

Physics guided machine learning for multi-material decomposition of tissues from dual-energy CT scans of simulated breast models with calcifications

Muralikrishnan Gopalakrishnan Meena^{1,†}, Amir K. Ziabar², Singanallur V. Venkatakrishnan², Isaac R. Lyngaas¹, Matthew R. Norman¹, Balint Joo¹, Thomas L. Beck¹, Charles A. Bouman³, Anuj J. Kapadia⁴, and Xiao Wang⁴

¹National Center for Computational Sciences, ²Electrification and Energy Infrastructure Division, ⁴Computational Sciences and Engineering Division, Oak Ridge National Laboratory, Oak Ridge, TN, USA

³School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA

[†]Email address for correspondence: gopalakrishm@ornl.gov

Abstract

We introduce a physics guided data-driven method for image-based multi-material decomposition for dual-energy computed tomography (CT) scans. The method is demonstrated for CT scans of virtual human phantoms containing more than two types of tissues. The method is a physics-driven supervised learning technique. We take advantage of the mass attenuation coefficient of dense materials compared to that of muscle tissues to perform a preliminary extraction of the dense material from the images using unsupervised methods. We then perform supervised deep learning on the images processed by the extracted dense material to obtain the final multi-material tissue map. The method is demonstrated on simulated breast models with calcifications as the dense material placed amongst the muscle tissues. The physics-guided machine learning method accurately decomposes the various tissues from input images, achieving a normalized root-mean-squared error of 2.75%.

Introduction

Computed Tomography (CT) imaging comprises passing X-rays through an opaque object at different orientations to obtain images depicting the internal structure of the object [2]. The X-ray projections are detected by the X-ray detector, and these X-ray projections from different view angles of CT rotations are arranged into a 2D array, and we call such array the “sinogram”. Then, the CT images are reconstructed from the sinograms using different CT image reconstruction algorithms [1], such as filtered back-projection (FBP) and Model-Based Iterative Reconstruction (MBIR) [23, 25, 27]. CT imaging has a wide array of applications ranging from medical and security [3, 10, 12, 27, 23, 25] to industrial and scientific [9, 6, 19, 26]. As the objects being scanned comprise various materials, it is important to identify the location and density of each material in the object. One important task in CT imaging is to identify the underlying structure of an object as various applications will have objects with various materials. Thus, it is important to extract the material composition (density and location of materials) of the reconstructed image. Since each pixel (or voxel) in the CT image consists of the total radio-density of all materials in the object, the task of extracting the material composition at each pixel becomes challenging when there exists a mixture of materials at a given pixel. This problem is called multi-material decomposition in CT imaging.

If the object only consists of 2 materials, one solution to quantitatively decompose the 2 materials is through the use of dual-energy CT (DECT) scans [18]. In DECT, the object is incident with X-rays from two different energies, which correspond to the characteristic absorption capabilities of the two materials. This results in an inverse problem, allowing 2 inversions to extract the material composition of the 2 materials. The task becomes challenging when the object comprise of more than 2 materials, which cannot be solved using 2 inversions. The use of DECT is the standard in current clinical practices to decompose up to 2 materials [7], but practical realistic application of material decomposition requires the capability to decompose multiple materials.

Traditional multi-material decomposition in DECT has been performed through projection-based techniques concurrent with the iterative methods for reconstructing the images [14, 15, 20]. While these methods are highly accurate, they require access to sinograms, which have limited public access. In addition, projection-based approaches are limited to 2D images and cannot be applied to 3D image volumes [15]. Thus, projection-based multi-material decomposition is not currently used in clinical practice and is limited to academic research [18, 7, 8, 29].

There have been recent trends in image-based decomposition which are performed solely in the image domain [17, 18, 8, 29]. Thus, they can be considered as a post-processing step with respect to the image reconstruction procedure. These methods have the advantage of not needing access to sinograms. Nonetheless, they have poor accuracy [20] and are not guided by physics. In this work, we use machine learning (ML) methods guided by physics information to perform image-based multi-material decomposition of tissue maps from DECT scans of subjects with multiple materials.

Image data

The current study expands on the authors’ participation in the 2022 American Association of Physicists in Medicine (AAPM) Grand Data Challenge, DL-spectral CT (www.aapm.org/GrandChallenge/DL-spectral-CT/) [24]. The data challenge broadly involved decomposition of tissue maps from simulated dual-energy X-ray CT scans of virtual human subjects with more than two tissues. The dataset in the study comprised 1000 cases of simulated 2D breast CT scans [21] with four materials: three tissues (adipose, fibroglandular, and calci-

