IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received May 3, 2019, accepted May 30, 2019, date of publication June 5, 2019, date of current version June 26, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2921013

MCDNet – A Denoising Convolutional Neural Network to Accelerate Monte Carlo Radiation Transport Simulations: A Proof of Principle With Patient Dose From X-Ray CT Imaging

ZHAO PENG[1,3], HONGMING SHAN[1], TIANYU LIU[2], XI PEI[3], GE WANG[1], X. GEORGE XU[1,2,3]

[1]Department of Biomedical Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180, USA
[2]Department of Mechanical, Aerospace and Nuclear Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180, USA
[3]Department of Engineering and Applied Physics, University of Science and Technology of China, Hefei, Anhui 230026, China,

Corresponding author: X. George Xu (xug2@rpi.edu)

This work was supported in part by the National Institutes of Health/National Institute of Biomedical Imaging and Bioengineering (NIH/NIBIB) under Grant R42EB019265-01A1, in part by the NIH/NIBIB under Grant U01EB017140, in part by the National Natural Science Foundation of China under Grant 11575180, and in part by an International Training Grant from the American Association of Physicists in Medicine (AAPM).

ABSTRACT X-ray computed tomography (CT) is a popular diagnostic imaging tool that has caused public concern over potential radiation risks to the patient. Monte Carlo (MC) simulations are the most accurate methods to calculate x-ray interactions with the patient's body and voxel-wise dose distributions, but the statistical methods suffer from extremely long computing time that is required to achieve necessary statistical precision. In this paper, we propose and demonstrate the Monte Carlo Denoising Net (MCDNet), a convolutional encoder-decoder neural network, for the purpose of accelerating the MC radiation transport simulations for patient CT dosimetry. A unique set of full-body anatomically realistic adult voxel phantoms of various sizes and a GPU-based parallel MC code were used to produce adequate training and testing data for supervised learning. Gamma index passing rate (GIPR) was used to evaluate the performance of predicted dose maps. For CT scan protocols considered in this study, MCDNet is found to have the ability of predicting dose maps of 9.9×10^7 photons from corresponding dose maps of 1.3×10^6 photons, yielding 76× speed-up in terms of photon numbers used in the MC simulations. MCDNet is the first CNN-based method to speed-up MC radiation transport simulations involving 3D and heterogeneous patient anatomies for x-ray CT. Future studies will test the feasibility of applying the deep-learning based denoising strategy to other MC radiation transport applications.

INDEX TERMS CT dose, Monte Carlo simulation, dose map, machine learning, deep neural network, Gamma Index Passing Rate.

I. INTRODUCTION

Worldwide, the number of diagnostic x-ray CT examinations has increased more than 2 dozen folds in the past 3 decades, owing to rapidly improving multi-detector CT (MDCT) technologies [1], [2]. Today, the CT procedure is replacing conventional radiography as the initial diagnostic exam in many emergency rooms and CT scanners are integrated with modalities in nuclear medicine and in image-guided radiotherapy where a patient can receive multiple CT scans [3]. The International Commission on Radiological Protection (ICRP) has reported that CT doses can approach or exceed radiation levels that have been shown to result in an increase in cancer incidence [1]. The radiation risk is even greater to pregnant and pediatric patients [2]. The American College of Radiology (ACR) urgently calls for more effective methods to evaluate and manage such imaging doses, recognizing that existing dose computational tools are insufficient for patient-specific dose quantification, scanner optimization, and protocol comparison [4].

CT scanner design has been evolving continuously since the early 1970s and, currently, helical motion and multi-slice

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

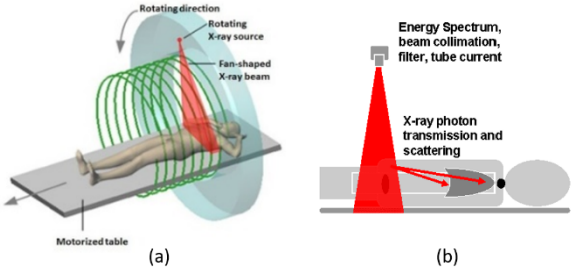

FIGURE 1. X-ray CT scans cause patients to be exposed to ionizing radiation. (a) In helical CT, the table moves as the x-ray source and detectors rotate. (b) X-rays scattering irradiates health tissues outside the scan range (red arrows).

capabilities are available. Fig. 1 (a) depicts a patient being scanned by a CT system and Fig. 1 (b) shows that a portion of the x-ray photons undergoes radiation interactions, depositing radiation energies in the organs inside and outside the beam. In the current CT organ dose paradigm, a ''whole-body'' phantom is needed for calculating organ doses from CT. Such calculations would consider typical x-ray source parameters such as x-ray energy spectrum, beam collimation, filter, x-ray tube current and irradiation time.

For CT, x-ray photons are less than 160 keV and, in such energies, they interact with the tissues primarily via the ''photoelectric effects'' and ''Compton scattering'' [5]. The probability of an interaction occurring within an organ or tissue depends on energy, tissue electron density and composition. Although radiation transport problems can be solved by several methods, only ''Monte Carlo (MC) radiation transport calculation'', a simulation method originally developed and refined for nuclear weapons research in the 1940s at Los Alamos [6]–[8], is able to account for all aspects of radiological physics within 3D heterogeneous media such as the human body. The inherent statistical uncertainty can be controlled to less than 1% which is often more precise than an experimental result. Today, MC methods are integral to nuclear engineering and radiological medical physics owing to powerful and affordable computers. The MC radiation transport community has made available a number of well-tested, large-scale MC code packages [8]–[13]. However, the MC calculation for CT doses using a production code can take several hours to achieve acceptable statistical uncertainty, preventing the MC methods from wider clinic applications [14]. At present, the traditional wisdom of accelerating MC simulation is to achieve a parallel calculation by the GPU. However, to achieve acceptable accuracy, the number of simulative photons cannot be reduced. Here, we proposed an out of box method that can reduce the simulative photons while maintaining the comparable accuracy. Therefore, whatever MC codes are based on, either CPU or GPU, the proposed method has the advantage of reducing the simulative photons so as to achieve acceleration, which means that we can achieve further acceleration on the advanced GPU-based parallel MC method. The potential limitation is that the proposed network may not produce an acceptable dose map if the input dose map is very noisy caused by insufficient simulative photons. In this study, we demonstrated that our method can reduce the number of photons by about 76 times while maintaining acceptable accuracy.

TABLE 1. Parameters of different adult phantoms.

Phantom	Sex	Weight (kg)	Dimensions (voxels)
1	Male	73	114 × 93 × 509
2	Male	86	114 × 99 × 509
3	Female	63	112 × 90 × 469
4	Female	74	115 × 99 × 469
5	Female	89	116 × 107 × 469

This study was inspired by ideas proposed by Wang regarding the use of phantoms data in deep learning for medical imaging [15]. Here we propose a method to accelerate the MC calculations using a convolutional neural network (CNN) that is trained to predict dose distribution characteristics learned from MC simulation data. Fig.2 shows the overall flow chart of the proposed method. First, we generate a database of dose maps for low- to high-photon fluences using MC simulations of CT scans. Then we train the CNN to map the low-photon dose maps to high-photon dose maps. At the core of this method is a new Monte Carlo ''denoising'' CNN algorithm – called Monte Carlo Denoising Net (MCDNet) – that is based on our previous study [16]. Finally, we evaluate the quality of predicted dose distributions by comparing voxel-wise dose levels with the ground truth. A quick literature survey shows that MCDNet is the first to demonstrate a CNN-based method for the purposes of accelerating MC radiation transport dose simulations for 3D and heterogeneous patient anatomies, thus opening the door for many potential applications.

