

LEARNING-BASED MATERIAL DECOMPOSITION IN DUAL ENERGY CT USING AN UNROLLED ESTIMATOR

Megan Lantz, Greg Ongie

Marquette University
Department of Mathematics and Statistical Sciences
Milwaukee, WI, USA

ABSTRACT

Reconstructing multiple material-specific images in dual-energy CT is a challenging non-linear inverse problem. Traditionally, the reconstruction process consists of two steps: material decomposition and tomographic reconstruction. Model-based iterative reconstruction methods that combine material decomposition and tomographic reconstruction into a unified “one-step” framework provide improved estimates, but require longer reconstruction times. To address this, we propose a supervised machine learning technique to accelerate one-step iterative dual-energy CT reconstruction. Specifically, we train a deep neural network embedded in an “unrolling” of a model-based iterative algorithm. We demonstrate this approach on the problem of identifying three materials (adipose tissue, fibroglandular tissue, and calcifications) from simulated breast dual-energy CT data. Empirically, we find that the unrolling approach gives accurate material map estimates in few iterations, and outperforms a baseline image-domain learning approach.

Index Terms— dual-energy CT, material decomposition, model-based iterative reconstruction, machine learning

1. INTRODUCTION

Current medical practice relies heavily on mammography for breast imaging. However, as many as 43% of women ages 40-74 have breast tissue that contains a large proportion of fibroglandular tissue (dense breast tissue), and mammography is known to be less effective in detecting calcifications for these women [1]. Conventional single-energy computed tomography (CT) is similarly limited in its ability to differentiate between adipose (fatty) tissue, fibroglandular tissue, and calcifications, even with the use of dedicated breast CT systems [2]. Using dual-energy CT (DECT) to reconstruct distinct material maps for adipose tissue, fibroglandular tissue, and calcifications offers potential diagnostic improvement.

However, reconstructing material maps from DECT transmission data is a challenging non-linear – and often ill-posed – inverse problem. Traditionally, reconstruction is performed

by a two-step procedure that separates the problems of tomographic reconstruction and materials decomposition, but this approach can introduce significant artifacts in the presence of noise and incomplete data acquisition. Instead, several “one-step” model-based iterative reconstruction methods have also been proposed [3, 4, 5]. These approaches require optimizing a large-scale non-convex (and often non-smooth) cost function, which presents practical challenges. Recently, extensions of the alternating direction method of multipliers (ADMM) to non-convex objectives have been investigated for one-step multi-energy CT reconstruction [6, 7, 8, 9]. These approaches have shown success in material map estimation, but require many iterations for accurate reconstructions, leading to long reconstruction times.

We investigate a novel deep learning approach inspired by the non-convex ADMM algorithm proposed in [9]. Specifically, we train a network that estimates material maps from noisy/undersampled dual-energy transmission data, which is then embedded inside the ADMM iterations as a regularization step. This estimator (initialization layer plus a fixed number of unrolled iterations of ADMM) is trained end-to-end in a supervised manner. Our main technical contribution is to show that the non-convex ADMM updates can be back-propagated through using implicit differentiation, allowing for end-to-end training. The proposed approach is shown to improve the accuracy of material maps obtained using a baseline image-domain learning-based approach on simulated data.

Related work Several recent works have investigated machine learning techniques for materials decomposition in dual-energy CT. Image-domain approaches include training a convolutional neural network to directly estimate materials decomposition from attenuation images [10, 11], and sparse transform learning [12]. Our approach is similar to other dual-domain approaches, such as [13], which proposed training separate networks to restore data in projection and image domain, and [14], which uses a learned generative model as a prior. Different from these works, our approach is based on approximating a maximum likelihood estimate using an unrolling of non-convex ADMM [9].

