

Advanced Prior Modeling for 3D Bright Field Electron Tomography

Suhas Sreehari^a, S. V. Venkatakrishnan^b, Lawrence F. Drummy^c, Jeffrey P. Simmons^c, and Charles A. Bouman^a

^aSchool of ECE, Purdue University, West Lafayette, IN, USA;

^bLawrence Berkeley National Laboratory, Berkeley, CA, USA;

^cAir Force Research Laboratory, Dayton, OH, USA

ABSTRACT

Many important imaging problems in material science involve reconstruction of images containing repetitive non-local structures. Model-based iterative reconstruction (MBIR) could in principle exploit such redundancies through the selection of a log prior probability term. However, in practice, determining such a log prior term that accounts for the similarity between distant structures in the image is quite challenging. Much progress has been made in the development of denoising algorithms like non-local means and BM3D, and these are known to successfully capture non-local redundancies in images. But the fact that these denoising operations are not explicitly formulated as cost functions makes it unclear as to how to incorporate them in the MBIR framework.

In this paper, we formulate a solution to bright field electron tomography by augmenting the existing bright field MBIR method to incorporate any non-local denoising operator as a prior model. We accomplish this using a framework we call *plug-and-play priors* that decouples the log likelihood and the log prior probability terms in the MBIR cost function. We specifically use 3D non-local means (NLM) as the prior model in the plug-and-play framework, and showcase high quality tomographic reconstructions of a simulated aluminum spheres dataset, and two real datasets of aluminum spheres and ferritin structures. We observe that streak and smear artifacts are visibly suppressed, and that edges are preserved. Also, we report lower RMSE values compared to the conventional MBIR reconstruction using qGGMRF as the prior model.

Keywords: Non-local means, plug-and-play, bright field, 3D tomography.

1. INTRODUCTION

Structural biology^{1,2} and material science³ are examples of research areas that require the reconstruction of images containing many similar or identical particles. The considerable data redundancy in such images points toward under-exploited potential of prior modeling in image reconstruction of these particles. As an important example, cryo electron microscopy (EM) tomography involves single particle reconstructions using several views of the same particle.¹ However, in the more general area of 3D transmission electron microscopy (TEM) tomography, no solution currently exists to fully exploit the redundancy in images constituted by many similar or identical particles.

Model-based iterative reconstruction (MBIR) is a general framework for solving inverse imaging problems by computing the maximum a posteriori (MAP) estimate of the unknown.⁴⁻⁹ Traditionally, MBIR solves a single optimization problem that tightly couples the log likelihood term (based on the data) and the log of the prior probability. This constrains the choice of prior models we can use – making it challenging to model non-local dependencies in biological and materials images.

Non-local means (NLM) is among the most promising and reliable methods for denoising images due to its ability to capture non-local redundancies.¹⁰⁻¹⁴ There have been several efforts to weave in non-local regularizers within the Bayesian setting. Chen et al. proposed an MRF-style prior, albeit with non-local spatial dependencies,

Further author information: (Send correspondence to S.S.)
S.S.: E-mail: ssreehar@purdue.edu, Telephone: 1 646 812 0677

to perform Bayesian tomographic reconstruction.¹⁵ The authors adopt a two-step optimization involving non-local weight update, followed by the image update. However, the cost function related to the prior model is a function of time, and there is no single fixed cost function that is minimized. Chun et al. proposed non-local regularizers for emission tomography¹⁶ based on alternating direction method of multipliers (ADMM),^{17–20} using Fair potential²¹ as the non-local regularizer, instead of non-local means. This model is restricted to convex potential functions, which in practice is a very strong constraint, and severely limits how expressive the model can be. Yang et al. proposed a unifying energy minimization framework for non-local regularization,²² resulting in a model that captures the intrinsically non-convex behavior required for modeling distant particles with similar structure. However, there is no single fixed cost function that is minimized.

In this paper, we present a formulation to incorporate 3D non-local means as a prior model in 3D bright field electron tomographic reconstruction. We encapsulate 3D non-local means in what we call *plug-and-play priors* framework,²³ which is based on ADMM. Plug-and-play decouples the log likelihood and the log prior probability terms in the optimization procedure, thereby allowing us to use independent software modules for implementing the forward and prior models. Though NLM is not explicitly formulated as the solution to an optimization problem, in Sec 4 we empirically demonstrate convergence of the ADMM iterations.

