

Deep Learning Based Sinogram Correction for Metal Artifact Reduction

Muhammad Usman Ghani and W. Clem Karl; Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA.

Abstract

Computed Tomography (CT) is a non-invasive imaging technique that reconstructs cross-sectional images of scenes from a series of projections acquired at different angles. In applications such as airport security luggage screening, the presence of dense metal clutter causes beam hardening and streaking in the resulting conventionally formed images. These artifacts can lead to object splitting and intensity shading that make subsequent labeling and identification inaccurate. Conventional approaches to metal artifact reduction (MAR) have post-processed the artifact-filled images or interpolated the metal regions of the sinogram projection data. In this work, we examine the use of deep-learning-based methods to directly correct the observed sinogram projection data prior to reconstruction using a fully convolutional network (FCN). In contrast to existing learning-based CT artifact reduction work, we work completely in the sinogram domain and train a network over the entire sinogram (versus just local image patches). Since the information in sinograms pertaining to objects is non-local, patch-based methods are not well matched to the nature of CT data. The use of an FCN provides better computational scaling than historical perceptron-based approaches. Using a poly-energetic CT simulation, we demonstrate the potential of this new approach in mitigating metal artifacts in CT.

Introduction

X-ray computed tomography (CT) allows visualizing the interior of objects in a non-destructive fashion. A source transmits X-rays which are detected by a receiver after passing through the region of interest (ROI). This configuration provides projection data of material properties of the scene at different offsets and angles. The resulting observed projection data is commonly referred as the sinogram. Reconstruction algorithms then couple the observed sinogram with an imaging model to reconstruct cross-sectional attenuation images of the original scene. The reconstructed images represent X-ray attenuation properties of objects in the ROI, which are dependent on X-ray energy.

The Filtered back projection (FBP) algorithm is the standard reconstruction method used in most CT scanners. FBP produces excellent reconstructions when we have access to high-quality sinogram data with complete angular coverage and scenes with constrained density ranges. In the presence of highly attenuating metallic objects, however, FBP can produce severe streak artifacts resulting from inconsistencies in the underlying assumptions on which FBP is based [1, 2]. These image artifacts produce intensity shading and object splitting which makes the subsequent image analysis steps of object segmentation and identification inaccurate. These challenges are especially strong in security applications, where scenes are more likely to contain a large spectrum

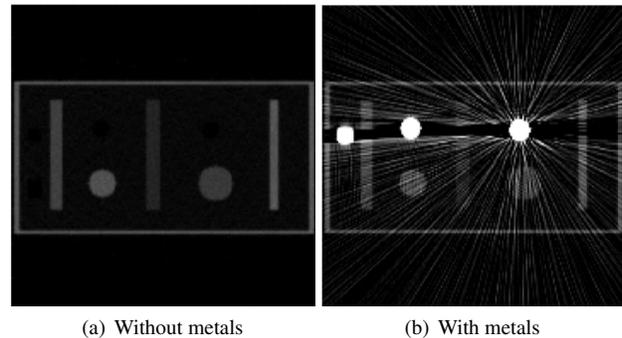


Figure 1. Reconstruction of an example scene without and with metallic objects using FBP method. This illustrates the challenging nature of metal artifacts which can severely impact CT reconstructions.

of materials. The problem is illustrated in Figure 1 for a synthetic example, where reconstructions of the same scene with and without highly-absorbing metals are presented.

X-ray CT has applications in a vast range of areas including medical imaging [3], materials sciences [4], and security [5]. Metal artifacts are a major challenge in CT image reconstructions. In this paper, we propose a deep learning-based approach for metal artifact reduction (MAR). We apply a convolutional neural network (CNN) that learns to perform sinogram correction using a large simulated dataset. In particular, we use a fully convolutional network (FCN) over the entire sinogram to correct the sinogram data, and finally use a conventional FBP algorithm for reconstruction of the corrected sinogram data. In this we are removing problems in the projection data rather than in the subsequent image. Our proposed approach is simple, efficient, and in initial experiments seems to outperform existing popular MAR techniques.

Prior Work

Many MAR techniques have been proposed in literature but the problem still remains the focus of research. The existing MAR methods can be categorized into four groups: physics-based correction, iterative reconstruction, image post-processing, and sinogram correction methods. Physics-based methods involve pre-processing of the projection data using different filtering operations to reduce effects of metal artifact inducing mechanisms. Some examples include adaptive filtering [6], edge-preserving blur filter [6].

Iterative reconstruction methods generally use a model-based approach to minimize a well-defined objective function subject to defined constraints. De Man et al. [7] proposed