II. MATERIAL AND METHOD

MC dose calculations for patients receiving CT imaging were performed on full-body computational phantoms [17]. Previously, we have developed a library of 5 male and 5 female adult phantoms with different body weights ranging from 5th to 95th percentile [18], as shown in Fig 3. These phantoms have detailed anatomical information, each having more than 100 well-defined organs and tissues for the purposes of MC simulations. The geometries were originally modeled using triangular surfaces, known as boundary representation, and were then converted into voxels to allow for fast Monte Carlo simulations [18]. These phantoms provide the basis for the generation of the ground truth for CNN learning. In this study, 2 RPI adult male and 3 female phantoms were used, each having a uniform voxel size of 0.35 cm × 0.35 cm × 0.35 cm, as summarized in Table 1. Four phantoms were used for training, whereas the others for testing.

A model representing a GE Lightspeed Pro 16 multi-detector CT was modeled and validated in our previous studies [19], [20]. The simulated scanning protocol is 100 kVp, 20 mm beam collimation, axial body scan, involving a total of 120 rotations to cover from the thorax to abdomen of

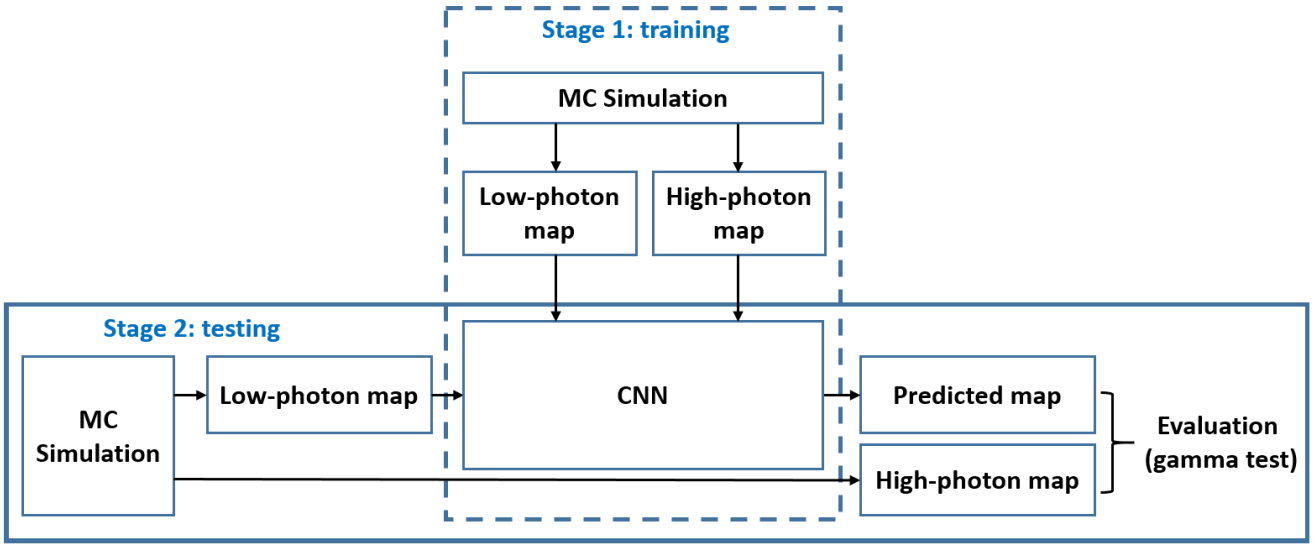

FIGURE 2. Overall flow chart of the MCDNet training and testing. "Low-photon map" means a dose map from MC simulations involving a low photon number, and "High-photon map" means a dose map from MC simulations involving a high photon number. The dashed box represents the training process for MCDNet involving a database of dose maps from various MC simulations derived in human phantoms. The solid box represents the testing process for the MCDNet-predicted true dose maps by comparing voxel-wise dose accuracy.

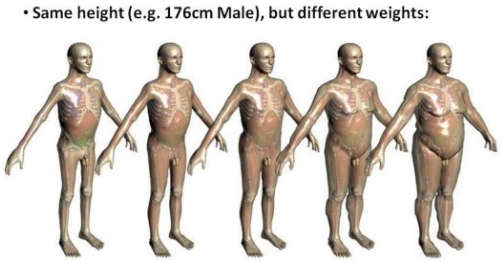

FIGURE 3. Deformable phantoms representing adult males at 5th, 25th, 50th, 75th, 95th weight-percentiles of the population [18].

the patient phantom. CT scanner's continuous rotational motion is simulated using the step-and-shoot pattern, with each rotation approximated by 16 discrete positions [20]. A GPU-accelerated MC code, ARCHER, previously developed by us was used in this study [21]. ARCHER is an accurate and fast MC simulation tool that supports both Nvidia and AMD GPU accelerators. In this study, the Nvidia Titan X card was used.

For each phantom, a ground truth dose distribution map was first obtained using 1.6×10^9 photons per CT scanner rotation. A sequence of MC simulations using progressively fewer photons (from 1.6×10^8 to 1.6×10^4) were then performed, taking less MC simulation time but resulting in an increasingly greater statistical uncertainty. Fig. 4 shows the dose maps of varying statistical noise for different photon numbers.

To accelerate the Monto Carlo radiation transport simulations, the proposed MCDNet is essentially a deep learning method to predict the denoising mapping from low-photon-fluence dose maps to high-photon-fluence dose maps [16]. Assuming that M_{lp} and M_{hp} denote the dose distribution maps from low photons and high photons, respectively, the goal of the denoising process is then to seek a function F that can predict a high-photon dose distribution map for a given low-photon dose distribution map:

$$F : M_{lp} \rightarrow M_{hp}.$$

Fig. 5 illustrates the structure of the proposed MCDNet, for learning the mapping function F. MCDNet is designed as an encoder-decoder network, with the encoder having 5 convolutional layers and the decoder having 5 deconvolutional layers. The number of (de)convolutional layers is empirically chosen. Four dashed arrow lines in the figure indicate four conveying paths [16], [22] that copy and reuse early feature-maps as the input to later layers having the same feature-map size using a concatenation operation to preserve high-resolution features. The solid line in the figure indicates a residual skip connection [23] that sums up the input and output of the MCDNet to reduce the searching space of the network output. After the summation operation, there is a ReLU activation layer, then the network outputs the final results. Each layer has 32 filters, each of size 3×3, except for the final layer that has only 1 filter. One convolutional layer of 32 filters of kernel size 1×1 is used after each concatenation operation to reduce the total number of feature-maps from 64 to 32. Each (de)convolutional layer is followed by a ReLU activation function [24].

