Published in final edited form as:

Proc SPIE Int Soc Opt Eng. 2021 February ; 11595: . doi:10.1117/12.2582268.

Manifold Reconstruction of Differences: A Model-Based Iterative Statistical Estimation Algorithm With a Data-Driven Prior

Matthew Tivnan, J. Webster Stayman

Biomedical Engineering, Johns Hopkins University, Baltimore, MD

Abstract

Manifold learning using deep neural networks been shown to be an effective tool for building sophisticated prior image models that can be applied to noise reduction in low-dose CT. A manifold is a low-dimensional space that captures most of the variation in a class of data (e.g. chest CT images). This nonlinear reduction in dimensionality is one way to build up a sophisticated model of the features and structures present in image data. This nonlinear dimensionality reduction tends to learn features associated with the signal rather than the noise when provided with a large training dataset. However, the internal workings operation are typically not easy to understand. The results can be highly nonlinear and object-dependent. After training the model is fixed and during reconstruction is typically no check for fitness between the image estimates and the measured data. We propose a new iterative CT reconstruction algorithm, called Manifold Reconstruction of Differences (MRoD), which combines physical and statistical models with a data-driven prior based on manifold learning. The MRoD algorithm involves estimating a *manifold component*, approximating common features among all patients, and the *difference component* which has the freedom to fit the measured data. By applying a sparsity-promoting penalty to the difference image rather than a hard constraint to the manifold, the MRoD algorithm is able to reconstruct features which are not present in the training data. The difference component itself may be independently useful. While the manifold captures typical patient features (e.g. healthy anatomy), the difference image highlights patient-specific elements (e.g. pathology). In this work, we present a description of an optimization framework which combines trained manifold-based modules with physical modules. We present a simulation study using anthropomorphic lung data showing that the MRoD algorithm can both isolate differences between a particular patient and the typical distribution, but also provide significant noise reduction with less bias than a typical penalized likelihood estimator in composite manifold plus difference reconstructions.

1. INTRODUCTION

In the last 50 years, the state-of-the-art in CT reconstruction has progressed from analytical methods, to model-driven iterative methods, and, in recent years, data-driven methods with deep learning models have gained in popularity. In a continuous noiseless mathematical framework, filtered back-projection is the closed-form mathematical inversion of the radon transform, and there are equally exact extensions to fan-beam and cone-beam geometries. The transition to model-driven methods such as model-based iterative reconstruction (MBIR) or iterative image denoising was motivated, in large part, by the non-idealities of CT image formation caused by measurement noise. Low-dose CT, in particular, benefits

greatly from the use of a statistical estimation framework such as maximum likelihood (ML) because of the incorporation of a noise model via statistical weighting. Another important benefit of MBIR is the ability to incorporate prior information about the estimates via regularization such as a roughness penalty in penalized likelihood (PL) reconstruction. However, common regularizers such as quadratic and huber penalties can effectively lower the spatial resolution of the final estimates. This bias towards smooth images is introduced by the estimator, not the physical data acquisition, and it is characterized by the local impulse response. We claim that these roughness penalties are intuitively a poor approximation of the prior distribution. For example, an image containing only one constant value in all voxels will have a penalty of zero, since it is perfectly smooth, even though it bears little resemblance to the anatomic features that are relevant to medical imaging. There is another way to incorporate prior information in the special case where there is access to a previous scan of the same subject. In this case, MBIR can be used for reconstruction of differences (ROD) where the estimates represent a difference with respect to a reference image. Then, the difference image can be encouraged via a penalty term to be small in magnitude and sparse, thereby utilizing the previous image as prior information. In the past, a more advanced regularization which encourages anatomical image features without the need for a reference image seemed out of reach. However, with the advent of deep-learning-based data processing, in particular manifold learning, there is hope for a more sophisticated model of the prior distribution.