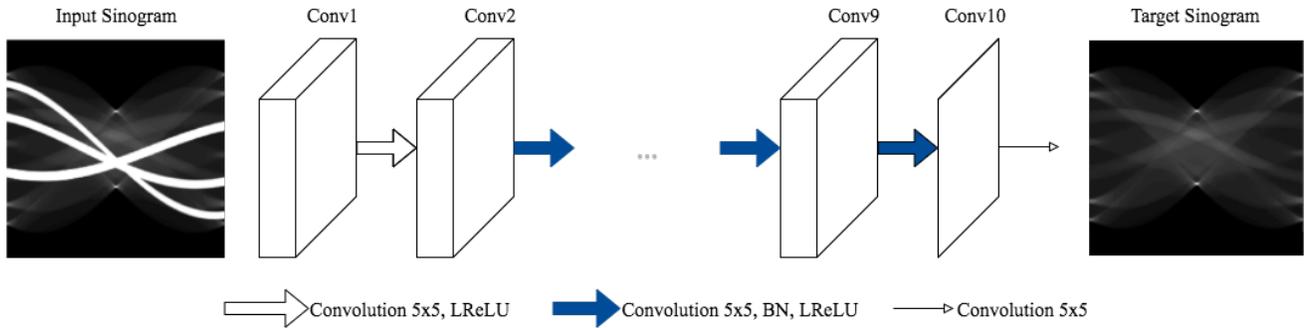


Figure 2. FCN architecture used in this paper to learn sinogram correction.

a maximum-likelihood-based MAR approach by using poly-energetic model. Hamelin et al. [8] approximated data as Gaussian distribution and used the poly-energetic model to reduce metal artifacts. Zhang et al. [9] suggested a penalized least squares methods subject to non-negativity constraint and discard metal traces data.

Image post-processing approaches rely on image processing methods to reduce artifacts in the image domain. Soltanian-Zadeh et al. [10] exploit the high-frequency nature of metal-induced streaks and estimate streaks as the difference between original reconstruction and its low-pass-filtered version. Bal et al. [11] applied an adaptive filter in the metal affected regions of the image. Gjesteby et al. [12] used a simple CNN to learn a mapping from metal affected image to artifact-free image.

Sinogram correction methods are the most popular MAR techniques, the idea is to replace metal-corrupted sinogram data with estimated values in an effort to perform sinogram correction. The values for metal-corrupted sinogram regions are typically estimated using neighboring information. Kalender et al. [13] consider metal-corrupted data as missing data and estimate the missing data using a one-dimensional linear interpolation method (referred as LI-MAR in the rest of the paper). LI-MAR is considered a benchmark and is so far the most cited MAR technique [2]. Morin and Raeside [14] use nearest neighbor search to estimate the missing data. Mahnken et al. [15] suggested a two-dimensional interpolation that replaces the metal-corrupted data by a weighted sum of 16-nearest neighbors (referred as WNN-MAR in the rest of the paper). Zhao et al. [16] perform sinogram interpolation using wavelet coefficients. Normalized MAR (NMAR) [17] uses a prior image to normalize the corrupted sinogram and a simple interpolation method to estimate metal-corrupted data in the normalized sinogram. NMAR heavily relies on a prior image and is not suitable for security or other applications where a correct prior image cannot be robustly found. A comprehensive review of MAR techniques can be found in Gjesteby et al. [2].

Deep Learning Based Sinogram Correction

In this paper, we exploit the structured nature of sinograms and learn a CNN-based function to identify and correct the corrupted data in sinogram domain. We propose a completely automated approach to perform sinogram correction that can handle diverse sizes and numbers of metals placed on arbitrary locations. Our approach exploits a large simulated dataset to learn to per-

form sinogram correction using a deep CNN. As opposed to other learning based methods, we work in the sinogram domain and use complete sinogram data to learn a mapping function to correct the sinogram data. We use a fully convolutional network (FCN) architecture in which convolutional kernels of all layers learn representations at same scale. As compared to fully connected networks these structures have fewer learning parameters and due to convolutional nature of operations, exploits local neighborhood information.

Fully Convolutional Network (FCN)

With AlexNet [18] winning the ImageNet challenge in 2012 [19] has resulted in renewed interest in theoretical and applications aspects of neural networks. The availability of different machine learning frameworks and hardware computing capabilities have enabled researchers to rapidly develop new algorithms and achieve state-of-the-art performance using neural networks-based methods which in many cases even surpass human performance [20, 21, 22]. These successes have attracted researchers in many related areas to focus on developing machine learning-based algorithms to solve interesting problems in their domain. Apart from image classification, CNNs have shown tremendous performance in various image restoration tasks such as image denoising [23], artifact reduction [24], and image super-resolution [25].

Our CNN architecture is inspired by VDSR [25], which is a 20-layer FCN used for image super-resolution. Our network is different in that it uses 10 convolutional layers with 5×5 convolutional kernels whereas VDSR used 3×3 kernels. The network architecture used in this paper is presented in Figure 2. All the layers in our network have 64 channels except the last layer which has only 1 channel. We use single-pixel stridden convolutions and use zero-padding so that convolutions do not change the image size. We use the leaky rectified linear unit (LReLU) non-linearity with 0.2 slope at the first 9 layers. Additionally, we use the batch normalization (BN) [26] operation at each layer except the first and last layer. This architecture results in 41×41 theoretical effective receptive field, so our networks can potentially use information from a large area to estimate the missing data. The input to our FCN is the metal-corrupted sinogram and corresponding target is the artifact-free sinogram. The artifact-free sinogram is generated by simulating the same scene with no metallic materials in the field of view. We use the mean-square-error loss function to train the network.