2. PROBLEM FORMULATION

In this work, we assume a rapid kV switching DECT acquisition where high energy ($80kV$) and low energy ($50kV$) projection views are interleaved. We model the acquired DECT transmission counts $c_{k,\ell}$ $k \in \{80kV, 50kV\}$, along ray ℓ as $c_{k,\ell} \sim \text{Poisson}(\hat{c}_{k,\ell})$ where

$$\hat{c}_{k,\ell} = \sum_{i=1}^{E_k} s_{k,i} \exp \left(- \sum_{m=1}^M \mu_{k,i,m} y_{m,\ell} \right), \quad (1)$$

where $y_{m,\ell}$ is the projection of the m th material map along ray ℓ , $\mu_{k,i,m}$ is the attenuation coefficient of material m for photons at energy i and energy window k , and $s_{k,i}$ is the incident photon spectral density at energy i . To make the dependence on the material maps $\mathbf{x} \in \mathbb{R}^{MN}$ explicit, we define $\hat{\mathbf{c}}(\mathbf{y})$ to be the vector of transmission counts $\hat{c}_{k,\ell}$ at all energies k and all rays ℓ , and let $\mathbf{y} = \mathbf{P}\mathbf{x}$, where $\mathbf{P}$ is the discrete fan-beam projection applied separately to each material map.

The maximum likelihood estimate of the material maps given dual-energy transmission counts $\mathbf{c}$ is the solution of the optimization problem

$$\min_{\mathbf{x}} L(\mathbf{P}\mathbf{x}) \quad (2)$$

where $L(\mathbf{y}) = L(\mathbf{y}; \mathbf{c})$ is the Poisson negative log-likelihood:

$$L(\mathbf{y}) = \sum_{k,\ell} \left(\hat{c}_{k,\ell}(\mathbf{y}) - c_{k,\ell} - c_{k,\ell} \log \left(\frac{\hat{c}_{k,\ell}(\mathbf{y})}{c_{k,\ell}} \right) \right)$$

Note that we may write $L(\mathbf{y}) = L_c(\mathbf{y}) + L_d(\mathbf{y}) + C$ where

$$L_c(\mathbf{y}) = \sum_{k,\ell} \hat{c}_{k,\ell}(\mathbf{y}) \text{ and } L_d(\mathbf{y}) = - \sum_{k,\ell} c_{k,\ell} \log(\hat{c}_{k,\ell}(\mathbf{y})).$$

with L_c convex, and L_d non-convex but smooth, and C is a constant independent of $\mathbf{y}$.

3. METHODS

Our goal is to enhance an ADMM algorithm designed to optimize (2) with supervised learning techniques. Here we briefly discuss the form the ADMM updates take for this problem, and how these updates may be backpropagated through. Then we discuss how to train an initialization/regularization network for the ADMM algorithm.

3.1. One-step iterative reconstruction via ADMM

Here we summarize the steps involved in applying the non-convex ADMM algorithm proposed in [9] to minimize the non-convex loss (2). First, we recast (2) as a constrained optimization problem:

$$\min_{\mathbf{x}, \mathbf{y}} L(\mathbf{y}) \text{ s.t. } \mathbf{P}\mathbf{x} = \mathbf{y}.$$

Consider the associated augmented Lagrangian:

$$F(\mathbf{x}, \mathbf{y}, \mathbf{u}) = L(\mathbf{y}) + \langle \mathbf{P}\mathbf{x} - \mathbf{y}, \mathbf{u} \rangle + \frac{1}{2} \|\mathbf{P}\mathbf{x} - \mathbf{y}\|_{\Sigma}^2$$

where $\mathbf{u}$ is a vector of Lagrange multipliers, and $\|\mathbf{z}\|_{\Sigma} := \mathbf{z}^{\top} \Sigma \mathbf{z}$ where Σ is a diagonal ‘‘step-size’’ matrix. Following [9], we choose $\Sigma = \text{diag}(\mathbf{P}\mathbf{1}/\sigma)^{-1}$ where $\sigma > 0$ is a user-set parameter, and $\mathbf{1}$ is a vector of all ones. The non-convex ADMM algorithm of [9] proceeds by cycling between minimizing F with respect to $\mathbf{x}$ and $\mathbf{y}$, with the other variables fixed, followed by an update of the dual variable $\mathbf{u}$. However, when updating the $\mathbf{y}$ variable, the non-convex function $L(\mathbf{y})$ is replaced with a convex surrogate given by linearizing the smooth part about the current iterate $\mathbf{y} = \mathbf{y}^{(k)}$:

$$L(\mathbf{y}) \approx L_c(\mathbf{y}) + L(\mathbf{y}^{(k)}) + \langle \nabla L_d(\mathbf{y}^{(k)}), \mathbf{y} - \mathbf{y}^{(k)} \rangle + C.$$