Typically, these methods are formulated as a maximum *a posteriori* (MAP) estimator with a negative log-posterior objective function which, according to Bayes rule, can be broken into a negative log-likelihood term, which represents data fidelity, summed with a negative log-prior term which promotes image properties which are considered to be likely *a priori* without reference to the data. The data fidelity term requires a physical forward model which relates the image estimates to the data, as well as a noise model. In this way the goodness-of-fit between image estimates and projection-domain CT data can be evaluated as a probability and optimized iteratively. If no prior information is enforced, that is, all potential image results are considered equally likely *a priori*, the MAP estimator reduces to a maximum likelihood (ML) estimator. Assuming a perfectly matched physical model, ML is an unbiased estimator. However, ML suffers greatly from noise, especially in low-dose CT or irregular sampling geometries. The prior model for the image estimates is often substituted for a crude approximation which promotes certain image properties such as a roughness penalty. This introduction of prior information tends to be effective increasing the signal to noise ratio by means of noise suppression, but it also introduces some bias. In the case of a roughness penalty, this bias is towards smooth images and, accordingly, this type of MBIR is typically labeled as penalized likelihood (PL) to distinguish it from a true MAP estimator for which the penalty term should be based on a prior probability distribution of the estimates. Until recently, a true characterization of the prior distribution of image estimates seemed intractable. However, this task may now be within reach using powerful new deep learning tools. Specifically, models for manifold learning are particularly well suited to this task.

Manifold learning is a form of data-driven nonlinear dimensionality reduction which has shown great promise in the task of low-dose CT reconstruction. Such learning involves

finding a latent representation of the image data in a lower-dimensional space than the original voxelized basis. This latent space, or manifold, can encode for a large portion of the variations within a certain class of images, and is coupled with a nonlinear mapping from the manifold back to the original image space. The transformation between the image space and the latent space is trained using, e.g., deep neural networks (DNN).¹ A convolutional architecture has found use in image denoising,² and has also been incorporated into full deep learning CT reconstruction models.³ The concept of operation is that the features which contribute to the signal are, in some sense, consistent for a certain class of images of interest, whereas the contribution from noise is inconsistent for different images. Therefore, the network will tend to learn to extract features associated with the signal, rather than the noise, since the features associated with the signal will be more useful, on average, for the full ensemble of images. This is analogous to repeated noisy acquisitions of the same measurement which would theoretically have the same signal, but different different noise realizations. Of course, the signal portion of the image is not identical for two CT images of different patients, so the goal is to find a manifold which can represent many images. For this reason, this data-driven approach requires a large training dataset which sufficiently samples the variety of images that make up the class of interest. First, there is an encoder stage, which transforms and compresses the image inputs to a low-dimensional representation which we call the manifold. Next, there is a decoder stage which maps the manifold representation back to the original image space. This model can be viewed as a lossy nonlinear compressor. Training the weights of the network amounts to finding the most efficient low-dimensional representation of the images which minimizes error due to compression. Typically, training cases are provided as both inputs and outputs to the model, so the goal is to train a process which makes the minimal change to the training images, despite the compressed representation at an intermediate layer. One of the major concerns regarding deep learning CT reconstruction is that the processing may be highly object-dependent and can have a highly nonlinear response to specific stimuli.⁴ Moreover, there is also a risk of over-fitting to the training data. For example, if a new patient has a lesion that was not represented in the training data, it may be difficult or impossible to accurately represent that feature if the image result is constrained to lie on the manifold. The net effect is that certain features that are present may not be visualized in the data, or vice-versa, false features may be injected into the image, but the overall image has the appearance of “good” image quality. Another critique of this type of hard constraint to the manifold is that there is no explicit check to ensure goodness of fit between image estimates and measured data. However, some efforts have been made to combine manifold-based noise reduction with a physics-based statistical estimators by alternating between a maximum likelihood reconstruction and a DNN-autoencoder-based denoising module.⁵