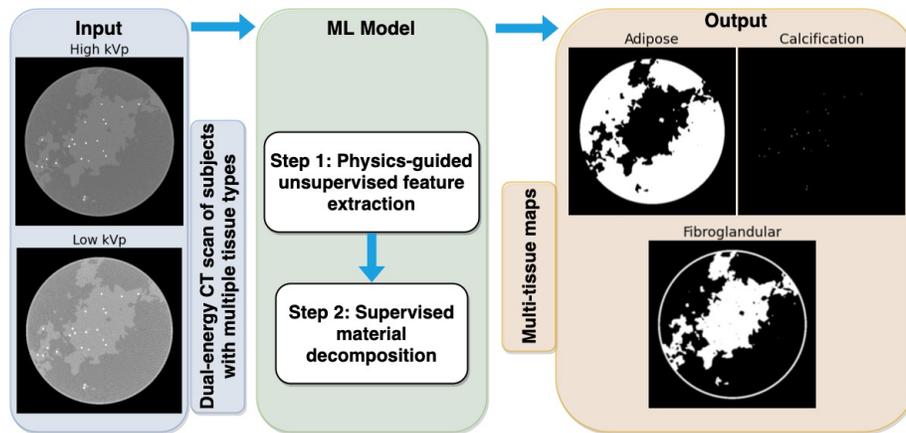


Figure 1. Sample portrayal of the image-based machine learning framework for multi-material decomposition using dual-energy CT scans.

fication) and air. Each case contains $256 \text{ views} \times 1024$ detector-pixel X-ray transmission data from 50 and 80 kVp dual-energy acquisitions. The dataset also contained reconstructed images of 512×512 pixel resolution for the high- and low-kVp acquisitions. These images were reconstructed using standard negative logarithm processing and FBP. However, they exhibited streaking and cupping (non-uniform background gray level) artifacts.

The objective of this work is to decompose the three types of tissues and obtain tissue maps using only the CT reconstructed images of the breast tissue phantoms. We note that the AAPM data challenge provided three options to obtain the tissue maps: using the transmission data, the FBP reconstructed images, or a combination of both. However, in the current study, we chose to use only the FBP images provided by the AAPM data challenge for obtaining the tissue maps. We did not use other algorithms to reconstruct the images from the provided transmission data.

For this work, we implemented a combination of unsupervised and supervised learning techniques to decompose the 3 tissues and obtain the respective tissue maps by only using the FBP images as the inputs, as shown in Fig. 1. Note that while the teams who participated in the AAPM data challenge attained near floating-point zero accuracy, none of the top 10 teams utilized only the FBP reconstructed images to perform the decomposition [24]. They relied on different variants of image reconstruction techniques to recover the tissue maps.

Method

The ML framework comprise of two steps as shown in Fig. 1. Step 1 involves using the input images (high- and low-kVp) to perform unsupervised feature extraction guided by physics knowledge. Results from step 1 are fed into step 2 to perform supervised multi-material decomposition. Step 1 broadly extracts certain features which are used to aid step 2. We note that while the current multi-material decomposition framework is demonstrated on the simulated 2D breast CT scans from the AAPM data challenge, our ML framework can be generalized to other mammography datasets.

Step 1: Physics-guided unsupervised feature extraction

The objective of step 1 is to extract the dense material(s) from the input images. The motivation of step 1 is to reduce the number of materials to be extracted by step 2, and thus, reduce the complexity of the multi-material decomposition problem. Moreover, we use physics guidance to improve the overall quality of the decomposition.

The physics information we use is that high-density lesions have significantly different X-ray mass attenuation coefficients compared to human tissues at different X-ray energies. X-ray mass attenuation coefficient is the amount of X-ray absorbed by a material at a given X-ray energy. In the current dataset, the calcification regions are assumed to be composed of hydroxyapatite [5], a form of calcium. We take advantage of the existing physics knowledge in the literature that calcium and the materials of the other tissues, adipose and fibroglandular, absorb X-rays at different rate at different X-ray energies (for a depiction of the difference in X-ray attenuation coefficients, see Fig. 1 of [28]). Thus, taking the difference between the high- and low-kVp images reveals the calcification regions. To generalize, the high-density lesions in a DECT scan can be extracted by taking the difference of the low- and high-kVp images. We use this physics information as a guidance to perform unsupervised feature extraction of the dense material (calcification) from the dataset, as shown in Fig. 2. Step 1 can be summarized broadly into the following:

1. Extract an image populated with high-density lesion regions (calcification regions in the current dataset) by subtracting the high- and low-kVp images. We call this image the high-low kVp difference image. (Fig. 2(a))
2. Use unsupervised image segmentation to extract the high-density lesion regions from the difference image. (Fig. 2(b))
3. Replace the high-density lesion regions from the high- and low-kVp images with the background values (air in the current dataset). (Fig. 3)
4. Provide the three images: high-calcification, low-calcification, and calcification images, to step 2.