The proposed MCDNet is a modified version of the Conveying-Path Convolutional Encoder-decoder (CPCE) used in our previous study for CT image denoising [16]. Compared to CPCE, the proposed MCDNet has two more layers and a residual skip connection from the input to the output for the network. The skip connection enables the network to infer the noise from the input image directly. It should be noted

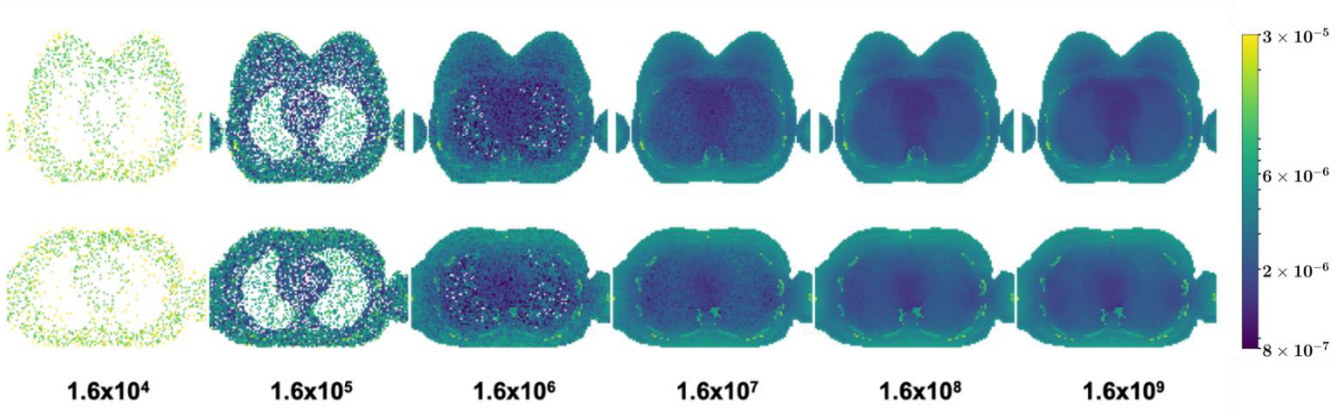

FIGURE 4. The MC simulated dose maps showing different statistical noise levels. The first row is for dose maps of the female phantom and the second row is for a male phantom. Different columns are for the MC simulated dose maps for different photon numbers that are listed at the bottom.

that the output space of noise is much smaller than the truth dose space. Therefore, the skip connection can reduce the searching space of the network output, making the network converge faster. In addition to this, the parameters of the network are usually initialized by a Gaussian distribution with zero mean, which indicates that the initialized network without training tends to output an image that all values close to zero. In this case, the skip connection enables the network to output the input itself, which is similar to ground-truth in structures. Therefore, at the starting point, the network with skip connection can produce better output than that without skip connection, which also can boost the training speed. Furthermore, MCDNet is similar to the well-known U-Net for biomedical imaging segmentation [22] but without the down-sampling operation in U-Net that can lead to the loss of details. Such theoretical analysis ensures that the MCDNet is able to achieve the reasonable performance.

III. EXPERIMENT

A. DATA PROCESSING

All dose maps from five phantoms are normalized into the range of [0, 1]. In order to get sufficient training data, we trained our network on patches with a data augmentation technique. More specifically, we extracted the patches of size 64 × 64 from each dose map with a moving stride being 16. Then, we randomly split the patches into the training set and validation set with a 4:1 ratio. Finally, each patch in the training set is augmented by rotating 90, 180 and 270 degrees, respectively, and flipping the patches left and right as well as up and down. Table 2 shows the data sizes in the training and validation set.

B. LOSS FUNCTION

The parameters in MCDNet are optimized by minimizing the Mean Squared Error between the outputs of the network and the reference high-photon dose map patches according to,

$$\sum_{i=1}^{N} \left\| F\left(M_{lp}^{i}\right) - M_{hp}^{i} \right\|_2^2$$

TABLE 2. The number of patches (64 × 64) of different adult phantoms in the training and validation set.

Phantom	Training set	Validation set
1	4608	192
2	6912	288
3	4608	192
4	6912	288
5	6912	288

where N is the total number of training samples. This minimization problem can be solved by various algorithms; in this work, we adopt the Adam algorithm to update the parameters [25]. The gradients of the parameters are computed using a back-propagation algorithm [26].

C. TRAINING AND TESTING

At the training stage, the inputs of the network are the low-photon dose patches of size 64 × 64, the ground truth data are the high-photon dose patches with the same size. The initial learning rate is 0.001, and the initial number of training epoch is 1000. The validation loss is calculated for every epoch and the training process is terminated when the validation loss is no longer decreased after 20 consecutive epochs. In the testing stage, the inputs of the network are the low-photon dose maps with original size, the outputs of the network are the predicted high-photon dose maps with the same size.

For the CT imaging cases considered in this experiment, the highest number of photons is 1.6×10^9, the number of photons required to yield an MC dose map with acceptable statistic uncertainty. The lower photon numbers include 1.6×10^4, 1.6×10^5, 1.6×10^6 and 1.6×10^7. The dose maps of 1.6×10^8 photons are found to be very close to that of the 1.6×10^9 photons, so they were not used as training data. We trained a CNN model for each of the low-number dose maps to produce a total of 4 dose prediction models. For each

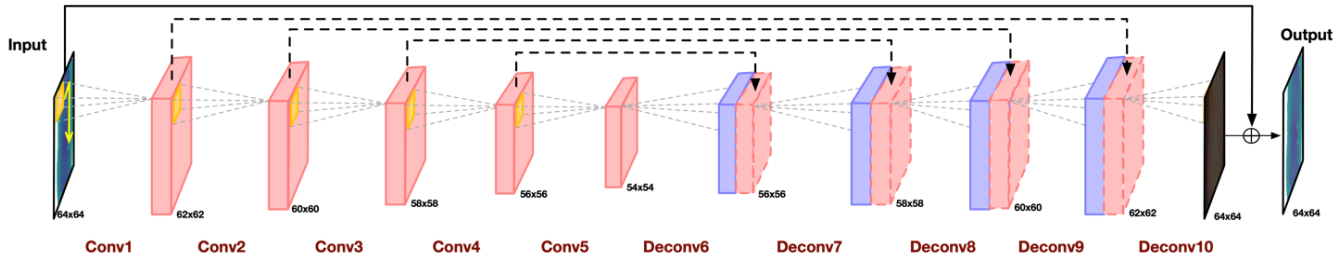

FIGURE 5. The proposed MCDNet structure consisting of 5 convolutional and 5 deconvolutional layers. Four dashed arrow lines indicate four conveying paths that copy and reuse early feature-maps as the input to later layers with the same feature-map size via concatenation operation to preserve high-resolution features. The solid line indicates a residual skip connection summing up the input and output of the MCDNet. Each layer has 32 filters except for the final layer that has only 1 filter. One convolutional layer of 32 1 × 1 filters is used after each concatenation operation to reduce the number of feature-maps from 64 to 32. Each (de)convolutional layer is followed by a rectified linear unit (ReLU) activation function [24]. The numbers next to each feature-map represent its spatial size, given the training patch size of 64 × 64.

model, we performed 5-fold cross-validation to obtain testing results.

All experiments were performed on a Linux operation system. Keras was used to design and train our neural network, the backend is TensorFlow [27]. The hardware includes (1) GPU: Nvidia GeForce Titan X Graphics Card with 12GB memories, and (2) CPU: Intel Xeon Processor X5650 with 16GB memories.

D. EVALUATION STANDARD

In this experiment, we performed a Gamma Test which is a gold standard in comparing dose maps by the radiation therapy community [28]. The distance-to-agreement is 3 mm, and the dose-difference is 3%. The Gamma Index Passing Rate (GIPR) is calculated for a predicted dose map which is rated by how close the value is to unity.

