R4-B.1: Toward Advanced Baggage Screening: Reconstruction & Automatic Target Recognition

I. PARTICIPANTS INVOLVED FROM JULY 1, 2019 TO JUNE 30, 2020

Faculty/Staff			
Name	**Title**	**Institution**	**Email**
Charles Bouman	Co-PI	Purdue University	bouman@purdue.edu
Graduate, Undergraduate and REU Students			
Name	**Degree Pursued**	**Institution**	**Month/Year of Graduation**
Wenrui Li	PhD	Purdue University	9/2022
Venkatesh Sridhar	PhD	Purdue University	9/2020

II. PROJECT DESCRIPTION

A. Project Overview

3D computed tomography (CT) has been widely deployed in air transportation security applications such as check-baggage screening because of its ability to better detect threats. Based on this success, there is great interest in deploying CT in other areas of transportation security, such as checkpoint security for carry-on baggage and cargo inspection. However, for applications such as checkpoint screening and cargo, it is likely not practical to support the traditional rotating gantry structure or even rotating tables that can collect a full set of projection views assumed by traditional Fourier-based reconstruction methods. This means that this new generation of scanners will need to depend on a new generation of image reconstruction algorithms that go beyond the limitations of Fourier-based reconstruction methods, such as filtered back-projection (FBP).

The following are major challenges faced by next-generation CT systems for Transportation Security Administration (TSA) applications.

- **Large volumes:** Cargo scanning at the resolution required to detect threats will drive systems to reconstruct enormous numbers of voxels. Using existing methods, this has two undesirable consequences. First, it requires more views, which slows down physical scanning; and second, the increase in the amount of data and voxels dramatically slows reconstruction.
- **Dense materials:** In cargo applications, large volumes of dense material will make scanning difficult since photon penetration will be limited.
- **Sparse views:** Physical and practical constraints will likely limit the number of views that can be acquired by fieldable scanning systems for both baggage and checkpoint scanning. This means that reconstruction algorithms will need to be able to reconstruct accurate 3D volumes from sparse views.
- **Beam hardening and scatter due to metal:** Heterogeneous materials such as metals create a combination of beam hardening and scatter effects that break traditional models of reconstruction. While beam hardening for a single material without scatter can be compensated for, it is nearly impossible to accurately compensate for the beam hardening in a single-energy system when it is combined with

scatter effects and multiple materials. This creates large-scale artifacts that can cause erroneous segmentation and missed detection of threats.

B. State of the Art and Technical Approach

Over the last decade, model-based iterative reconstruction (MBIR) and other regularized iterative reconstruction methods have changed how people think about 3D reconstruction from CT data [1–4]. By using MBIR, we have demonstrated that 3D reconstruction from sparse views is possible [5, 6]. However, major challenges still remain. While MBIR reconstruction methods offer some solutions to these problems [7, 8], these algorithms suffer from very high computational requirements, continued (but reduced) susceptibility to beam-hardening/scatter artifacts due to metal, and continued (but reduced) degradation in quality from sparse view data.

A new generation of reconstruction algorithms that blend physical models with artificial intelligence (AI) or machine learning (ML) methods hold great promise in addressing these challenges. The breakthroughs of ML and in particular deep learning (DL) methods have opened new directions for both dramatically reducing computation and dramatically improving quality of reconstructions. The challenge of ML-based methods is to integrate the well-understood models of physics with the information that can be extracted from large corpuses of data. In some cases, this blending of data and physics models can be subtle, such as in the modeling of metal artifacts that requires the integration of well-understood physics models with empirical models of scatter distribution in the presence of unknown physical scattering parameters.

Following is a list of advanced technologies that will likely be critical to future CT imaging technologies needed by the Department of Homeland Security (DHS).

- **Integrated data, physics, and sensor models:** Next generation CT algorithms will need to integrate multiple sources of information from the physical sensor measurements, the known physics of the materials, and the empirical data distributions in order to create accurate 3D reconstructions with the limitations of sparse view data, measurement distortion due to scattering and beam hardening, heterogeneous materials, large and dense objects, and detector and photon counting noise sources. To meet this challenge, a new generation of algorithms should be developed for the integration of heterogeneous information that we call multi-agent consensus equilibrium (MACE) [9]. MACE methods are based on the plug-and-play (PnP) methods [10, 11], but they generalize to multiple sources of information that may not be formulated in an optimization framework. This framework can be used to integrate traditional physics-based models with ML models based on DL [12, 13].
- **Integrated beam-hardening and scatter models:** The complex interactions of beam-hardening and scatter are extremely difficult to accurately model for heterogeneous materials scanned with single energy CT. To address this, one can use a blend of physics and ML modeling to integrate the known physics of CT sensors with the higher order data-driven sensor models. By integrating this type of model with prior image models, our goals was to reduce metal artifacts and improve segmentation accuracy.
- **Integrated segmentation models:** The ultimate goal of reconstruction is to accurately detect targets, and this is critically dependent upon accurate segmentation. ML methods provide a formal basis for incorporating segmentation into the reconstruction process either through training or ML-based prior models. Our goal was to use the methods such as PnP to integrate in segmentation models, or to directly train ML methods to incorporate segmentation into the reconstruction process.
- **Fast reconstruction:** ML methods can be used in a variety of ways to speed up reconstruction [14, 15, 16]. One approach is to train DL methods to remove noise, but this is not effective for artifact reduction. An alternative approach that we have pioneered is to use DL methods to accurately approximate MBIR

reconstructions in a method we call DL-MBIR [14, 15]. One particular extension of this is 2.5D DL-MBIR, which uses convolutional neural networks to accurately approximate MBIR reconstructions. A goal was to use the 2.5D processing structure to achieve quality near full 3D processing, with computation that is very close to 2D.

- **Incorporation of side information:** In many applications, such as cargo scanning, there is side information available, such as manifests, that describe cargo contents that can be potentially used to improve reconstruction quality or improve segmentation or detection accuracy. This will require the formulation of general methods for incorporating side information in the reconstruction, segmentation, and threat detection process.
- **Dynamic sparse-view selection:** Large dense objects present special challenges in 3D imaging due to the difficulty of achieving sufficient penetration. While some of these can be addressed using high Kv X-ray sources and photon counting detectors, these hardware upgrades add cost and also may not be sufficient. There is a need to develop methods for dynamic acquisition of sparse views [17] that can achieve higher signal-to-noise ratio (SNR) given a fixed time budget by selecting views that are more informative based on dynamic measures of object content and structure. We have based this approach on our previous research in statistical learning approach to dynamic sampling (SLADS) [18], which has proven to be useful in dynamic 2D scanning of materials using scanning electron microscopy and other point-probe measurement methods.

We worked with ALERT to obtain data from previous task orders, and we also used simulated data to perform experiments.

B.1. Research on Direct Sparse-View Reconstruction Using Long Short-Term Memory Processing of Stacked Back-Projections

Over the past few years, sparse-view CT has captured the attention of the CT community since sparse-view data acquisition can drastically reduce scan time and dosage. However, in sparse-view applications, conventional CT reconstruction algorithms will suffer from severe streak artifacts and image distortion since projection data is undersampled. MBIR can solve sparse-view CT reconstruction problems. However, the high computational cost of MBIR makes it unsuitable for calculating reconstruction in real time.

Also over the past five years, DL has emerged as a fundamentally new approach to image reconstruction. The two major advantages of deep neural networks (DNNs) are that they can be trained to produce high-quality reconstructions from sparse and noisy data and that they are much more computationally efficient than iterative reconstruction methods.

We proposed stacked back-projection (SBP) as a general structure for direct tomographic reconstruction that preserves all the sinogram information, avoids the need for computationally intractable dense networks, and allows for the natural introduction of convolutional neural network (CNN) structures in reconstruction. We also proposed a DNN based on long short-term memory (LSTM), with SBP as an efficient computational structure that yields high quality reconstructions. In this approach, our model has a convolutional LSTM followed by a U-Net. Further, to avoid a high dependency of the output on the last view projection of SBP, we developed a rotational LSTM that can train two parallel LSTMs using different projections as the last view. Compared to CNN-based methods, our approach incorporates a convolutional LSTM that improves algorithmic convergence and leads to better reconstruction quality.

We train our networks on simulated data and validate our approach on both simulated and real data. We demonstrate that the SBP-LSTM architecture produces higher quality images as compared to post-processing of FBP images and SBP-CNN architecture.

B.1.a. Stacked Back-Projection

Figures 1 and 2 illustrate the procedure of the SBP generation. As shown in Figure 1, all the sinogram views acquired from a single 2D slice are first filtered by a ramp filter. Then, for each single-view projection, we back-project it as listed in the outer ring in Figure 1.

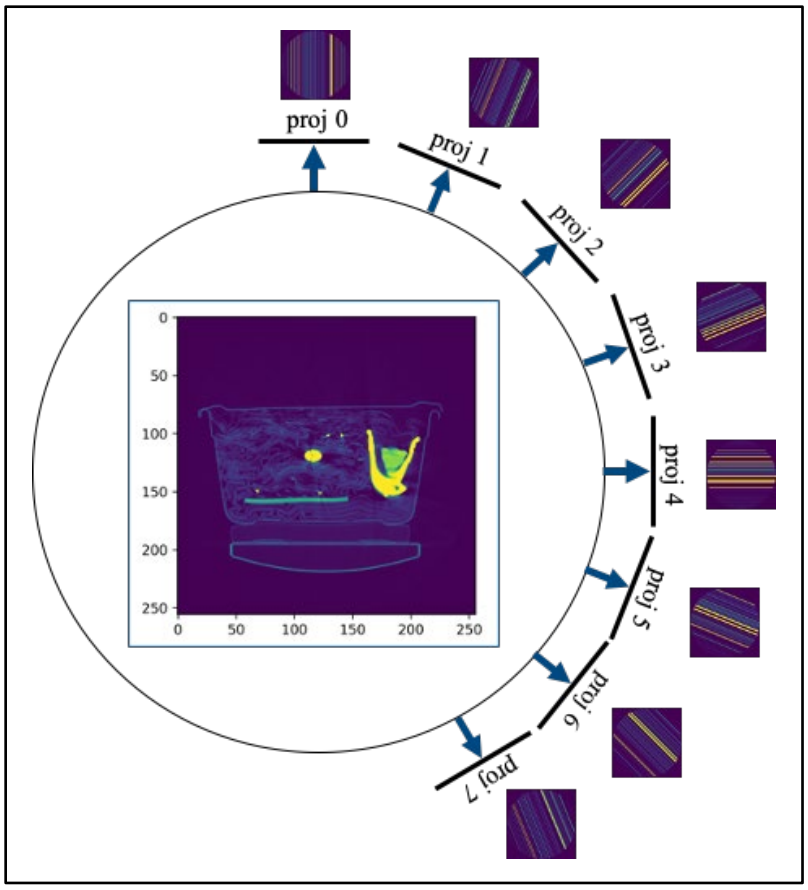

Figure 1: Single-view back-projection generation. In this example, sinogram is projected from sixteen equispaced views $[0, \pi]$.