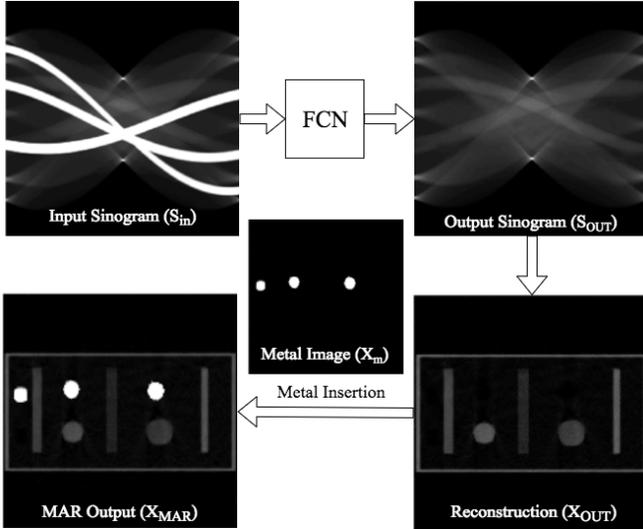


Figure 3. FCN-MAR algorithm is presented, once FCN is trained completely it is used to correct sinogram data which produces reconstructions without metals. The metal image is then used to insert metals back to produce final output of the algorithm.

FCN-MAR Algorithm

Our FCN-MAR algorithm has three major steps: (i) metal(s) segmentation in the reconstructed image, (ii) sinogram correction, and (iii) reconstruction using FBP and inserting segmented metal(s). These steps are illustrated in Figure 3. This paper focuses on simulated data only where we can use oracle information for the metal segmentation step. So we use perfect knowledge of metals for metal image X_m . We train FCN on a large dataset which is later used to perform sinogram correction. Once sinogram is corrected, we use FBP to reconstruct the image and insert the metal(s) back in the reconstructed image. While we use oracle information for metal segmentation step in this work, in practice, we do not have access to perfect information about the metals. One of the ways to automate the metals segmentation step is to reconstruct the image using incorrect sinogram data and apply thresholding followed by morphological operations.

Experiments

This section discusses our experiments. Our aim was to understand the potential for sinogram-based deep learning. We describe our accurate poly-energetic CT simulation setup, MAR dataset generation, FCN training strategy, reconstruction results, and compare the performance of our approach to two popular MAR methods. We consider LI-MAR and WNN-MAR for comparison and use mean-square-error to compare results of different methods.

We used a physically accurate poly-energetic CT simulation to generate data for training and testing. An X-ray source transmits multi-energetic X-rays which are partially absorbed by scene objects with attenuation $\mu(E)$ and subsequently received by an energy integrating detector. The energy-dependent absorption is in accordance with the X-ray absorption properties of the object, $\mu(E)$, described by linear attenuation coefficients (LACs). In this work we ignore X-ray scatter for simplicity. The detectors perform a weighted sum received photons over the range of the ener-

gies in the beam. For ideal scenarios, the resulting sinogram can be specified by Equation 1.

$$I = - \sum_{i=1}^N I_0 \cdot \eta(E_i) \cdot e^{-\sum_l \mu(\vec{x}, E_i)} \quad (1)$$

where I represents intensity of received ray, I_0 represents source intensity factor, N denotes the number of energies used to estimate the discretized version of continuous energy dependence of the source beam, $\eta(E_i)$ denotes the CT system weighting function, $\mu(\vec{x}, E_i)$ represents the LAC value at energy E_i at the scene location \vec{x} , and \sum_l denotes the projection operation commonly referred as the Radon Transform.

In real systems, however, the sinogram is subjected to electronic and data-dependent noise. We model electronic noise as a Gaussian distribution with zero-mean and σ_e^2 variance. The mean of electronic noise is generally non-zero, but can be estimated and subtracted to result in a zero-mean effect. The data-dependent noise is modeled as the Poisson distribution with mean I . Further, it is common to use as data a log-normalized sinogram. The resulting log-normalized observed sinogram, \tilde{I}_n , is given by Equation 2.

$$\tilde{I}_n = - \ln \left(\frac{\sum_{i=1}^N \text{Poisson} \left(I_0 \cdot \eta(E_i) \cdot e^{-\sum_l \mu(\vec{x}, E_i)} \right) + \mathcal{N}(0, \sigma_e^2)}{\sum_{i=1}^N I_0 \cdot \eta(E_i)} \right) \quad (2)$$

MAR Dataset

We create a MAR dataset by generating pairs of artifact-free and metal-corrupted scenes and their corresponding sinograms. Our simulation setup is based on Equation 2. The weighting function corresponds the Imatron C300 scanner by Crawford et al. [27]. We use $N = 86$ uniformly sampled energy bins between 10 KeV and 95 KeV, source flux $I_0 = 1.8 \times 10^5$, and electronic noise variance σ_e^2 equivalent to the statistics of 170 photons received at 70 KeV (suggested by [28]). We use LAC values from the NIST XCOM database [29] to simulated suitcases of size $179mm \times 179mm$. We used 255 detector channels and 256 projection angles uniformly sampled between 0^0 and 180^0 .

In order to create MAR datasets, we prepared a stochastic suitcase simulator which inserts the objects at different locations in the suitcase. Here, object material, shape, and location is selected randomly. For training, we generated 10,000 suitcases with non-metal materials, and their corresponding sinograms with up to 3 metals. For each suitcase, we generate a 1 metal, a 2 metal, and a 3 metal example. This results in 30,000 pairs of suitcases which we used to train our network. For testing, we generated 3,000 similar pairs and compared the performance of our method to LI-MAR and WNN-MAR.

