad (3)$$

where N is the batch size, w , h are width and height of the 2D projection image, x_i is the network output at pixel location i , σ is the sigmoid function, and y_i is the corresponding ground truth pixel value.

F. EVALUATION

We used 100 manually segmented metals from projections of 10 clinical scans as ground truths for evaluation of different metal segmentation methods. We evaluated U-Net-based models trained on noisy and clean simulations. For each of the noisy and clean simulations, the model was trained on full size, crops, and their combination. For forward projection-based conventional MAR (CMAR), we segmented the metals in 3D images using connected component region growing segmentation similar to [3]. Metal seeds were started at the voxels with intensities of more than 7000 HU. The metal region was grown starting from the seeds, by checking the intensities of the connected pixels in the neighborhood. If the neighboring pixel had an intensity of more than 3000 HU, it was included in the metal region. The segmented 3D metals were forward projected to create metal traces in the 2D projections. We also experimented with higher values than 3000 HU for lower (region growing) threshold, but then some metals were missed and undersegmented.

We also compared our results with moving metal artifact reduction (MMAR) [28]. We used MMAR to refine the coarse metal traces obtained by CMAR. For the patch-based training (Noisy (crops) and Clean (crops)), we followed [31] and used random crops of size 512×512 during the training. The inference was done on full size projections. For a fair comparison, we used the same U-Net network with the same parameters for all compared methods. For the comparison with the consistency check, we followed [31] to reconstruct 3D binary metal from the initial segmentation of Noisy (full+crops) model and forward project 3D binary metal to get updated metal traces.

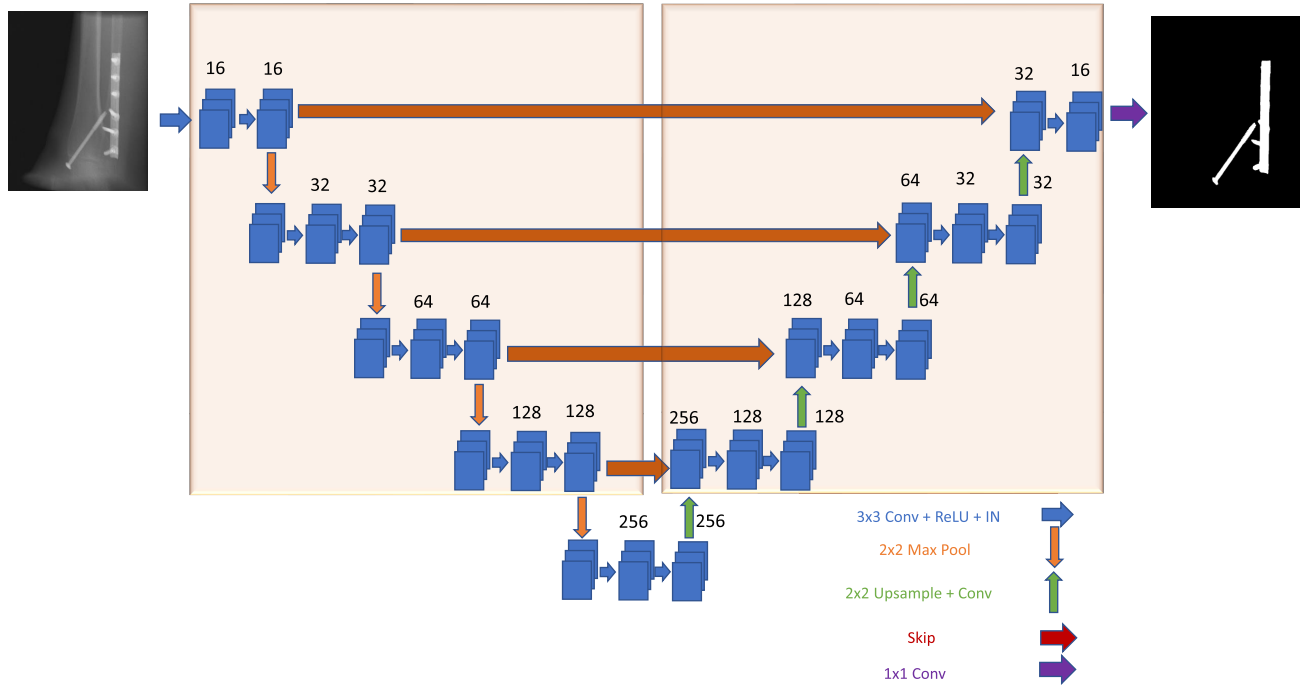


FIGURE 4. Network architecture used for metal segmentation. Simulated projections are given as input to the network. A sigmoid function and a threshold of 0.5 was applied at the output of the network to get binary metal trace.

The quantitative segmentation performance of the compared methods was evaluated using Dice Similarity Coefficient (DSC), Intersection over union (IOU) and False positive rate (FPR) [42]. These metrics are defined as

$$DSC = \frac{2TP}{2TP + FP + FN}, \quad (4)$$

$$IOU = \frac{TP}{TP + FP + FN}, \quad (5)$$

$$FPR = \frac{FP}{TN}, \quad (6)$$

where TP is true positive, TN is true negative, FP is false positive, and FN is false negative.

