



DeepDRR – A Catalyst for Machine Learning in Fluoroscopy-Guided Procedures

Mathias Unberath^{1(✉)}, Jan-Nico Zaech^{1,2}, Sing Chun Lee¹, Bastian Bier^{1,2}, Javad Fotouhi¹, Mehran Armand³, and Nassir Navab¹

¹ Computer Aided Medical Procedures, Johns Hopkins University, Baltimore, USA
unberath@jhu.edu

² Pattern Recognition Lab, Friedrich-Alexander-Universität Erlangen-Nürnberg, Erlangen, Germany

³ Applied Physics Laboratory, Johns Hopkins University, Baltimore, USA

Abstract. Machine learning-based approaches outperform competing methods in most disciplines relevant to diagnostic radiology. Interventional radiology, however, has not yet benefited substantially from the advent of deep learning, in particular because of two reasons: (1) Most images acquired during the procedure are never archived and are thus not available for learning, and (2) even if they were available, annotations would be a severe challenge due to the vast amounts of data. When considering fluoroscopy-guided procedures, an interesting alternative to true interventional fluoroscopy is *in silico* simulation of the procedure from 3D diagnostic CT. In this case, labeling is comparably easy and potentially readily available, yet, the appropriateness of resulting synthetic data is dependent on the forward model. In this work, we propose DeepDRR, a framework for fast and realistic simulation of fluoroscopy and digital radiography from CT scans, tightly integrated with the software platforms native to deep learning. We use machine learning for material decomposition and scatter estimation in 3D and 2D, respectively, combined with analytic forward projection and noise injection to achieve the required performance. On the example of anatomical landmark detection in X-ray images of the pelvis, we demonstrate that machine learning models trained on DeepDRRs generalize to unseen clinically acquired data without the need for re-training or domain adaptation. Our results are promising and promote the establishment of machine learning in fluoroscopy-guided procedures.

Keywords: Monte Carlo simulation · Volumetric segmentation
Beam hardening · Image-guided procedures

M. Unberath and J.-N. Zaech—Both authors contributed equally and are listed in alphabetical order.

1 Introduction

The advent of convolutional neural networks (ConvNets) for classification, regression, and prediction tasks, currently most commonly referred to as deep learning, has brought substantial improvements to many well studied problems in computer vision, and more recently, medical image computing. This field is dominated by diagnostic imaging tasks where (1) all image data are archived, (2) learning targets, in particular annotations of any kind, exist traditionally [1] or can be approximated [2], and (3) comparably simple augmentation strategies, such as rigid and non-rigid displacements [3], ease the limited data problem.

Unfortunately, the situation is more complicated in interventional imaging, particularly in 2D fluoroscopy-guided procedures. First, while many X-ray images are acquired for procedural guidance, only very few radiographs that document the procedural outcome are archived suggesting a severe lack of meaningful data. Second, learning targets are not well established or defined; and third, there is great variability in the data, e.g. due to different surgical tools present in the images, which challenges meaningful augmentation. Consequently, substantial amounts of clinical data must be collected and annotated to enable machine learning for fluoroscopy-guided procedures. Despite clear opportunities, in particular for prediction tasks, very little work has considered learning in this context [4–7].

A promising approach to tackling the above challenges is *in silico* fluoroscopy generation from diagnostic 3D CT, most commonly referred to as digitally reconstructed radiographs (DRRs) [4,5]. Rendering DRRs from CT provides fluoroscopy in known geometry, but more importantly: Annotation and augmentation can be performed on the 3D CT substantially reducing the workload and promoting valid image characteristics, respectively. However, machine learning models trained on DRRs do not generalize to clinical data since traditional DRR generation, e.g. as in [4,8], does not accurately model X-ray image formation. To overcome this limitation we propose DeepDRR, an easy-to-use framework for realistic DRR generation from CT volumes targeted at the machine learning community. On the example of view independent anatomical landmark detection in pelvic trauma surgery [9], we demonstrate that training on DeepDRRs enables direct application of the learned model to clinical data without the need for re-training or domain adaptation.