To perform the unsupervised feature extraction, we use the multi-class Otsu thresholding [13] to separate the pixels at the regions of high-density lesions in the high-low kVp difference im-

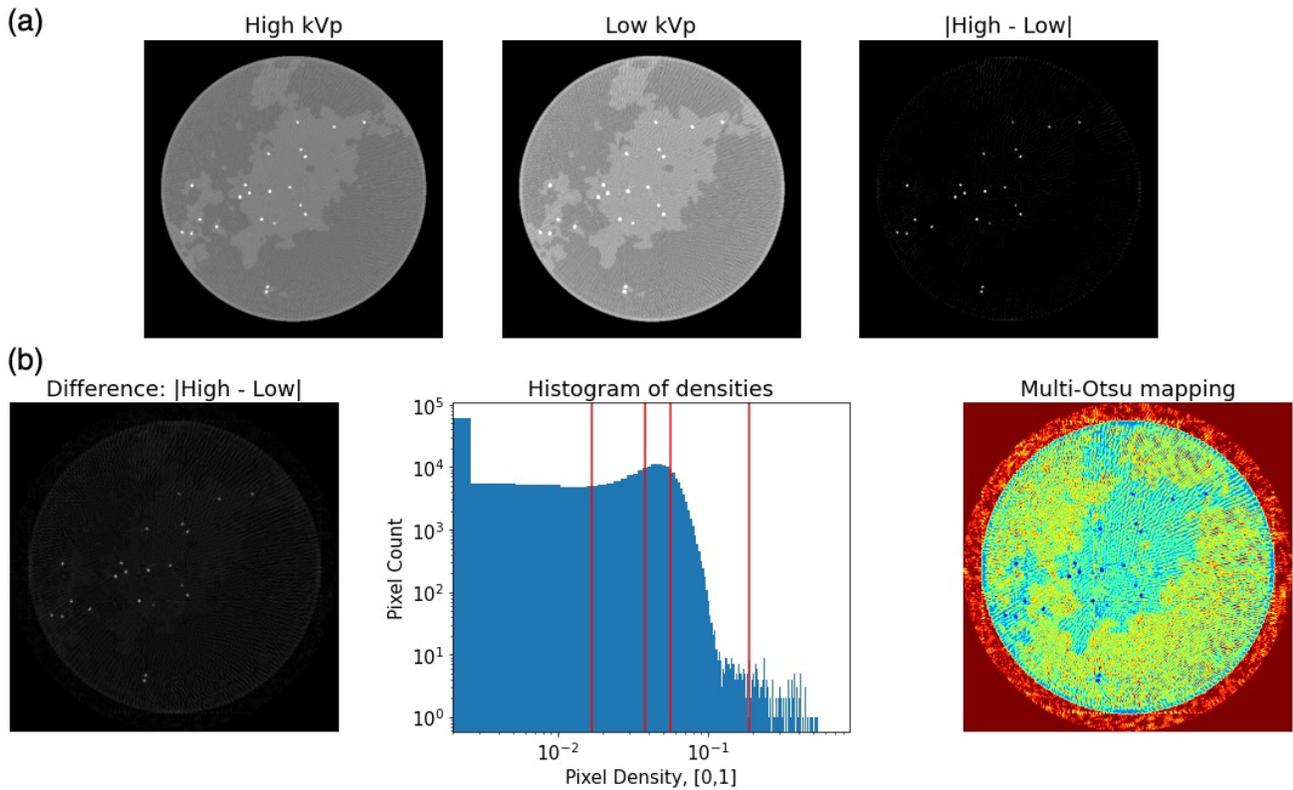


Figure 2. Step 1 of the machine learning framework. (a) The FBP reconstructed high-kVp (left) and low-kVp (center) images. The difference between the high- and low-kVp images (right) highlights the high-density lesion regions (calcification in the current dataset). (b) The multi-Otsu image thresholding routine is applied to the difference image. The input high-low kVp difference image (left), histogram of pixel densities (center), and output labels (right) are shown.