Figure 2 depicts a pipeline from sinogram to SBP. Sixteen single-view back-projections are stacked together following the order of the equispaced views $[0, \pi]$.

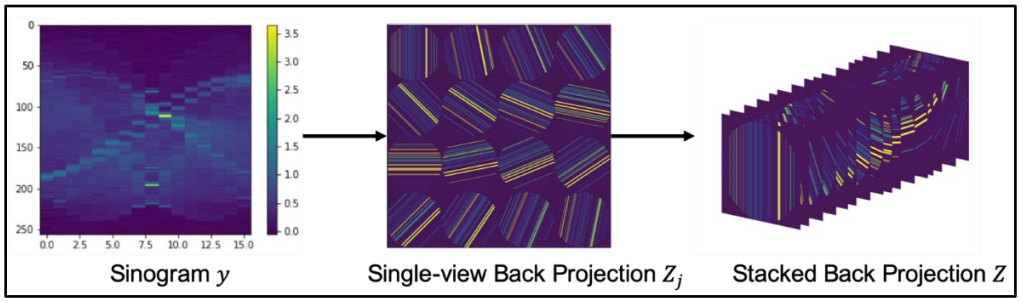

Figure 2: SBP generation. In this example, sinogram is projected from sixteen equispaced views $[0, \pi]$.

B.1.b. Deep Neural Network Models

Since we want to demonstrate the effect of using SBP as input to the deep neural network, we first implement a baseline model using FBP as input. According to [1], many deep neural networks have been proposed to learn the artifacts of the image generated from the measurement by some analytic methods. For example, Wang et al. propose several different neural networks [2-6] to recover a normal dose CT image from a noisy low dose CT image.

The architecture of our learning model is shown in Figures 3 and 4. We use both the FBP image (256 × 256 × 1) and SBP image (256 × 256 × 16) as input to our CNN model to demonstrate the effect of using SBP. We base our CNN model on the U-Net [7] by adding three convolutional blocks before the original U-Net.

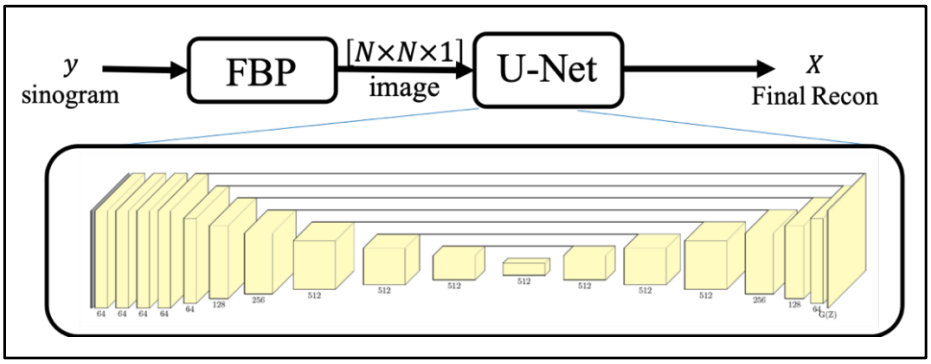

Figure 3: FBP + U-Net. The first block consists of a 3 × 3 convolution layer (with 64 filter, stride = 1) and a rectified linear unit (ReLU) activation function. We do not use activation functions in the final block. The remaining blocks consist of a 3 × 3 convolution layer (with 64 filter, stride = 1), ReLU activation function, and batch normalization.

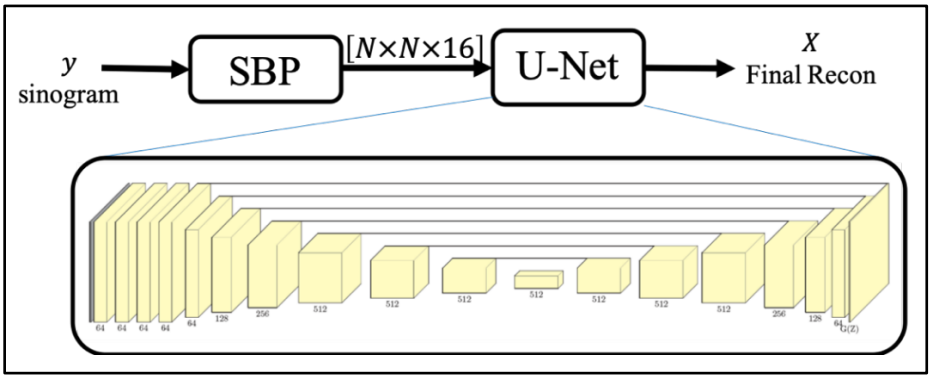

Figure 4: SBP + U-Net. Other layers remain the same as in Figure 3.

LSTM model for SBP: When using a CNN as above, we treat each single-view back projection as a channel, like the RGB channels in a color image, and stack them together in a manner by which the order of the channel does not affect the reconstruction. However, projections of a sinogram generated under a sorted angle sequence have order. Therefore, we introduce a sequential processing network to better handle the SBP data. This is reminiscent of Le et al. [8], who use LSTM to correlate adjacent views in 3D mesh segmentation.

The architecture of our proposed model is shown in Figure 5. We treat multiple views as a temporal sequence. In practical training, we need to expand and rearrange the original SBP into a 4D tensor consisting of time, rows, columns, and channels. After rearrangement, the number of views of the corresponding sinogram is consistent with the dimensionality of time, while the dimensionality of channels is one.

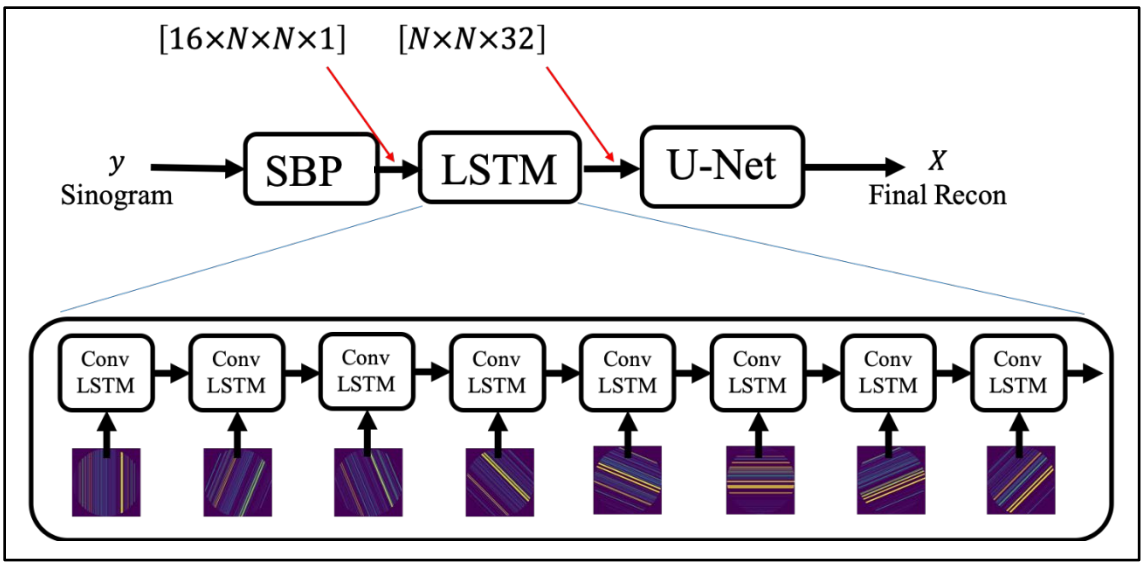

Figure 5: SBP + LSTM + U-Net. The first three blocks are replaced by the convolutional LSTM layer (kernel size = 3, strides = 1, padding = "same," and number of output filters in convolution = 32). Other layers remain the same as in Figure 3.

We base our recurrent layer on convolutional LSTM. Convolutional LSTM replaces a fully connected operation with a convolutional operation, which leads to a novel combination of LSTM and CNN. A traditional CNN-LSTM model uses CNN to extract features, followed by an LSTM and a dense layer, as in [8]. Our proposed model first uses the rearranged SBP tensor as input to the convolutional LSTM; the output is the hidden state output of the last sequence step, which is a 3D tensor shown in Figure 5. The output of the convolutional LSTM will plug into the U-Net. The convolutional LSTM can correlate adjacent views and can better separate low-noise projections from noisy projections. Also, the U-Net can further remove streak artifacts.

Rotational stride: Since we output only the last sequence step of the convolutional LSTM, we need to introduce a half rotational stride to avoid a high dependency of the output on the last view projection. As shown in Figure 6, by using a rotational stride of $\pi/2$, we actually train two parallel convolutional LSTMs. Then we concatenate the last hidden states ($N \times N \times 32$) of two LSTMs as the final output ($N \times N \times 64$) of our LSTM.

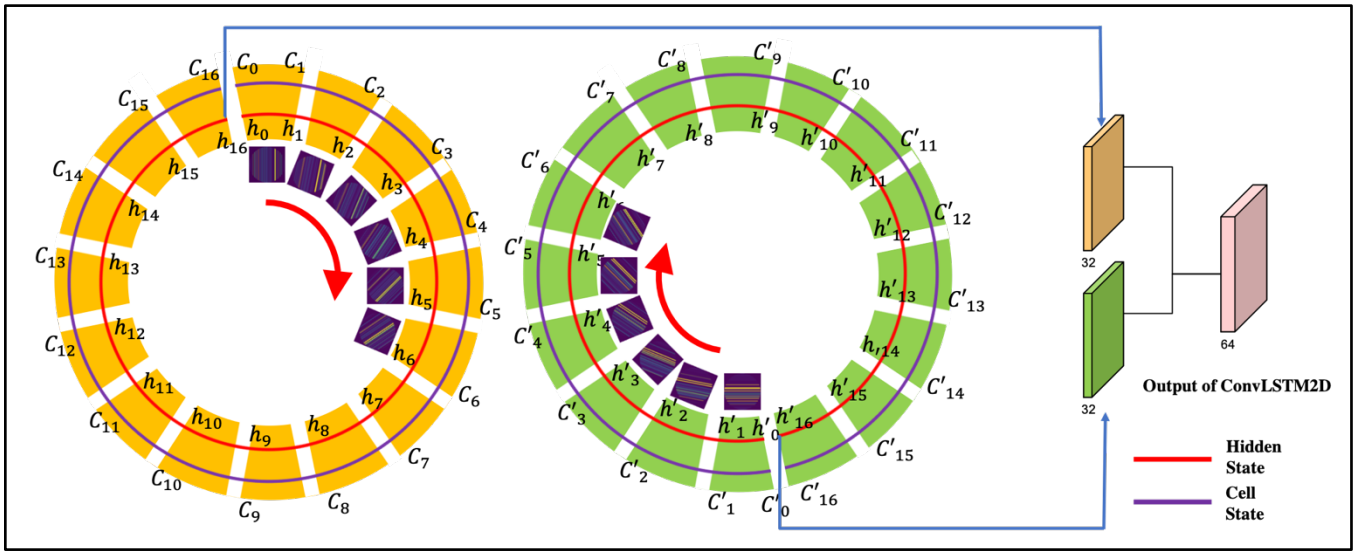

Figure 6: Half rotational stride convolutional LSTM.