Another important consideration is the computational tractability of the denoising algorithm. Any 3D patch-based denoising is a computationally daunting endeavor. Fortunately, in our application, the samples abound with similar or identical structures. This allows us to restrict the NLM patch weighting to voxels within a smaller search window, thereby speeding up our 3D NLM implementation, while not compromising on the availability of good patches.

Due to the generality of the plug-and-play technique, this work results in an MBIR framework that is compatible with any denoising algorithm as a prior model, and thus opens up a huge opportunity to adopt any spatial constraint in the reconstructions. This enhancement is presented as an augmentation of the existing MBIR framework for bright field electron tomography,²⁴ which features Bragg scatter modeling and anomaly detection.

2. PLUG-AND-PLAY FRAMEWORK

We introduce a methodology for incorporating a general class of denoising operations to function as a prior model for our reconstruction problem. This approach, which we call plug-and-play priors,²³ has the advantage of having two separate software modules – one to solve the inversion and the other to implement the denoising operator. Plug-and-play is an application of ADMM,¹⁷ wherein we split the state variable to decouple the forward and prior model terms of MAP estimation. Plug-and-play solves the MAP estimation problem even when the prior term does not have an explicit cost function.

In the MBIR framework, reconstruction is typically formulated as the maximum *a posteriori* (MAP) estimate of the unknowns given the measurements, which is given by,

$$\hat{x} = \underset{x \geq 0}{\operatorname{argmin}} \{l(y; x) + \beta s(x)\}, \quad (1)$$

where y represents the data, x represents the unknown pixels, $l(y; x) = -\log p(y|x)$ is the forward model term, $s(x) = -\log p(x)$ is the prior model term, and β controls the regularization applied.

In,²³ the authors split the variable x (of Eq. (1)) into two variables x and v , and reformulate Eq. (1) as,

$$(\hat{x}, \hat{v}) = \underset{\substack{x \geq 0 \\ v=x}}{\operatorname{argmin}} \{l(y; x) + \beta s(v)\}. \quad (2)$$

Then, we solve Eq. (2) by applying the ADMM technique on the augmented Lagrangian given by,

$$L_\lambda(x, v, u) = l(y; x) + \beta s(v) + \frac{1}{2\sigma_\lambda^2} \|x - v + u\|_2^2, \quad (3)$$

where u is a scaled dual variable (initialized to zero) and σ_λ is the augmented Lagrangian parameter.

Each iteration of ADMM algorithm consists of the following steps.

$$\hat{x} = \mathbb{F}(y, \tilde{x}; \sigma_\lambda), \quad (4)$$

$$\hat{v} = \mathbb{H}(\tilde{v}; \sigma_n), \quad (5)$$

$$u = u + (\hat{x} - \hat{v}), \quad (6)$$

where $\tilde{x} = \hat{v} - u$, and $\tilde{v} = \hat{x} + u$.

We define the inversion operator, \mathbb{F} , used in Eq. (4) as,

$$\mathbb{F}(y, \tilde{x}; \sigma_\lambda) = \operatorname{argmin}_{x \geq 0} \left\{ l(y; x) + \frac{\|x - \tilde{x}\|_2^2}{2\sigma_\lambda^2} \right\}. \quad (7)$$

Further, we define the denoising operator, \mathbb{H} , used in Eq. (5) as,

$$\mathbb{H}(\tilde{v}; \sigma_n) = \operatorname{argmin}_{v \geq 0} \left\{ \frac{\|\tilde{v} - v\|_2^2}{2\sigma_n^2} + s(v) \right\}, \quad (8)$$

where $\sigma_n = \sqrt{\beta}\sigma_\lambda$ is the assumed noise standard deviation used in the denoising algorithm.

We use the 3D non-local means as the denoising operator \mathbb{H} , though it does not necessarily solve an optimization problem as in Eq. (8). In Sec. 4, however, we show the convergence of \hat{x} and \hat{v} . In theory, the value of σ_λ does not affect the reconstruction in case of convex optimization, but in practice, a well-chosen value of σ_λ can speed up the ADMM convergence. Our approach is to choose the value of σ_λ to be approximately equal to the amount of variation in the reconstruction, as below.

$$\sigma_\lambda^2 \approx \operatorname{var}[x|y]. \quad (9)$$

Since we do not know x , we instead use an initial reconstruction of x using a baseline algorithm. The unitless parameter β controls the regularization. Increasing the value of β leads to greater regularization, while decreasing it reduces the regularization.