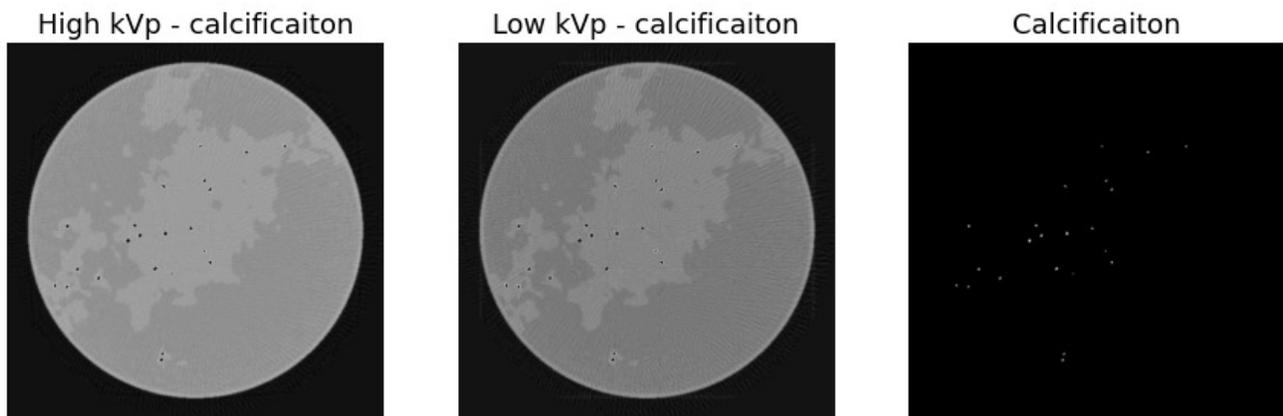


Figure 3. Input to step 2. The high-density lesion regions (calcification in the current dataset) identified from multi-Otsu algorithm are subtracted from the high-kVp (left) and low-kVp (center) images, and the corresponding pixels are replaced by the background values (air in the current dataset). The high-density lesion regions are also added as an input to step 2 (right).

age as shown in Fig. 2(b). For the current dataset, we found that extracting the calcification regions require 5 classes to be identified using the multi-Otsu thresholding. Intuitively, the thresholding algorithm would identify 4 classes - the pixels corresponding to calcification, adipose, fibroglandular, and the background air. Nonetheless, the region around the breast model exhibits a halo effect, which the multi-Otsu thresholding algorithm identifies as an additional class, as shown in Fig. 2(b)-right. The algorithm is unable to accurately extract the adipose and fibroglandular tissues, even by solely using either the high- or low-kVp images.

Although the multi-Otsu thresholding algorithm broadly identifies the calcification regions, these regions does not accurately match the true calcification map. We thus use the extracted calcification regions as a guidance for step 2, the supervised ML model to perform multi-material decomposition.

Step 2: Supervised material decomposition

We use deep learned convolutional neural networks, specifically the U-Net model [22], to perform supervised multi-material decomposition. The U-Net model is a well established deep learning model that performs biomedical image segmentation. We use U-Net to perform supervised regression to map the pixels in the images resulting from step 1 to the material maps of the tissues at the respective pixels. The input to the U-Net model is a 3D image with 3 channels (or a 3D matrix of size $512 \times 512 \times 3$): (1) the high-kVp image, (2) the low-kVp image, and (3) the high-density lesion regions extracted using step 1. The high-density lesion regions (calcification in the current dataset) in the high- and low-kVp images are replaced with the values of the background, which is air in the current dataset, as shown in Fig. 3. The details of the U-Net model architecture for the current dataset are summarized below:

- Input: 3 channels of 2D images with 512×512 resolution
 1. high kVp images with calcification regions replaced by air
 2. low kVp images with calcification regions replaced by air
 3. calcification regions extracted using step 1
- Output: 3 channels of 2D images with 512×512 resolution representing mapping of the material composition of the 3 tissues.
- Hidden layer(s): U-Net model with mostly the default architecture used in [22]
 - The encoder and decoder have layers of channel sizes $\{64, 128, 512, 1024\}$
 - ReLU activation function is used in all layers.
 - We added a sigmoid function at the output layer to enforce the outputs to be in the range of $[0, 1]$.

We call our physics-guided ML model as UNetPhy. For comparison of the UNetPhy model with a baseline image-based ML model, we construct a U-Net model without the physics-guidance from step 1. This U-Net only model, herein called the UNet model, takes only 2 inputs, the raw high- and low-kVp images (without the calcification regions replaced by air), and the tissue maps as the outputs. The UNet model represents canonical image-based multi-material ML models which map the DECT

images to the material densities without any physics information embedded in the modeling framework.

Training the U-Net models and hyperparameter search

We use the same training setup and model hyperparameters for both U-Net models, UNet and UNetPhy. For training the U-Net models, we use the NAdam optimizer [11, 4, 16] for optimizing the weights of the neural network based on an L_2 loss of the predicted material maps against the ground truth maps. The L_2 loss function (mean-squared error, MSE) corresponds to that used for solving regression problems. We also tested other loss functions such as cross entropy, binary cross-entropy (BCE), BCE with logits, Dice coefficient, and a combination of Dice and BCE. For the regression problem, the MSE results in the best optimization of the network parameters. The dataset (both inputs and outputs) is not scaled as the pixel values are in the range of $[0, 1]$. From the total number of images available, 95% is used for training and the rest for testing the model. The training set is split at a 70 : 30 ratio for training and validation. The training images are randomly shuffled during each epoch of training.