FCN Training

We applied the Adam optimizer [30] to train FCN with a mini-batch size of 16. We used initial learning rate of 0.01 with momentum parameters $\beta_1 = 0.9, \beta_2 = 0.999$. The network was trained for 200 epochs and the final state of the network was later used on test data. We used Tensorflow [31] to implement and train this network. The network training took approximately 48 hours on a Nvidia Tesla P100 GPU. Once the training is

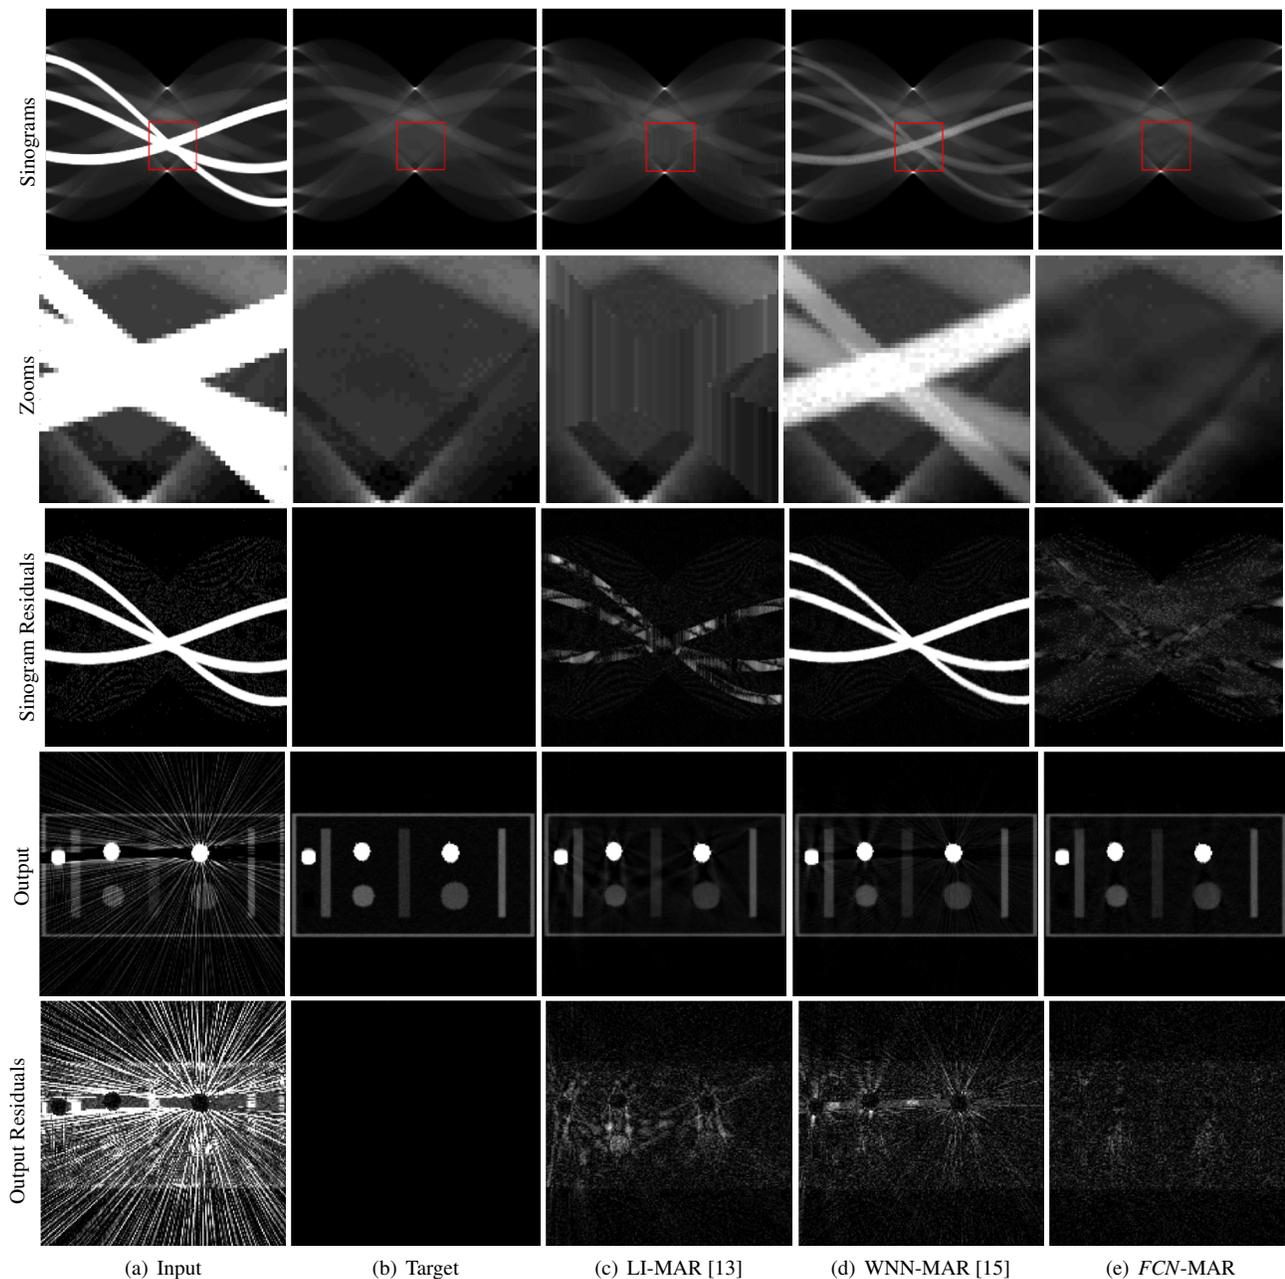


Figure 4. Input and output sinograms, reconstructions and residuals corresponding to considered MAR methods are shown for an example scene. FCN-MAR evidently outperforms other methods by producing sinograms which are very similar to target sinograms. Zoomed patches displayed in range $[0.42 \ 1.5]$ accentuate the differences among different sinogram correction methods. Reconstruction results show the effectiveness of our deep learning based FCN-MAR approach in artifact reduction. FCN-MAR suppresses most of the metal-induced streaks while preserving the structure.

completed, it takes less than a second to process each test sample.

Results and Discussion

We present a comparison of MAR methods, showing their impact in both the sinogram and the image domain for two examples of our larger set chosen as particularly challenging examples. In each case the input and target sinogram pairs along with the outputs of all three methods considered are presented.

The process of sinogram correction consists of two steps: identification of metal corrupted traces, and replacing those data points with estimate values. For LI-MAR and WNN-MAR, we use oracle information for identification step and use corresponding algorithms to estimate new values from the reliable neighboring data. However, our approach is completely automated in that it automatically identifies the corrupted data and estimates new data. This is possible due to the learning-based nature of our approach that relies on deep-learning and big-data to learn to

