

Plug-and-Play Priors for Bright Field Electron Tomography and Sparse Interpolation

Suhas Sreehari*, *Student Member, IEEE*, S. V. Venkatakrishnan, *Member, IEEE*, Brendt Wohlberg, *Senior Member, IEEE*, Lawrence F. Drummy, Jeffrey P. Simmons, *Member, IEEE*, and Charles A. Bouman, *Fellow, IEEE*

Abstract—Many material and biological samples in scientific imaging are characterized by non-local repeating structures. These are studied using scanning electron microscopy and electron tomography. Sparse sampling of individual pixels in a 2D image acquisition geometry, or sparse sampling of projection images with large tilt increments in a tomography experiment, can enable high speed data acquisition and minimize sample damage caused by the electron beam.

In this paper, we present an algorithm for electron tomographic reconstruction and sparse image interpolation that exploits the non-local redundancy in images. We adapt a framework, termed plug-and-play (P&P) priors, to solve these imaging problems in a regularized inversion setting. The power of the P&P approach is that it allows a wide array of modern denoising algorithms to be used as a “prior model” for tomography and image interpolation. We also present sufficient mathematical conditions that ensure convergence of the P&P approach, and we use these insights to design a new non-local means denoising algorithm. Finally, we demonstrate that the algorithm produces higher quality reconstructions on both simulated and real electron microscope data, along with improved convergence properties compared to other methods.

Index Terms—Plug-and-play, prior modeling, bright field electron tomography, sparse interpolation, non-local means, doubly-stochastic gradient non-local means, BM3D.

I. INTRODUCTION

TRANSMISSION electron microscopes are widely used for characterization of material and biological samples at the nano-meter scale [1]–[3]. In many cases, these electron microscopy samples contain many repeating structures that are similar or identical to each other. High quality reconstruction of these samples from tomographic projections is possible by exploiting the redundancy caused by repeating structures. As an important example, cryo-electron microscope (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.

Another important imaging problem is that raster scanning an electron beam across a large field of view is time consuming and can damage the sample. For this reason, there is growing interest in reconstructing full resolution images from sparsely sampled pixels [4], [5]. The redundancy in material and biological samples suggests that it is possible to reconstruct such images with sufficient fidelity by acquiring only a few random samples in the image and using an advanced image reconstruction algorithm that exploits non-local redundancies.

Conventionally, model-based iterative reconstruction (MBIR) solves a single optimization problem that tightly couples the log likelihood term (based on the data) and the log of the prior probability [6]–[14]. MBIR can, in principle, exploit redundancy in microscope images for tomographic reconstruction. This requires selection of the appropriate log prior probability, which is very challenging in practice. Patch-based denoising algorithms such as non-local means (NLM) [15]–[17] and BM3D [18] have been very successful in exploiting non-local redundancy in images. However, since NLM and BM3D are not explicitly formulated as cost functions, it is unclear how to use them as prior models in the MBIR framework. Venkatakrishnan et al. [19] developed a semi-empirical framework termed plug-and-play priors to incorporate such algorithms into general inverse problems, but limited results were presented and the convergence of the algorithm was not discussed. Chen et al. [20] proposed an MRF-style prior, but with non-local spatial dependencies, to perform Bayesian tomographic reconstruction. The authors adopted a two-step optimization involving non-local weight update, followed by the image update. However, the cost function changes every iteration, so that there is no single fixed cost function that is minimized. Chun et al. [21] proposed non-local regularizers for emission tomography based on alternating direction method of multipliers (ADMM) [22]–[25], using Fair potential [26] 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 [27], resulting in a model that captures the intrinsically non-convex behavior required for modeling distant particles with similar structure. However, it is not clear under what conditions their method converges. Non-local regularizers using PDE-like evolutions

S. Sreehari and C. A. Bouman are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, 47907 USA
*E-mail: ssreehar@purdue.edu

S. V. Venkatakrishnan is with Lawrence Berkeley National Laboratory, Berkeley, CA.

B. Wohlberg is with Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM.

L. F. Drummy and J. P. Simmons are with Air Force Research Laboratory, Dayton, OH.

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.

and total variation are proposed to solve inverse problems [28], [29].