We propose a new iterative CT reconstruction algorithm, called Manifold Reconstruction of Differences (MRoD), which estimates images using a physics-based statistical data-fidelity term, and a manifold-based penalty term. These terms are combined into one penalized-likelihood estimator which is optimized directly. The MRoD algorithm involves jointly estimating two images: the *manifold component*, which contains features which are typical among all patients, and the *difference component* which contains the changes needed to fit the measured data. A final composite image estimate may be formed via sum of

these two components. A penalty is applied to minimize the difference component and promote sparsity. This effectively encourages image estimates which very close to manifold-constrained case without invoking a hard constraint. These small changes in the difference image are guided by a data-fidelity term which uses physics-based signal and noise models to evaluate the goodness-of-fit between the final images (a summation the two components) and the measured data. Therefore, the MRoD algorithm has the ability to reconstruct image features which may be rare or irregular and did not appear in the manifold training cases. Difference images themselves may independent, clinically useful outputs. For example, while the manifold component captures features which are typical across all patients (e.g. healthy anatomy), the difference highlights patient-specific atypical elements (e.g. pathology or rare variations within the patient population). It is often these differences that are of greatest interest to radiologists.

In this work, we present a mathematical description of the algorithm as well a description of the software implementation. This includes the combination of trained deep DNN modules, forward/back-projectors, and other physics-based modules in a unified iterative optimization framework. In a simulated lung imaging study, we characterize the effect of the penalty weight on noise and bias. Comparisons to maximum-likelihood and roughness-penalized likelihood estimators are also provided. Finally, we present a numerical experiment which involves reconstructing pathological cases. This demonstrates that the MRoD algorithm is capable of imaging clinically meaningful irregularities which were not present in the training data.

2. METHODS

2.1 Manifold Learning with a DNN Encoder-Decoder

We consider a DNN encoder-decoder model consisting of an encoder which compresses the image data to the manifold representation and a decoder which maps that manifold back to the original image space. If the N voxels comprising the image data are represented as a single vector $\mathbf{x} \in \mathfrak{R}^N$, then the encoder is a nonlinear operation $\mathbf{m} = \mathcal{E}(\mathbf{x}) \in \mathfrak{R}^M$ which maps the image data onto an M -dimensional manifold where $M < N$. The decoder is another nonlinear operation $\tilde{\mathbf{x}} = \mathcal{D}(\mathbf{m}) \in \mathfrak{R}^N$ which maps the manifold back to the image space. Trained weights for the DNN are optimized to minimize the L1-norm of the residual, $\|\mathbf{x} - \mathcal{D}(\mathcal{E}(\mathbf{x}))\|$. Fig. 1 shows the DNN encoder-decoder consisting of convolutional and fully connected layers with leaky ReLu activation functions. The input and output images have $256 \times 256 \times 1$ voxels and the intermediate manifold representation has $1 \times 1 \times 8192$ features. The encoder has 2×2 max pooling layers and the decoder has 2×2 up-sampling layers. In the following sections, we will incorporate this data-driven manifold as a penalty in an iterative CT reconstruction algorithm to encourage image features which were common in the training data.

Manifold learning requires a large dataset of images which sufficiently sample a class of interest (e.g. chest CT images). The weights of this deep learning model are then trained by providing sample images as both inputs and outputs to this network. The goal is to train the full model to behave approximately as an identity operator when applied to images from

this class of interest. This is challenging because there model is fundamentally structured as a lossy compressor with a low-dimensional representation at an intermediate layer halfway through the network.