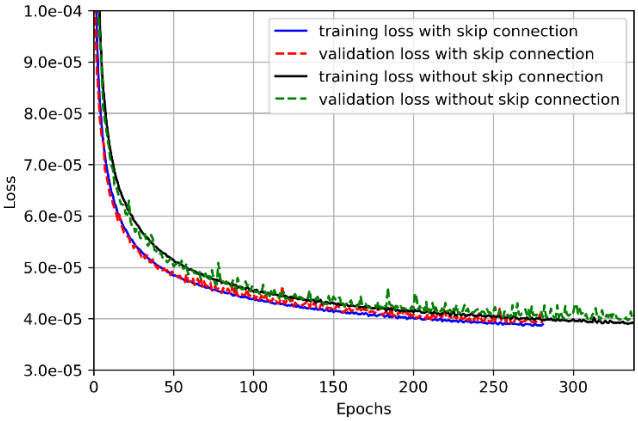

FIGURE 6. The training loss and validation loss in the model with and without skip connection trained by the input dose maps of 1.6 × 106 photons.

IV. RESULTS

For all 4 models based on MCDNet, the one trained with dose maps of 1.6×10^4 photons failed to produce desirable results, suggesting a lower bound for the proposed deep learning method. The other 3 successful models have a similar training process, and we only show one of them here. Fig. 6 presents the training process of the model for the dose map of 1.6×10^6 photons. In the model with skip connection, the training loss and validation loss decrease as the epoch number increase until the 282nd epoch when the training process is stopped. The training, validation and testing procedure we used is typical in the deep learning community, which effectively avoids either overfitting or underfitting. To better validate our analysis about skip connection, we trained a model without the skip connection based on the input dose maps of 1.6×10^6 photons. It can be seen from the Gamma Index Passing Rate that they can achieve comparable performance with skip connection (0.9962 ± 0.0033) and without skip connection (0.9962 ± 0.0034), but the model with the skip connection converges faster and requires fewer epochs. Fig. 6 shows the difference between the two models in the training process.

Table 3 and Fig. 7 present the statistical data in the 5-fold cross-validation test. In the MC simulations, the mean computational time, and the mean and the standard deviation of GIPR are calculated for dose maps. In the MCDNet, the mean and the standard deviation of GIPR in the 5-fold cross-validation are calculated, but the computational time is not showed because it is less than 3s in each model, which is negligibly small. The dose maps of 1.6×10^8 photons in MC simulations that were not used as training data are also listed in the table to compare the performance of our network. It can be seen that GIPRs obtained from the proposed MCDNet are improved to values that extremely close to unity fairly quickly, even for photon numbers as low as 1.6×10^5. When the number of photons is 1.6×10^4, there are too many blank voxels in the original dose distribution maps, as shown in the first column of Fig. 4 previously. At that level of photon number, MCDNet is not convergent due to insufficient training information, and finally the output of the network is blank in the testing, this is related to that the last layer of MCDNet is a ReLU activation function. What needs to be explained is that the GIPRs are still far more than zero because there is the blank background in the all dose distribution maps. For photon number of 1.6×10^5, the MCDNet improves the GIPR from 0.5676 to 0.9678, approaching to the GIPR value of 0.9767 from 1.6×10^7 photons in MC simulations. For the photon number of 1.6×10^6, the MCDNet improves the GIPR from 0.7862 to 0.9962. And similarly, for 1.6×10^7 photons,

TABLE 3. The computational time and GIPR for MC simulations and MCDNet predictions.

Number of photons	MC simulations			MCDNet	
	Time (min)	GIPR		GIPR	
		Mean	Std	Mean	Std
1.6×10^4	0.185	0.4448	0.0676	0.4447	0.0659
1.6×10^5	0.235	0.5676	0.0658	0.9678	0.0285
1.6×10^6	0.275	0.7862	0.0529	0.9962	0.0033
1.6×10^7	0.465	0.9767	0.0173	0.9996	0.0005
1.6×10^8	2.515	0.9999	0.0003		
1.6×10^9	21.92				

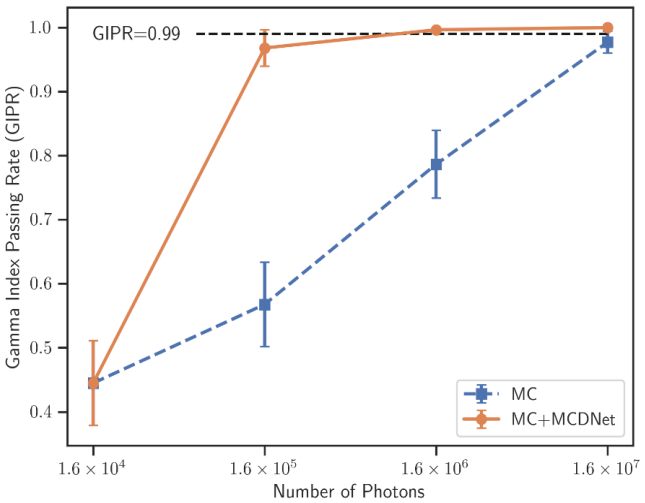

FIGURE 7. The mean GIPR of each model in the 5-fold cross-validation test which includes 600 dose maps in total. The squares in the dashed line indicate the mean GIPR of original dose maps from MC simulations. The circles in the solid line indicate the mean GIPR of predicted dose maps after using the MCDNet.

the MCDNet improves the GIPR from 0.9767 to 0.9996. Fig. 8 offers further visual evidence of these results. As can be seen in the first column of Fig. 8, there is still too much noise in MC simulations of 1.6×10^5 photon number, but the corresponding MCDNet results exhibit very little noise. Second and Third columns of Fig. 8 display nearly perfect dose maps predicted by MCDNet.

As shown in the row of MC simulations in Table 3, the GPU-based parallel MC computing time using ARCHER does not increase linearly as the number of photons increases, due to the heterogeneous CPU-GPU nature of the algorithm in ARCHER [21]. From these data, to obtain a low-noise dose map with a GIPR that is greater than 0.99, it is estimated that about 9.9×10^7 photons are needed for the original MC simulations using ARCHER which took 1.640 minutes. In contrast, the MCDNet can predict a dose map with GIPR of 0.99 using only 1.3×10^6 photons in the MC simulations using ARCHER which only took 0.266 minutes, yielding 76× speed-up in terms of photon numbers used in the MC simulations. Please note that ARCHER was run under the CPU-GPU architecture and, as such, the speedup in terms of wall clock time is not linear – in fact the speedup of the wall clock time for ARCHER code is found to be about 6 showing the ARCHER code is extremely fast and efficient in utilizing the parallel GPU device (i.e., if a CPU-based traditional code such as MCNP or GEANT4 were used, the speedup of wall clock time would be about 76×).

Fig. 9 shows the output of the network before the summation operation (learned noise). As a reference, we also provide the input image, the ground truth, the prediction image, and the real noise image (ground truth – input image). it shows that the network learns the noise from the input image, and produces the final result after summing. It is noted that the background of the learned noise image is negative, but passing through a ReLU activation layer after the summation operation, the background of the prediction image is the same with the ground truth as shown.