Algorithm 1 Plug-and-Play Priors

```

1: procedure  $\hat{x} \leftarrow \text{RECONSTRUCT}(y, \beta, \sigma_\lambda)$ 
2:   (Inputs: Measurement  $y$ , regularization  $\beta$ , augmented Lagrangian parameter  $\sigma_\lambda$ .)
3:   (Output: Estimate  $\hat{x}$  of the unknown image  $x$ )
4:   initialize  $\hat{x}; \hat{v} \leftarrow \hat{x}; u \leftarrow 0$ 
5:    $\sigma_n^2 \leftarrow \beta\sigma_\lambda^2$  // Variance for denoising algorithm
6:   Set MAX_ITER // Default value: 20
7:   for  $i = 1$  to MAX_ITER do
8:     if stopping criteria are met then
9:       return  $\hat{x}$ 
10:    else
11:       $\tilde{x} \leftarrow \hat{v} - u$ 
12:       $\hat{x} \leftarrow \mathbb{F}(y, \tilde{x}; \sigma_\lambda)$  //  $\mathbb{F}$ : inversion operator
13:       $\tilde{v} \leftarrow \hat{x} + u$ 
14:       $\hat{v} \leftarrow \mathbb{H}(\tilde{v}; \sigma_n)$  //  $\mathbb{H}$ : denoising with prior
15:       $u \leftarrow u + (\hat{x} - \hat{v})$  // Update of  $u$ 
16:    end if
17:  end for
18:  return  $\hat{x}$ 
19: end procedure

```

3. BRIGHT FIELD ELECTRON TOMOGRAPHY USING PLUG-AND-PLAY FRAMEWORK

Let $S \subset \mathbb{Z}^3$ be the set of all lattice points, and $s = (s_1, s_2, s_3) \in S$ be the s -th voxel. Also, let y be the tomographic measurement, x be the unknown, and \hat{x} be the reconstruction. For denoising operations, let \tilde{v} be the noisy image, and \hat{v} be the denoised image.

3.1 The denoising operator, \mathbb{H} : 3D non-local means

In this section, we present 3D non-local means which we use as our denoising operator, $\mathbb{H}(\tilde{v}; \sigma_n)$, in Eq. (8). Non-local means is a denoising algorithm that filters each voxel by a weighted mean of all pixels in the image. Non-local means leads to greater sharpness in the filtered image, as compared to local smoothing-based denoising methods.²⁵

Let us define the cubic window, W , to window out cubic patches as we pass through the image.

$$W = \{s \in \mathbb{Z}^3 : \|s\|_\infty \leq R\}, \quad (10)$$

where $R \in \mathbb{N}$ defines the size of the cubic patches. The gray-level of the j -th voxel of the cubic patch P_s around voxel s is then given by,

$$P_{s,j} = \tilde{v}_{s+j}, \quad (11)$$

where $j \in W$. To maintain computational tractability, we make use of a search window – to which we restrict the weighting of candidate voxels. The search window of voxel s is given by,

$$\Omega_s = \{r \in S : \|r - s\|_\infty \leq n\} \subset S, \quad (12)$$

where $n \in \mathbb{N}$ defines the size of the search window.

We define the weight exerted by voxel r on voxel s as,

$$w_{s,r} = \exp \left\{ \frac{-\|P_r - P_s\|_2^2}{\sigma_n^2} \right\}. \quad (13)$$

For every voxel s , we then normalize $\{w_{s,r}\}_{r \in \Omega_s}$ as below.

$$\tilde{w}_{s,r} = \frac{w_{s,r}}{\sum_{r \in \Omega_s} w_{s,r}}. \quad (14)$$

The denoised image is then computed as,

$$\hat{v}_s = \sum_{r \in \Omega_s} \tilde{w}_{s,r} \tilde{v}_s. \quad (15)$$

3.2 Computing the inversion operator, \mathbb{F}

The likelihood function that models the bright field TEM data is given by,²⁴

$$p(y|x, d, \sigma) = \frac{1}{Z} \exp \left\{ -\frac{1}{2} \sum_{k=1}^K \sum_{i=1}^M \beta_{T,\delta} \left((y_{k,i} - A_{k,i,*}x - d_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma} \right) \right\}, \quad (16)$$

where K is the number of tilts, $\lambda_{k,i}$ is the electron counts corresponding to the i -th measurement at the k -th tilt, $y_{k,i} = -\log \lambda_{k,i}$, $\lambda_{D,k}$ is the blank scan value at the k -th tilt, $d_k = -\log \lambda_{D,k}$, A_k is the $M \times N$ tomographic forward projection matrix associated with the k -th tilt, $A_{k,i,*}$ is the i -th row of A_k , σ^2 is a proportionality