2.2 Penalized-Likelihood Estimation Framework

The MRoD algorithm can be described mathematically as a penalized-likelihood estimator which includes a physics-based data-fidelity term and a manifold-based penalty. A general physical model for x-ray CT measurements is given by $\bar{\mathbf{y}}(\mathbf{x}) = \mathbf{B} \exp(-\mathbf{A}\mathbf{x})$, where $\mathbf{A}$ is a forward projector used to model the system geometry and $\mathbf{B}$ is a diagonal scaling operator used to model the gain. A penalized weighted least-squares objective function may be written as

$$\Phi(\mathbf{x}) = L(\mathbf{x} | \mathbf{y}) + R(\mathbf{x}), \quad L(\mathbf{x} | \mathbf{y}) = (\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x}))^T \Sigma_{\mathbf{y}}^{-1} (\mathbf{y} - \bar{\mathbf{y}}(\mathbf{x})) \quad (1)$$

The objective function $\Phi(\mathbf{x})$ has a weighted least-squares term, $L(\mathbf{x} | \mathbf{y})$, which quantifies the data fidelity with respect to the measured CT data, $\mathbf{y}$, and a penalty term $R(\mathbf{x})$ which serves as a regularizer. In this work, we will compare the new proposed manifold-regularized case with a standard quadratic roughness penalty given by $R(\mathbf{x}) = \beta \mathbf{x}^T \mathbf{R} \mathbf{x}$ where $\mathbf{R}$ is a discrete Laplace operator and β is the penalty strength. In this case, the penalty applies to the magnitude of the 2D gradient and thereby encourages smooth images. This prior can be effective at noise reduction by suppressing high frequencies. However, it can also introduce bias in the form of spatial blur at the interfaces between objects which lowers the effective resolution. The quadratic penalized likelihood estimator is the minimizer of the objective function, $\Phi(\mathbf{x} | \mathbf{y})$, using the quadratic roughness penalty. The next section proposes a second design for $R(\mathbf{x})$ and a corresponding estimator which utilizes a data-driven prior trained using manifold learning.

2.3 Manifold Reconstruction of Differences

The MRoD algorithm is inspired by the Reconstruction of Difference (RoD) algorithm⁶ which involves estimating the difference with respect to some reference image (e.g. from a previous scan). However, the MRoD algorithm does not require an *a priori* reference image. Instead, the reference image is replaced with a manifold-based approximation image which is jointly estimated together with the difference component. MRoD is a penalized-likelihood estimator with an L1-norm penalty to promote small sparse difference estimates. The proposed estimator can be formulated as follows.

$$\Phi(\mathbf{m}, \mathbf{x}_D) = L(\mathcal{D}(\mathbf{m}) + \mathbf{x}_D | \mathbf{y}) + \gamma |\mathbf{x}_D|, \quad \hat{\mathbf{m}}, \hat{\mathbf{x}}_D = \underset{\mathbf{m}, \mathbf{x}_D}{\operatorname{argmin}} \Phi(\mathbf{m}, \mathbf{x}_D), \quad \hat{\mathbf{x}} = \mathcal{D}(\hat{\mathbf{m}}) + \hat{\mathbf{x}}_D \quad (2)$$

The nonlinear operator $\mathcal{D}$ is the decoder portion of the trained encoder-decoder model. Therefore, the manifold component, $\mathcal{D}(\hat{\mathbf{m}})$, is an approximate image which lies on the manifold. The likelihood term is evaluated using the final composite image estimate, $\hat{\mathbf{x}} = \mathcal{D}(\hat{\mathbf{m}}) + \hat{\mathbf{x}}_D$, where $\hat{\mathbf{x}}_D$ is the difference component. The regularization strength is controlled by the parameter γ . As γ approaches zero, the difference image is unconstrained

so the estimator approaches a maximum-likelihood estimator. As γ approaches infinity, the the solution will be constrained to lie on the manifold.