In order to better understand how the proposed method works on the feature learning, we randomly select a testing dose map from 1.6×10^6 photons in MC simulations and fed it into MCDnet, then we visualized the feature maps learned by the proposed network on the first, the fifth, and the ninth convolutional layers as shown in Fig. 10. The display range is the same for the feature maps in the same layer. In the red box of each subplot, we provide the input image, the ground-truth image, the real noise image (ground-truth – input image), and the learned noise image (the output of network before summation operation) as the references to facilitate the understanding of the feature maps shown in the blue box. Note that the ninth convolutional layer is before the last deconvolutional layer. All the feature maps refer to as the output of ReLU activation function. In addition, we briefly gave the explanation about the feature learning as follows. In Fig. 10 (a), the feature maps at an early layer tend to be roughly similar to the input image in terms of the structures, which indicates that the first convolutional layer tries to extract the noise but also preserve the structures. In contrast, the feature maps at a late layer shown in Fig. 10 (c) tend to be more similar to the final noise. The value scale of different feature-maps is discrepant, and each feature map looks fairly uniform, which indicates that the ninth convolutional layer learned the different scales of noise. As a trade-off, the feature maps at the middle layer shown in Fig. 10 (b) are mostly local with much coarse texture, and the original image features are not obvious. It means that the fifth convolutional layer probably infers more noise features from the input image. In conclusion, the network is trying to infer the noise from the input image layer by layer.

V. DISCUSSION

In this study, we considered the dose map simulated from 1×10^9 photons to be the ground truth, i.e., noise free. We produced dose maps of varying degree of noise from 1×10^4 photons to 1×10^9 photons, each involving a full-body anatomically 3D voxel phantom irradiated by x-rays from

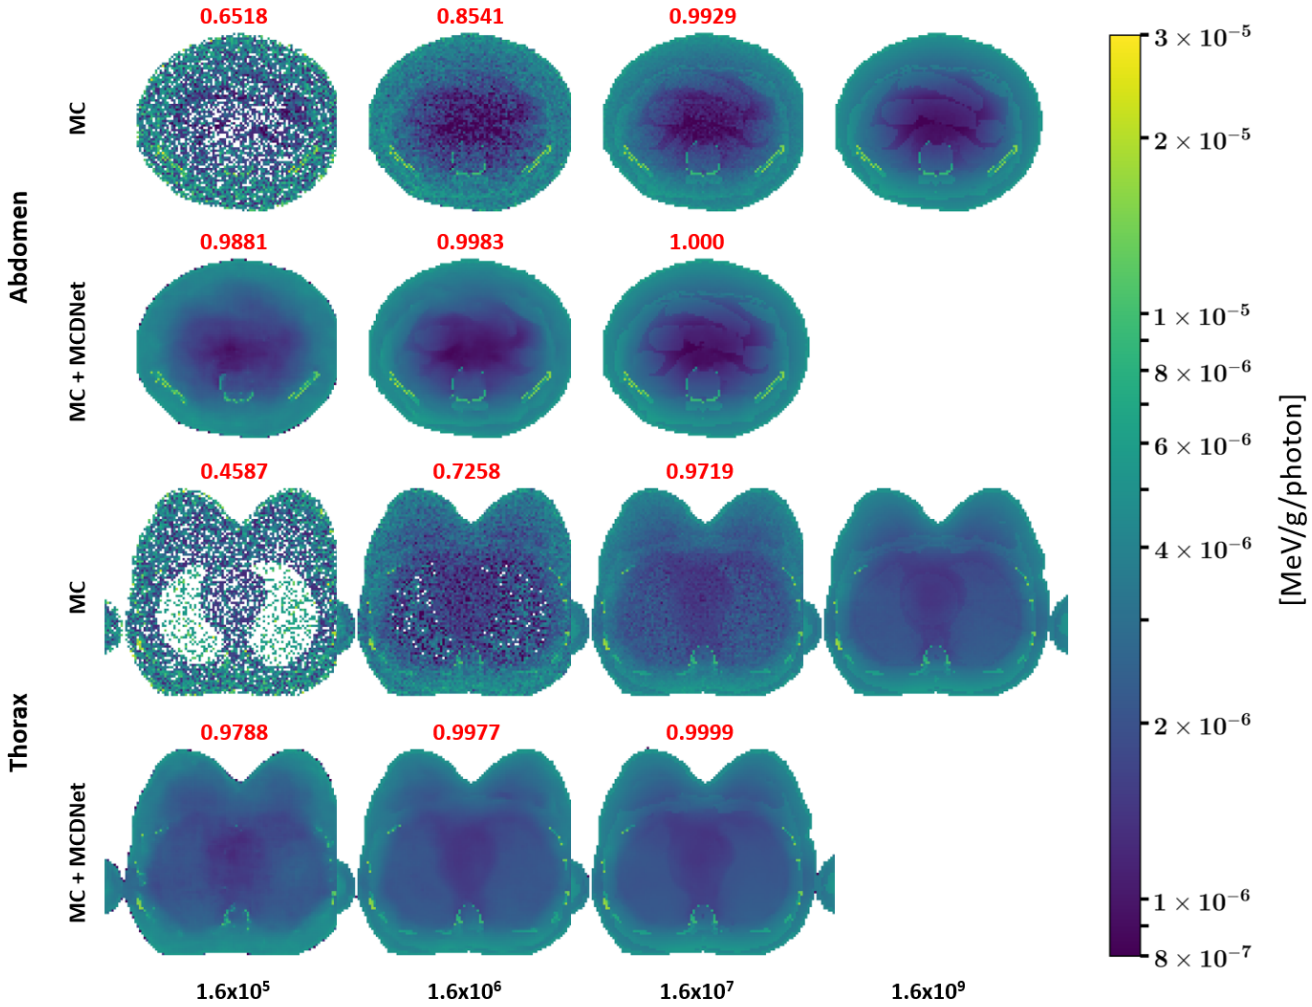

FIGURE 8. The original MC simulated dose maps and MCDNet predicted dose maps using the RPI female phantom as an example. The two top rows compare dose maps for the abdominal CT scans obtained from MC and MCDNet, and the bottom two rows compare those for the thorax CT scans. Each column represents dose maps in different number of photons (denoted by the blacknumbers at the bottom). The pixels are normalized for all dose maps with yellow pixels 3×10^{-5} MeV / ghoton dark blue referring to 8×10^{-7} MeV / g / photon particularly, the white pixels referring to zero.The red number at the top of each dose map represents the corresponding GIPR for 1.6×10^{9} photons(i.e., the ground truth).