2 Methods

2.1 Background and Requirements

DRR generation considers the problem of finding detector responses given a particular imaging geometry according to Beer-Lambert law [10]. Methods for *in silico* generation of DRRs can be grouped in analytic and statistical approaches, i.e. ray-tracing and Monte Carlo (MC) simulation, respectively. Ray-tracing algorithms are computationally efficient since the attenuated photon fluence of a detector pixel is determined by computing total attenuation along a 3D line

that then applies to all photons emitted in that direction [8]. Commonly, ray-tracing only considers a single material in the mono-energetic case and thus fails to model beam hardening. In addition and since ray-tracing is analytic, statistical processes during image formation, such as scattering, cannot be modeled. Conversely, MC methods simulate single photon transport by evaluating the probability of photon-matter interaction, the sequence of which determines attenuation [11]. Since the probability of interaction is inherently material and energy dependent, MC simulations require material decomposition in CT that is usually achieved by thresholding of CT values (Hounsfield units, HU) [12] and spectra of the emitter [11]. As a consequence, MC is very realistic. Unfortunately, for training-set-size DRR generation on conventional hardware, MC is prohibitively expensive. As an example, accelerated MC simulation [11] on an NVIDIA Titan Xp takes ≈ 4 h for a single X-ray image with 10^{10} photons. To leverage the advantages of MC simulations in clinical practice, the medical physics community provides further acceleration strategies if prior knowledge on the problem exists. A well studied example is variance reduction for scatter correction in cone-beam CT, since scatter is of low frequency [13].

Unfortunately, several challenges remain that hinder the implementation of realistic *in silico* X-ray generation for machine learning applications. We have identified the following fundamental challenges at the interface of machine learning and medical physics that must be overcome to establish realistic simulation in the machine learning community: (1) Tools designed for machine learning must seamlessly integrate with the common frameworks. (2) Training requires many images so data generation must be fast and automatic. (3) Simulation must be realistic: Both analytic and statistic processes such as beam-hardening and scatter, respectively, must be modeled.

2.2 DeepDRR

Overview: We propose DeepDRR, a Python, PyCUDA, and PyTorch-based framework for fast and automatic simulation of X-ray images from CT data. It consists of 4 major modules: (1) Material decomposition in CT volumes using a deep segmentation ConvNet; (2) A material- and spectrum-aware ray-tracing forward projector; (3) A neural network-based Rayleigh scatter estimation; and (4) Quantum and electronic readout noise injection. The individual steps of DeepDRR are visualized schematically in Fig. 1 and explained in greater detail in the remainder of this section. The fully automated pipeline is open source available for download¹.

Material Decomposition: Material decomposition in 3D CT for MC simulation is traditionally accomplished by thresholding, since a given material has a characteristic HU range [12]. This works well for large HU discrepancies, e. g. air ($[-1000]$ HU) and bone ($[200, 3000]$ HU), but may fail otherwise, particularly

¹ Github link: <https://github.com/mathiasunberath/DeepDRR>.

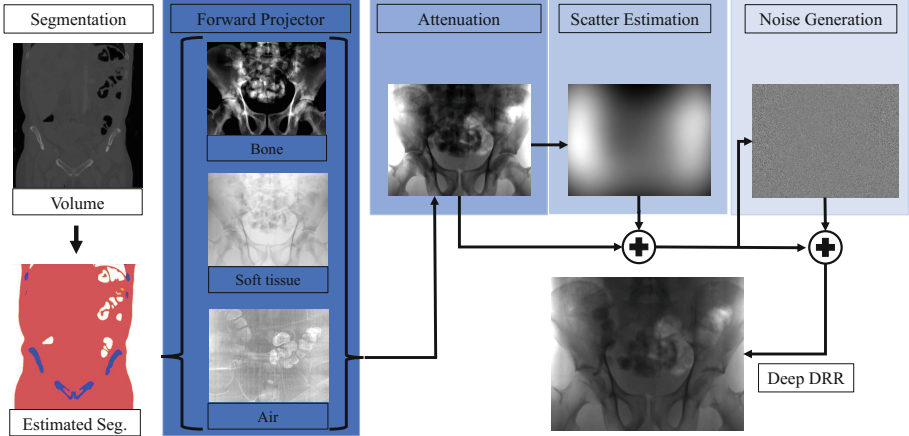


Fig. 1. Schematic overview of DeepDRR.

between soft tissue ($[-150, 300]$ HU) and bone in presence of low mineral density. This is problematic since, despite similar HU, the attenuation characteristic of bone is substantially different of soft tissue [10]. Within this work, we use a deep volumetric ConvNet adapted from [3] to automatically decompose air, soft tissue, and bone in CT volumes. The ConvNet is of encoder-decoder structure with skip-ahead connections to retain information of high spatial resolution while enabling large receptive fields. The ConvNet is trained on patches with $128 \times 128 \times 128$ voxels with voxel sizes of $0.86 \times 0.86 \times 1.0$ mm yielding a material map $M(\mathbf{x})$ that assigns a candidate material to each 3D point \mathbf{x} . We used the multi-class Dice loss as the optimization target. 12 whole-body CT data were manually annotated, and then split: 10 for training, and 2 for validation and testing. Training was performed over 600 epochs until convergence where, in each epoch, one patch from every volume was randomly extracted. During application, patches of $128 \times 128 \times 128$ voxels are fed-forward with stride of 64 since only labels for the central $64 \times 64 \times 64$ voxels are accepted.