As mentioned in the model setup description, we mostly use the default setting of the U-Net model from [22]. Moreover, the first layer of the default model has 64 channels, which might be too large of a convolution from the input data having only 3 channels. We customized the layers in the U-Net to study the effect of making the model shallower. We varied the channel size of the first hidden layer from 4 to 64 and the total number of layers from 3 to 6. We noted the accuracy to increase with the size of the first layers channel size as well as the total number of hidden layers. We choose a model with channels $\{64, 128, 512, 1024, 2048, 4096\}$ to perform the modeling.

Choosing the regression loss and the sigmoid activation function at the output layer enables us to use the U-Net model for material decomposition, which is generally used as an image segmentation model. We also use the same architecture to perform segmentation of the images to the 3 tissues. The only change that is required is to convert the material density values of the output maps to binary values for each tissue at each pixel. This is done by simply rounding each density value to the nearest integer, 0 or 1. Thus, using the same model architecture we can perform both multi-material decomposition and image segmentation. We tested the optimization with the different loss functions mentioned above. The loss functions MSE and BCE gave the best results.

Since U-Net architectures can get prohibitively expensive to train with deeper layers, there is a need to accelerate training using parallel techniques and high-performance computing. This is more pertinent when expanding the current modeling framework to 3D images. The current 2D image dataset and U-Net model takes about 32 secs per epoch to train on a single NVIDIA V100 GPU. Regardless of our current image dataset and model being computationally tractable in a single GPU and not requiring parallel training routines, we demonstrate the capability to scale the training algorithm in the current study. We use the Distributed Data-Parallel training technique by PyTorch to perform multi-GPU, multi-node training of the U-Net model on the Summit supercomputer hosted by the Oak Ridge Leadership Computing Facility. We are able to achieve strong scaling of the training algorithm up to 128 nodes on Summit (6 NVIDIA V100 GPUs per node) as demonstrated in Fig. 4.

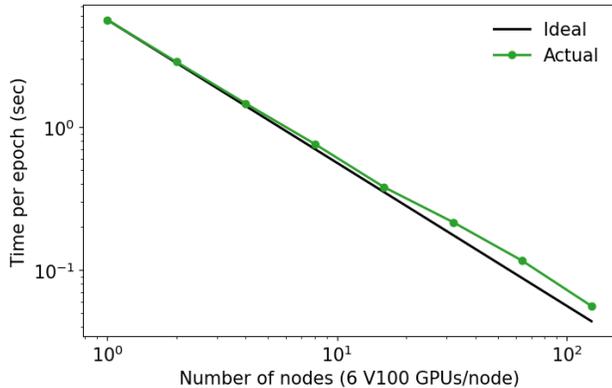


Figure 4. Strong scaling of the training algorithm for U-Net using Distributed Data-Parallel technique of PyTorch.

The batch size (or mini-batch size) is another key hyperparameter which determines the speed as well as accuracy of training. We observe that, as expected, increasing the batch size increases the number of epochs required to converge to model parameters with good accuracy. The number of nodes used in the data-parallel training also contributes to this convergence. For all the results in the current study, we used 8 nodes and a batch size of 1.

Lastly, the amount of training sample is also crucial for training deeper complex models. During training, we observe artifacts in the training and validation curves such as high-fluctuations, large gaps between the two curves, and validation loss lower than training loss. These suggest the amount of samples might not be representative of the problem. Since the data is isotropic in nature without a preferred direction, more training data can be generated by augmenting the images through rotations. For the current study, we augmented the dataset by simple 90° rotations, generating a total of 2000 images in the dataset. Although this helped reduce the artifacts in training curve, the overall accuracy of the model did not increase dramatically (less than 7% increase in accuracy).

Results

The results for evaluating a sample test case from the dataset are shown in Fig. 5. This test case was not used for training the U-Net models. We see from Figs. 5(a-b) that qualitatively, the physics-guided ML model UNetPhy can effectively decompose the tissue maps. The total root mean square error (RMSE) is of the order $\mathcal{O}(10^{-2})$ and the normalized root mean square error (NRMSE) $\|y - \hat{y}\|_2 / \|y\|_2 = 2.75\%$. The difference between the ground truth map and the UNetPhy model shown in Fig. 5(c) highlights the boundary pixels of each tissues map to accumulate the largest error.

The comparison of the ground truth and the UNet model (UNet without physics guidance) are shown in Fig. 5(d). The difference images depict the higher accumulation of error by the UNet model compared to the UNetPhy model, particularly in the tissue boundaries and within the tissue regions. Many dense “spot” like artifacts are found within the adipose tissue (left and top parts of Fig. 5(d)-left) and the fibroglandular tissue (center part of Fig. 5(d)-center). We find a relative increase in the NRMSE er-

ror by 18% for the UNet model compared to the UNetPhy model. These results demonstrate the need and advantage of the framework of the UNetPhy model, that adds the dense material information into the inputs.