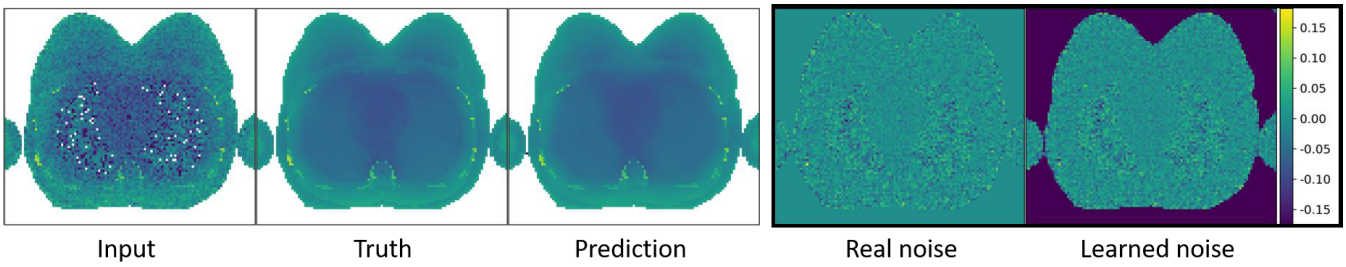

FIGURE 9. The output of the network before the summation operation (learned noise). As a reference, the input image, the ground truth, the prediction image, and the real noise image (ground truth -input image) are also provided. It is noted that the background of the learned noise image is negative, but passing through a ReLU activation layer after the summation operation, the backgroundof the prediction image is the same with the ground truth as shown.

the CT scanner. To generate such as large amounts of necessary data for training and testing, a home-grown GPU-based parallel Monte Carlo code, ARCHER, was used in this study. ARCHER is both high accurate and efficient; for Monte Carlo CT imaging dose calculations, results from ARCHER are in good agreement with those from MCNP, but the GPU-accelerated ARCHER code is over 200 times faster [29]. In the other words, if a CPU-based MCNP code were used in this study instead, we would have to spend about 2,000 days using a CPU cluster equivalent of a Tian

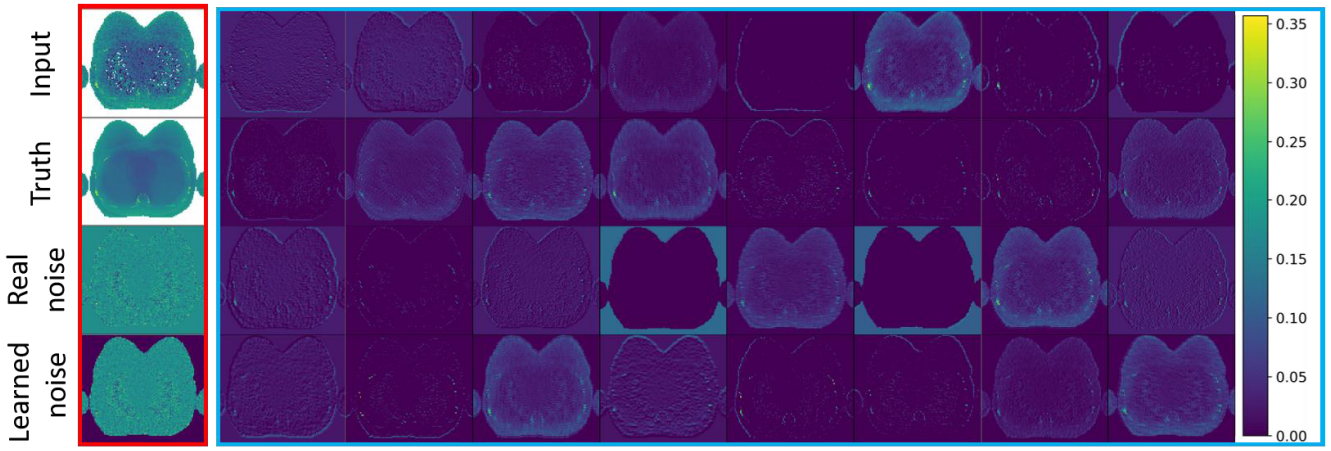

(a)

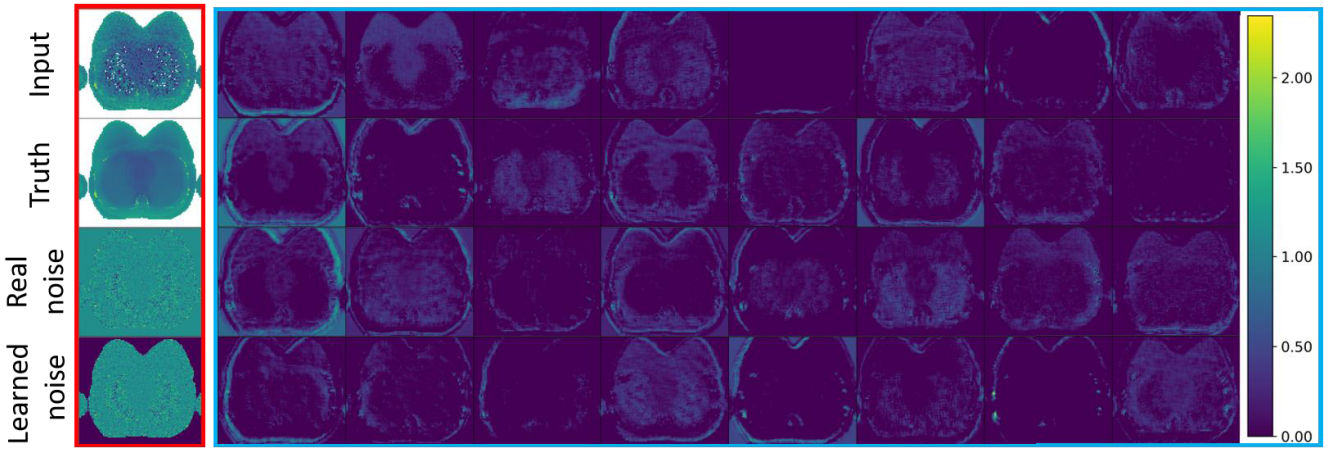

(b)

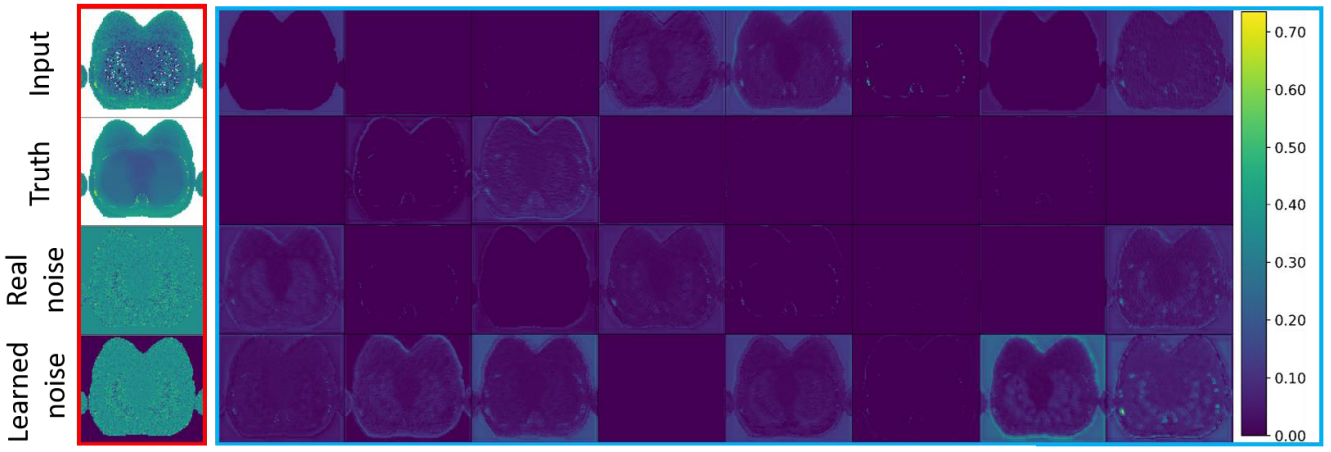

(c)

FIGURE 10. **The visualization of feature maps learned by the proposed network on the first, the fifth, and the ninth convolutional layers, respectively. In the red box of each subplot, we provide the input image, the ground-truth image, the real noise image (ground-truth - input image), and the learned noise image (the output of network before summation operation) as the references to facilitate the understanding of the feature maps shown in the blue box. Note that the ninth convolutional layer is before the last deconvolutional layer. All the feature maps refer to as the output of ReLU activation function after one convolutional layer. (a) The feature maps after the first convolutional layer. (b) The feature maps after the fifth convolutional layer. (c) The feature maps after the ninth convolutional layer.**

X GPU device to produce a similar amount of training data. Therefore, All the results of MC simulation in this paper are the state-of-the-art.