Additionally, the $\mathbf{x}$ -update can be simplified by adding in a proximal term of the form $\|\mathbf{x} - \mathbf{x}^{(k)}\|_{\mathbf{H}}^2$ for a specific choice of positive definite $\mathbf{H}$ (see [9] for details). Together, this yields the non-convex ADMM updates

$$\mathbf{y}^{(k+1)} = \arg \min_{\mathbf{y}} \{ L_c(\mathbf{y}) + \langle \mathbf{y}, \nabla L_d(\mathbf{y}^{(k)}) - \mathbf{u}^{(k)} \rangle + \frac{1}{2} \|\mathbf{P}\mathbf{x}^{(k)} - \mathbf{y}\|_{\Sigma}^2 \} \quad (3)$$

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \mathbf{Q}^{-1} \mathbf{P}^{\top} (\Sigma(\mathbf{y}^{(k+1)} - \mathbf{P}\mathbf{x}^{(k)}) - \mathbf{u}^{(k)}) \quad (4)$$

$$\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} + \Sigma(\mathbf{P}\mathbf{x}^{(k+1)} - \mathbf{y}^{(k+1)}) \quad (5)$$

where $\mathbf{Q} = \text{diag}(\sigma \mathbf{P}^{\top} \mathbf{1})$ (the role of the $\mathbf{Q}$ matrix is discussed further in [9]). The $\mathbf{y}$ -update (3) has no closed-form solution, but the objective is smooth, strongly convex, and separable over all energy/ray pairs (k, ℓ) , and is optimized efficiently by running Newton iterations in parallel [9].

As an additional regularization step, we propose adding an intermediate $\mathbf{x}$ -update following the $\mathbf{y}$ -update (3) of the form

$$\bar{\mathbf{x}}^{(k)} = R_{\theta}(\mathbf{y}^{(k+1)}) \quad (6)$$

where R_{θ} is a learned regularization network (described below), such that $\bar{\mathbf{x}}^{(k)}$ is used in place of $\mathbf{x}^{(k)}$ in (4).

3.2. Backpropagating through ADMM updates

The ADMM $\mathbf{y}$ -update (3), which is defined implicitly as the solution to an optimization problem, does not have a simple closed-form expression. Despite this fact, we show how to differentiate the $\mathbf{y}$ -update with respect to the ADMM state variables $(\mathbf{x}, \mathbf{y}, \mathbf{u})$. First, observe that the $\mathbf{y}$ -update requires solving an optimization problem of the form

$$\min_{\tilde{\mathbf{y}}} \underbrace{L_c(\tilde{\mathbf{y}}) + \langle \tilde{\mathbf{y}}, \mathbf{v} \rangle + \frac{1}{2} \|\tilde{\mathbf{y}}\|_{\Sigma}^2}_{F(\tilde{\mathbf{y}}, \mathbf{v})}$$

where $\mathbf{v} = \nabla L_d(\mathbf{y}) - \mathbf{u} - \Sigma \mathbf{P} \mathbf{x}$, with $(\mathbf{x}, \mathbf{y}, \mathbf{u})$ being the current ADMM state variables. This optimization problem is smooth and strongly convex, thus has a unique minimizer $\mathbf{y}^*$ characterized by the optimality condition $\nabla_{\mathbf{y}} F(\mathbf{y}^*, \mathbf{v}) = \mathbf{0}$. This is equivalent to the fixed-point equation

$$\mathbf{y}^* = -\Sigma^{-1} \mathbf{v} - \Sigma^{-1} \nabla L_c(\mathbf{y}^*) := \mathbf{f}(\mathbf{y}^*, \mathbf{v}). \quad (7)$$

Using this fact, we may compute the derivatives of $\mathbf{y}^*$ with respect to $\mathbf{v}$ (and hence with respect to the ADMM state variables) using implicit differentiation. In particular, backpropagation requires vector-Jacobian products (VJPs) of the form $\mathbf{w}^\top \left(\frac{\partial \mathbf{y}^*}{\partial \mathbf{v}} \right)$ where $\mathbf{w}$ is an arbitrary vector. Differentiating both sides of the fixed-point equation (7) we may deduce

$$\frac{\partial \mathbf{y}^*}{\partial \mathbf{v}} = \left(\mathbf{I} - \left[\frac{\partial \mathbf{f}}{\partial \mathbf{y}^*} \right] \right)^{-1} \frac{\partial \mathbf{f}}{\partial \mathbf{v}}$$

and so the (transpose of) VJPs with any vector $\mathbf{w}$ is given by

$$\left(\frac{\partial \mathbf{y}^*}{\partial \mathbf{v}} \right)^\top \mathbf{w} = \left(\frac{\partial \mathbf{f}}{\partial \mathbf{v}} \right)^\top \left(\mathbf{I} - \left[\frac{\partial \mathbf{f}}{\partial \mathbf{y}^*} \right]^\top \right)^{-1} \mathbf{w}. \quad (8)$$