Metal traces found from the compared methods were inpainted from the nearby pixels using a 2D interpolation method. The inpainted projections were reconstructed using a modification of FDK [34] algorithm. The 3D metals obtained from the region growing-based method were put back in the reconstructions.

III. RESULTS

In this section we compare the quantitative metal segmentation performance of U-Net-based models with the forward projection-based CMAR, MMAR, and the consistency check. Then we show the qualitative impact of metal trace segmentation on the reconstructed images.

A. QUANTITATIVE ANALYSIS

1) CLINICAL DATA WITH METALS

As performance measures, we calculated the mean DSC and IOU for 10 test scans. Table 2 shows the segmentation

TABLE 2. Mean DSC (%) and mean IOU (%) (\pm standard error) calculated for metal trace segmentation on 10 test scans. \uparrow indicates that higher values are better.

Method	mean DSC \uparrow	mean IOU \uparrow
CMAR	90.5 \pm 2.0	84.4 \pm 3.2
MMAR	89.1 \pm 2.2	81.5 \pm 3.4
Noisy (full)	93.7 \pm 0.9	88.6 \pm 1.6
Noisy (crops)	94.4 \pm 0.7	89.6 \pm 1.2
Noisy (full+crops)	94.8 \pm 0.6	90.2 \pm 1.1
Consistency check	91.0 \pm 3.4	85.0 \pm 4.8
Clean (full)	93.0 \pm 1.0	87.4 \pm 1.6
Clean (crops)	92.2 \pm 1.6	86.5 \pm 2.1
Clean (full+crops)	93.8 \pm 0.8	88.6 \pm 1.4

performance of CMAR, MMAR, six U-Net-based models, and the consistency check-based segmentation applied on the model Noisy (full+crops). Compared to the CMAR and MMAR, all six U-Net-based models had higher mean DSCs and IOUs. This shows that the deep learning methods segmented metals more accurately than the forward projection-based CMAR and MMAR. The model Noisy (full+crops) gave the highest mean DSC=94.8 (SE=0.6) and IOU = 90.2 (SE=1.1). When MMAR was applied on the dilated metal traces obtained from CMAR, the mean DSC decreased from 90.5 (SE=2.0) to 89.1 (SE=2.2) and the mean IOU decreased from 84.4 (SE=3.2) to 81.5 (3.4). Furthermore, when consistency check [31] was applied on the best model Noisy (full+crops), it decreased the mean DSC from 94.8 (SE=0.6) to 91.0 (SE=3.4) and the mean IOU from 90.2 (SE=1.1) to 85.0 (SE=4.8). The mean DSC and IOU scores decreased due to the challenging cases of metals which




FIGURE 5. Boxplot visualization of DCE and IOU scores for CMAR, MMAR, U-Net models, and consistency check. The rectangles contain data within the first and the third quartiles. The endpoints of the lower and upper whiskers are selected as the first quartile - 1.5 times the interquartile range (IQR) and third quartile + 1.5 IQR, respectively. The medians are visualized as red lines and the means as green triangles. The outliers are the points that are outside the interval defined by the whiskers.

is explained further in the qualitative analysis in Section III-B. Among the three models trained on clean simulations, Clean (full+crops) had the highest mean DSC=93.8 (SE=0.8) and IOU=88.6 (SD=1.4), which is similar to the finding that the Noisy (full+crops) model had the highest mean DSC and IOU among the three models trained on noisy simulations.

We present the DSC and IOU results in Fig. 5. The box plots of DSCs and IOUs of MMAR had the largest spread over values. The median values of DSCs and IOUs of Noisy (full + crops) model were close to the median values of Noisy (crops). Boxplots of Noisy (full+crops) did not have any outliers as opposed to one outlier in Noisy (crops). The boxplots of Clean (crops) and Clean (full+crops) had more outliers than Noisy (crops) and Noisy (full+crops). Furthermore, the boxplots of consistency check had more vertical spread and showed one outlier in comparison to the boxplots of Noisy (full + crops).

2) CLINICAL DATA WITHOUT METALS

When comparing the segmentation performance of the models in the presence of the metals, it is also important to know how the models behave when the CBCT projections do not contain any metal. This comparison might show the model's generalization. This comparison does not require ground truth metal traces. Table 3 shows the FPR of the 6 models calculated for metal-free projections from 6 clinical scans. The mean FPR for the Noisy (full+crops) was 0.51×10^{-3} (SE= 0.16×10^{-3}) which was less than the mean FPR given by Noisy (crops), i.e., 0.85×10^{-3} (SE= 0.36×10^{-3}) and the mean FPR given by Noisy (full),

TABLE 3. Mean FPR (10^{-3}) (\pm standard error) calculated for 6 scans without any metals. \downarrow indicates that lower values are better.