MCDNet has trained about the noise patterns involving 3D, heterogeneous patient anatomies in the context of x-ray CT imaging. The results show that MCDNet is quite capable

of predicting dose maps of acceptable statistical precision without suffering from overfitting – a problem with lacking training data. The limitations of this study are that all phantoms are adults, and we only considered the 100kVp x-ray energy spectrum. However, these limitations do not change the conclusion of this study and can be easily remedied in the future. Based on the MCDNet, we can further improve the denoising performance by considering the local structure and the similarity between low-photon dose and high-photon dose [30], [31]. In addition, to better understand the way of feature learning, we will further analyze the learned feature maps [32]. It should be also noted that, to derive dose to each of the radiosensitive organs, image analysis towards multi-organ segmentation can also be done [33].

VI. CONCLUSION

Using a unique set of full-body anatomically realistic voxel phantoms and a GPU-based parallel MC code, we proposed and tested a convolutional neural network named MCDNet to predict dose maps having acceptably small statistical noise using MC-simulated low-photon-fluence, high-noise dose map as input. For CT scan protocols considered in this study, MCDNet is found to have the ability of predicting dose maps of 9.9×10^7 photons from corresponding dose maps of 1.3×10^6 photons, yielding 76× speed-up in terms of photon numbers used in the MC simulations. Although there is perhaps still room for network optimization, the performance of the proposed methods is already impressive. Future work will further improve the denoising performance and apply to patient-specific clinical applications of CT dose and other procedures such as radiation treatment planning, radiation shielding design, and nuclear reactor physics that are governed by the same statistical and denoising principles.

ACKNOWLEDGEMENTS

(Zhao Peng and Hongming Shan contributed equally to this work.)

REFERENCES

[1] J. Valentin, *Managing Patient Dose in Multi-Detector Computed Tomography (MDCT)*. New York, NY, USA: Elsevier, 2007.
[2] D. J. Brenner and E. J. Hall, "Computed tomography—An increasing source of radiation exposure," *New England J. Med.*, vol. 357, no. 22, pp. 2277–2284, 2007.
[3] A. Ding, J. Gu, A. V. Trofimov, and X. G. Xu, "Monte Carlo calculation of imaging doses from diagnostic multidetector CT and kilovoltage cone-beam CT as part of prostate cancer treatment plans," *Med. Phys.*, vol. 37, no. 12, pp. 6199–6204, 2010.
[4] E. S. Amis, Jr., P. F. Butler, K. E. Applegate, S. B. Birnbaum, L. F. Brateman, J. M. Hevezi, F. A. Mettler, R. L. Morin, M. J. Pentecost, G. G. Smith, K. J. Strauss, and R. K. Zeman, "American College of Radiology white paper on radiation dose in medicine," *J. Amer. College Radiol.*, vol. 4, no. 5, pp. 272–284, 2007.
[5] F. H. Attix, *Introduction to Radiological Physics and Radiation Dosimetry*. Hoboken, NJ, USA: Wiley, 1986.
[6] *MCNP—A General Monte Carlo N-Particle Transport Code, Version 5*, X-5 Monte Carlo Team, Los Alamos Nat. Lab., Los Alamos, NM, USA, 2003, vol. 1.
[7] F. B. Brown, "Recent advances and future prospects for Monte Carlo (selected papers of the joint international conference of supercomputing in nuclear applications and Monte Carlo: SNA+ MC 2010)," *Prog. Nucl. Sci. Technol.*, vol. 2, pp. 1–4, 2011.
[8] J. Hammersley and D. J. Handscomb, *Monte Carlo Methods*. London, U.K.: Methuen, 1964.
[9] W. R. Nelson, D. W. Rogers, and H. Hirayama, *The EGS4 Code System*. Stanford, CA, USA: SLAC, 1985.
[10] R. E. Prael, "High-energy particle Monte Carlo at los alamos," in *Monte-Carlo Methods and Applications in Neutronics, Photonics and Statistical Physics*. Berlin, Germany: Springer, 1985, pp. 196–206.
[11] H. Hughes, R. Prael, and R. Little, "MCNPX—The LAHET/MCNP code merger, X-division research note XTM-Rn (U) 97-012," Los Alamos Nat. Lab., Los Alamos, NM, USA, Tech. Rep. LA-UR-97-4891, 1997.
[12] S. Agostinelli, J. Allison, K. A. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, and F. Behner, "Geant4—A simulation toolkit," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 506, no. 3, pp. 250–303, 2003.
[13] F. Salvat, J. M. Fernández-Varea, and J. Sempau, "PENELOPE-2008: A]code system for Monte Carlo simulation of electron and photon transport," in *Proc. Workshop*, 2006, vol. 4, no. 6222, p. 7.
[14] H. Zaidi and G. Sgouros, *Therapeutic Applications of Monte Carlo Calculations in Nuclear Medicine*. Boca Raton, FL, USA: CRC Press, 2002.
[15] G. Wang, "A perspective on deep imaging," *IEEE Access*, vol. 4, pp. 8914–8924, 2016.
[16] H. Shan, Y. Zhang, Q. Yang, U. Kruger, M. K. Kalra, L. Sun, W. Cong, and G. Wang, "3-D convolutional encoder-decoder network for low-dose CT via transfer learning from a 2-D trained network," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1522–1534, Jun. 2018.
[17] X. G. Xu, "An exponential growth of computational phantom research in radiation protection, imaging, and radiotherapy: A review of the fifty-year history," *Phys. Med. Biol.*, vol. 59, no. 18, p. R233, 2014.
[18] A. Ding, M. M. Mille, T. Liu, P. F. Caracappa, and X. G. Xu, "Extension of RPI-adult male and female computational phantoms to obese patients and a Monte Carlo study of the effect on CT imaging dose," *Phys. Med. Biol.*, vol. 57, no. 9, p. 2441, 2012.
[19] A. Ding, "Development of a radiation dose reporting software for X-ray computed tomography (CT)," Rensselaer Polytech. Inst., Troy, NY, USA, Tech. Rep., 2012.
[20] J. Gu, B. Bednarz, P. F. Caracappa, and X. G. Xu, "The development, validation and application of a multi-detector CT (MDCT) scanner model for assessing organ doses to the pregnant patient and the fetus using Monte Carlo simulations," *Phys. Med. Biol.*, vol. 54, no. 9, p. 2699, 2009.
[21] X. G. Xu, T. Liu, L. Su, X. Du, M. Riblett, W. Ji, D. Gu, C. D. Carothers, M. S. Shephard, F. B. Brown, M. K. Kalra, and B. Liu, "ARCHER, a new Monte Carlo software tool for emerging heterogeneous computing environments," in *Proc. SNA + MC Joint Int. Conf. Supercomput Nucl. Appl.*, Monte Carlo, CA, USA: EDP Sciences, 2014, Art. no. 06002.
[22] 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.* Springer, 2015, pp. 234–241.
[23] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 770–778.
[24] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proc. 27th Int. Conf. Mach. Learn. (ICML)*, 2010, pp. 807–814.
[25] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, *arXiv:1412.6980*. [Online]. Available: https://arxiv.org/abs/1412.6980
[26] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *Nature*, vol. 323, p. 9, Oct. 1986.
[27] M. Abadi "Tensorflow: A system for large-scale machine learning," in *Proc. 12th USENIX Symp. Oper. Syst. Design Implement. (OSDI)*, 2016, pp. 265–283.
[28] D. A. Low, W. B. Harms, S. Mutic, and J. A. Purdy, "A technique for the quantitative evaluation of dose distributions," *Med. Phys.*, vol. 25, no. 5, pp. 656–661, 1998.
[29] T. Liu, X. G. Xu, and C. D. Carothers, "Comparison of two accelerators for Monte Carlo radiation transport calculations, NVIDIA Tesla M2090 GPU and Intel Xeon Phi 5110p coprocessor: A case study for X-ray CT imaging dose calculation," *Ann. Nucl. Energy*, vol. 82, pp. 230–239, Aug. 2015.
[30] N. Liu, L. Wan, Y. Zhang, T. Zhou, H. Huo, and T. Fang, "Exploiting convolutional neural networks with deeply local description for remote sensing image classification," *IEEE Access*, vol. 6, pp. 11215–11228, 2018.