Analytic Primary Computation: Once segmentations of the considered materials $M = \{\text{air, soft tissue, bone}\}$ are available, the contribution of each material to the total attenuation density at detector position \mathbf{u} are computed using a given geometry (defined by projection matrix $\mathbf{P} \in \mathbb{R}^{3 \times 4}$) and X-ray spectral density $p_0(E)$ via ray-tracing:

$$\begin{aligned}
 p(\mathbf{u}) &= \int p(E, \mathbf{u}) dE \\
 &= \int p_0(E) \exp \left(\sum_{m \in M} \delta(m, M(\mathbf{x})) (\mu/\rho)_m(E) \int \rho(\mathbf{x}) d\mathbf{l}_{\mathbf{u}} \right) dE, \quad (1)
 \end{aligned}$$

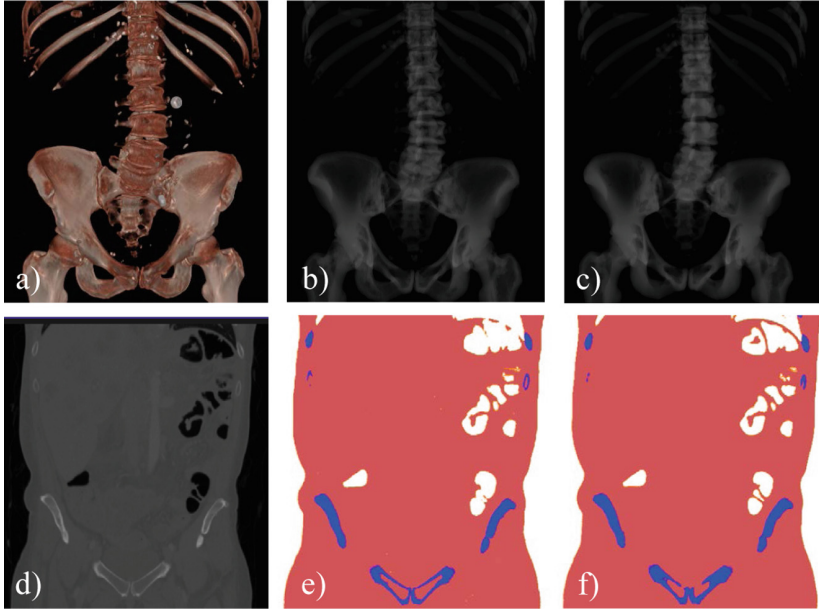


Fig. 2. Representative results of the segmentation ConvNets. From left to right, the columns show input volume, manual segmentation, and ConvNet result. The top rows shows volume renderings of the bony anatomy and respective label, while the bottom row shows a coronal slice through the volumes.

where $\delta(\cdot, \cdot)$ is the Kronecker delta, \mathbf{l}_u is the 3D ray connecting the source position and 3D location of detector pixel \mathbf{u} determined by \mathbf{P} , $(\mu/\rho)_m(E)$ is the material and energy dependent linear attenuation coefficient [10], and $\rho(\mathbf{x})$ is the material density at position \mathbf{x} derived from HU values. The projection domain image $p(\mathbf{u})$ is then used as input to our scatter prediction ConvNet.

Learning-Based Scatter Estimation: Traditional scatter estimation relies on variance-reduced MC simulations [13], which requires a complete MC setup. Recent approaches to scatter estimation via ConvNets outperform kernel based methods [14] while retaining the low computational demand. In addition, they inherently integrate with deep learning software environments. We define a ten layer ConvNet, where the first six layers generate Rayleigh scatter estimates and the last four layers, with 31×31 kernels and a single channel, ensure smoothness. The network was trained on 330 images generated via MC simulation [11], augmented by random rotations and reflections. The last three layers were trained using pre-training of the preceding layers. The input to the network is down-sampled to 128×128 pixels.