Method	FPR \downarrow
Noisy (full)	14 ± 5
Noisy (crops)	0.85 ± 0.36
Noisy (full+crops)	0.51 ± 0.16
Clean (full)	23 ± 5
Clean (crops)	8 ± 2
Clean (full+crops)	0.44 ± 0.16

i.e., 14×10^{-3} (SE= 5×10^{-3}). Similarly, the model trained on simulations with 5000 photons per pixel (Clean (full+crops)) had least mean FPR i.e., 0.44×10^{-3} (SE= 0.16×10^{-3}) compared to the mean FPR given by Clean (crops), i.e., 8×10^{-3} (SE= 2×10^{-3}) and the mean FPR given by Clean (full), i.e., 23×10^{-3} (SE= 5×10^{-3}). Overall, this analysis on the metal-free projections showed that the models trained on the full size and cropped projections together were more robust the false positives compared to the models trained on only full size or only cropped projections.

3) COMPARISON OF DIFFERENT TRAINING RUNS

The comparison of DSC scores on test data for the 6 models obtained from 3 independent training runs is shown in Table 4. The mean DSC score for each run is shown in the columns named, seed1 (2060), seed2 (12060), and seed3 (22060). The last column named delta is the difference between maximum and minimum DSC scores across three runs. The model Noisy (full+crops), trained on the noisy full size and cropped projections, had the best mean DSC scores in each independent run, i.e., 94.8 (SE=0.6), 94.5 (SE=0.7), and 94.5 (SE=0.7). The delta for the model Noisy (full+crops) was 0.3, which was less than the delta of model Noisy (full), i.e., 2.3 and the delta of model Noisy (crops), i.e., 1.9. Similarly, the delta for the model Clean (full+crops) trained on clean full size and crops was 0.4, which was less than the delta of model trained on Clean (full), i.e., 2.1 and the delta of the model Clean (crops), i.e., 2.4. It can be deduced that the models trained on full size and crops were giving more stable mean DSC scores across three independent runs. Furthermore, the model Noisy (full+crops) had a slightly smaller delta (0.3) than the delta of the model Clean (full+crops), i.e., 0.4. These results show that the model trained on noisy full size and cropped projections was most robust model in terms of DSC scores.

B. QUALITATIVE ANALYSIS

For the qualitative analysis, we show the metal segmentation and the effect of accurate metal segmentation on the corresponding reconstructions. For all qualitative visualizations we compared Noisy (full+crops) reconstructions with the reconstructions from CMAR, MMAR, consistency check step [19] applied on the segmentation from Noisy (full+crops), and Clean (full+crops). We analyzed

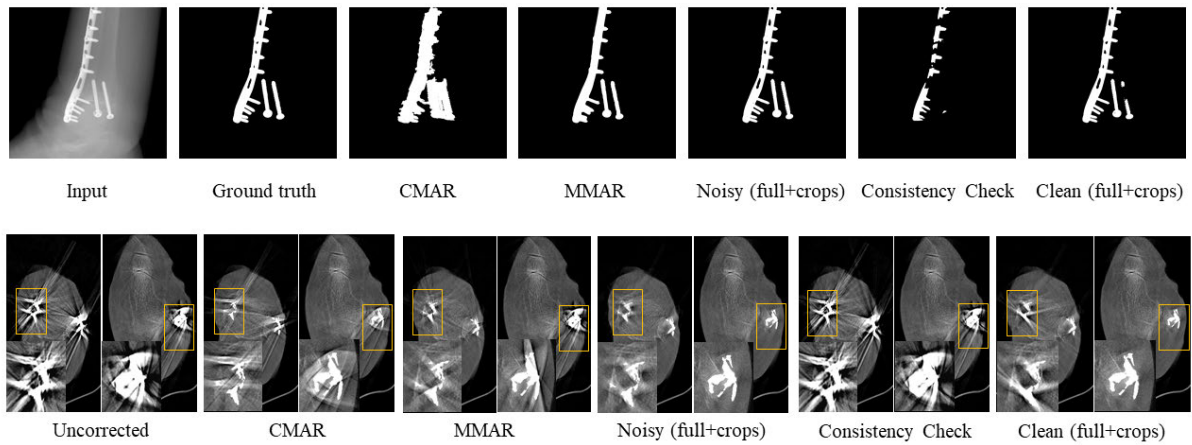


FIGURE 6. Test scan 4 with large motion and complex metal implants. The first row shows segmentation of metals in an input projection by various methods. The second row shows two axial image slices from the reconstructions (slice location 370 and 511, respectively) and the corresponding magnified region of interest. Yellow boxes show the location of the region of interest in the axial image. Our method, Noisy (full+crops) segmented metals in the projection most accurately which reduced most of the artifacts in the reconstructed axial slices.