[31] Y. Zhang, H. Zhang, X. Chen, M. Liu, X. Zhu, S.-W. Lee, and D. Shen, "Strength and similarity guided group-level brain functional network construction for MCI diagnosis," *Pattern Recognit.*, vol. 88, pp. 421–430, Apr. 2019.
[32] G. Zhou, Q. Zhao, Y. Zhang, T. Adalı, S. Xie, and A. Cichocki, "Linked component analysis from matrices to high-order tensors: Applications to biomedical data," *Proc. IEEE*, vol. 104, no. 2, pp. 310–331, Feb. 2016.
[33] Z. Peng, X. Fang, H. Shan, T. Liu, X. Pei, P. Yan, G. Wang, B. Liu, M. Kalra, and X. G. Xu, "Multi-Organ segmentation of CT images using deep-learning for instant and patient-specific dose reporting," in *Proc. 61st Annu. Meeting Exhib. Amer. Assoc. Med. Physicists (AAPM)*, San Antonio, TX, USA, Jul. 2019.

ZHAO PENG received the B.S. degree from the University of Science and Technology of China (USTC), in 2017, where he is currently pursuing the degree. He is currently a Visiting Scholar with the Rensselaer Polytechnic Institute, Troy, NY, USA. His research interests include the application of deep learning for medical images analysis and radiation treatment planning.

HONGMING SHAN received the B.S. degree from Shandong University of Technology, China, in 2011, and the Ph.D. degree from Fudan University, China, in 2017. He is currently a Postdoctoral Scholar with the Rensselaer Polytechnic Institute, Troy, New York, USA. His research interests include machine/deep learning, computer vision, dimension reduction, and biomedical imaging.

TIANYU LIU received the Ph.D. degree from Rensselaer Polytechnic Institute (RPI), Troy, New York, USA, in 2014. He is currently a Computer Scientist with RPI, with the expertise in developing high-performance computing software as well as computational phantoms for radiation dosimetry. He is the Lead Developer of ARCHERTM, a commercial GPU-accelerated Monte Carlo dose engine.

XI PEI received the Ph.D. degree from the University of Chinese Academy of Sciences, in 2011. He is currently an Associate Professor with the Department of Engineering and Applied Physics, University of Science and Technology of China, Hefei, China, where he conducts research on computational and optimization methods of medical imaging and radiotherapy treatment planning. He has authored 30 peer-reviewed papers, 40 conference abstracts, and 4 patents. He is a member of the Radiation and Environment Committee of the Chinese Biophysics Physics Society (BPS) and the Accurate Radiotherapy Technical Branch of the Chinese Biomedical Medical Engineering Society (BMES).

GE WANG received the Ph.D. degree from the State University of New York at Buffalo, New York, in 1992. He is the Clark & Crossan Endowed Chair Professor of the Department of Biomedical Engineering and Director of Biomedical Imaging Center, Rensselaer Polytechnic Institute, Troy, NY, USA. He has published the first spiral/helical cone-beam/multi-slice CT algorithm in 1991 and since then more than 100 papers systematically covering theory, algorithms, artifact reduction, and biomedical applications in this area. Currently, there are more than 100 million medical CT scans yearly with a majority in the spiral/helical cone-beam/multi-slice mode. His group developed interior tomography to solve the long-standing "interior problem," and enable omni-tomography ("all-in-one") with CT-MRI as an example. He initiated the area of bioluminescence tomography. His results were featured in Nature, Science, and PNAS, and recognized with academic awards. He wrote more than 450 journal publications, receiving a high number of citations. In 2016, he gave the first perspective on neural-network-based tomographic imaging as the new frontier of machine learning. His team has been in collaboration with world-class groups, and continuously well-funded by federal agencies and academic-industrial partnerships translating machine learning techniques into imaging products. His research interests include X-ray CT, MRI, optical molecular tomography, multimodality fusion, and machine learning. He is the Lead Guest Editor of the five IEEE TRANSACTIONS ON MEDICAL IMAGING SPECIAL ISSUES, Founding Editor-in-Chief of the International Journal of Biomedical Imaging, Outstanding Associate Editor of the IEEE TRANSCATIONS ON MEDICAL IMAGING, a Board Member of IEEE ACCESS, and Associate Editor of the IEEE TRANSACTIONS ON RADIATION AND PLASMA MEDICAL SCIENCES, and Medical Physics. He is a Fellow of the SPIE, OSA, AIMBE, AAPM, and AAAS.

X. GEORGE XU received the Ph.D. degree in nuclear engineering from Texas A&M University, in 1994. He has been on the Faculty of Rensselaer Polytechnic Institute, Troy, NY, USA, since 1995, and currently holding the Edward E. Hood Endowed Chair of Engineering with joint appointment with the Department of Mechanical, Aerospace, and Nuclear Engineering and Department of Biomedical Engineering. He also holds an adjunct faculty appointment with the College of Physics, University of Science and Technology of China, Hefei, China. He has mentored more than 50 Ph.D. and M.S. students, covering computational and experimental methods of "radiation dosimetry" for radiation protection, medical imaging, and radiotherapy. In the past 24 years, their research has been funded by NSF, DOE, NIST, and NIH, as well as by the private sector industry. He has authored more than 550 archival journal papers and conference abstracts, as well as 130 invited talks. He is the Co-Editor of "Handbook of Anatomical Models for Radiation Dosimetry" and author of several comprehensive review articles. He is a Fellow of ANS, HPS, and AAPM, as well as an Elected Council Member of the NCRP. He is currently an Editorial Board Member of many journals including Medical Physics and Physics in Medicine and Biology. Among his professional recognitions are the NSF/ CAREER Award, ANS/RPSD Professional Excellence Award, and CIRMS Randal S. Caswell Award for Distinguished Achievements, and HPS Distinguished Scientific Achievement Award. He is the Co-Founder of Virtual Phantoms Inc. that commercializes VirtualDoseTM (a patient radiation dose tracking software for CT and interventional radiology procedures) and ARCHERTM (a GPU-based Monte Carlo dose computing software for radiotherapy treatment planning and verification).