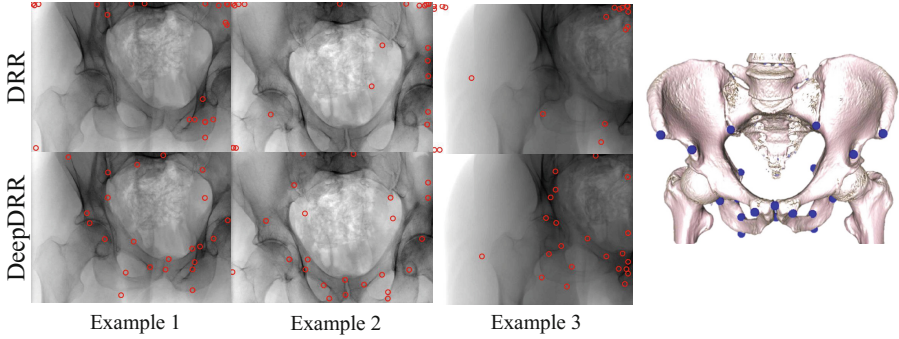


Fig. 3. Anatomical landmark detection on real data using the method detailed in [9]. Top row: Detections of a model trained on conventional DRRs. Bottom row: Detections of a model trained on the proposed DeepDRRs. No domain adaption or re-training was performed. Right-most image: Schematic illustration of desired landmark locations.

Noise Injection: After adding scatter, $p(\mathbf{u})$ expresses the energy deposited by a photon in detector pixel \mathbf{u} . The number of photons is estimated as:

$$N(\mathbf{u}) = \sum_E \frac{p(E, \mathbf{u})}{E} N_0, \quad (2)$$

to obtain the number of registered photons $N(\mathbf{u})$ and perform realistic noise injection. In Eq. 2, N_0 (potentially location dependent $N_0(\mathbf{u})$, e.g. due to bow-tie filters) is the emitted number of photons per pixel. Noise in X-ray images is a composite of uncorrelated quantum noise due to photon statistics that becomes correlated due to pixel crosstalk, and correlated readout noise [15]. Due to beam hardening, the spectrum arriving at any detector pixel differs. To account for this fact in the Poisson noise model, we compute a mean photon energy for each pixel by $\bar{E}(\mathbf{u})$ and estimate quantum noise as $\bar{E}/N_0(p_{Poisson}(N) - N)$, where $p_{Poisson}$ is the Poisson generating function. Since real flat panel detectors suffer from pixel crosstalk, we correlate the quantum noise of neighboring pixels by convolving the noise signal with a blurring kernel [15]. The second major noise component is electronic readout noise. Electronic noise is signal independent and can be modeled as additive Gaussian noise with correlation along rows due to sequential readout [15]. Finally, we obtain a realistically simulated DRR.

3 Experiments and Results

3.1 Framework Validation

Since forward projection and noise injection are analytic processes, we only assess the prediction accuracy of the proposed ConvNets for volumetric segmentation and projection domain scatter estimation. For volumetric segmentation of air, soft tissue, and bone in CT volumes, we found a misclassification rate of

$(2.03 \pm 3.63) \%$ which is in line with results reported in previous studies using this architecture [3]. Representative results on the test set are shown in Fig. 2. For scatter estimation, the evaluation on a test set consisting of 30 image yielded a normalized mean squared error of 9.96%. For 1000 images with 620×480 px, the simulation per image took 0.56 s irrespective of number of emitted photons.

3.2 Task-Based Evaluation

Fundamentally, the goal of DeepDRR is to enable the learning of models on synthetically generated data that generalizes to unseen clinical fluoroscopy without re-training or other domain adaptation strategies. To this end, we consider anatomical landmark detection in X-ray images of the pelvis from arbitrary views [9]. The authors annotated 23 anatomical landmarks in CT volumes of the pelvis (Fig. 3, last column) and generated DRRs with annotations on a spherical segment covering 120° and 90° in RAO/LAO and CRAN/CAUD, respectively. Then, a sequential prediction framework is learned and, upon convergence, used to predict the 23 anatomical landmarks in unseen, real X-ray images of cadaver studies. The network is learned twice: First, on conventionally generated DRRs assuming a single material and mono-energetic spectrum, and second, on DeepDRRs as described in Sect. 2.2. Images had 615×479 pixels with 0.616^2 mm pixel size. We used the spectrum of a tungsten anode operated at 120 kV with 4.3 mm aluminum and assumed a high-dose acquisition with $5 \cdot 10^5$ photons per pixel. In Fig. 3 we show representative detections of the sequential prediction framework on unseen, clinical data acquired using a flat panel C-arm system (Siemens Cios Fusion, Siemens Healthcare GmbH, Germany) during cadaver studies. As expected, the model trained on conventional DRRs (upper row) fails to predict anatomical landmark locations on clinical data, while the model trained on DeepDRRs produces accurate predictions even on partial anatomy. In addition, we would like to refer to the comprehensive results reported in [9] that were achieved using training on the proposed DeepDRRs.