constant, Λ_k is a diagonal matrix whose entries are set such that $\frac{\sigma^2}{\Lambda_{k,ii}}$ is the variance of $y_{k,i}$, $d = [d_1, \dots, d_K]$ is the offset parameter vector, and Z is the partition function, and $\beta_{T,\delta}(\cdot)$ is the generalized Huber function defined as,

$$\beta_{T,\delta}(x) = \begin{cases} x^2 & \text{if } |x| < T \\ 2\delta T|x| + T^2(1 - 2\delta) & \text{if } |x| \geq T \end{cases} \quad (17)$$

and the surrogate function to the generalized Huber function defined as,

$$Q_{T,\delta}(x; x') = \begin{cases} x^2 & \text{if } |x'| < T \\ \frac{\delta T}{|x'|} x^2 + \delta T|x'| + T^2(1 - 2\delta) & \text{if } |x'| \geq T \end{cases} \quad (18)$$

We use the generalized Huber function to account for outliers in the measurement due to Bragg scatter.

To perform the inversion based on the plug-and-play inversion operator (given in Eq. (7)), we must minimize the cost function below.

$$c(x, d, \sigma; \tilde{x}, \sigma_\lambda) = \frac{1}{2} \sum_{k=1}^K \sum_{i=1}^M \beta_{T,\delta} \left((y_{k,i} - A_{k,i,*}x - d_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma} \right) + MK \log(\sigma) + \frac{\|x - \tilde{x}\|_2^2}{2\sigma_\lambda^2}. \quad (19)$$

Next we implement the inversion operator, $\mathbb{F}(y, \tilde{x}; \sigma_\lambda)$, by minimizing the cost function given in Eq. (19).

$$(\hat{x}, \hat{d}, \hat{\sigma}) \leftarrow \underset{x \geq 0, d, \sigma}{\operatorname{argmin}} c(x, d, \sigma; \tilde{x}, \sigma_\lambda). \quad (20)$$

We use the majorization-minimization technique along with a coordinate descent strategy to solve the joint-optimization of the non-convex cost function in Eq. (19). The surrogate function we use for the majorization-minimization algorithm is given by,

$$Q(x, d, \sigma; x', d', \sigma'; \tilde{x}, \sigma_\lambda) = \frac{1}{2} \sum_{k=1}^K \sum_{i=1}^M Q_{T,\delta}(h_{k,i}(x, d, \sigma); h_{k,i}(x', d', \sigma')) + MK \log(\sigma) + \frac{1}{2\sigma_\lambda^2} \|x - \tilde{x}\|_2^2, \quad (21)$$

where $h_{k,i}(x, d, \sigma)$ is defined as,

$$h_{k,i}(x, d, \sigma) = (y_{k,i} - A_{k,i,*}x - d_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma}. \quad (22)$$

We now repeatedly construct and minimize Q one voxel at a time. We can write the surrogate cost function with respect to a single voxel, j , as below.

$$\tilde{c}(t) = \theta_1 t + \frac{\theta_2}{2} (t - x'_j)^2 + \frac{(t - \tilde{x}_j)^2}{2\sigma_\lambda^2}, \quad (23)$$

where, θ_1 and θ_2 are defined²⁴ as,

$$\theta_1 = - \sum_{k=1}^K \sum_{i=1}^M A_{k,i,j} \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} \left(b'_{k,i} e'_{k,i} \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} + (1 - b'_{k,i}) \delta T \frac{e'_{k,i}}{|e'_{k,i}|} \right),$$

$$\theta_2 = \sum_{k=1}^K \sum_{i=1}^M A_{k,i,j}^2 \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} \left(b'_{k,i} \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} + (1 - b'_{k,i}) \delta T \frac{1}{|e'_{k,i}|} \right), \quad (24)$$

$e'_{k,i} = y_{k,i} - A_{k,i,*}x' - d'_k$, and $b'_{k,i}$ is a binary indicator variable to simplify the update equations, defined as,

$$b'_{k,i} = \begin{cases} 1 & \text{if } \left| (y_{k,i} - A_{k,i,*}x' - d'_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} \right| < T \\ 0 & \text{if } \left| (y_{k,i} - A_{k,i,*}x' - d'_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} \right| \geq T \end{cases} \quad (25)$$