Loss function: For reconstructing an X-ray CT image for security applications, the range from 0 to 2,000 Hounsfield units (HU) (scaled so that air = 0 HU and water = 1,000 HU) is critical. Therefore, we used a modified MSE loss to train our models:

$$L_{MSE} = \|f(x) - f(\hat{x})\|^2$$

$$\text{where, } f(x) = \frac{x}{|x|+2000}$$

B.1.c. Experimental Results

To test our approach, we used 980 synthetic sinograms generated from the ALERT Task Order 4 (Advances in Automatic Target Recognition for CT-Based Object Detection Systems) CT dataset by Python as well as 200 sinograms from the ALERT Task Order 3 (Research and Development of Reconstruction Advances in CT-Based Object Detection Systems) CT dataset, using an Imatron Model C300 scanner X-ray CT system.

Figure 7 together with Table 1 compare both qualitative and quantitative results of the image reconstruction algorithms using 16 view simulated data. Both root mean squared error (RMSE) and structural image similarity metric (SSIM) are used as quantitative measures of image quality. For Figure 7, a display range of 0 to 2,000 HU is used. Notice that that proposed SBP+LSTM+U-Net algorithm produces the best overall results and reduces striking artifacts such as those highlighted in the red box. Perhaps surprisingly, SBP+LSTM+U-Net even performs better than the much more computationally expensive MBIR algorithm with a simple Q-GGMRF prior model. Also notice that for our example the introduction of CGAN training does not significantly improve the result relative to more a more standard MSE loss function.

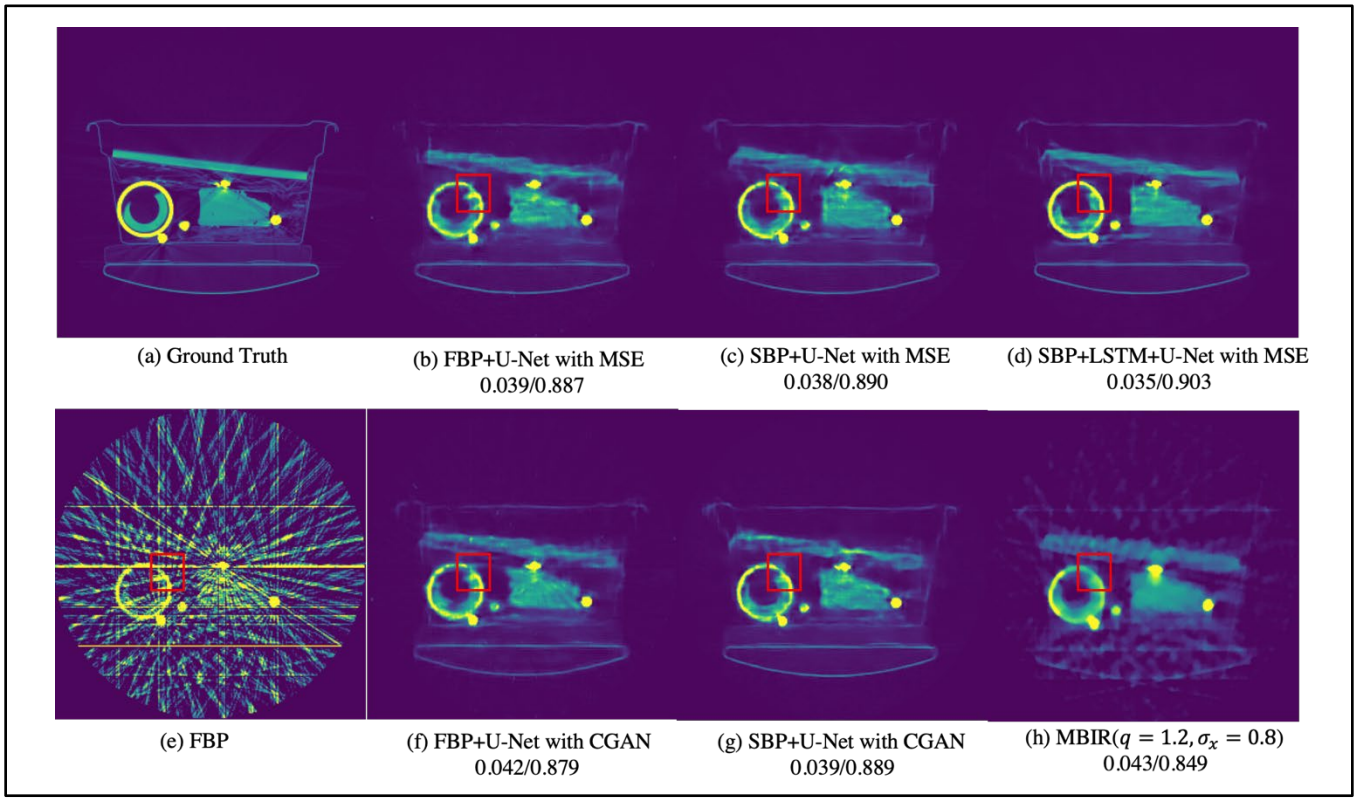

Figure 7: Results on a simulated image. Modified RMSE/SSIM. Display range is 0 to 2,000 HU. Both DL-based methods can successfully suppress most of the streak artifacts. However, FBP + U-Net fails to suppress the remaining streak artifacts shown in the red frame. SBP + LSTM + U-Net can sharpen the reconstruction image and recover important information from low-noise projections.

Model	Training Loss	RMSE	SSIM
FBP + U-Net	MSE	0.0331	0.9028
SBP + U-Net	MSE	0.0321	0.9074
FBP + U-Net	CGAN	0.0339	0.8957
SBP + U-Net	CGAN	0.0320	0.9068
SBP + LSTM + U-Net (unidirectional with rotational stride)	MSE	0.0303	0.9153
MBIR	N/A	0.0374	0.8686

Table 1: Quantitative evaluation on simulated data. SBP + LSTM + U-NET has the best modified RMSE and SSIM metric among all the methods.

Figure 8 together with Table 2 compare both qualitative and quantitative results of the image reconstruction algorithms using 16 view real data subsampled from a full set of views. Again, the RMSE and SSIM are used as metrics of image quality, but since this is real data ground truth is not available to use as a reference image. Therefore, we use the reconstruction form the full set of views as our reference, and this is done using and MBIR and FBP reference for fairness. Perhaps not surprisingly, the MBIR algorithm results in the best image quality when an MBIR reference image is used. However, when an FBP reference image is used, then SBP+LSTM+U-Net algorithm results in the best image quality. Also, visually, the SBP+LSTM+U-Net algorithm results in comparable image quality to MBIR, with less streaking and better overall sharpness, but perhaps more artifacts in some locations. So on the real data, the SBP+LSTM+U-Net algorithm results in comparible image quality to MBIR but with much lower computation.