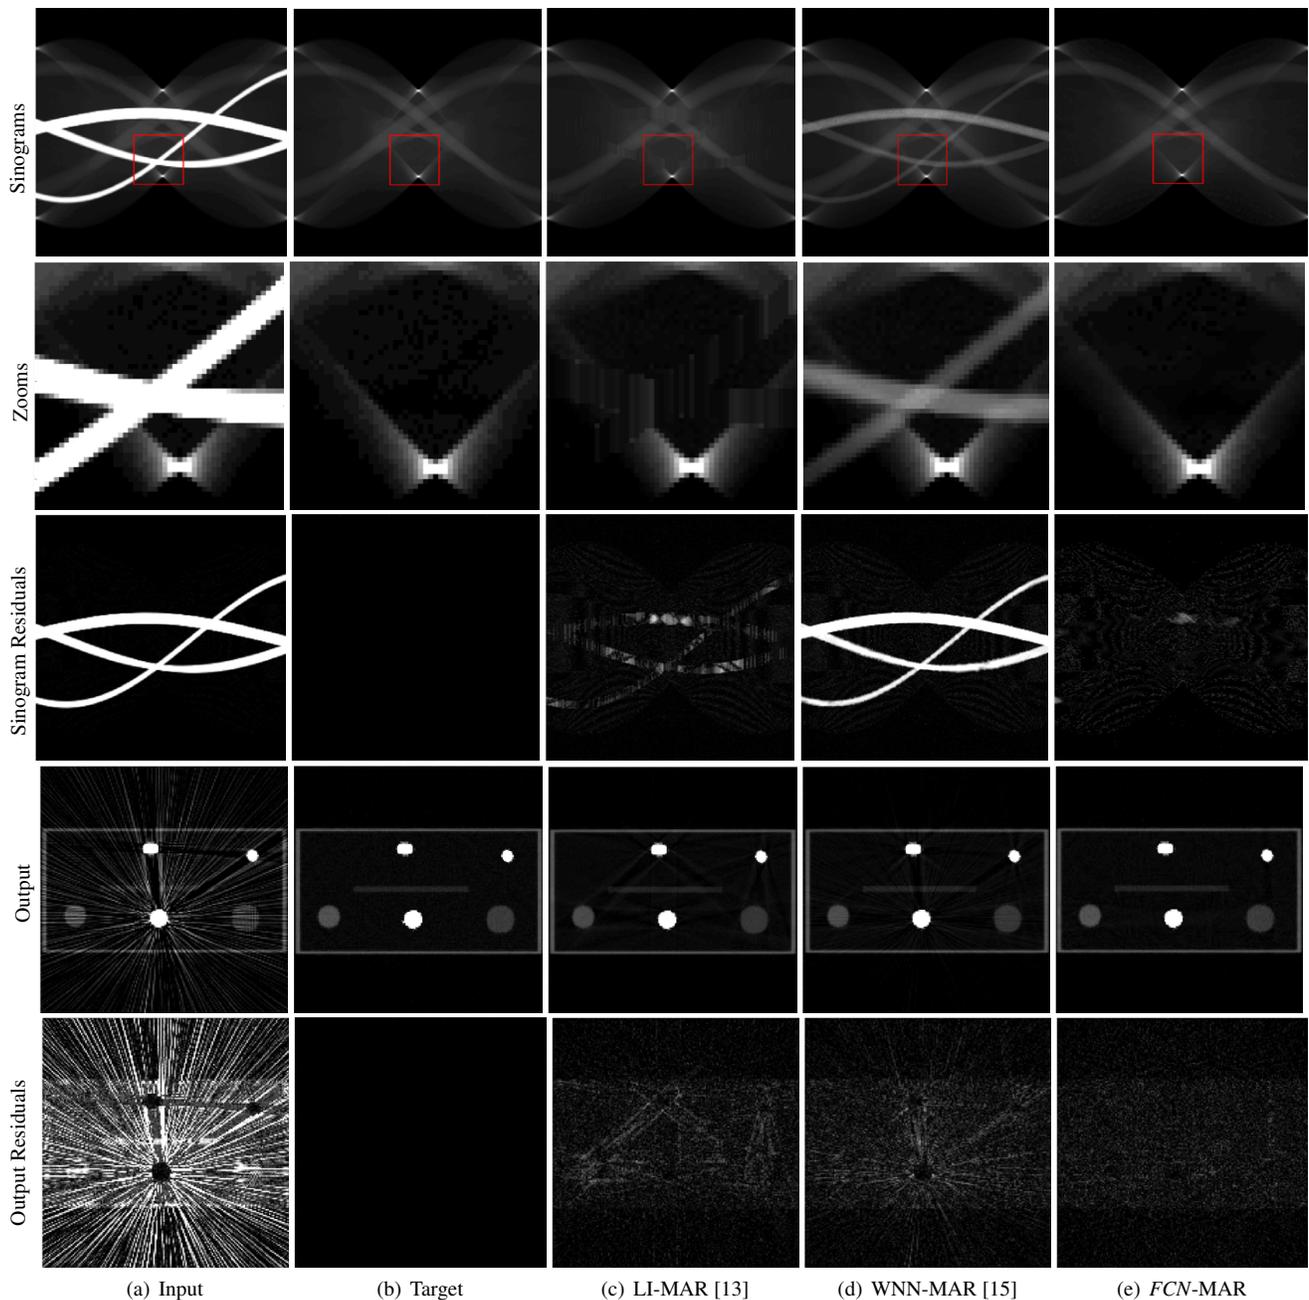


Figure 5. Input and outputs sinogram, resulting reconstructions, and residuals corresponding to different MAR methods are presented for another example scene. FCN-MAR evidently outperforms other methods by producing sinograms which are very similar to target sinograms and therefore suppress most of the streaks and preserves structure.

identify and correct the corrupted data.

Example 1 is shown in Figure 4 and contains three metal objects in a challenging configuration in which they happen to align. The first row presents the original sinogram, target sinogram, and the results of the various methods. Bright traces in the input sinogram depict the incorrect data which are not present in the target sinogram. LI-MAR uses a 1-D linear interpolation scheme to estimate new values, it uses information only from the data acquired at the same angle. This is why values estimated by LI-MAR are a

good approximation in areas with less content and when diameter of metal traces is small. WNN-MAR uses a 2-D weighted nearest neighbor scheme to estimate new values, however, it does not completely discard the original data and assigns a small weight to original data. This is why we still see dimmed traces corresponding to the corrupted regions. Our FCN-MAR approach relies on a CNN to correct the corrupted data, and resulting sinogram exhibit the effectiveness of our approach in accurately correcting the corrupted data.

We show zoomed patches of sinograms in second row of Figure 4 to accentuate the differences among outputs of different techniques. It is clear that LI-MAR results have a linear structure due to the inherent assumptions behind its interpolation scheme that data is only correlated along the detector channels. Additionally, it can be noticed that estimated values along the border of traces are more accurate as compared to central pixels in the corrupted regions. WNN-MAR uses 2-D information and assigns small weight to original data, this is why dimmed traces are still visible. Additionally, since WNN-MAR uses a rule-based approach to estimate new values, it will not be able to accurately estimate values in areas with high content. On the other hand, zoomed patches of FCN processed sinograms exhibit structure and intensity values which are very close to the target sinograms.