Since Eq. (23) is quadratic in t , we can write a closed form expression for the update of t ,

$$t^* = \frac{\tilde{x}_j + \theta_2 x'_j \sigma_\lambda^2 - \theta_1 \sigma_\lambda^2}{1 + \theta_2 \sigma_\lambda^2}. \quad (26)$$

The updated value of the j -th voxel (enforcing positivity constraint on x) is given by,

$$x'_j \leftarrow \max \{0, t^*\}. \quad (27)$$

By setting the gradient of $Q(x', d, \sigma'; x', d', \sigma'; \tilde{x}, \sigma_\lambda)$ in Eq. (21) with respect to d to zero, we get the optimal update for d as,

$$d'_k \leftarrow d'_k + \frac{\sum_{i=1}^M \sqrt{\Lambda_{k,ii}} \left[e'_{k,i} b'_{k,i} \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} + \delta T \frac{e'_{k,i}}{|e'_{k,i}|} (1 - b'_{k,i}) \right]}{\sum_{i=1}^M \sqrt{\Lambda_{k,ii}} \left[b'_{k,i} \frac{\sqrt{\Lambda_{k,ii}}}{\sigma'} + \frac{\delta T}{|e'_{k,i}|} (1 - b'_{k,i}) \right]}. \quad (28)$$

In a similar manner, we can write the update of σ as,

$$\sigma' \leftarrow \sqrt{\frac{\sum_{k=1}^K \sum_{i=1}^M e'^2_{k,i} \Lambda_{k,ii} b'_{k,i} + \sum_{k=1}^K \sum_{i=1}^M (1 - b'_{k,i}) \delta T |e'_{k,i}| \sigma' \sqrt{\Lambda_{k,ii}}}{MK}}. \quad (29)$$

We approximate the minimization in Eq. (7) by one iteration of the above optimization.

4. RESULTS

In this section, we present the results of bright field tomographic reconstruction of (1) a simulated dataset of aluminum spheres of varying radii, (2) a real dataset of aluminum spheres, and (3) a real dataset of ferritin structures. We compare three reconstruction methods – filtered backprojection, MBIR with qGGMRF prior,²⁶ and the plug-and-play reconstruction with 3D NLM as the prior model. We used the OpenMBIR software (www.openmbir.org) for reconstructions using qGGMRF as the prior model, and 3D non-local means reconstruction was performed using the plug-and-play framework that we incorporated onto the OpenMBIR codebase. For NLM-based reconstructions, the qGGMRF-reconstructions were used as the initialization. All the reconstruction results shown below are $x - z$ slices, and have been plotted in the range $[0, 19.8 \times 10^{-3}] \text{ nm}^{-1}$.

Table 1: Plug-and-play algorithm parameters

	(Simulated) Al spheres	(Real) Al spheres	(Real) ferritin
$\sigma_\lambda \text{ (nm}^{-1}\text{)}$	8.66×10^{-3}	8.66×10^{-3}	8.66×10^{-3}
β	3.68	4.77	4.25

The qGGMRF parameters used for all reconstructions are as follows: $q = 1$, $p = 1.2$, and $c = 0.001$. The NLM patch size used for all reconstructions was $5 \times 5 \times 5$.

For each experiment, we include a convergence plot (see Figs. 3, 6, and 9) that shows the percentage normalized RMSE between \hat{x} and \hat{v} . Specifically, the y -axis represents the quantity $\frac{\|\hat{x}^{(k)} - \hat{v}^{(k)}\|}{\|\hat{x}^{(\infty)}\|}$, where $\hat{x}^{(k)}$ and $\hat{v}^{(k)}$ are the values of \hat{x} and \hat{v} after the k -th iteration of the plug-and-play algorithm, respectively, and $\hat{x}^{(\infty)}$ is the final value of the reconstruction, \hat{x} . This quantity drops to as low as 0.2% for each of the three experiments, demonstrating empirical convergence of the plug-and-play algorithm.

4.1 Aluminum spheres (simulated) dataset

The aluminum spheres simulated dataset contains 47 equally-spaced tilts about the y -axis, spanning $[-70^\circ, +70^\circ]$. Dosage $\lambda_{D,k} = 1865$ counts using Beer's law model. The noise model is Gaussian, with variance set equal to the mean. The phantom also contains effects that resemble Bragg scatter. The dimensions of the phantom are 256 nm, 512 nm, and 512 nm – along z , x , and y axes, respectively.