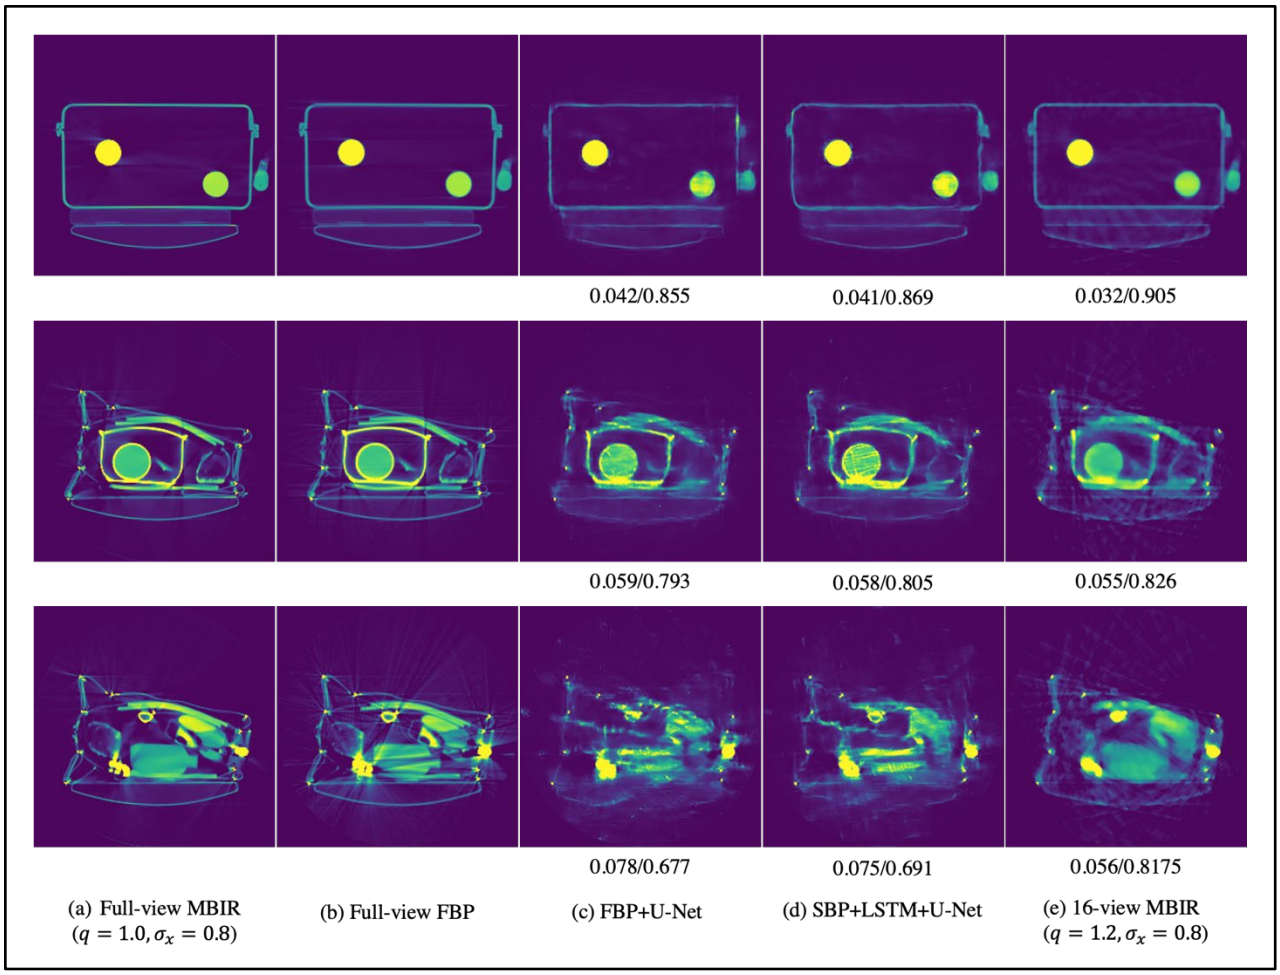

Figure 8: Sparse-view reconstruction on real sinogram. Modified RMSE/SSIM compared with full-view MBIR. Display range is from 0 to 2,000 HU. The MBIR reconstruction on sparse-view sinogram has a more accurate result than our proposed method on real data. However, our proposed method still outperforms other deep models on real data.

Model	Training Loss	RMSE (Relative to Full-View MBIR)	SSIM (Relative to Full-View MBIR)	RMSE (Relative to Full-View FBP)	SSIM (Relative to Full-View FBP)
FBP + U-Net	MSE	0.0606	0.7695	0.0493	0.8110
SBP + U-Net	MSE	0.0613	0.7740	0.0501	0.8112
FBP + U-Net	CGAN	0.0609	0.7781	0.0493	0.8124
SBP + U-Net	CGAN	0.0605	0.7758	0.0502	0.8123
SBP + LSTM + U-Net (unidirectional with rotational stride)	MSE	0.0589	0.7870	0.0470	0.8242
MBIR	N/A	0.0505	0.8404	0.0514	0.7983

Table 2: Quantitative evaluation on real data. Calculating the RMSE and SSIM between the sparse view reconstruction and a reference image computed using the full set of views and either MBIR or FBP. Notice that SBP+LSTM+U-Net and MBIR have the best overall image quality.

B.2. Source Code and Licensing for High-Performance Iterative CT Reconstruction

The parallel super-voxel iterative coordinate optimization (PSV-ICD) method [12] is a high-performance implementation of MBIR for multicore CPU systems. This approach has fast algorithmic convergence since it is based on the ICD method [9] and is roughly 187 times faster on a 20-core CPU system as standard serial ICD reconstruction on a single CPU core [12].

Figure 9 illustrates how the PSV-ICD accelerates the computation MBIR by reorganizing the projection data and system-matrix entries for each super-voxel (SV, a group of spatially close voxels) into a compact contiguous memory layout called the SV-buffer. This approach has two major benefits. First, the data access for the voxel-wise ICD update in a given SV is faster due to sequential memory access and superior cache-locality. Second, the computation of each voxel-wise ICD update can be accelerated using SIMD vector processing units. To efficiently utilize each CPU core, the update of multiple SVs that are spatially far apart is parallelized across CPU cores.

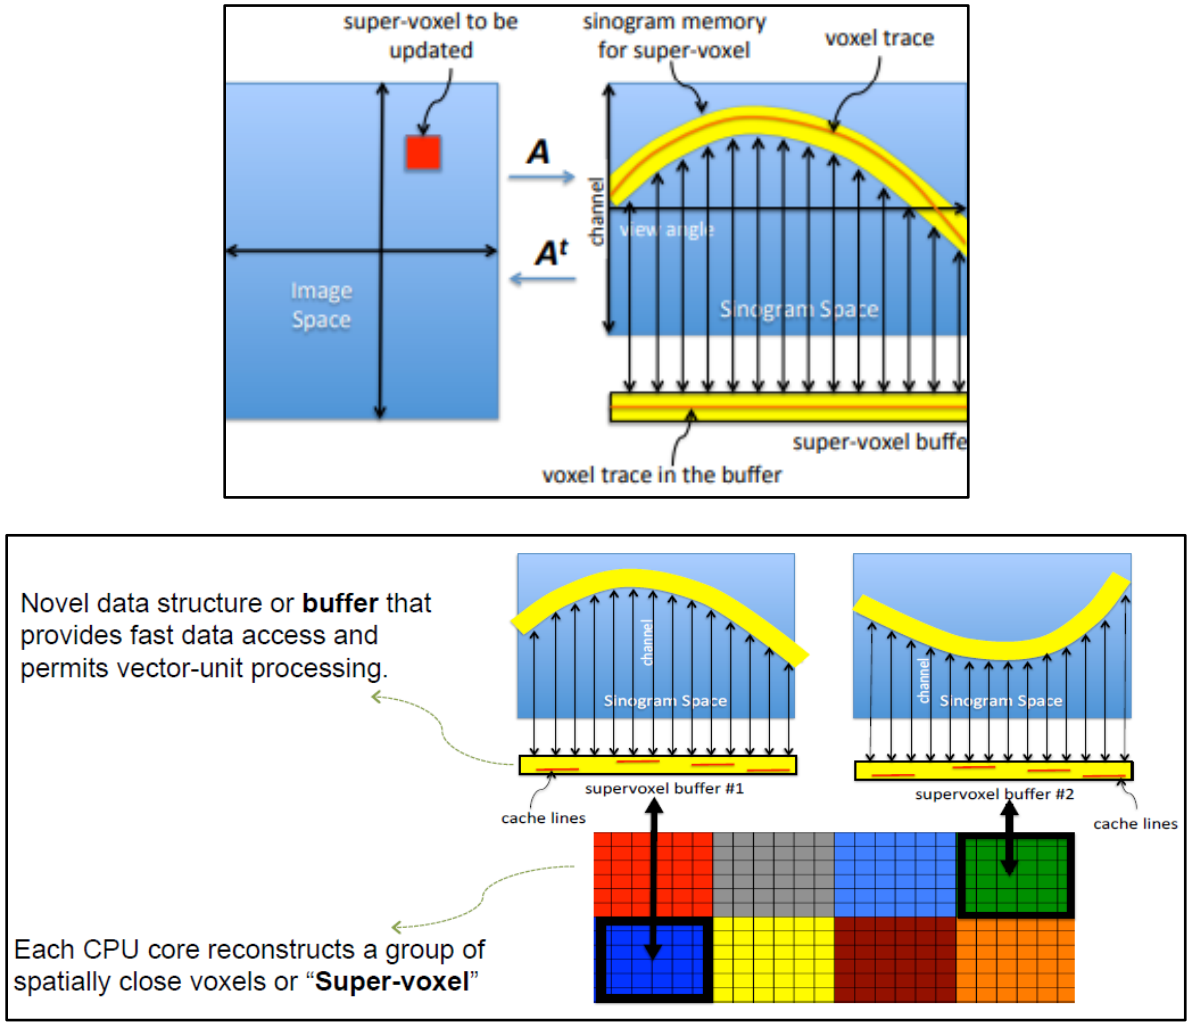

Figure 9: Illustration of high-performance CT reconstruction using the parallel SV-ICD approach; (top) SV buffer formation and (bottom) parallelization of SV-ICD algorithm across CPU cores for reconstructing single slice. The projection data for each SV is reorganized into a more compact data structure called the SV-buffer that permits fast sequential memory access and high data-caching efficiency. Further, the ICD update of different SVs that are spatially far apart is parallelized across multiple CPU cores.

In collaboration with High Performance Imaging LLC (HPI), we publicly released an open-source version of the PSV-ICD algorithm suitable for use on a high-performance multicore compute node. The software, which is available on GitHub (https://github.com/HPImaging), is licensed under the BSD-3-Clause License.

We also built an open-source Python package, svmbir, on GitHub (https://github.com/cabouman/svmbir) that makes the SV MBIR code easy to use and integrate into user software. Enhanced versions of the software suitable for CPU and GPU clusters are available from HPI for commercial licensing.

B.3. Research on Distributed Iterative High-Resolution CT Reconstruction

Tomographic reconstruction algorithms can be roughly divided into two categories: analytical reconstruction methods and regularized iterative reconstruction methods, such as MBIR. MBIR methods have the advantage that they can improve reconstructed image quality particularly when projection data are sparse and/or the X-ray dosage is low. This is because MBIR integrates a model of both the sensor and object being imaged into the reconstruction process. However, the high computational cost of MBIR often makes it less suitable for solving large reconstruction problems in real time.

We proposed a MACE algorithm for distributing both the computation and memory of MBIR reconstruction across a large number of parallel nodes [11, 16, 17]. In MACE, each node stores only a sparse subset of views and a small portion of the system matrix, and each parallel node performs a local sparse-view reconstruction, which—on repeated feedback from other nodes—converges to the global optimum.

Our distributed approach can also incorporate advanced denoisers as priors to enhance reconstruction quality. In this case, we obtain a parallel solution to the serial framework of PnP priors [15, 18], which we call MACE-PnP. Further, a direct implementation of MACE is not practical since it involves repeated application of proximal operators that are themselves iterative. To make MACE practical, we introduced a partial-update method [11, 16] that eliminates nested iterations and proves that it converges to the same global solution.

Finally, we validated our approach on a distributed memory system with real CT datasets taken from security and synchrotron imaging applications, respectively. We also demonstrate an implementation of our approach on a massive supercomputer that can perform large-scale 3D reconstruction in real time.

B.3.a. MACE and MACE-PnP Algorithms

Figure 10 illustrates our two approaches to this distributed CT reconstruction problem. While both the approaches integrate multiple sparse-view reconstructions across a compute cluster into a high-quality reconstruction, they differ based on how the prior model is implemented.

Figure 10 (top) depicts our basic MACE approach that utilizes conventional edge-preserving regularization as a prior model and converges to the maximum a posteriori (MAP) estimate. Figure 10 (bottom) shows our second approach called MACE-PnP which allows for distributed CT reconstruction using PnP priors. These PnP priors [15] substantially improve reconstructed image quality by implementing the prior model using a denoising algorithm based on methods such as BM3D [19] or deep residual CNNs [20]. We prove that MACE-PnP provides a parallel algorithm for computing the standard serial PnP reconstruction of [15].