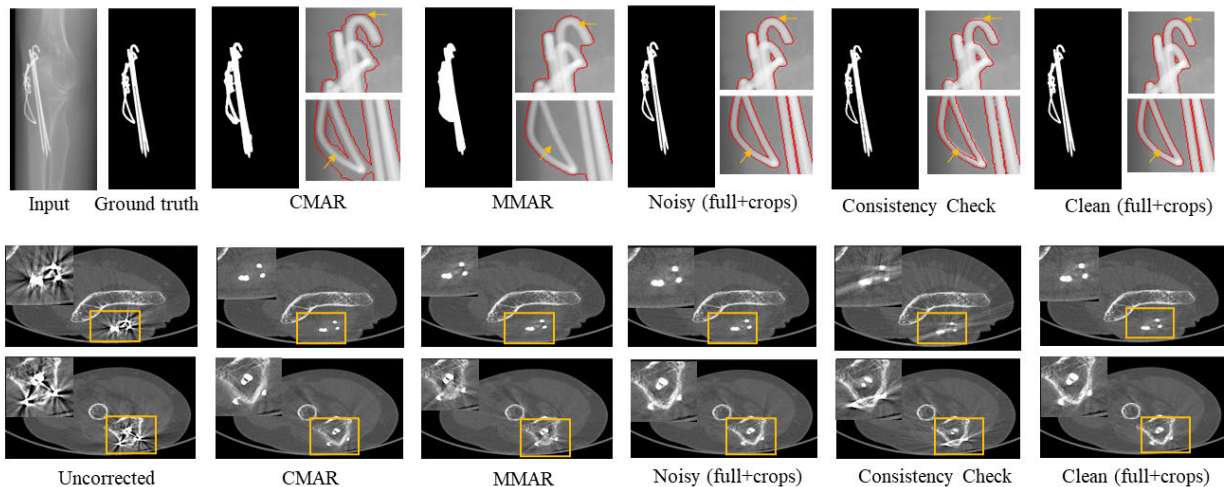


FIGURE 7. Test scan 9 with a complex metal shape. The first row shows an input projection, ground truth, and segmentation of the metal by several methods. Parts of the metal are enlarged to show the minute segmentation differences on the right of the segmented projections. The second and third rows show two axial image slices from the reconstructions with the corresponding magnified regions of interest. The axial slices are from slice location 169 and 396, respectively. Yellow boxes show the location of the region of interest in the axial image. Noisy (full+crops) and Clean (full+crops) segmented metals well which reduced most of the artifacts and preserved more details in the reconstructions.

TABLE 4. Mean DSC (%) (\pm standard error) calculated for three training runs for metal trace segmentation on 10 test scans. Delta is difference between maximum and minimum mean DSC for the corresponding method.

Method	seed 1	seed 2	seed 3	delta
Noisy (full)	93.7 \pm 0.9	92 \pm 1.4	94.2 \pm 0.8	2.2
Noisy (crops)	94.4 \pm 0.7	92.5 \pm 1.6	92.5 \pm 1.3	1.9
Noisy (full+crops)	94.8 \pm 0.6	94.5 \pm 0.7	94.5 \pm 0.7	0.3
Clean (full)	93.0 \pm 1.0	90.9 \pm 2.2	91.9 \pm 1.4	2.1
Clean (crops)	92.2 \pm 1.6	91.2 \pm 2.0	93.4 \pm 1.0	2.4
Clean (full+crops)	93.8 \pm 0.8	93.4 \pm 0.9	93.4 \pm 1.1	0.4

improvements in the image quality in complex cases such as the motion affected and out-of-FOV metal cases.

1) MOTION AFFECTED VOLUME

The qualitative results on the challenging motion case are shown in the Fig. 6. The segmentation of metals and two axial images from the reconstructions are visualized. The uncorrected images contained large artifacts from the presence of the large metals as well as the motion during the scan. The CMAR method over-segmented the metals in the projection. The conventional method reduced some artifacts in the reconstructed images but still there were many remaining streaks and darkening around the metal. The over-segmentation of metals in CMAR reduced some bone details in the axial image as well. MMAR refined the segmentations from CMAR but the holes inside the

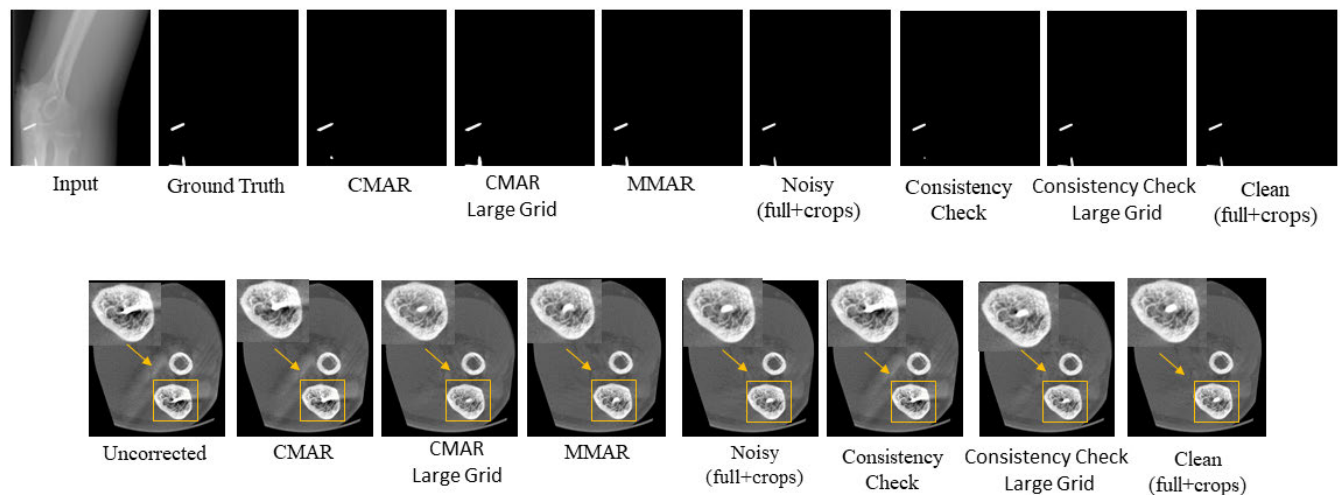


FIGURE 8. Test scan 6 containing out-of-FOV metal. The first row shows segmentation of the metals in the input projection by different methods and corresponding ground truth. The second row shows an axial slice from the reconstruction with the corresponding magnified region of interest. MMAR, Noisy (full+crops), and Clean (full+crops) segmented metals accurately without the need of larger reconstruction.