Figure 1: 0° tilt of the aluminum spheres (simulated) dataset.

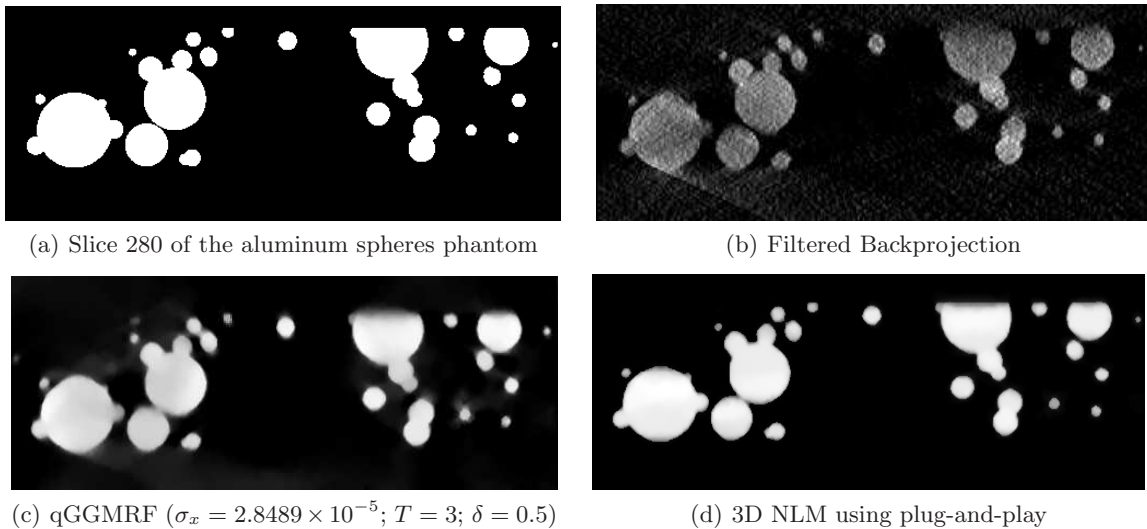


Figure 2: Tomographic reconstruction of slice 280 of the simulated dataset. For visually equivalent regularization, NLM-based reconstruction is clearer and relatively artifact-free.

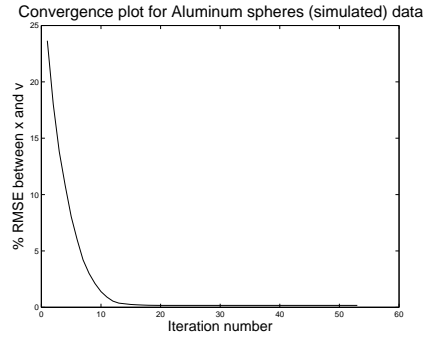


Figure 3: The convergence rate of the plug-and-play algorithm with 3D NLM as the prior model.

Fig. 1 shows a 0° tilt image of the simulated TEM data. Since this is a bright-field image, the aluminum spheres appear dark against a bright background. Fig. 2 shows the ground truth along with three reconstructions of slice 280 along the x - z plane. The NLM-based reconstruction has no smear artifacts, and also has a lower RMSE value of $0.2531 \times 10^{-3} \text{ nm}^{-1}$ compared to $0.4581 \times 10^{-3} \text{ nm}^{-1}$ for qGGMRF, and $8.3988 \times 10^{-3} \text{ nm}^{-1}$ for filtered backprojection. The edges are also sharper in the NLM reconstruction.

4.2 Aluminum spheres (real) dataset

The aluminum spheres dataset (see Fig. 4) has 67 equally-spaced tilts about the y -axis, spanning $[-70^\circ, +60^\circ]$.





Figure 4: 0° tilt of the very noisy aluminum spheres (real) dataset.



(a) Filtered Backprojection



(b) qGGMRF ($\sigma_x = 1.2 \times 10^{-4}$; $T = 3$; $\delta = 0.5$)



(c) 3D NLM using plug-and-play

Figure 5: Tomographic reconstruction of slice 307 of the aluminum spheres dataset. For equivalent regularization, NLM-based reconstruction is clearer and has less smear artifacts.

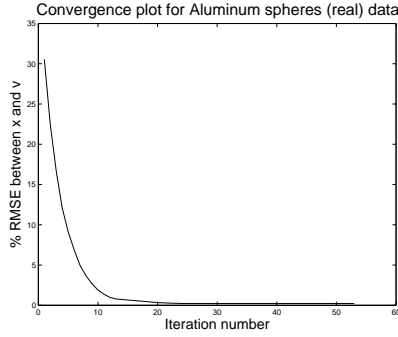


Figure 6: The convergence rate of the plug-and-play algorithm with 3D NLM as the prior model.