Further, we present the sinogram residuals of considered methods in third row of Figure 4. We see very bright residual traces for the input data. LI-MAR reduces the residuals to some extent, but they are still visible in most of regions. WNN-MAR dims the residuals but they are still present all over the sinogram. Residual traces are minimal in the FCN processed sinograms and there is no structure associated with residuals.

Output images generated by different methods are presented in fourth row of 4. Notice that output images are generated by using FBP to reconstruct images and inserting metals back again to the reconstructions using metal image X_m . The presented results show that if we do not use any MAR technique, the produced reconstructions suffer from severe streaking artifacts which limit the usefulness of CT. The objects nearby metals are broken into multiple pieces, and there are several streaks. LI-MAR output image shows that it suppresses existing streaks but it creates new streaks and results in the splitting of rubber sheet at a different location. Additionally, intensity shading is still visible. WNN-MAR does not create new streaks, but it fails to completely suppress existing streaks. Rubber sheets are still split into two pieces and intensity shading is also visible. FCN-MAR output shows that it suppresses most of the existing streaks and recover the structure hidden under the streaks. Additionally, rubber sheets are not split into pieces. So FCN-MAR not only suppresses most of the streaks and recover the structure hidden under streaks but also keep rubber sheets intact.

We present output image residuals in the fifth row of Figure 4. There are very bright streaks present in the residual image corresponding to original data. LI-MAR reduces some of the streaks but there is structure associated in the LI-MAR residual image. WNN-MAR residual image show both structure and residual streaks. FCN-MAR suppresses most of the streaks and there is not structure associated with the residual image which shows that FCN-MAR successfully suppresses most of the streaking and structural artifacts.