The right-hand side of (8) reduces to finding the solution $\mathbf{q}$ of the linear system $(\Sigma + \nabla^2 L_c(\mathbf{y}^*)) \mathbf{q} = \mathbf{w}$ where $\nabla^2 L_c(\mathbf{y}^*)$ denotes the Hessian of L_c at $\mathbf{y}^*$. This system decouples over the ray index ℓ resulting in a collection of n_{rays} independent $M \times M$ linear systems (where M is the number of materials) that are solvable in parallel. Therefore, once the update $\mathbf{y}^*$ is obtained, computing VJPs as needed for backpropagation involves a straightforward linear solve.

3.3. Transmissions-to-Materials (T2M) Network

In order to initialize and regularize the proposed ADMM unrolling, we train a “transmissions-to-materials” (T2M) network $M_\theta(\cdot)$ that takes as input the DECT transmission counts $\mathbf{c}$ and outputs an estimate of the material maps $\mathbf{x}$. Specifically, the T2M network performs two operations: first, it computes an estimate of the high/low energy attenuation images by applying filtered back projection to the negative log of the high/low transmission data $\mathbf{c}$ (zero-filling any missing views in projection domain), then it applies a trainable image-to-image fully convolutional neural network $U_\theta(\cdot)$ to estimate a materials decomposition; in our experiments we use a U-net architecture [15] for U_θ . Put in symbols, the T2M network has the form: $M_\theta(\mathbf{c}) = U_\theta(\text{FBP}(-\log(\mathbf{c})))$.

We convert the T2M network into a regularization network R_θ (as used in the ADMM $\bar{\mathbf{x}}$ -update (6)) by simulating transmission data from the current projection-domain estimate $\mathbf{y}$, and passing this into the T2M network, i.e., we set $R_\theta(\mathbf{y}) = M_\theta(\hat{\mathbf{c}}(\mathbf{y}))$, where $\hat{\mathbf{c}}$ is defined in (1). This is similar to the “reuse-and-regularize” approach proposed in [16], where a pre-trained reconstruction network is reused as a regularizer in an iterative reconstruction algorithm.

3.4. End-to-end Training

Our unrolled ADMM estimator is constructed as follows: we initialize with $\mathbf{x}^{(0)} = M_\theta(\mathbf{c})$, $\mathbf{y}^{(0)} = \mathbf{P} \mathbf{x}^{(0)}$, and $\mathbf{u}^{(0)} = \mathbf{0}$, then run K iterations of ADMM with the embedded T2M network. The returned estimate of the material maps is $\hat{\mathbf{x}} = \mathbf{x}^{(K)}$, i.e., the final ADMM $\mathbf{x}$ -update after K iterations. We write this estimator as $\hat{\mathbf{x}} = \text{ADMM}_\theta^{(K)}(\mathbf{c})$, where θ are the trainable parameters of the embedded T2M network.

Given a training set $\{(\mathbf{x}_i, \mathbf{c}_i)\}_{i=1}^n$ where $\mathbf{x}_i$ is a set of ground truth material maps $\mathbf{x}_i$, and $\mathbf{c}_i$ is the corresponding DECT transmission data, we train the unrolled estimator by attempting to minimize the mean-squared error over the training set:

$$\min_{\theta} \sum_{i=1}^n \|\text{ADMM}_\theta^{(K)}(\mathbf{c}_i) - \mathbf{x}_i\|_2^2$$

Using the implicit differentiation routine outlined in Sec. 3.2, the unrolled ADMM estimator is fully differentiable with respect to the T2M network parameters θ , and gradients of the above loss function can be computed using standard autodifferentiation tools.

4. EXPERIMENTS

In our experiments, we use the simulated alternating view DECT transmission data based on breast CT phantoms provided with the AAPM Spectral Computed Tomography Image Reconstruction Challenge [17]. The goal of the AAPM Spectral CT challenge was to reconstruct three material maps (adipose tissue, fibroglandular tissue, and calcifications) from alternating view DECT transmission data in the absence of noise. Our work extends this by additionally considering simulated Poisson noise in the transmission counts.