metals could not be recovered. The corresponding image slices for MMAR still had some remaining artifacts. The Noisy (full+crops) model segmented metals in projections most accurately. Most of the artifacts were reduced and the bones near the metal were clearly visible in the axial images. Because of the large motion, the consistency check method missed some of the blurred metals in the reconstructions and the forward projection of these metals removed most of the metal trace in the projections. This caused the artifacts to reappear in the images and affected image quality adversely. The model Clean (full+crops) missed some part of metals in the projection, this caused artifacts in one of the axial images.

2) UNDER- AND OVER-SEGMENTATION

Another interesting case is shown in Fig. 7. The metals were over-segmented by the region growing-based CMAR. The over-segmentation reduced some details near the metal such as the bone is blurred in the reconstructed axial image. MMAR could not deal with the complex shape of the metal and dilated the CMAR segmentation more. On the other hand, Noisy (full+crops) and Clean (full+crops) models segmented the metal traces well. As a result, the axial images from both of those methods had more details visible around the metal and the bone is better preserved. The consistency check missed only small part of the metal traces in the projection. The small under-segmentation caused significant artifacts in the reconstructed images. The results clearly showed that the under-segmentation may cause large artifacts in complex cases and over-segmentation may blur crucial details.

3) OUT-OF-FOV METAL

Fig. 8 shows when metal was in the out-of-FOV region, forward projections of image-domain metal failed to segment metal traces in the projection domain (CMAR and consistency check), which, in turn, caused artifacts in the

reconstructions. MMAR was able to recover the metals from the CMAR and did not need larger reconstruction. The original volume size for the reconstruction of the scan was $801 \times 801 \times 601$. To account for out-of-FOV metals in the forward projections, image reconstruction needed to be done on a larger size grid. When reconstruction was done for a volume size of $801 \times 801 \times 661$, the forward projection included all of the metals in the projections (CMAR large grid and consistency check large grid). Noisy (full+crops) and Clean (full+crops) models segmented the metal traces directly in the projections without the need of larger reconstruction. As we used the same metal inpainting and reconstruction procedure for all methods, it is clear from Fig. 8 that most of the artifacts were reduced in the reconstructed images if the metals were segmented well in the projections. In this case, the out-of-FOV metal was relatively close to the reconstruction edge. In some clinical cases where the out-of-FOV metal is far from the reconstruction FOV, the segmentation of metal is not possible. In such situations, our method will still segment metals in the projections without large reconstructions.

4) EFFECT OF VOXEL SIZE

The effect of small voxel sizes on the forward projections of metals is shown in Fig. 9. The images were reconstructed with a voxel size of 0.6 mm. The forward projection did not consider the voxel size and sampling artifacts are visible in the metal projections of the CMAR. Furthermore, CMAR missed some part of the metal, which was recovered by consistency check and MMAR. However, the MMAR segmentation resulted in over-segmentation of the metals and the consistency check segmentation still shows some sampling artifacts. Additional artifact can be seen in the reconstruction of CMAR, MMAR, and the consistency check. The model trained on the clean simulations missed

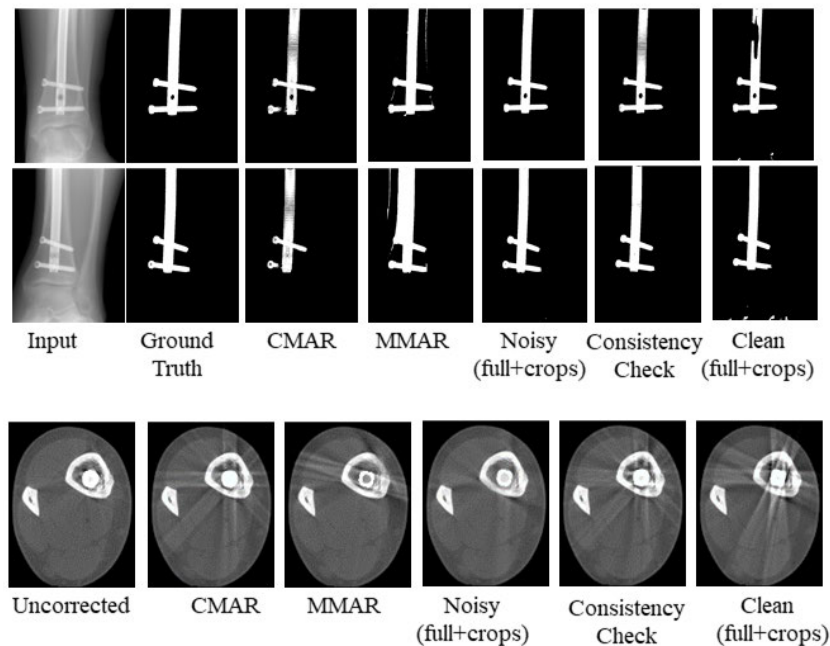


FIGURE 9. Test scan 8 with 0.6 mm voxel size. The first and second rows show segmentation of metals in the input projections by different methods and the corresponding ground truths. The third row shows an axial slice from the reconstructions. Our method, Noisy (full+crops) segmented metal traces best.

some metal part in the projections which caused artifacts in the reconstructed images. The model trained on noisy simulations segmented metals well which did not cause any additional artifacts. This visualization shows, that our method does not need to account for the voxel size as it segments metals directly in the projections.