Sinogram and output image results for example 2 are shown in and Figure 5. First row presents input, target and output sinograms produced by different methods. Similar to example 1, LI-MAR and WNN-MAR fail to accurately correct the corrupted data in most regions of the sinogram and FCN-MAR produces sinogram values and structure which are very similar to the target sinogram. This difference becomes prominent in the zoomed patches shown in second row. LI-MAR has a linear structure, WNN-MAR dims the metal traces, and FCN accurately corrects the structure and intensity values. This claim is further supported

Comparison of considered methods on test data in terms of average MSE.

Without MAR	LI-MAR	WNN-MAR	FCN-MAR
7.22×10^{-5}	1.22×10^{-6}	1.43×10^{-6}	5.83×10^{-7}

by the sinogram residuals presented in third row, where LI-MAR and WNN-MAR still have high residuals, and FCN have very small residual values. Output results are generated by applying FBP to the resulting sinograms and inserting metals back using metal image X_m . We present a qualitative comparison of different methods using reconstructed images in fourth row which shows that original sinogram results in very bright streaks and structural artifacts including object splitting. LI-MAR suppresses original streaks but creates new streaks and structural artifacts including splitting of rubber sheet at two different points. WNN-MAR fails to completely suppress the streaking and structural artifacts. However, FCN-MAR suppresses most of the streaks and structural artifacts. The rubber sheet split into pieces in other outputs appears as a single object in the FCN-MAR output image. The superiority of FCN-MAR approach becomes prominent when we look at the output image residuals presented in the last row. While LI-MAR and WNN-MAR both still have residual structural and streaking artifacts, FCN-MAR residual image exhibit no structure and errors are reduced to minimal values.

Further, we use average mean-square-error (MSE) to perform a quantitative comparison of considered methods in Table 1. The quantitative analysis reveals that *FCN-MAR* reduces 99% of MSE as compared to standard FBP method. Furthermore, it also reduces 52% and 59% relative MSE as compared to LI-MAR and WNN-MAR methods respectively. It shows that our proposed *FCN-MAR* method suppresses most of the streaks present in the original reconstructed produced using FBP method. Both qualitative and quantitative analysis show that our proposed method greatly suppresses the metal-related artifact and outperforms two popular MAR techniques considered in this paper.

LI-MAR and WNN-MAR perform better in areas with less content and suffer in areas with high content. It is clear from both examples that output sinograms produced by LI-MAR and WNN-MAR fail to completely correct the sinograms. However, both structure and intensity values of sinograms produced by *FCN-MAR* are very similar to the target sinograms. Zoomed patches of sinograms show that our deep-learning-based method not only recover the sinogram structure but also produces intensity values which are very close to the target the sinograms. The sinogram residuals also show that sinogram error for our method is minimal. Further, both considered methods LI-MAR and WNN-MAR reduce streaks, however, residual streak level is still quite high and objects nearby metals are still broken. On the other hand, *FCN-MAR* produces excellent reconstructions by suppressing most of the streaks and recovering the objects nearby metals.

Conclusion

Metal-induced artifacts severely degrade the quality of CT reconstructions and limit their usability in subsequent steps of object segmentation and identification. Although many MAR techniques have been proposed in literature, however, their effective-

ness is limited. Even the most cited technique, LI-MAR [13], fails to suppress streaks completely and introduces new streaks in some cases. In this paper, we presented a deep learning based method to reduce metal artifacts. Our method lies in the sinogram correction category of MAR techniques, we train a CNN on a large CT dataset to perform sinogram correction on complete sinogram data. We show that our method effectively corrects the sinograms. i.e., their structure and intensity values are very similar to target sinograms. Our analysis shows that FCN-MAR produces excellent reconstructions which are far superior as compared to LI-MAR and WNN-MAR. It also shows that FCN-MAR reduces 99% of average MSE as compared to FBP, the technique which is used in most of the CT scanners today. The learning-based nature of our algorithm makes it powerful that rather than using a rule-based approach it learns the sinogram correction function from a large dataset. A future extension of this work could be a similar network that can be trained in an unsupervised setting.

Acknowledgments

This material is based upon work supported by the U.S. Department of Homeland Security, Science and Technology Directorate, Office of University Programs, under Grant Award 2013-ST-061-ED0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