For the segmentation problem using the same modeling framework as UNetPhy but the modified outputs, we obtained similar results with a normalized RMSE of 6.75%. Similar observations regarding the hyperparameter adjustments as with the regression problem are also observed for the segmentation problem.

We note that the L_2 errors of both the material decomposition ($\sim 2\%$) and segmentation ($\sim 6\%$) ML models are slightly higher compared to other projection-based methods. The source of error could be on the lack of sinogram-based physics information and the high artifacts in the FBP reconstructed images. Nonetheless, we emphasize the advantage of our method to only rely on the CT reconstructed images and not on the sinograms, which enables accessible usage of our method by the general public.

Concluding remarks

We introduce a physics-guided ML framework to perform image-based multi-material decomposition of dual-energy CT scans comprising of more than 2 types of materials. The methodology involves a combination of unsupervised and supervised techniques, which is novel for image-based material decomposition methods in medical images. The physics guidance using the difference of the dual-energy images to extract high-density materials has been very valuable to the current framework. Moreover, we also use the same framework to demonstrate image segmentation, thus solving 2 problems with the same ML framework.

One crucial aspect that the accuracy of image-based methods depends on is the quality of the reconstructed images [24], which is hindered by the artifacts from the FBP reconstruction. We are exploring other image reconstruction algorithms to reconstruct the images from the transmission data provided by the organizers of the AAPM challenge. We hope that images with reduced artifacts would improve the accuracy of the overall ML multi-material decomposition framework.

Acknowledgments

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

References

- [1] C. A. Bouman. *Foundations of Computational Imaging: A Model-Based Approach*. SIAM, 2022.
- [2] T. M. Buzug. *Computed tomography*. Springer, 2011.