IV. CONCLUSION

In this paper, we proposed to use noisy Monte Carlo simulations to train a U-Net architecture for the segmentation of metals in CBCT projections. We experimentally demonstrated that synthetic data could substitute real data for metal segmentation training. We showed that the model trained with noisy simulations outperformed the model trained on time-consuming cleaner simulations for the metal segmentation task. We also showed that adding the crops to the full size projections during the training helped to get more robust metal segmentation on the real clinical test scans. The forward projection based methods, such as CMAR, MMAR, and consistency checks are prone to errors and require tuning of parameters. However, segmenting the metals directly in projections is more robust solution.

The model trained on full size projections and crops, significantly reduced the false positives in the non-metal projections. Although, we trained our model on limited simulations, we have shown the applicability of the model on a diversified real dataset acquired from multiple anatomies and scanner sites. We have discussed noticeable improvements in the reconstruction image quality of motion-affected

and out-of-FOV metal affected scans by obtaining better segmentation traces of metals. Recently, a Fourier dual-domain restoration network has been shown to be more robust to metal trace defects [43] which is another approach to alleviate the artifacts coming from wrong segmentation of metals. In our future work, we will consider to test the application of our method on more diversified and larger dataset including dental scans.

V. STUDY LIMITATIONS

While we have demonstrated excellent results for metal artifact correction using our method, we have used only simulated data for the training. This approach helps to address the challenges associated with obtaining enough clinical data for training and manually segmented ground truth. However, simulated projections might not fully capture all the variations and complexities inherent in the clinical scans. This could potentially impact the generalizability and effectiveness of our proposed method in real-world scenarios. To address this limitation, we incorporated noisy simulations and included 10 clinical scans for testing. Although we have included complex cases, such as those with motion artifact and out-of-fov metals, there is still a necessity to include a larger dataset to fully establish the model's generalizability and effectiveness.

As we have noted earlier, metal artifact reduction methods often require precise segmentation of metal traces in CBCT projections. The efficacy of the trained model's performance depends upon the accuracy of this ground truth metal

segmentation. If the metals in the simulated training data are not segmented accurately, the performance of the trained model may degrade on the real clinical cases.