2.4 Lung Imaging Simulation Study

To demonstrate the MRoD algorithm, we conducted a simulated lung imaging study. Data were generated using the XCAT 2.0 anthropomorphic digital phantom. The phantom was parameterized by 97 values which represent volumes and sizes of various organs and other anatomical structures as well as stretches and rotations of the entire body. For each realization, the patient had a 50% chance to be male or female, and the remaining parameters were randomly generated according to a multivariate Gaussian distribution with a standard deviation of 3% for size/volume parameters, 2 degrees for rotation parameters, and a correlation coefficient of 0.75 between size/volume parameters and 0.00 for all other combinations. These parameters were initialized for 100 patients in the training data. Digital phantoms were $256 \times 256 \times 100$ voxels with voxel size $2 \times 2 \times 2$ mm. Voxel values represented attenuation coefficients at 100 keV. Image volumes were cropped along the inferior-superior axis to include slices between the diaphragm and the superior extent of the lungs. Slices were randomly shuffled across all data since training and data processing was conducted on a 2D slice-by-slice basis. The DNN encoder-decoder shown in Figure 1 was trained by using each training sample as a $256 \times 256 \times 1$ input and output with a mean absolute error loss function. The trained decoder was then used as the $\mathcal{D}$ operator as described in a previous section. Synthetic data were generated based on circular fan-beam CT system with 1100 mm source-to-detector distance, 830 mm source-to-axis distance, and 1 mm detector pixels. Uncorrelated Poisson noise was added to the synthetic measured data according to a dose of 10^5 photons per pixel per view.

Both MRoD and QPL estimators were implemented using the Keras deep learning framework.⁷ Custom static layers were designed for forward- and back-projectors, and other operations needed for the “physics-based” modules. Thus, iterative reconstruction uses the same optimization framework and software that is typically used to train a DNN. The only variable quantities during optimization were the image and manifold estimates. For the MRoD case, the pre-trained decoder was included in the estimation model, but the weights are treated as fixed during the iterative CT reconstruction. The estimation was conducted using 1,000 iterations of the Adam optimizer with a learning rate of 0.001.

Three additional simulated patients (distinct from training data) were generated using the above process and a slice at the center of the lungs was extracted for evaluation. The first image was used to compare the noise-bias relationship for the MRoD and QPL algorithms by conducting a sweep of the penalty strength parameters β and γ . A noiseless version of the data was also generated for this case. Bias is defined as the Root-Mean-Squared Error (RMSE) between the image estimates for the noiseless case and the ground truth, and noise is defined as the RMSE between estimates for the noisy and noiseless cases. The second and third images are used to simulate pathological cases. For the second image, we manually added a spiculated lung nodule, and for the third image we manually added a 4mm crack in a rib. These pathological features were not present in the training data.

3. RESULTS

The results from the first study are shown in Figure 2. Each point on the noise vs bias plot represents a different reconstruction for which the level of noise and bias were evaluated. For low penalty strengths, the two algorithms have very similar behavior, since the limiting behavior for both algorithms is a maximum-likelihood estimator. However, for the range of practically useful penalty strengths, the MRoD algorithm has significantly less bias for a matched noise level. We select the case for each method with a noise level closest to 30 HU, which corresponds to relatively heavy regularization, and show the image estimates alongside the ground truth. As expected, QPL bias appears as spatial blur. For the MRoD case, we show the manifold component, the difference component, and the final image estimate. The manifold component captures the general shape of the patient and consists of extremely smooth regions since it was trained with noiseless ground-truth phantoms. Much soft tissue variation in the lungs is absent from the manifold component. However, such variations are represented in the difference image where the data-fidelity term of the objective function may compensate for this residual density in the lungs. The final MRoD image appears to be a close match to the ground truth with much less bias due to spatial blur than the QPL image.