Fig. 4 shows a 0° tilt image of the real aluminum spheres TEM data. Fig. 5 shows three reconstructions of slice 307 along the x-z plane. The NLM-based reconstruction has less smear artifacts than the qGGMRF reconstruction, and far more clarity than the filtered backprojection reconstruction.

4.3 Ferritin samples (real) dataset

The ferritin spheres dataset (see Fig. 7) has 57 equally-spaced tilts about the y -axis, spanning $[-70^\circ, +70^\circ]$.






Figure 7: 0° tilt of the contrast-stretched version of the ferritin (real) dataset.



(a) Filtered Backprojection



(b) qGGMRF ($\sigma_x = 3.77 \times 10^{-4}$; $T = 3$; $\delta = 0.5$)



(c) 3D NLM using plug-and-play

Figure 8: Tomographic reconstruction of a portion of slice 1149 of the ferritin dataset. For equivalent regularization, NLM-based reconstruction is clearer and has less smear artifacts.

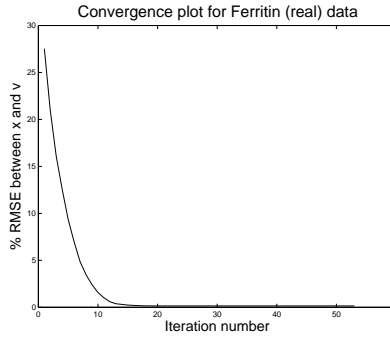


Figure 9: The convergence rate of the plug-and-play algorithm with 3D NLM as the prior model.

Fig. 7 shows a 0° tilt image of the real ferritin structures TEM data. Fig. 8 shows three reconstructions of slice 1149 along the x-z plane. The NLM-based reconstruction has less blur compared the qGGMRF reconstruction, and is far more clear than the filtered backprojection reconstruction. Also, the core of the ferritin structures are sharp and distinct in the NLM reconstruction.

5. CONCLUSIONS

In this paper, we augmented the current MBIR solution to bright field electron tomography by incorporating plug-and-play capabilities, thus enabling the application of a wide variety of spatial constraints as prior models, including those models that are not necessarily expressed as solutions to optimization problems. We implemented 3D non-local means as a prior model, and achieved high quality tomographic reconstructions in all of the three experiments involving both real and simulated datasets. NLM resulted in visibly suppressed smear artifacts and greater clarity (for all three datasets), and lower RMSE values (for the simulated dataset) than the qGGMRF and filtered backprojection reconstructions.

ACKNOWLEDGMENT

This work was supported by an AFOSR/MURI grant #FA9550-12-1-0458, by UES Inc. under the Broad Spectrum Engineered Materials contract, and by the Electronic Imaging component of the ICMD program of the Materials and Manufacturing Directorate of the Air Force Research Laboratory, Andrew Rosenberger, program manager.

We thank Marc DeGraef, Professor of Materials Science and Engineering at Carnegie Mellon University, for the simulated aluminum spheres dataset.