REFERENCES

- [1] R. Schulze, U. Heil, D. Gross, D. Bruellmann, E. Dranischnikow, U. Schwanecke, and E. Schoemer, "Artefacts in CBCT: A review," *Dentomaxillofacial Radiol.*, vol. 40, no. 5, pp. 265–273, Jul. 2011.
- [2] W. A. Kalender, R. Hebel, and J. Ebersberger, "Reduction of CT artifacts caused by metallic implants," *Radiology*, vol. 164, no. 2, pp. 576–577, Aug. 1987.
- [3] E. Meyer, R. Raupach, M. Lell, B. Schmidt, and M. Kachelrieß, "Normalized metal artifact reduction (NMAR) in computed tomography," *Med. Phys.*, vol. 37, no. 10, pp. 5482–5493, Sep. 2010.
- [4] J. Wang, S. Wang, Y. Chen, J. Wu, J.-L. Coatrieux, and L. Luo, "Metal artifact reduction in CT using fusion based prior image," *Med. Phys.*, vol. 40, no. 8, Jul. 2013, Art. no. 081903.
- [5] M. Meilinger, C. Schmidgunst, O. Schütz, and E. W. Lang, "Metal artifact reduction in cone beam computed tomography using forward projected reconstruction information," *Zeitschrift für Medizinische Physik*, vol. 21, no. 3, pp. 174–182, Sep. 2011.
- [6] X. Duan, L. Zhang, Y. Xiao, J. Cheng, Z. Chen, and Y. Xing, "Metal artifact reduction in CT images by sinogram TV inpainting," in *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, Oct. 2008, pp. 4175–4177.
- [7] C. Peng, B. Qiu, M. Li, Y. Guan, C. Zhang, Z. Wu, and J. Zheng, "Gaussian diffusion sinogram inpainting for X-ray CT metal artifact reduction," *Biomed. Eng. OnLine*, vol. 16, no. 1, pp. 1–17, Dec. 2017.
- [8] H. Liao, W.-A. Lin, Z. Huo, L. Vogelsang, W. J. Sehnert, S. K. Zhou, and J. Luo, "Generative mask pyramid network for CT/CBCT metal artifact reduction with joint projection-sinogram correction," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*. Cham, Switzerland: Springer, 2019, pp. 77–85.
- [9] M. U. Ghani and W. C. Karl, "Fast enhanced CT metal artifact reduction using data domain deep learning," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 181–193, 2020.
- [10] C. Peng, B. Li, P. Liang, J. Zheng, Y. Zhang, B. Qiu, and D. Z. Chen, "A cross-domain metal trace restoring network for reducing X-ray CT metal artifacts," *IEEE Trans. Med. Imag.*, vol. 39, no. 12, pp. 3831–3842, Dec. 2020.
- [11] H. Agrawal, A. Hietanen, and S. Särkkä, "Metal artifact reduction in cone-beam extremity images using gated convolutions," in *Proc. IEEE 18th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2021, pp. 1087–1090.
- [12] J. Ahn, Y. Choi, and J. Baek, "Convolutional neural network based metal artifact reduction method in dental CT image," *Proc. SPIE*, vol. 11595, pp. 715–720, Feb. 2021.
- [13] L. Gjestebj, Q. Yang, Y. Xi, H. Shan, B. Claus, Y. Jin, B. De Man, and G. Wang, "Deep learning methods for CT image-domain metal artifact reduction," *Proc. SPIE*, vol. 10391, pp. 147–152, Sep. 2017.
- [14] K. Liang, L. Zhang, H. Yang, Y. Yang, Z. Chen, and Y. Xing, "Metal artifact reduction for practical dental computed tomography by improving interpolation-based reconstruction with deep learning," *Med. Phys.*, vol. 46, no. 12, Dec. 2019.
- [15] D. J. Moseley, J. H. Siewerdsen, and D. A. Jaffray, "High-contrast object localization and removal in cone-beam CT," *Proc. SPIE*, vol. 5745, pp. 40–50, Apr. 2005.
- [16] C. Peng, B. Li, M. Li, H. Wang, Z. Zhao, B. Qiu, and D. Z. Chen, "An irregular metal trace inpainting network for X-ray CT metal artifact reduction," *Med. Phys.*, vol. 47, no. 9, pp. 4087–4100, Sep. 2020.
- [17] L. Yu, Z. Zhang, X. Li, and L. Xing, "Deep sinogram completion with image prior for metal artifact reduction in CT images," *IEEE Trans. Med. Imag.*, vol. 40, no. 1, pp. 228–238, Jan. 2021.
- [18] W.-A. Lin, H. Liao, C. Peng, X. Sun, J. Zhang, J. Luo, R. Chellappa, and S. K. Zhou, "DuDoNet: Dual domain network for CT metal artifact reduction," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2019, pp. 10512–10521.
- [19] M. D. Ketcha, M. Marrama, A. Souza, A. Uneri, P. Wu, X. Zhang, P. A. Helm, and J. H. Siewerdsen, "Sinogram + image domain neural network approach for metal artifact reduction in low-dose cone-beam computed tomography," *J. Med. Imag.*, vol. 8, no. 5, Mar. 2021, Art. no. 052103.
- [20] M. Stille, B. Kratz, J. Müller, N. Maass, I. Schasiepen, M. Elter, I. Weyers, and T. M. Buzug, "Influence of metal segmentation on the quality of metal artifact reduction methods," *Proc. SPIE*, vol. 8668, pp. 902–907, Mar. 2013.
- [21] S. Choi, S. Moon, and J. Baek, "A metal artifact reduction method for small field of view CT imaging," *PLoS One*, vol. 16, no. 1, Jan. 2021, Art. no. e0227656.
- [22] A. P. Candemil, B. Salmon, D. Q. Freitas, G. M. Ambrosano, F. Haiter-Neto, and M. L. Oliveira, "Metallic materials in the exomass impair cone beam CT voxel values," *Dentomaxillofacial Radiol.*, vol. 47, no. 6, Jul. 2018, Art. no. 20180011.
- [23] A. P. Candemil, B. Salmon, D. Q. Freitas, F. Haiter-Neto, and M. L. Oliveira, "Distribution of metal artifacts arising from the exomass in small field-of-view cone beam computed tomography scans," *Oral Surgery, Oral Med., Oral Pathol. Oral Radiol.*, vol. 130, no. 1, pp. 116–125, Jul. 2020.
- [24] C. Nardi, R. Molteni, C. Lorini, G. G. Taliani, B. Matteuzzi, E. Mazzoni, and S. Colagrande, "Motion artefacts in cone beam CT: An in vitro study about the effects on the images," *Brit. J. Radiol.*, vol. 89, no. 1058, Feb. 2016, Art. no. 20150687.
- [25] H. Agrawal, "Deep learning methods for visual fault diagnostics of dental X-ray systems," M.S. thesis, Dept. Neurosci. Biomed. Eng., Aalto Univ., Sch. Sci., Espoo, Finland, 2018.
- [26] B. De Man, J. Nuyts, P. Dupont, G. Marchal, and P. Suetens, "Metal streak artifacts in X-ray computed tomography: A simulation study," *IEEE Trans. Nucl. Sci.*, vol. 46, no. 3, pp. 691–696, Jun. 1999.
- [27] J. Toftgaard, W. Fledelius, D. Seghers, M. Huber, M. Brehm, E. S. Worm, U. V. Elström, and P. R. Poulsen, "Moving metal artifact reduction in cone-beam CT scans with implanted cylindrical gold markers," *Med. Phys.*, vol. 41, no. 12, Nov. 2014, Art. no. 121710.
- [28] A. Hahn, M. Knaup, M. Brehm, S. Sauppe, and M. Kachelrieß, "Two methods for reducing moving metal artifacts in cone-beam CT," *Med. Phys.*, vol. 45, no. 8, pp. 3671–3680, Aug. 2018.
- [29] A. Amirkhanov, C. Heinzl, M. Reiter, J. Kastner, and E. Grollier, "Projection-based metal-artifact reduction for industrial 3D X-ray computed tomography," *IEEE Trans. Vis. Comput. Graph.*, vol. 17, no. 12, pp. 2193–2202, Dec. 2011.
- [30] M. A. A. Hegazy, M. H. Cho, M. H. Cho, and S. Y. Lee, "U-net based metal segmentation on projection domain for metal artifact reduction in dental CT," *Biomed. Eng. Lett.*, vol. 9, no. 3, pp. 375–385, Aug. 2019.
- [31] T. M. Gottschalk, A. Maier, F. Kordon, and B. W. Kreher, "View-consistent metal segmentation in the projection domain for metal artifact reduction in CBCT—An investigation of potential improvement," *Mach. Learn. Biomed. Imag.*, vol. 1, pp. 1–10, Dec. 2021.
- [32] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. (MICCAI)*. Cham, Switzerland: Springer, 2015, pp. 234–241.
- [33] J. Maier, E. Eulig, T. Vöth, M. Knaup, J. Kuntz, S. Sawall, and M. Kachelrieß, "Real-time scatter estimation for medical CT using the deep scatter estimation: Method and robustness analysis with respect to different anatomies, dose levels, tube voltages, and data truncation," *Med. Phys.*, vol. 46, no. 1, pp. 238–249, Jan. 2019.
- [34] L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone-beam algorithm," *J. Opt. Soc. Amer. A, Opt. Image Sci.*, vol. 1, no. 6, pp. 612–619, 1984.
- [35] F. Salvat, J. M. Fernández-Varea, and J. Sempau, "PENELOPE-2006: A code system for Monte Carlo simulation of electron and photon transport," in *Proc. Workshop*, vol. 4, no. 6222. Paris, France: Nucl. Energy Agency, 2006, p. 7.
- [36] A. Badal and A. Badano, "Accelerating Monte Carlo simulations of photon transport in a voxelized geometry using a massively parallel graphics processing unit," *Med. Phys.*, vol. 36, no. 11, pp. 4878–4880, Oct. 2009.
- [37] D. E. Cullen, J. H. Hubbell, and L. Kissel, "EPDL97: The evaluated photo data library 97 version," Law. Liv. Nat. Lab. (LLNL), Livermore, CA, USA, Tech. Rep. UCRL-50400-Vol.6-Rev.5, 1997.
- [38] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," 2016, *arXiv:1607.08022*.
- [39] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32. Red Hook, NY, USA: Curran Associates, 2019, pp. 8024–8035.
- [40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*.