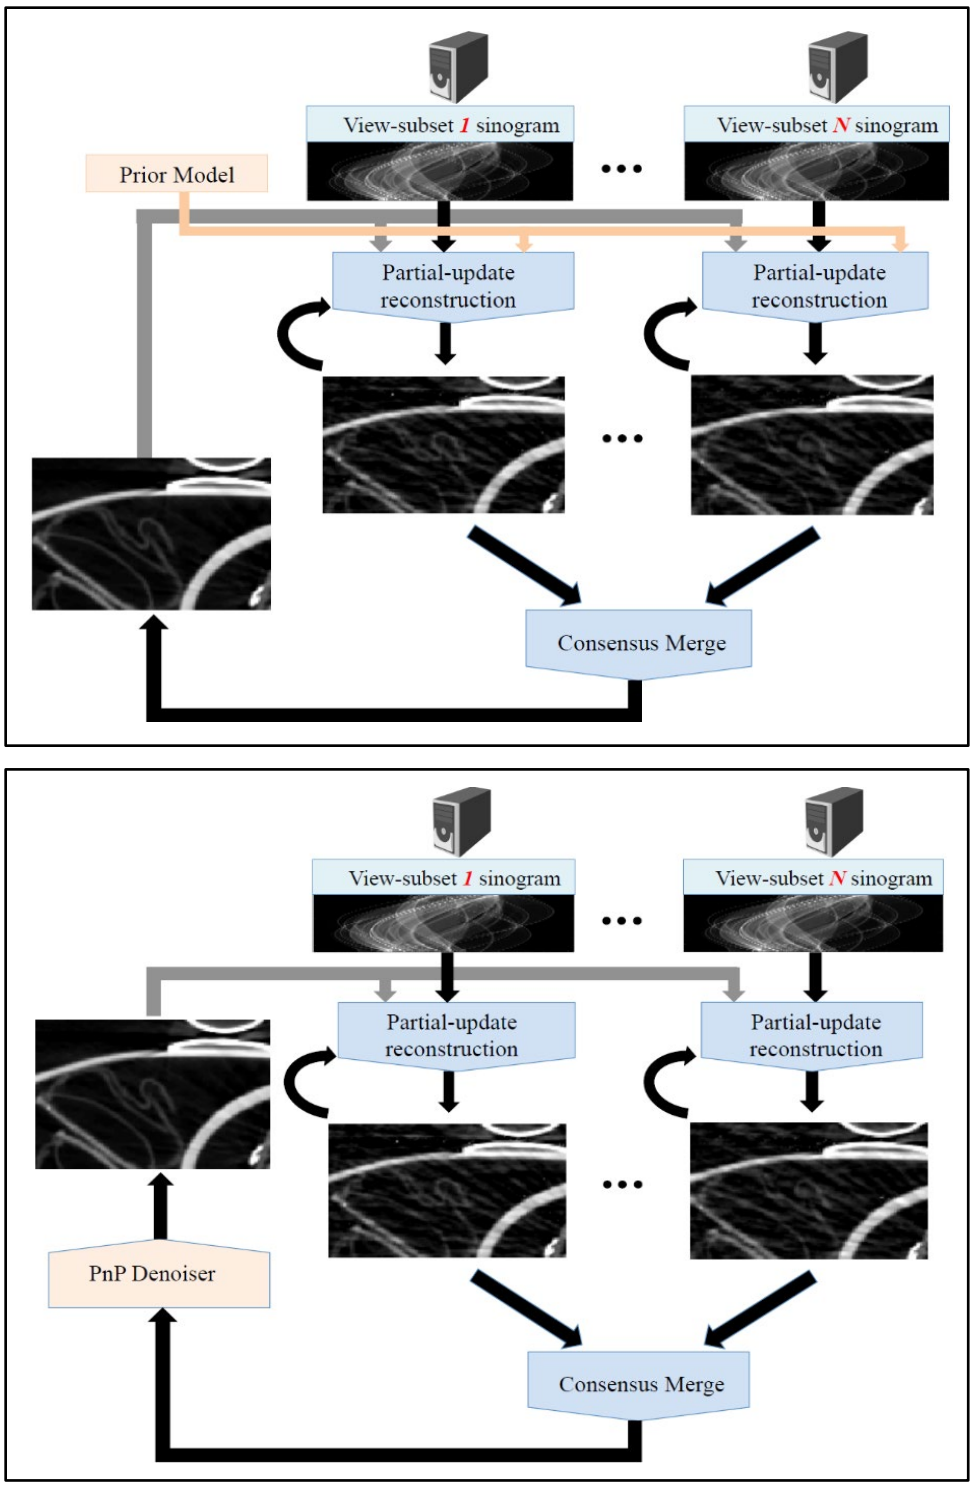

Figure 10: Illustration of the MACE algorithm for distributing CT reconstruction across a parallel cluster of compute nodes. The MACE algorithm works by splitting data into view subsets and reconstructing them in parallel. The individual reconstructions are then merged in an iterative loop that results in (top) the MAP reconstruction in the case of conventional prior models or (bottom) standard PnP reconstruction in the case of advanced prior models.

B.3.b. Mathematical Framework for MACE and MACE-PnP Algorithms

B.3.b.i. CT Reconstruction as a MAP Estimation Problem

To parallelize CT reconstruction, we first split the global MAP estimation problem into a sum of N auxiliary functions given by the following:

$$\text{Minimize } (x;\beta) = \sum_{i=1}^{N} f_i(x;\beta)\,,$$

where for each i, $f_i(x)$ represents a sparse-view reconstruction problem associated with the i^{th} view subset. More specifically, we define

$$f_i(x;\beta) = \frac{1}{2}\|y_i - A_i x\|_{\Lambda}^2 + \frac{1}{N}h(x)\,,$$

where y_i is the projection data from the i^{th} view subset, A_i is the projection operator for the i^{th} view subset, $\beta \geq 0$ is the regularization level, and $h(x)$ is the regularizing prior model. So, the MAP estimate x^* is given by the minimization of the objective function $f(x;\beta)$.

B.3.b.ii. MACE Framework

We apply the MACE framework [10] to solve the above reconstruction problem. The key idea of the MACE framework is to integrate N different data models or *agents*, represented by $F_i\,, i = 1, \cdots, N$ and generate a single coherent fit x^*. The MACE framework [10] specifies x^* through the equilibrium conditions

$$F_i(x^* + u_i^*) = x^*, i = 1, \ldots, N$$

$$\sum_{i=1}^{N} u_i^* = 0.$$

To solve the above CT MAP estimation problem using the MACE framework, we define each agent $F_i\,, i = 1, \cdots, N$ as a proximal map associated with the sparse-view reconstruction problem f_i given by

$$F_i(v_i) = \arg\min_x \left\{ f_i(x) + \frac{1}{2\sigma_\lambda}\|x - v_i\|^2 \right\},$$

where $\sigma_\lambda > 0$ is a parameter that can be selected arbitrarily, but that effects convergence speed of our algorithm. In this case, it can be shown that the MACE solution x^* is exactly the MAP estimate.

To represent the MACE conditions more compactly, we define the parallel-agent operator F and consensus operator G as

$$F(v) = \begin{bmatrix} F_1(v_1) \\ \vdots \\ F_N(v_N) \end{bmatrix}, \text{where } v = \begin{bmatrix} v_1 \\ \vdots \\ v_N \end{bmatrix} \text{and}$$

$$G(v) = \begin{bmatrix} \bar{v} \\ \vdots \\ \bar{v} \end{bmatrix}, \text{where } \bar{v} = \frac{1}{N}\sum_{i=1}^{N} v_i\,.$$

Then, the MACE equilibrium conditions are compactly represented as

$$F(v^*) = G(v^*),$$

where v^* is then the equilibrium solution to the MACE equations. Once v^* is known, then the desired solution to our CT MAP estimation problem is given by the component-wise average of v^*,

$$x^* = \frac{1}{N}\sum_{i=1}^{N} v_i^*$$

It can be shown that the solution to the MACE equations is also the solution to the following fixed-point problem [10]:

$$(2F - I)(2G - I)w^* = w^* ,$$

where w^* is defined through a change of coordinates given by $v^* = (2G - I)w^*$. Consequently, the fixed-point solution can be computed using the Mann iteration given by

$$w^{k+1} = (1 - \rho)w^k + \rho(2F - I)(2G - I)w^k$$

where again ρ is a user-selectable parameter between 0 and 1 that affects convergence speed of the algorithm. Importantly, note that the above approach parallelizes CT reconstruction across N view subsets, since the proximal operators $F_1,\ldots,F_N$ in F that compute the individual sparse-view reconstructions can be applied independently.

B.3.b.iii. MACE-PnP Framework

In this section, we generalize our approach to incorporate PnP priors implemented with advanced denoisers. Since we will be incorporating the prior as a denoiser, for this section we drop the prior terms in the equations of section B.3.b.i by setting $\beta = 0$. So let $f(x) = f(x;\beta)$ denote the CT log likelihood function with $\beta = 0$ and no prior term, and let $F(x)$ denote its corresponding proximal map.

Then Buzzard et al. in [10] show that the PnP framework of Sreehari et al. in [15] can be specified by the following equilibrium conditions:

$$F(x^* - \alpha^*; \sigma) = x^*$$

$$H(x^* + \alpha^*) = x^*$$

where H is the PnP denoiser used in place of a prior model. This framework supports a wide variety of denoisers including BM3D [19] and residual CNNs [20] that can be used to improve reconstruction quality as compared to conventional prior models [15, 18].

Let $f_i(x) = f_i(x;\beta)$ denote the CT log likelihood function with $\beta = 0$ and no prior term, and let $F(x)$ denote its corresponding proximal map. Then we can show that the below equilibrium conditions are exactly equivalent to that of the PnP framework [16].

$$F_i(x^* + u_i^*; \sigma) = x^*, i = 1, \ldots, N$$

$$H(x^* + \alpha^*) = x^*$$

$$\alpha^* + \textstyle\sum_{i=1}^{N} u_i^* = 0.$$

Once again, we formulate a fixed-point method to solve the above problem. One specific way is to use an approach similar to the one in section B.3.b.i, where all agents are applied in parallel. However, we take a slightly different approach in which the denoising is applied in series rather than in parallel. To do so, we define the parallel-operator F and consensus operator G_{H} as

$$F(v) = \begin{bmatrix} F_1\left(v_1\,;\,\sqrt{N}\sigma\right) \\ \vdots \\ F_N\left(v_N;\,\sqrt{N}\sigma\right) \end{bmatrix}, \text{where } v = \begin{bmatrix} v_1 \\ \vdots \\ v_N \end{bmatrix} \text{and}$$

$$G_H(v) = \begin{bmatrix} H\bar{v} \\ \vdots \\ H\bar{v} \end{bmatrix}, \text{where } \bar{v} = \frac{1}{N}\sum_{i=1}^{N} v_i\,.$$

Then we show that x^* in the above equilibrium conditions can be computed by finding the fixed-point of map T_H defined by $T_H = (2F - I)(2G_H - 1)$ [16]. More specifically, to compute x^* we first find w*, the fixed-point of T_H by using the Mann iteration

$$w^{k+1} = (1-\rho)w^k + \rho T_H w^k,$$

where $0<\rho<1$, and then applying the denoiser H to the component-wise average of w^*. Note that this method is parallelizable across N view subsets, since the parallel operator F typically dominates the computation in comparison with consensus operator G_H.