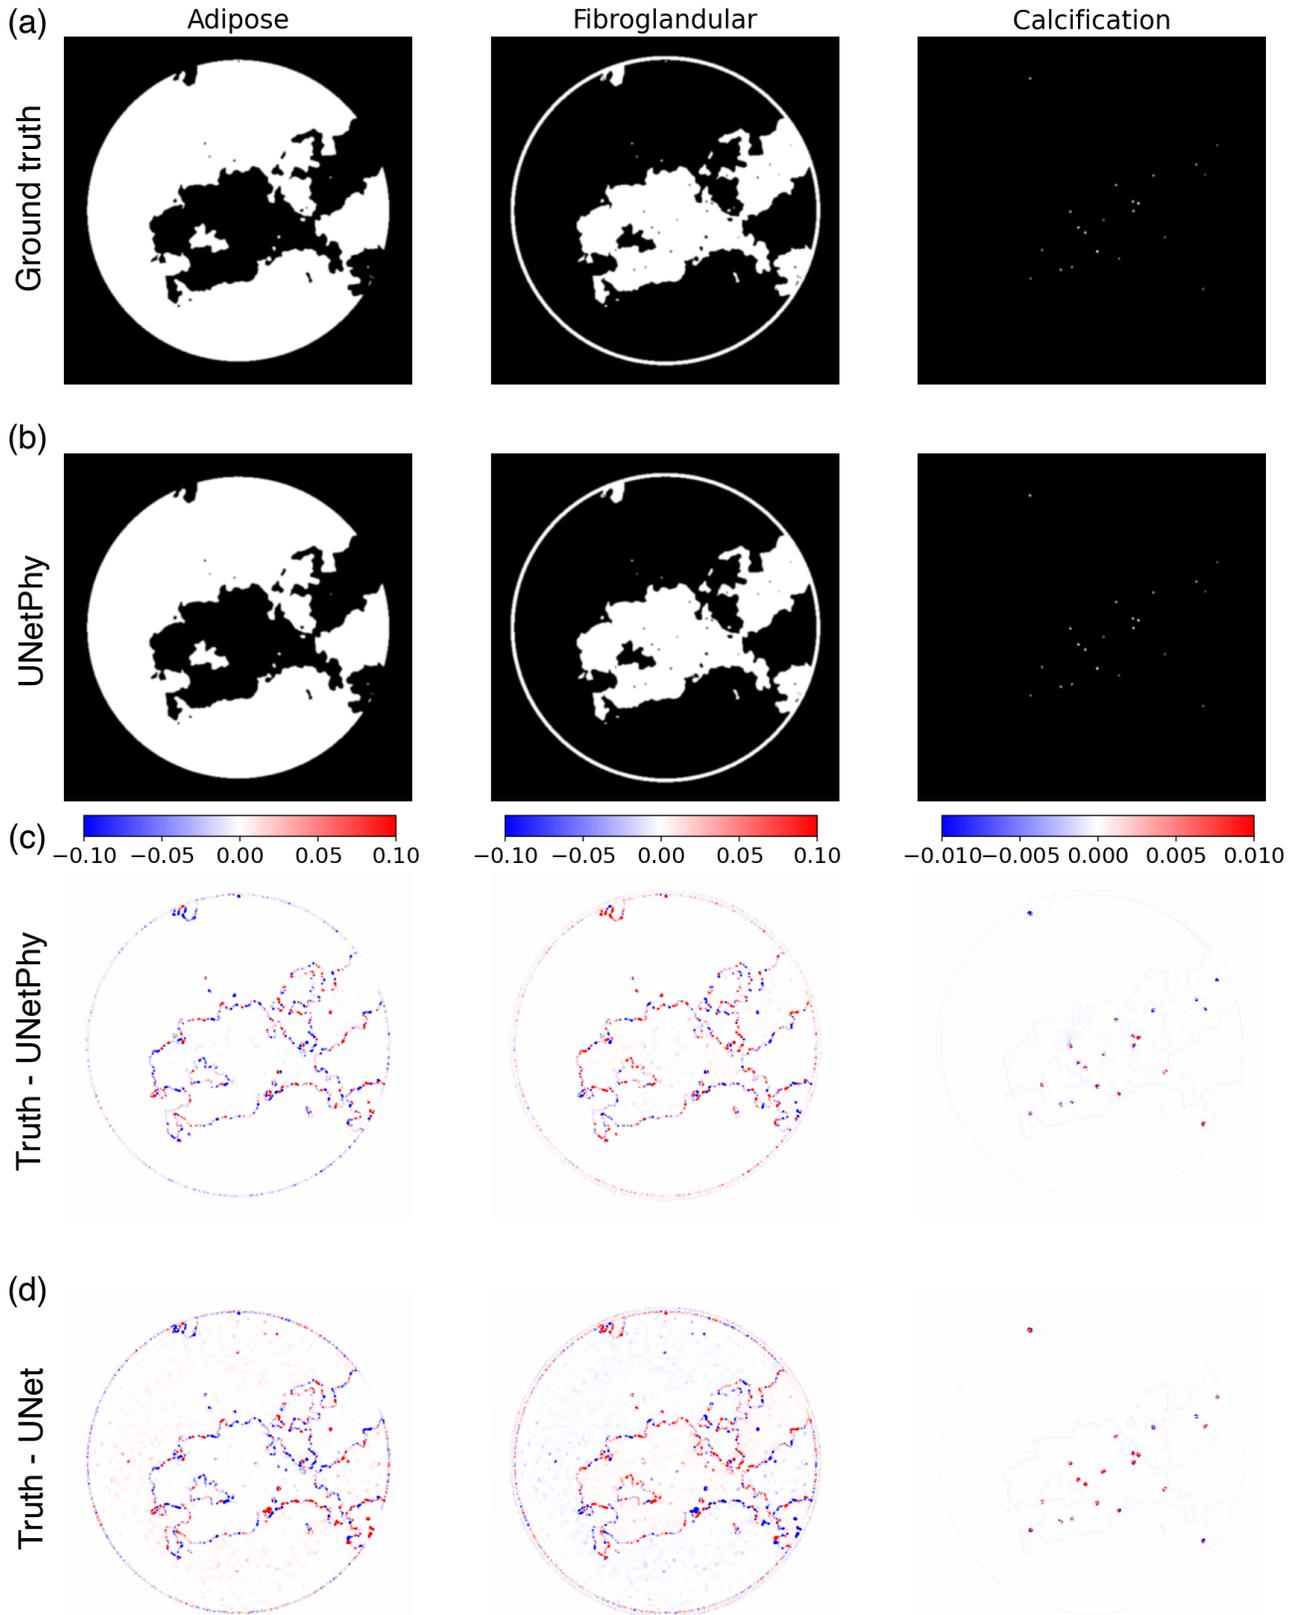


Figure 5. Sample results of evaluating a test case from the dataset. (a) Ground truth tissue maps. (b) Tissue maps decomposed using the current framework, the physics-guided UNet model (UNetPhy). (c) Difference between the ground truth and UNetPhy maps. (d) The difference computed between the ground truth and the output of the UNet model with no physics guidance, which uses only the high- and low-kVp images as inputs.