References

- [1] F Edward Boas and Dominik Fleischmann. Ct artifacts: causes and reduction techniques. *Imaging Med*, 4(2):229–240, 2012.
- [2] Lars Gjestebj, Bruno De Man, Yannan Jin, Harald Paganetti, Joost Verburg, Drosoula Giantsoudi, and Ge Wang. Metal artifact reduction in ct: where are we after four decades? *IEEE Access*, 4:5826–5849, 2016.
- [3] Idris A Elbakri and Jeffrey A Fessler. Statistical image reconstruction for polyenergetic x-ray computed tomography. *IEEE transactions on medical imaging*, 21(2):89–99, 2002.
- [4] Ahmet Tuysuzoglu, Yuehaw Khoo, and W Clem Karl. Variable splitting techniques for discrete tomography. In *Image Processing (ICIP), 2016 IEEE International Conference on*, pages 1764–1768. IEEE, 2016.
- [5] Limor Martin, Ahmet Tuysuzoglu, W Clem Karl, and Prakash Ishwar. Learning-based object identification and segmentation using dual-energy ct images for security. *IEEE Transactions on Image Processing*, 11(24):4069–4081, 2015.
- [6] F Edward Boas and Dominik Fleischmann. Evaluation of two iterative techniques for reducing metal artifacts in computed tomography. *Radiology*, 259(3):894–902, 2011.
- [7] Bruno De Man, Johan Nuyts, Patrick Dupont, Guy Marchal, and Paul Suetens. An iterative maximum-likelihood polychromatic algorithm for ct. *IEEE transactions on medical imaging*, 20(10):999–1008, 2001.
- [8] Benoit Hamelin, Yves Goussard, David Gendron, Jean-Pierre Dussault, Guy Cloutier, Gilles Beaudoin, and Gilles Soulez. Iterative ct reconstruction of real data with metal artifact reduction. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, pages 1453–1456. IEEE, 2008.
- [9] Xiaomeng Zhang, Jing Wang, and Lei Xing. Metal artifact reduction in x-ray computed tomography (ct) by constrained optimization. *Medical physics*, 38(2):701–711, 2011.
- [10] Hamid Soltanian-Zadeh, Joe P Windham, and Jalel Soltanianzadeh. Ct artifact correction: an image-processing approach. In *Medical Imaging 1996*, pages 477–485. International Society for Optics and Photonics, 1996.
- [11] Matthieu Bal, Hasan Celik, Krishna Subramanyan, Kai Eck, and Lothar Spies. A radial adaptive filter for metal artifact reduction. In *Medical Imaging 2005: Image Processing*, volume 5747, pages 2075–2083. International Society for Optics and Photonics, 2005.
- [12] Lars Gjestebj, Qingsong Yang, Yan Xi, Hongming Shan, Bernhard EH Claus, Yannan Jin, Bruno De Man, and Ge Wang. Deep learning methods for ct image-domain metal artifact reduction. In *Developments in X-Ray Tomography XI*, volume 10391, page 103910W. International Society for Optics and Photonics, 2017.
- [13] Willi A Kalender, Robert Hebel, and Johannes Ebersberger. Reduction of ct artifacts caused by metallic implants. *Radiology*, 164(2):576–577, 1987.
- [14] Richard L Morin and DE Raeside. A pattern recognition method for the removal of streaking artifact in computed tomography. *Radiology*, 141(1):229–233, 1981.
- [15] Andreas H Mahnken, Rainer Raupach, Joachim E Wildberger, Bettina Jung, Nicole Heussen, Thomas G Flohr, Rolf W Günther, and Stefan Schaller. A new algorithm for metal artifact reduction in computed tomography: in vitro and in vivo evaluation after total hip replacement. *Investigative radiology*, 38(12):769–775, 2003.
- [16] Shiyong Zhao, DD Robeltson, Ge Wang, Bruce Whiting, and Kyongtae T Bae. X-ray ct metal artifact reduction using wavelets: an application for imaging total hip prostheses. *IEEE transactions on medical imaging*, 19(12):1238–1247, 2000.
- [17] Esther Meyer, Rainer Raupach, Michael Lell, Bernhard Schmidt, and Marc Kachelrieß. Normalized metal artifact reduction (nmar) in computed tomography. *Medical physics*, 37(10):5482–5493, 2010.
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [19] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [20] Benjamin Graham. Fractional max-pooling. *arXiv preprint arXiv:1412.6071*, 2014.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [22] Chen-Yu Lee, Patrick W Gallagher, and Zhuowen Tu. Generalizing pooling functions in convolutional neural networks: Mixed, gated, and tree. In *Artificial Intelligence and Statistics*, pages 464–472, 2016.
- [23] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE Transactions on Image Processing*, 2017.
- [24] Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017.
- [25] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Pro-*

ceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1646–1654, 2016.

- [26] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- [27] C Crawford, H Martz, and WC Karl. Research and development of reconstruction advances in ct-based object detection systems—final report, dept. homeland security center excellence, alert, boston. Technical report, MA, Tech. Rep. HSHQDC-12-J-00056, 2013.
- [28] Patrick J La Riviere. Penalized-likelihood sinogram smoothing for low-dose ct. *Medical physics*, 32(6):1676–1683, 2005.
- [29] Martin J Berger and JH Hubbell. Xcom: Photon cross sections on a personal computer. Technical report, National Bureau of Standards, Washington, DC (USA). Center for Radiation Research, 1987.
- [30] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [31] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale ma-

chine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.

Author Biography

Muhammad Usman Ghani received the B.S. degree in Electrical Engineering from COMSATS Institute of Information Technology, Lahore, Pakistan, in 2013, and the M.S. degree in Computer Science from Sabanci University, Istanbul, Turkey, in 2016. He is currently a Ph.D. candidate at the Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA. His research interests include machine learning, inverse problems, computer vision, and biomedical image processing.

W. Clem Karl received the Ph.D. degree in Electrical Engineering and Computer Science from MIT in 1991. He is currently the Chair of the Electrical and Computer Engineering Department at Boston University. His research interests are in the areas of statistical signal and image processing, estimation, detection, and medical signal and image processing. He has served as the Editor-in-Chief of the IEEE Transactions on Image Processing and the IEEE Transactions on Computational Imaging.