The accuracy of each material map reconstructions is assessed using the root mean square error (RMSE), defined as $\|\hat{\mathbf{x}} - \mathbf{x}\|_2 / \sqrt{N}$, where $\hat{\mathbf{x}}$ is the reconstructed material maps, $\mathbf{x}$ is the ground truth material maps, and N is the total number of pixels in one set of three material maps. Additionally, the Poisson negative log-likelihood cost (as defined in (2)) is also computed for each reconstruction. To compare experimental results across algorithms, the mean RMSE on a test set of 100 simulated ground truth phantoms is used.

We compare our unrolled estimator against two baseline approaches: a T2M Network (U-net) trained to map from high/low energy attenuation maps to three material maps, and an iterative ADMM algorithm that uses the T2M Network output as the initialization for an ADMM algorithm with generalized total variation (gTV) regularization (similar to [18]).

For the T2M network, a training set of 800 phantom images was used to train the network using a basic U-net architecture. For each training phantom, we generated a set of simulated dual energy transmissions with 512 views, 1024 detector pixels, and 50,000 photons/detector pixel. The transmission images were then used as input to the U-net model




Fig. 1. Material map estimates for a single test phantom (best viewed digitally).

	Test RMSE (mean $\pm$ std. dev.)	Avg. Recon Time (sec.)
T2M network	0.0216 ± 0.0025	0.02
ADMM w/TV reg. (100 iter)	0.0213 ± 0.0025	9.09
Unrolled ADMM (proposed)	0.0209 ± 0.0025	0.82

Table 1. Quantitative results on test set

which was trained over 50 epochs to minimize the MSE loss between the model generated material map reconstruction and the original phantom image. Applying the trained T2M network model to our test set of 100 phantom images yielded a test RMSE of 0.0216 ± 0.0025 and an average reconstruction time of 0.02 seconds.

The second baseline builds off of the pre-trained T2M network by using the network output as the initialization for 100 iterations of ADMM with gTV constraints. For the test set, this approach yields a mean test RMSE of 0.0213 ± 0.0025 , a slight improvement over the T2M network alone. The mean reconstruction time for this approach is substantially higher at 9.09 seconds per image.

For the proposed approach, we unrolled $K = 8$ iterations of ADMM with the embedded T2M network. Using pre-trained T2M network weights as the initialization, we re-trained the unrolled ADMM estimator on the same 800 image training set. After re-training, the unrolled ADMM estimator resulted in a mean test RMSE of 0.0209 ± 0.0025 and an average reconstruction time of 0.82 seconds. Using this approach we see a decrease in RMSE from both the T2M

network and the T2M initialized ADMM with gTV regularization with only a small increase in reconstruction time from the T2M network approach.

Figure 1 shows each of the reconstruction methods applied to a sample breast phantom from the test set. The difference images show that the proposed approach shows less error at the boundaries of the tissues. We also observe that some spurious dark spots present in the Adipose map of the two baselines approaches are suppressed with the proposed approach.

5. CONCLUSION

We propose training an unrolled iterative estimator for one-step DECT material decomposition based on a non-convex ADMM algorithm. The unrolled estimator contains a trainable “transmissions-to-materials” network acting as an initialization layer and regularizer. Our experiments on synthetic breast CT data show that end-to-end training of the unrolled estimator yields improved accuracy over the estimates provided by the “transmissions-to-materials” network alone and its iterative refinement with total variation constrained ADMM. In future work, we will explore alternative types of regularization networks inside the iterations of the non-convex ADMM algorithm, as well as the application of transfer learning approaches to enable the application of our methods to real data, which has recently shown to be effective in a spectral CT context [19].

6. COMPLIANCE WITH ETHICAL STANDARDS

This is a numerical simulation study for which no ethical approval was required.