Figures 3 and 4 show the ground truth and MRoD reconstructions of the pathological cases for the lung nodule and the cracked rib. In both cases, the manifold component shows an approximation which does not capture these features entirely since they did not appear in the training data. The manifold component for the nodule is empty space. The manifold component for the cracked rib looks like an intact rib that extends through the cracked region. This is the danger of over-fitting to the training data faced by many deep learning CT reconstruction algorithms. However, the MRoD algorithm is guided by physics-based models to compensate for these errors in the difference image. For the nodule case, the spiculated lesion is clearly shown in the difference image. For the cracked rib case, the difference component contains both the negative signal needed to represent the crack, as well as the positive signal needed to represent the full extent of the rib. In both cases, the MRoD image estimate appears to be a close match to the ground truth. This is promising because it indicates that the MRoD algorithm is capable of both representing features that were not present in the training data, but also highlighting the presence of these differences from the manifold of typical patients.

4. CONCLUSION

The MRoD algorithm combines advantages of physics-based statistical estimation and deep-learning CT. It uses sophisticated physical models, statistical weights, and prior knowledge about the estimates via the manifold-based penalty. Unlike many other deep learning CT methods, the final image estimates are not constrained lie on the manifold. The estimates are guided by physics-based models which allow for deviation from prior expectations when needed. The experimental results show that the MRoD algorithm is highly effective at reducing noise with small levels of bias. The results from from the pathological examples show that this method is capable of reconstructing image features that never appeared in the training data. Finally, both the composite manifold plus different images and difference

component alone appear to be clinically useful outputs, where the latter element can highlight differences from typical anatomy showing what is unique about this patient. In the future we are interested in applying task-driven fluence modulation,⁸ source spectral modulation^{9,10} or non-standard source trajectories¹¹ to mold the data acquisition patterns according the content of the difference component.

ACKNOWLEDGMENTS

The authors thank Dr. Paul Segar who provided access to the XCAT 2.0 anthropomorphic digital phantom software. This work is supported, in part, by NIH grant R01CA249538.

REFERENCES

- [1]. Bengio Y et al., *Neural Computation*, vol. 18, no. 10, pp. 2509–2528, 2006. [PubMed: 16907635]
- [2]. Hein M and Maier M, in *Advances in neural information processing systems*, 2007, pp. 561–568.
- [3]. Zhu B et al., *Nature*, vol. 555, no. 7697, pp. 487–492, 2018. [PubMed: 29565357]
- [4]. Gang GJ et al., in *Medical Imaging 2019: Physics of Medical Imaging*, vol. 10948. International Society for Optics and Photonics, 2019, p. 109480L.
- [5]. Ma G et al., *arXiv preprint arXiv:1810.12255*, 2018.
- [6]. Pourmorteza A et al., *Physics in Medicine & Biology*, vol. 61, no. 5, p. 1986, 2016. [PubMed: 26894795]
- [7]. Chollet F et al. (2015). [Online]. Available: <https://github.com/fchollet/keras>
- [8]. Stayman JW et al., in *Medical Imaging 2016: Physics of Medical Imaging*, vol. 9783. International Society for Optics and Photonics, 2016, p. 97830X.
- [9]. Tivnan M et al., in *Medical Imaging 2019: Physics of Medical Imaging*, vol. 10948. SPIE, 2019, p. 109481A.
- [10]. _____, *arXiv preprint arXiv:2010.07483*, 2020.
- [11]. Ouadah S et al., in *Medical Imaging 2017: Physics of Medical Imaging*, vol. 10132. International Society for Optics and Photonics, 2017, p. 101320H.

Figure 1:

The DNN encoder-decoder model which is used for manifold learning. The encoder portion (green arrow) transforms a $256 \times 256 \times 1$ image of to a $1 \times 1 \times 8192$ manifold representation and the decoder portion (yellow arrow) transforms the manifold back to the image space. For each training sample, the same image is provided as both input and output to the model so this lossy nonlinear compressor can be trained to extract the most important image features.

Figure 2: Noise vs bias and example reconstructions for QPL and MRoD. Violet circles indicates image examples shown.

Figure 3:
Spiculated lung nodule example. Ground truth, manifold, difference, and final images are shown.

Figure 4:
Cracked rib example. Ground truth, manifold, difference, and final images are shown.