B.3.b.iv. Partial-Update Method for Speeding Up MACE and MACE-PnP

A direct application of the MACE and MACE-PnP algorithms is not practical, since the proximal operators F_{i} = 1,…,N that are evaluated in each Mann iteration require iterative optimization on their own. This results in nested loops of optimization that slow down the convergence speed.

To overcome this issue, we proposed a partial-update method that replaces the proximal operator with fast noniterative updates. In [16] we show that this method significantly speeds up the MACE and MACE-PnP algorithms without sacrificing convergence. We refer to this modified version of MACE with partial updates as partial-update MACE.

B.3.c. Experimental Results

We compare the image quality and convergence of our distributed MACE approach against a standard single-node reconstruction method. The single-node method is conventional ICD reconstruction from all views on a single node [9]. In the MACE approach, the views are distributed across multiple nodes, and we compute each partial update using one pass of ICD optimization.

To measure algorithmic convergence, we specify computation in units called equits. For the case of the conventional serial algorithm, we define one equit to be the computation associated with the update of all the pixels within the region of interest using the entire set of sinogram views. Since this definition accounts for the number of views being processed by the node, note that on a distributed implementation, one equit is equivalent to having each node perform one full ICD iteration using its subset of views.

Figure 11 (left) shows the distributed MACE reconstruction using N = 16 nodes for a single slice of the ceramic-composite dataset acquired by the Advanced Light Source (ALS) synchrotron at Lawrence Berkeley National Labs (LBNL). For this dataset, the CT measurements per slice consist of 2,560 sinogram views × 1,024 detector channels, and our reconstruction size is 512 × 512 voxels. From Figure 11 (middle) and

(right), we notice that MACE reconstruction with N = 16 nodes has quality equivalent to the single-node reconstruction, even though each node utilizes <7% of the sinogram views for its local reconstruction.

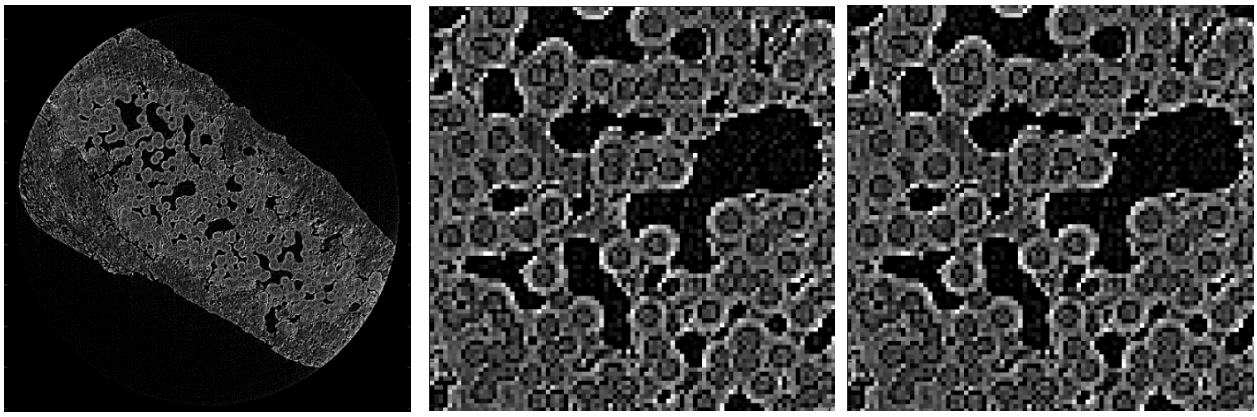

Figure 11: Image quality comparison using the LBNL ceramic composite dataset: (left) MACE reconstruction method with N = 16 nodes; (middle) close-up of the MACE reconstruction method with N = 16 nodes; (right) close-up of single-node reconstruction method. Notice that MACE produces image quality equivalent to the single-node method.

Figure 12(a) shows the distributed MACE-PnP reconstruction using N = 16 nodes and CNN as a prior model for a single slice of the ALERT Task Order 3 baggage scan. For this dataset, the CT measurements per slice consist of 720 sinogram views × 1,024 detector-channels, and our reconstruction size is 512 × 512 pixels. For this experiment, our CNN prior is based on the deep res-net architecture of [20] and was trained using natural images. From Figure 12 (right) we notice that in comparison with conventional Q-GGMRF prior [13, 14], the CNN prior improves quality by reducing streaking artifacts, which can potentially enable better segmentation.

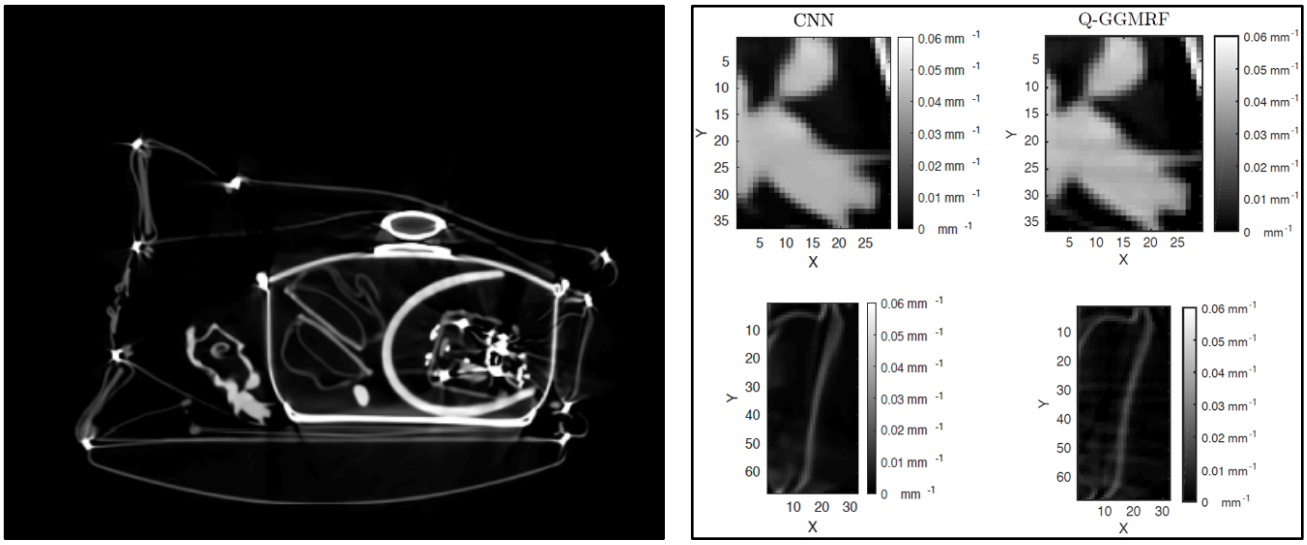

Figure 12: (left) MACE-PnP reconstruction with CNN prior and N = 16 nodes for the ALERT Task Order 3 baggage scan; (right) comparison of CNN prior versus Q-GGMRF prior. Notice that CNN prior reduces streaking artifacts in comparison to conventional edge-preserving regularization.

Figure 13 (left) shows that in the case of MACE with conventional Q-GGMRF prior, the number of equits required for convergence increases with the number of nodes, *N*, and consequently, the speedup is less than linear. However, from Figure 13 (right) we notice that in the case of MACE with CNN prior (MACE-PnP), the number of equits required for convergence does not change significantly from $N = 2$ to 8, and it is in fact less than that required by the serial approach when $N \leq 8$. Consequently, MACE with CNN prior provides a roughly linear speedup and has better parallel efficiency. Table 3 summarizes the relative speedup results for these experiments.

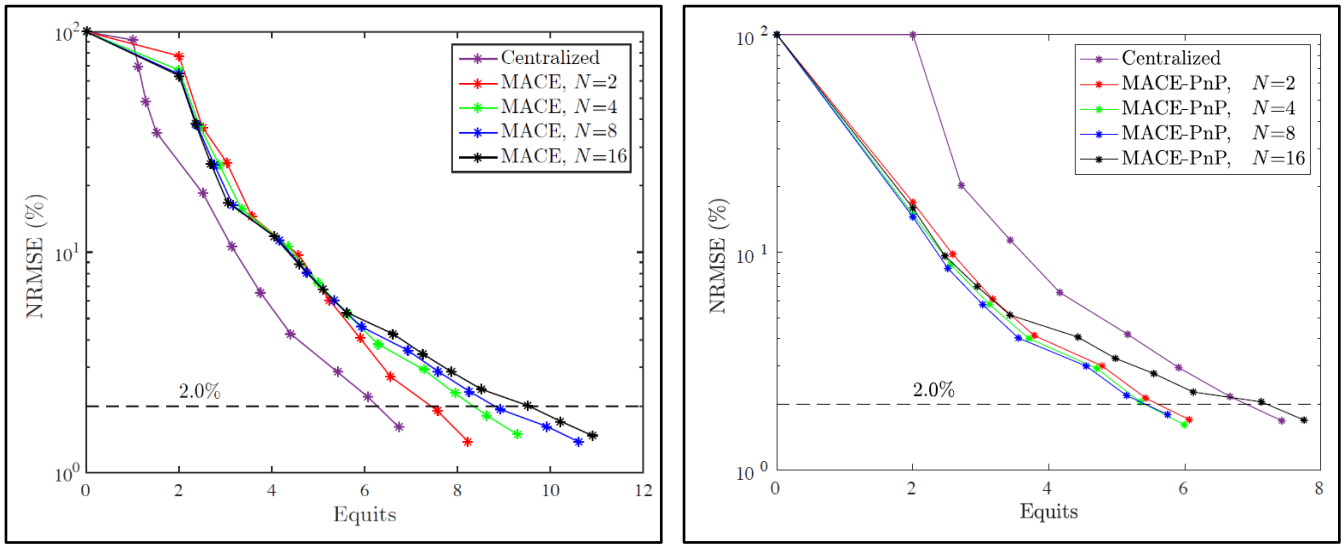

Figure 13: MACE algorithmic convergence using (left) Q-GGMRF prior and (right) res-net CNN prior (MACE-PnP) as a function of the number of nodes *N* for ALERT Task Order 3 baggage scan. In comparison with the Q-GGMRF prior, the CNN prior provides faster convergence and significantly better speedup over the serial reconstruction method.