- [3] S. Degirmenci, D. G. Politte, C. Bosch, N. Tricha, and J. A. O'Sullivan. Acceleration of iterative image reconstruction for x-ray imaging for security applications. In C. A. Bouman and K. D. Sauer, editors, *SPIE Proceedings*. SPIE, Mar. 2015.
- [4] T. Dozat. Incorporating Nesterov momentum into Adam. In *International Conference on Learning Representations*, 2016.
- [5] B. Ghamraoui et al. Fabrication of microcalcifications for insertion into phantoms used to evaluate x-ray breast imaging systems. *Biomedical Physics & Engineering Express*, 7(5):055021, 2021.
- [6] J. Gibbs, K. A. Mohan, E. Gulsoy, A. Shahani, X. Xiao, C. Bouman, M. De Graef, and P. Voorhees. The three-dimensional morphology of growing dendrites. *Scientific reports*, 5(1):11824, 2015.
- [7] J. R. Grajo, M. Patino, A. Prochowski, and D. V. Sahani. Dual energy ct in practice: basic principles and applications. *Appl Radiol*, 45(7):6–12, 2016.
- [8] T. Hyodo, N. Yada, M. Hori, O. Maenishi, P. Lamb, K. Sasaki, M. Onoda, M. Kudo, T. Mochizuki, and T. Murakami. Multimaterial decomposition algorithm for the quantification of liver fat content by using fast-kilovolt-peak switching dual-energy ct: clinical evaluation. *Radiology*, 283(1):108–118, 2017.
- [9] P. Jin, C. A. Bouman, and K. D. Sauer. A method for simultaneous image reconstruction and beam hardening correction. In *2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC)*, pages 1–5. IEEE, 2013.
- [10] P. Jin, E. Haneda, C. A. Bouman, and K. D. Sauer. A model-based 3d multi-slice helical ct reconstruction algorithm for transportation security application. In *Second International Conference on Image Formation in X-Ray Computed Tomography*, 2012.
- [11] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [12] S. J. Kisner, E. Haneda, C. A. Bouman, S. Skatter, M. Kourinny, and S. Bedford. Model-based ct reconstruction from sparse views. In *Second International Conference on Image Formation in X-Ray Computed Tomography*, pages 444–447, 2012.
- [13] P.-S. Liao et al. A fast algorithm for multilevel thresholding. *J. Inf. Sci. Eng.*, 17(5):713–727, 2001.
- [14] X. Liu et al. Quantitative imaging of element composition and mass fraction using dual-energy ct: Three-material decomposition. *Medical physics*, 36(5):1602–1609, 2009.
- [15] Y. Long and J. A. Fessler. Multi-material decomposition using statistical image reconstruction for spectral CT. *IEEE Transactions on Medical Imaging*, 33(8):1614–1626, Aug. 2014.
- [16] J. Ma and D. Yarats. Quasi-hyperbolic momentum and Adam for deep learning. In *International Conference on Learning Representations*, 2019.
- [17] C. Maaß, M. Baer, and M. Kachelrieß. Image-based dual energy ct using optimized precorrection functions: A practical new approach of material decomposition in image domain. *Medical physics*, 36(8):3818–3829, 2009.
- [18] P. R. Mendonça et al. A flexible method for multi-material decomposition of dual-energy ct images. *IEEE transactions on medical imaging*, 33(1):99–116, 2013.
- [19] K. A. Mohan, S. Venkatakrishnan, J. W. Gibbs, E. B. Gulsoy, X. Xiao, M. De Graef, P. W. Voorhees, and C. A. Bouman. Timbir: A method for time-space reconstruction from interlaced views. *IEEE Transactions on Computational Imaging*, 1(2):96–111, 2015.
- [20] R. A. Nasirudin et al. A comparison of material decomposition techniques for dual-energy ct colonography. In *Medical Imaging 2015: Physics of Medical Imaging*, volume 9412, pages 1102–1107. SPIE, 2015.
- [21] I. Reiser and R. Nishikawa. Task-based assessment of breast tomosynthesis: effect of acquisition parameters and quantum noise a. *Medical physics*, 37(4):1591–1600, 2010.
- [22] O. Ronneberger et al. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [23] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Transactions on Signal Processing*, 41(2):534–548, 1993.
- [24] E. Y. Sidky and X. Pan. Report on the aapm deep-learning spectral ct grand challenge. *arXiv preprint arXiv:2212.06718*, 2022.
- [25] J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh. A three-dimensional statistical approach to improved image quality for multislice helical CT. *Medical Physics*, 34(11):4526–4544, Oct. 2007.
- [26] X. Wang, K. A. Mohan, S. J. Kisner, C. Bouman, and S. Midkiff. Fast voxel line update for time-space image reconstruction. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1209–1213. IEEE, 2016.
- [27] X. Wang, A. Sabne, S. Kisner, A. Raghunathan, C. Bouman, and S. Midkiff. High performance model based image reconstruction. *ACM SIGPLAN Notices*, 51(8):1–12, Feb. 2016.
- [28] T. Xia et al. Dual energy ct for attenuation correction with pet/ct. *Medical physics*, 41(1):012501, 2014.
- [29] Y. Xue, Y. Jiang, C. Yang, Q. Lyu, J. Wang, C. Luo, L. Zhang, C. Desrosiers, K. Feng, X. Sun, et al. Accurate multi-material decomposition in dual-energy ct: A phantom study. *IEEE Transactions on Computational Imaging*, 5(4):515–529, 2019.