4 Discussion and Conclusion

We proposed DeepDRR, a framework for fast and realistic generation of synthetic X-ray images from diagnostic 3D CT, in an effort to ease the establishment of machine learning-based approaches in fluoroscopy-guided procedures. The framework combines novel learning-based algorithms for 3D material decomposition from CT and 2D scatter estimation with fast, analytic models for energy and material dependent forward projection and noise injection. On a surrogate task, i. e. the prediction of anatomical landmarks in X-ray images of the pelvis, we demonstrate that models trained on DeepDRRs generalize to clinical data without the need of re-training or domain adaptation, while the same model trained on conventional DRRs is unable to perform. Our future work will focus on improving volumetric segmentation by introducing more materials, in particular metal, and scatter estimation that could benefit from a larger training set

size. In conclusion, we understand realistic *in silico* generation of X-ray images, e. g. using the proposed framework, as a catalyst designed to benefit the implementation of machine learning in fluoroscopy-guided procedures. Our framework seamlessly integrates with the software environment currently used for machine learning and will be made open-source at the time of publication².

References

1. Kooi, T., et al.: Large scale deep learning for computer aided detection of mammographic lesions. *Med. Image Anal.* **35**, 303–312 (2017)
2. Roy, A.G., Conjeti, S., Sheet, D., Katouzian, A., Navab, N., Wachinger, C.: Error corrective boosting for learning fully convolutional networks with limited data. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 231–239. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_27
3. Milletari, F., Navab, N., Ahmadi, S.A.: V-Net: fully convolutional neural networks for volumetric medical image segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV), pp. 565–571. IEEE (2016)
4. Li, Y., Liang, W., Zhang, Y., An, H., Tan, J.: Automatic lumbar vertebrae detection based on feature fusion deep learning for partial occluded C-arm X-ray images. In: 2016 IEEE 38th Annual International Conference of the Engineering in Medicine and Biology Society (EMBC), pp. 647–650. IEEE (2016)
5. Terunuma, T., Tokui, A., Sakae, T.: Novel real-time tumor-contouring method using deep learning to prevent mistracking in X-ray fluoroscopy. *Radiol. Phys. Technol.* **11**, 43–53 (2017)
6. Ambrosini, P., Ruijters, D., Niessen, W.J., Moelker, A., van Walsum, T.: Fully automatic and real-time catheter segmentation in X-ray fluoroscopy. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10434, pp. 577–585. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66185-8_65
7. Ma, H., Ambrosini, P., van Walsum, T.: Fast prospective detection of contrast inflow in X-ray angiograms with convolutional neural network and recurrent neural network. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) MICCAI 2017. LNCS, vol. 10435, pp. 453–461. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-66179-7_52
8. Russakoff, D.B., et al.: Fast generation of digitally reconstructed radiographs using attenuation fields with application to 2D–3D image registration. *IEEE Trans. Med. Imaging* **24**(11), 1441–1454 (2005)
9. Bier, B., et al.: X-ray-transform invariant anatomical landmark detection for pelvic trauma surgery. In: Frangi, A.F., et al. (eds.) MICCAI 2018. LNCS, vol. 11073, pp. 55–63. Springer, Heidelberg (2018)
10. Hubbell, J.H., Seltzer, S.M.: Tables of X-ray mass attenuation coefficients and mass energy-absorption coefficients 1 keV to 20 MeV for elements $Z = 1$ to 92 and 48 additional substances of dosimetric interest. Technical report, National Institute of Standards and Technology (1995)

² The source code available at this link: <https://github.com/mathiasunberath/DeepDRR>.

11. Badal, A., Badano, A.: Accelerating Monte Carlo simulations of photon transport in a voxelized geometry using a massively parallel graphics processing unit. *Med. Phys.* **36**(11), 4878–4880 (2009)
12. Schneider, W., Bortfeld, T., Schlegel, W.: Correlation between CT numbers and tissue parameters needed for Monte Carlo simulations of clinical dose distributions. *Phys. Med. Biol.* **45**(2), 459 (2000)
13. Sisniega, A., et al.: Monte carlo study of the effects of system geometry and anti-scatter grids on cone-beam CT scatter distributions. *Med. Phys.* **40**(5) (2013)
14. Maier, J., Sawall, S., Kachelrieß, M.: Deep scatter estimation (DSE): feasibility of using a deep convolutional neural network for real-time X-ray scatter prediction in cone-beam CT. In: *SPIE Medical Imaging*, SPIE (2018)
15. Zhang, H., Ouyang, L., Ma, J., Huang, J., Chen, W., Wang, J.: Noise correlation in CBCT projection data and its application for noise reduction in low-dose CBCT. *Med. Phys.* **41**(3) (2014)