REFERENCES

1. Grünewald, K., Desai, P., Winkler, D. C., Heymann, J. B., Belnap, D. M., Baumeister, W., and Steven, A. C., “Three-dimensional structure of herpes simplex virus from cryo-electron tomography,” *Science* **302**(5649), 1396–1398 (2003).
2. Bárcena, M. and Koster, A. J., “Electron tomography in life science,” *Seminars in Cell & Developmental Biology* **20**(8), 920 – 930 (2009).
3. Midgley, P. A. and Dunin-Borkowski, R. E., “Electron tomography and holography in materials science.,” *Nature Materials* **8**(4), 271–280 (2009).
4. Mohammad-Djafari, A., “Joint estimation of parameters and hyperparameters in a bayesian approach of solving inverse problems,” in [*Image Processing, 1996. Proceedings., International Conference on*], **1**, 473–476 vol.2 (Sept. 1996).
5. Bouman, C. and Sauer, K., “A generalized Gaussian image model for edge-preserving MAP estimation,” *IEEE Trans. on Image Processing* **2**, 296–310 (July 1993).
6. Bouman, C. A., [*Model Based Image And Signal Processing*] (2013).
7. Sullivan, S. Z., Muir, R. D., Newman, J. A., Carlsen, M. S., Sreehari, S., Doerge, C., Begue, N. J., Everly, R. M., Bouman, C. A., and Simpson, G. J., “High frame-rate multichannel beam-scanning microscopy based on lissajous trajectories,” *Opt. Express* **22**, 24224–24234 (Oct 2014).
8. Oh, S., Milstein, A., Bouman, C., and Webb, K., “A general framework for nonlinear multigrid inversion,” *IEEE Trans. on Image Processing* **14**, 125–140 (Jan. 2005).
9. Mumcuoglu, E., Leahy, R., Cherry, S., and Zhou, Z., “Fast gradient-based methods for Bayesian reconstruction of transmission and emission PET images,” *IEEE Trans. on Medical Imaging* **13**, 687–701 (Dec. 1994).
10. Buades, A., Coll, B., and Morel, J.-M., “A review of image denoising algorithms, with a new one,” *Multiscale Modeling & Simulation* **4**(2), 490–530 (2005).
11. Wong, A. and Orchard, J., “A nonlocal-means approach to exemplar-based inpainting,” in [*Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*], 2600–2603, IEEE (2008).
12. Mairal, J., Bach, F., Ponce, J., Sapiro, G., and Zisserman, A., “Non-local sparse models for image restoration,” in [*Computer Vision, 2009 IEEE 12th International Conference on*], 2272–2279, IEEE (2009).
13. Coupé, P., Yger, P., Prima, S., Hellier, P., Kervrann, C., and Barillot, C., “An optimized blockwise nonlocal means denoising filter for 3-d magnetic resonance images,” *Medical Imaging, IEEE Transactions on* **27**(4), 425–441 (2008).
14. Tian, Z., Jia, X., Dong, B., Lou, Y., and Jiang, S. B., “Low-dose 4DCT reconstruction via temporal nonlocal means,” *Medical physics* **38**(3), 1359–1365 (2011).
15. Chen, Y., Ma, J., Feng, Q., Luo, L., Shi, P., and Chen, W., “Nonlocal prior bayesian tomographic reconstruction,” *Journal of Mathematical Imaging and Vision* **30**(2), 133–146 (2008).
16. Chun, S., Dewaraja, Y., and Fessler, J., “Alternating direction method of multiplier for emission tomography with non-local regularizers,” in [*Proc. Intl. Mtg. on Fully 3D Image Recon. in Rad. and Nuc. Med*], 62–5 (2013).
17. Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J., “Distributed optimization and statistical learning via the alternating direction method of multipliers,” *Foundations and Trends® in Machine Learning* **3**(1), 1–122 (2011).
18. Glowinski, R. and Marroco, A., “Sur l’approximation, par lments finis d’ordre un, et la rsolution, par pnalisation-dualit d’une classe de problemes de dirichlet non linaires,” *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique* **9**(R2), 41–76 (1975).
19. Gabay, D. and Mercier, B., “A dual algorithm for the solution of nonlinear variational problems via finite element approximation,” *Computers & Mathematics with Applications* **2**(1), 17–40 (1976).
20. Eckstein, J. and Bertsekas, D. P., “On the douglasrachford splitting method and the proximal point algorithm for maximal monotone operators,” *Mathematical Programming* **55**(1-3), 293–318 (1992).
21. Fair, R. C., “On the robust estimation of econometric models,” in [*Annals of Economic and Social Measurement, Volume 3, number 4*], 117–128, NBER (1974).

22. Yang, Z. and Jacob, M., "Nonlocal regularization of inverse problems: a unified variational framework," *Image Processing, IEEE Transactions on* **22**(8), 3192–3203 (2013).
23. Venkatakrishnan, S. V., Bouman, C. A., and Wohlberg, B., "Plug-and-play priors for model based reconstruction," in [*Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*], 945–948, IEEE (2013).
24. Venkatakrishnan, S. V., Drummy, L. F., De Graef, M., Simmons, J. P., and Bouman, C. A., "Model based iterative reconstruction for bright field electron tomography," *To appear in Computational Imaging, IEEE Transactions on* (2014).
25. Buades, A., Coll, B., and Morel, J.-M., "A non-local algorithm for image denoising," in [*Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*], **2**, 60–65, IEEE (2005).
26. Thibault, J.-B., Sauer, K. D., Bouman, C. A., and Hsieh, J., "A three-dimensional statistical approach to improved image quality for multislice helical ct," *Medical physics* **34**(11), 4526–4544 (2007).