Prior Model	*N* = 1	*N* = 2	*N* = 4	*N* = 8	*N* = 16
Q-GGMRF	1.00	1.80	3.27	6.00	10.29
BM3D	1.00	2.14	4.29	8.57	16.00

Table 3: MACE algorithmic speedup as a function of number of nodes *N*, for the LBNL ceramic composite dataset. Note that while MACE speedup is less than linear in the case of the Q-GGMRF prior, it is roughly linear in the case of the BM3D prior.

Table 4 illustrates the use of our distributed MACE approach for large-scale reconstruction on the massive National Energy Research Scientific Computing Center (NERSC) supercomputer in real time and analyzes its computational performance. For this experiment, we performed high-resolution reconstruction for a 3D ceramic-composite dataset acquired by the ALS synchrotron at LBNL. The CT measurements consist of 2,560 sinogram views × 1,024 detector channels × 1,200 detector rows, and our reconstruction size is 1280 × 1280 × 1,200 voxels. Each node in the cluster reconstructs 8 slices and utilizes the high-performance SV-ICD method [12] for efficient parallelization across multiple CPU cores. For this dataset, the conventional reconstruction method that does not distribute memory is not practical, since the overall CT system model is very large. From Table 4, we notice that our distributed MACE approach can achieve high-quality reconstructions in real time by reducing both the memory and computation (per iteration) on each node by a factor of *N*, where *N* denotes the number of view subsets. While the speedup of MACE increases with *N*, we notice that after $N = 4$, the parallel efficiency drops and the speedup based on real machine-time is less than the algorithmic speedup based on equits.

#View Subsets (*N*)	#Nodes	#Equits	#Time (seconds)	Algorithmic Speedup	Machine-Time Speedup
1	150	–	–	–	–
2	300	23.91	275	2.00	2.00
4	600	26.67	154	3.58	3.57
8	1200	34.03	121	5.62	4.55

Table 4: High-resolution reconstruction of a large 3D synchrotron dataset using the MACE approach on the NERSC supercomputer. MACE enables reconstruction of large datasets in real time by reducing both the computation and memory footprint per compute node by a factor of *N*, the number of view subsets. Note that while the speedup increases with *N*, the parallel efficiency drops after *N* = 4.

B.4. Major Contributions

Year 1 (2013–2014) outcomes:

- Developed MBIR algorithm for application in sparse view CT.
- Developed multislice helical scan geometry CT code for MBIR reconstruction.
- Developed an "implicit prior" method for image reconstruction. This method was the conceptual precursor to the PnP methods that were later developed and have come to be widely used in the computational imaging community.
- Published the first paper on PnP, which has 284 citations in Google Scholar as of August 25, 2020.
- Implemented 2D MBIR on reformatted Imatron data.

Year 2 (2014–2015) outcomes:

- Developed and implemented metal artifact reduction methods for MBIR reconstruction algorithm for application in sparse view CT.
- Implemented MBIR on Morpho and ALERT Task Order 3 data to demonstrate IQ improvements and improved Pd/Pfa (probability of detection/probability of false alarm) performance for baggage scanning data.
- Completed Secure Sensitive Information (SSI) report on improved performance of MBIR for baggage scanning applications.
- Developed novel segmentation-based beam hardening correction algorithm for MBIR. Resulted in P. Jin, C. A Bouman, and K. D Sauer, Trans. Computation Imaging, 2015, with 146 citations in Google Scholar as of August 25, 2020.

Year 3 (2015–2016) outcomes:

- Developed novel algorithm for joint segmentation and reconstruction using MBIR that won a best paper award from the *IEEE International Conference on Acoustics, Speech and Signal Processing*.
- Developed novel CT automatic target recognition (ATR) algorithm based on segmentation and classification method that was demonstrated in ALERT Task Order 4 competition.
- Invented SV MBIR reconstruction algorithm that has proven to be ~100–2,000 times faster than conventional MBIR and scalable to large-scale supercomputers. The SV-MBIR algorithm resulted in being a finalist for the Gorden Bell prize in 2017. It also was the basis for the formation of a small business,

High Performance Imaging, which currently has three ongoing projects with the Air Force and Air Force contractors as of July 1, 2020.

Year 4 (2016–2017) outcomes:

- Did initial research on MACE that is a follow-up to the PnP methods. This was used by others, including Professor Karl, for application in sparse-view CT reconstruction using DL.
- Implemented 3D SV algorithm on a large-scale supercomputer.

Year 5 (2017–2018) outcomes:

- Used the MACE algorithm to implement a highly parallel reconstruction algorithm for very large CT reconstruction on large-scale supercomputers.
- Demonstrated large-scale CT reconstruction on the NERSC supercomputer using MACE and SV-MBIR technologies.
- Created SBP architecture for direct reconstruction of sparse CT data using DNNs.

Year 6 (2018–2019) outcomes:

- Created SBP-LSTM algorithm for improved 2D image reconstruction quality.
- Developed partial-update PnP methods for large-scale parallel 3D image reconstruction.

Year 7 (2019–2020) outcomes:

- Extended the SBP-LSTM algorithm to 3D cone-beam geometry for fast and accurate CT image reconstruction.
- Published a TCI journal paper on large-scale parallel implementations of SV, MACE algorithms, with PnP priors.

B.5. Milestones

Our Year 7 activities focused on studying the potential of the methods described above to solve problems faced by TSA in baggage, checkpoint, and cargo scanning. In particular, we developed a novel algorithm for direct reconstruction from sinogram data using DNNs that is based on the SBP together with an LSTM architecture. These specific milestones were achieved:

- September 30, 2019—Performed preliminary study of methods and evaluated the potential of these methods for addressing TSA problems in cargo scanning. Identified mechanisms for obtaining simulated and real cargo-scanning data.
- December 31, 2019—Implemented ML-based reconstruction algorithms for real and simulated cargo-scanning data sets.
- March 31, 2019—Tested and optimized ML-based reconstruction algorithms using real and simulated cargo-scanning data sets. Assessed weaknesses and strengths. Formulated plan for optimizing performance in terms of computation speed and image quality.
- June 30, 2019—Evaluated image quality and computation speed of ML-based reconstruction algorithms using standard methods of image quality, such as normalized root mean squared error for simulated data, and artifact measures such as contrast ratio for real data sets.

In addition, we achieved the following milestone in parallel with high-speed MBIR reconstruction.

- June 30, 2019—Developed and implemented a novel algorithm based on the MACE framework for parallel reconstruction of large CT images using both traditional MBIR and PnP priors.

B.6 Final Results at Project Completion (Year 7)

Major project outcomes:

- Demonstrated the first practical methods for sparse view reconstruction. In [21] and [22], we had the first demonstrations that sparse view reconstruction was possible for practical transportation security problems such as cargo scanners. Practical sparse view reconstruction is of critical values since for many problems, including carry-on baggage and cargo, it may not be possible or practical to obtain a full set of views.
- Demonstrated the value of MBIR in reducing reconstruction artifacts that contribute to increased false alarm and missed detection rates [23].
- Introduced novel methods for metal artifact reduction in MBIR reconstruction [24,25]. These publications, done in collaboration with researchers at what was then Morpho Detection, demonstrate that MBIR can be used to reduce metal artifacts that cause false alarms and miss detections in baggage scanning.
- Developed the advanced prior methods for reconstruction that are generally known as plug-and-play (PnP) methods. These methods can dramatically improve reconstructed image quality by allowing the integration of physics models with AI and deep-learning–based prior models [18,15].
- Created a novel high-performance MBIR algorithm, known as super-voxel MBIR (SV-MBIR), that was demonstrated to speed reconstruction by well over a factor of 400 on CPU, GPU, and large-scale cluster computers [12,26,17,11,16]. The multicore CPU version of this algorithm is now available as open-source Python code and makes fast MBIR reconstruction practically accessible to a wide range of users.
- Created novel computationally efficient architecture known as stacked-back projection (SBP) for direct reconstruction of parallel and cone-beam data using deep neural networks [27, 28]. These emerging methods have the potential to achieve reconstruction quality that is near MBIR with a fraction of the computation time.
- Formed a startup company, High-Performance Imaging LLC (HPI), to translate state-of-the-art innovations in CT scanning technology to the problems of transportation security. It currently has a number of active research programs with the Air Force and is also working with US Department of Energy (DOE) researchers to translate the innovations of the ALERT research to important applications in manufacturing, material science, and sensing.

III. RELEVANCE AND TRANSITION

A. Relevance of Research to the DHS Enterprise

Metrics:

- Speed of computation—In the publications [12, 26, 17, 11, 16] we demonstrated reconstruction speedups of the svmbir algorithm relative to conventional parallel state-of-the-art implementations of MBIR of approximately 400 times. This relative speedup was achieved on multicore CPU implementations, multi-GPU implementations, and large-scale clusters such as the DOE's NERSC supercomputer.

- Image quality—Using MBIR with advanced priors and also using direct SBP reconstruction using deep neural networks we were able to demonstrate dramatic reconduction image artifacts for sparse view reconstruction, substantial image quality improvement using metal artifact reduction, and improved segmentation and object detection using integrated reconstruction/segmentation/classification methods.

B. Status of Transition at Project End

In collaboration with HPI, we publicly released an open-source version of the PSV-ICD algorithm suitable for use on a high-performance multicore compute node. The software, which is available on GitHub (https://github.com/HPImaging), is licensed under the BSD-3-Clause License. We are also building an open-source Python package, svmbir, on GitHub (https://github.com/cabouman/svmbir) that makes the SV MBIR code easy to use and integrate into user software. Enhanced versions of the software suitable for CPU and GPU clusters are available from HPI for commercial licensing.

In addition, we are planning on making open-source Python packages available for the SBP-LSTM reconstruction algorithm discussed. Our goal is to develop a version of the SBP-LSTM algorithm that can work for both parallel and cone beam geometry with arbitrary uniform view angle sampling and on-the-fly LSTM processing. With these innovations the package can be used on a wide variety of practical parallel and cone-beam tomography data sets without retraining.