Image interpolation is also a widely researched problem [30]. The approaches can be broadly classified into two categories - those based on local regularization and those on non-local regularization. In local approaches, the missing pixels are reconstructed from an immediate neighborhood surrounding the unknown values to encourage similarity between spatially neighboring pixels [31]. Spurred by the success of non-local means, there have been several efforts to solve the sparse interpolation problem using global patch based dictionary models [5], [32]–[35]. Li et al. [36] adapted a two stage approach similar to [37] to the problem of sparse image reconstruction using the BM3D denoising algorithm. However, this approach is not immediately applicable to denoising operators such as NLM and those formulated using a nonparametric point estimation framework [38]–[40]. The simplicity and success of NLM and BM3D has also led to the question of how these algorithms can be used to solve other inverse problems. In fact, Danielyan et al. [37] have adapted BM3D for image deblurring through the optimization of two cost functions balanced by the generalized Nash equilibrium.

In this paper, we present an algorithm for tomographic reconstruction and sparse image interpolation that exploits the non-local redundancies in microscope images. Our solution uses the plug-and-play (P&P) framework [19], [41] which is based on the alternating direction method of multipliers (ADMM) [23], [24] and decouples the forward model and the prior terms in the optimization procedure. This results in an algorithm that involves repeated application of two steps: an inversion step only dependent on the forward model, and a denoising step only dependent on the image prior model. The P&P takes ADMM one step further by replacing the prior model optimization by a denoising operator. However, while it is convenient to be able to use any denoising operator as a prior model, this new framework also begs the question as to whether P&P necessarily inherits the convergence properties of ADMM? We answer this important question by presenting a theorem that outlines the sufficiency conditions to be satisfied by the denoising operator in order to guarantee convergence of the P&P algorithm. We also present a proof for this convergence theorem partly based on the ideas presented by Moreau [42] and Williamson et al. [43]. Using this result, we then modify NLM to satisfy these sufficiency conditions and call it doubly-stochastic gradient NLM (DSG-NLM). We then apply DSG-NLM as a prior model to the tomographic reconstruction and sparse interpolation problems. This new DSG-NLM algorithm is based on symmetrizing the filter corresponding to the traditional NLM algorithm. Interestingly, Milanfar [44] has also discussed the benefit of symmetrizing the denoising operator, albeit in the context of improving the performance of heuristic denoising algorithms.

The plug-and-play electron tomography solution presented in this paper builds on the existing MBIR framework for bright field electron tomography [14], which models Bragg scatter and anomaly detection. We demonstrate that our proposed algorithm produces high quality tomographic reconstructions and interpolation on both simulated and real electron micro-

scope images. Additionally our method has improved convergence properties compared to using the standard NLM or the BM3D algorithm as a regularizer for the reconstruction. 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 a wide variety of spatial constraints to solve a wide variety of inverse problems.

II. PLUG-AND-PLAY FRAMEWORK

Let $x \in \mathbb{R}^N$ be an unknown image with a prior distribution given by $p(x)$, and let $y \in \mathbb{R}^M$ be the associated measurements of the image with conditional distribution given by $p(y|x)$. We will refer to $p(y|x)$ as the forward model for the measurement system. Then the maximum *a posteriori* (MAP) estimate of the image x is given by

$$\hat{x}_{MAP} = \underset{x \in \mathbb{R}^N}{\operatorname{argmin}} \{l(x) + \beta s(x)\}, \quad (1)$$

where $l(x) = -\log p(y|x)$, $\beta s(x) = -\log p(x)$, and β is a positive scalar used to control the level of regularization in the MAP reconstruction. In order to allow for the possibility of convex constraints, we will allow both $l(x)$ and $s(x)$ to take values on the extended real line, $\mathbb{R} \cup \{+\infty\}$. Using this convention, we can, for example, enforce positivity by setting $l(x) = +\infty$ for $x \leq 0$.

Splitting the variable x of equation (1) results in an equivalent expression for the MAP estimate given by

$$(\hat{x}, \hat{v}) = \arg \min_{\substack{x, v \in \mathbb{R}^N \\ x=v}} \{l(x) + \beta s(v)\}. \quad (2)$$

This contained optimization problem can then be computed by solving the following unconstrained augmented Lagrangian cost function given by

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

where u must be chosen to meet the constraint of $x = v$, and $\sigma_\lambda > 0$ is the augmented Lagrangian parameter¹.

It is well known that the solution to equation (3) may be computed using the ADMM algorithm. For this particular problem, the ADMM algorithm consists of iteration over the following steps:

$$\hat{x} \leftarrow \arg \min_{x \in \mathbb{R}^N} L_\lambda(x, \hat{v}; u) \quad (4)$$

$$\hat{v} \leftarrow \arg \min_{v \in \mathbb{R}^N} L_\lambda(\hat{x}, v; u) \quad (5)$$

$$u \leftarrow u + (\hat{x} - \hat{v}). \quad (6)$$

In fact, if $l(x)$ and $s(x)$ are both proper, closed, and convex functions, and a saddle point solution exists [23]–[25], then it is well known that the ADMM converges to the global minimum.