7. ACKNOWLEDGMENTS

G. Ongie was supported by NSF CRII award CCF-2153371.

8. REFERENCES

- [1] M. Rebolj, V. Assi, A. Brentnall, D. Parmar, and S. W. Duffy, "Addition of ultrasound to mammography in the case of dense breast tissue: systematic review and meta-analysis," *British journal of cancer*, vol. 118, no. 12, pp. 1559–1570, 2018.
- [2] K. K. Lindfors, J. M. Boone, T. R. Nelson, K. Yang, A. L. Kwan, and D. F. Miller, "Dedicated breast CT: initial clinical experience," *Radiology*, vol. 246, no. 3, pp. 725, 2008.
- [3] I. A. Elbakri and J. A. Fessler, "Statistical image reconstruction for polyenergetic X-ray computed tomography," *IEEE Transactions on Medical Imaging*, vol. 21, no. 2, pp. 89–99, 2002.
- [4] Y. Long and J. A. Fessler, "Multi-material decomposition using statistical image reconstruction for spectral CT," *IEEE Transactions on Medical Imaging*, vol. 33, no. 8, pp. 1614–1626, 2014.
- [5] A. Sawatzky, Q. Xu, C. O. Schirra, and M. A. Anastasio, "Proximal ADMM for multi-channel image reconstruction in spectral X-ray CT," *IEEE Transactions on Medical Imaging*, vol. 33, no. 8, pp. 1657–1668, 2014.
- [6] R. F. Barber, E. Y. Sidky, T. G. Schmidt, and X. Pan, "An algorithm for constrained one-step inversion of spectral CT data," *Physics in Medicine & Biology*, vol. 61, no. 10, pp. 3784, 2016.
- [7] R. F. Barber and E. Y. Sidky, "MOCCA: mirrored convex/concave optimization for nonconvex composite functions," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 5006–5056, 2016.
- [8] T. G. Schmidt, R. F. Barber, and E. Y. Sidky, "A spectral CT method to directly estimate basis material maps from experimental photon-counting data," *IEEE Transactions on Medical Imaging*, vol. 36, no. 9, pp. 1808–1819, 2017.
- [9] R. F. Barber and E. Y. Sidky, "Convergence for nonconvex ADMM, with applications to CT imaging," *arXiv preprint arXiv:2006.07278*, 2020.
- [10] D. P. Clark, M. Holbrook, and C. T. Badea, "Multi-energy CT decomposition using convolutional neural networks," in *Medical Imaging 2018: Physics of Medical Imaging*. SPIE, 2018, vol. 10573, pp. 415–423.
- [11] W. Zhang, H. Zhang, L. Wang, X. Wang, X. Hu, A. Cai, L. Li, T. Niu, and B. Yan, "Image domain dual material decomposition for dual-energy CT using butterfly network," *Medical Physics*, vol. 46, no. 5, pp. 2037–2051, 2019.
- [12] Z. Li, S. Ravishankar, Y. Long, and J. A. Fessler, "DECT-MULTRA: Dual-energy CT image decomposition with learned mixed material models and efficient clustering," *IEEE Transactions on Medical Imaging*, vol. 39, no. 4, pp. 1223–1234, 2019.
- [13] R. Zhang, J. Zhou, and Z. Yu, "A cascaded deep-learning reconstruction method for sparse-view kV-switching dual-energy CT," in *Medical Imaging 2020: Physics of Medical Imaging*. SPIE, 2020, vol. 11312, pp. 499–505.
- [14] Z. Shi, H. Li, Q. Cao, Z. Wang, and M. Cheng, "A material decomposition method for dual-energy CT via dual interactive Wasserstein generative adversarial networks," *Medical Physics*, vol. 48, no. 6, pp. 2891–2905, 2021.
- [15] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [16] D. Gilton, G. Ongie, and R. Willett, "Model adaptation for inverse problems in imaging," *IEEE Transactions on Computational Imaging*, vol. 7, pp. 661–674, 2021.
- [17] E. Sidky, X. Pan, and S. Armato, "Deep learning for inverse problems: Spectral computed tomography image reconstruction," 2022, URL: <https://www.aapm.org/GrandChallenge/DL-spectral-CT/>.
- [18] T. G. Schmidt, B. A. Sammut, R. F. Barber, X. Pan, and E. Y. Sidky, "Addressing CT metal artifacts using photon-counting detectors and one-step spectral CT image reconstruction," *Medical Physics*, 2022.
- [19] J. F. Abascal, N. Ducros, V. Pronina, S. Rit, P.-A. Rodesch, T. Broussaud, S. Bussod, P. C. Douek, A. Hauptmann, S. Arridge, et al., "Material decomposition in spectral CT using deep learning: a Sim2Real transfer approach," *IEEE Access*, vol. 9, pp. 25632–25647, 2021.