C. Transition Pathway and Future Opportunities

The software and algorithms described in this report and developed over the past seven years are being used to engage with a wide variety of government and commercial organizations including:

- Eli Lilly Corporation: We are using the cone-beam and parallel 3D and 4D MBIR algorithms to image medical devices, such as injectors, in collaboration with Eli Lilly. This collaboration also involves Canadian Light Source and Kinetic Vision.
- Los Alamos National Laboratory (LANL): We are using the software and algorithms developed in this effort as the core of our collaborations with LANL in the areas of hyperspectral neutron imaging, imaging of hydrodynamic events, and ptychographic imaging.
- Argonne National Laboratory's Advanced Photon Source (APS): We are using the parallel-beam geometry SV MBIR code as the basis of our collaborations with APS on topics such as 4D reconstruction form sparse view sampling of dynamic events.

D. Customer Connections

1. Morpho Detection: 2011–2016
    - Collaborated with Samit K Basu, Todd Gable, Walter Garms, Matthew Merzbacher, and Sondre Skatter on:
        - Astrophysics project for cargo scanner
        - Study of iterative reconstruction for checked baggage scanner
        - X-ray diffraction scanner for checked baggage scanner
2. Astrophysics: 2011–2013
    - Simon Bedford and Francois Zayek
        - Astrophysics project for cargo scanner

3. AS&E: 2017-2018
    - Aaron Couture and Jonathan Everett
        - Advanced X-ray diffraction scanner reconstruction done jointly with Tufts University.

IV. PROJECT ACCOMPLISHMENTS AND DOCUMENTATION

A. Peer Reviewed Journal Articles

Pending –

1. Sridhar, V., Wang, X., Buzzard, G.T., & Bouman, C.A., "Distributed Iterative CT Reconstruction Using Multi-Agent Consensus Equilibrium." *IEEE Transactions on Computational Imaging*, accepted 2020.

B. Peer Reviewed Conference Proceedings

1. Wang, X., Sridhar, V., Ronaghi, Z., Thomas, R., Deslippe, J., Parkinson, D., Buzzard, G.T., Midkiff, S.P., Bouman, C.A., & Warfield, S.K. "Consensus Equilibrium Method for Super-Resolution and Extreme-Scale CT Iterative Reconstruction." *SC 19: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, Denver, CO. 17–22 November 2019, pp. 86-1 to 86-13.

C. Other Presentations

1. Webinars
    a. "Plug and Play: A General Approach to AI and Sensor Model Fusion." Best Paper Prize, keynote lecture, *2020 SIAM Imaging Sciences Conference*. https://www.youtube.com/watch?v=GjCmxTqAJDo&feature=youtu.be.

D. Student Theses or Dissertations Produced from This Project

1. Sridhar, V., "Parallel Computational Methods for Model-Based Tomographic Reconstruction and Coherent Imaging." PhD, School of Electrical and Computer Engineering, Purdue University, May 2020.

E. Algorithms

See Section III.

V. REFERENCES

[1] S. Ravishankar, J. C. Ye, and J. A. Fessler, "Image Reconstruction: From Sparsity to Data-Adaptive Methods and Machine Learning,"arXiv:1904.02816 [cs, eess, stat], Apr. 2019, arXiv: 1904.02816.

[2] Q. Yang, P. Yan, Y. Zhang, H. Yu, Y. Shi, X. Mou, M. K. Kalra,Y. Zhang, L. Sun, and G. Wang, "Low-Dose CT Image Denoising Using a Generative Adversarial Network With Wasserstein Distance and Perceptual Loss," IEEE Transactions on Medical Imaging, vol. 37, no. 6,pp. 1348–1357, June 2018.

[3] C. You, L. Yang, Y. Zhang, and G. Wang, "Low-Dose CT via Deep CNN with Skip Connection and Network in Network," arXiv:1811.10564.

[4] C. You, Q. Yang, H. shan, L. Gjesteby, G. Li, S. Ju, Z. Zhang, Z. Zhao, Y. Zhang, W. Cong, and G. Wang,

"Structurally-Sensitive Multi-Scale Deep Neural Network for Low-Dose CT Denoising," IEEE Access, vol. 6, pp. 41 839–41 855, 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8416740/.

[5] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang, "Low-Dose CT with a Residual Encoder-Decoder Convolutional Neural Network (RED-CNN)," IEEE Transactions on Medical Imaging, vol. 36, no. 12, pp. 2524–2535, Dec. 2017, arXiv:1702.00288. [Online]. Available: http://arxiv.org/abs/1702.00288.

[6] H. Chen, Y. Zhang, W. Zhang, P. Liao, K. Li, J. Zhou, and G. Wang, "Low-dose CT denoising with convolutional neura network," arXiv:1610.00321 [physics], Oct. 2016, arXiv: 1610.00321. [Online]. Available: http://arxiv.org/abs/1610.00321.

[7] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Net-works for Biomedical Image Segmentation," in Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, ser. Lecture Notes in Computer Science, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Springer International Publishing, 2015, pp. 234–241.

[8] T. Le, G. Bui, and Y. Duan, "A Multi-View Recurrent Neural Network for 3D Mesh Segmentation," Computers & Graphics, vol. 66, pp. 103–112, Aug. 2017.

[9] C. A. Bouman and K. D. Sauer, "A Unified Approach to Statistical Tomography using Coordinate Descent Optimization," IEEE Trans. on Image Processing, vol. 5, no. 3, pp. 480-492, March 1996.

[10] G. T. Buzzard, S. H. Chan, S. Sreehari and C. A. Bouman, "Plug-and-Play Unplugged: Optimization Free Reconstruction Using Consensus Equilibrium," vol. 11, no. 3, SIAM Journal on Imaging Science, 2018.

[11] V. Sridhar, G.T. Buzzard and C. A. Bouman, "Distributed Framework for Fast Iterative CT Reconstruction from View-subsets", Proc. of IS&T Electronic Imaging Conference, 2018.

[12] X. Wang, A. Sabne, S. J. Kisner, A. Raghunathan, C. A. Bouman and S. P. Midkiff, "High Performance Model Based Image Reconstruction", 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), pp. 2:1-2:12, 2016.

[13] Bouman C.A. and Sauer K., "A Generalized Gaussian Image Model for Edge-Preserving MAP Estimation," IEEE Transacions on Image Processing, vol. 2, no. 3, pp. 296-310, July 1993.

[14] J. B. Thibault, K. D. Sauer, C. A. Bouman and J. Hsieh, "A Three-Dimensional Statistical Approach to Improved Image Quality for Multi-Slice Helical CT," Med. Phys., vol. 34, no. 11, pp. 4526-4544, 2007.

[15] S. Sreehari, S.V. Venkatakrishnan, B. Wohlberg, L. Drummy, J. P. Simmons, and C. A. Bouman, "Plug and Play Priors for Bright Field Electron Tomography and Sparse Interpolation", IEEE Transactions on Computational Imaging, vol. 2, no. 4, Dec. 2016.

[16] Sridhar V., Wang X., Buzzard G.T., and Bouman C.A., "Distributed Iterative CT Reconstruction using Multi-Agent Consensus Equilibrium", accepted to IEEE Transactions on Computational Imaging, 2020. (arXiv preprint:1911.09278).

[17] Wang, X., Sridhar, V., Ronaghi, Z., Thomas, R., Deslippe, J., Parkinson, D., Buzzard, G.T., Midkiff, S.P., Bouman, C.A., & Warfield, S.K. "Consensus Equilibrium Method for Super-Resolution and Extreme-Scale CT Iterative Reconstruction.", SC 19 the International Conference for High Performance Computing, Networking, Storage, and Analysis, pp. 86-1 to 86-13, 2019.

[18] Venkatakrishnan, S.V., Bouman, Charles A. and Wohlberg, Brendt, "Plug-n-play priors for model based reconstruction", IEEE Global Conference on Signal and Information Processing, 2013.

[19] Buades A., Coll, B. and Morel J.M., "A non-local algorithm for image denoising", IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), 2005.

[20] Zhang K., Zuo W., Chen Y., Meng D. and Zhang L., "Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising", IEEE Transactions on Image Processing.

[21] S. J. Kisner, E. Haneda, C. A. Bouman, S. Skatter, M. Kourinny, and S. Bedford, "Limited View Angle CT Reconstruction," Proceedings of SPIE-IS\&T Electronic Imaging -Computational Imaging X,} vol. 8296, January 23-24, 2012.

[22] Sherman J. Kisner, Eri Haneda, Charles A. Bouman, Sondre Skatter, Mikhail Kourinny, Simon Bedford, "Model-Based CT Reconstruction from Sparse Views," the proceedings of the Second International Conference on Image Formation in X-Ray Computed Tomography, pp. 444-447, Salt Lake City, Utah, June 24-27, 2012.

[23] Sherman J. Kisner, Pengchnog Jin, Charles A. Bouman, Ken Sauer, Walter Garms, Todd Gable, Seungseok Oh, Matthew Merzbacher, and Sondre Skatter, "Innovative Data Weighting for Iterative Reconstruction in Helical CT Security Baggage Scanner," 47th IEEE International Carnahan Conference on Security Technology, Medellin-Colombia, October 8-11, 2013.

[24] Pengchong Jin, Charles A. Bouman, and Ken D. Sauer, "A Model-Based Image Reconstruction Algorithm with Simultaneous Beam Hardening Correction for X-Ray CT," {\em IEEE Transactions on Computational Imaging}, vol. 1, no. 3, pp. 200-216, January 2016.

[25] Pengchong Jin, Dong Hye Ye, and Charles A. Bouman "Joint Metal Artifact Reduction and Segmentation of CT Images Using Dictionary-Based Image Prior and Continuous-Relaxed Potts Model," in the proceedings of the {\em IEEE International Conference on Image Processing}, September 2015.

[26] Amit Sabne, Xiao Wang, Sherman Kisner, Charles Bouman, Anand Raghunathan, and Samuel Midkiff, "Model-based Iterative CT Image Reconstruction on GPUs," 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '17), February 4-8, 2017.

[27] Dong Hye Ye, Gregery T. Buzzard, Max Ruby, and Charles A. Bouman, "Deep Back Projection for Sparse-View CT Reconstruction," IEEE Global Conference on Signal and Information Processing (GlobalSIP), pp. 1-5, Nov. 20, 2018.

[28] Wenrui Li, Gregery T. Buzzard, and Charles A. Bouman, "Direct Sparse-View CT Reconstruction using LSTM Processing of Stacked Back Projections," International Conference on Computational Photography (ICCP), Aprill 24-26, 2020.