- [41] L. Zhang, X. Wang, D. Yang, T. Sanford, S. Harmon, B. Turkbey, B. J. Wood, H. Roth, A. Myronenko, D. Xu, and Z. Xu, "Real-time scatter estimation for medical CT using the deep scatter estimation: Method and robustness analysis with respect to different anatomies, dose levels, tube voltages, and data truncation," *Med. Phys.*, vol. 39, no. 7, pp. 2531–2540, 2020.
- [42] A. A. Taha and A. Hanbury, "Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool," *BMC Med. Imag.*, vol. 15, no. 1, pp. 1–28, Dec. 2015.
- [43] Z. Li, Q. Gao, Y. Wu, C. Niu, J. Zhang, M. Wang, G. Wang, and H. Shan, "FD-MAR: Fourier dual-domain network for CT metal artifact reduction," 2022, *arXiv:2207.11678*.



HARSHIT AGRAWAL (Graduate Student Member, IEEE) received the B.Tech. degree (Hons.) from UPTU, India, in 2011, and the M.Sc. (Tech.) degree from Aalto University, Finland, in 2018, where he is currently pursuing the D.Sc. degree with the Department of Electrical Engineering and Automation Engineering. He was a Senior Engineer with GE HealthCare, India. He is also a Machine Learning Specialist with Planmeca Oy, Finland. His research interests include machine learning, especially application of deep neural networks in medical imaging.



ARI HIETANEN received the Ph.D. degree in theoretical physics from the University of Helsinki, Finland, in 2008. He is currently a Research Manager of algorithms with Planmeca Oy. His current research interests include machine learning and medical physics.



SIMO SÄRKKÄ (Senior Member, IEEE) is currently an Associate Professor with Aalto University, Finland. He has authored or coauthored more than 100 peer-reviewed scientific articles and three books. His research interests include multi-sensor data processing systems and machine learning methods with applications in medical and health technology, target tracking, inverse problems, and location sensing. He is a member of the IEEE Machine Learning for Signal Processing Technical Committee. He has been serving as an Associate Editor for the IEEE SIGNAL PROCESSING LETTERS.

...