¹ The augmented Lagrangian parameter, σ_λ , is related to the ADMM penalty parameter, λ , through a simple expression: $\sigma_\lambda = \frac{1}{\sqrt{\lambda}}$.

We can express the ADMM iterations more compactly by defining two operators. The first is an inversion operator F defined by

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

and the second is a denoising operator H given by

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

where $\sigma_n = \sqrt{\beta}\sigma_\lambda$ has the interpretation of being the assumed noise standard deviation in the denoising operator. Moreover, we say that H is the proximal mapping for the proper, closed, and convex function $s : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$.

Using these two operators, we can easily derive the plug-and-play algorithm shown in Algorithm 1 as an alternative form of the ADMM iterations. This formulation has a number of practical and theoretical advantages. First, in this form we can now “plug in” denoising operators that are not in the explicit form of the optimization of equation (8). So for example, we will later see that popular and effective denoising operators such as non-local means (NLM) [45] or BM3D [18], which are not easily represented in an optimization framework can be used in the plug-and-play iterations. Second, this framework allows for decomposition of the problem into separate software systems for the implementation of the inversion operator, F , and the denoising operator, H . In practice, as software systems for large inversion problems become more complex, the ability to decompose them into separate modules, while retaining the global optimality of the solution, can be extremely valuable.

The plug-and-play algorithm requires the selection of two parameters, β and σ_λ , and then the $\sigma_n = \sqrt{\beta}\sigma_\lambda$. The unitless parameter β can typically be chosen to be near 1, with larger or smaller values producing more or less regularization, respectively. In theory, the value of σ_λ does not affect the reconstruction for a convex optimization problem, but in practice, a well-chosen value of σ_λ can substantially speed up ADMM convergence [25], [46], [47]; so the careful choice of σ_λ is important. Our approach is to choose the value of σ_λ to be approximately equal to the amount of variation in the reconstruction. Formally stated, we choose

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

This choice for the value of σ_λ^2 is motivated by its role as the inverse regularizer in equation (7). In practice, this can be done by first computing an approximate reconstruction using some baseline algorithm, and then computing the sample variance in the approximate reconstruction.

Of course, for an arbitrary denoising algorithm, the question remains of whether the plug-and-play algorithm converges? The following section provides practical conditions for the denoising operator to meet that ensure convergence of the iterations.

Algorithm 1 Plug-and-play algorithm for implementation of a general forward model $F(\tilde{x}; \sigma_\lambda)$, and a prior model specified by the denoising operator in $H(\tilde{v}; \sigma_n)$.

```

initialize  $\hat{v}$ 
 $u \leftarrow 0$ 
while not converged do
   $\tilde{x} \leftarrow \hat{v} - u$ 
   $\hat{x} \leftarrow F(\tilde{x}; \sigma_\lambda)$ 
   $\tilde{v} \leftarrow \hat{x} + u$ 
   $\hat{v} \leftarrow H(\tilde{v}; \sigma_n)$ 
   $u \leftarrow u + (\hat{x} - \hat{v})$ 
end while

```

III. CONVERGENCE OF THE PLUG-AND-PLAY ALGORITHM

It is well known that the ADMM algorithm is guaranteed to converge under appropriate technical conditions. For example, if the optimization problem is convex and a saddle point solution exists then the iterations of ADMM converge [23]–[25]. However, in our plug-and-play approach, we will be using general denoising algorithms to implement the operator $H(\tilde{v}; \sigma_n)$, and therefore, the function $s(x)$ is not available to inspect. This raises the question of what conditions must $H(\tilde{v}; \sigma_n)$ and $l(y; x)$ must satisfy in order to ensure that the plug-and-play algorithm converges.

In the following theorem, we give conditions on both the log likelihood function, $l(x)$, and the denoising operator, $H(x)$, that are sufficient to guarantee convergence of the plug-and-play algorithm to the global minimum of some implicitly defined MAP cost function. This is interesting because it does not ever require that one know or explicitly specify the function $s(x)$. Instead, $s(x)$ is implicitly defined through the choice of $H(x)$.

Theorem III.1. *Let the negative log likelihood function $l : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$ and the denoising operator $H : \mathbb{R}^N \rightarrow \mathbb{R}^N$ meet the following conditions:*

- 1) $H(x)$ is a continuously differentiable function on \mathbb{R}^N ;
- 2) $\forall x \in \mathbb{R}^N$, $\nabla H(x)$ is a doubly stochastic matrix;
- 3) There exist a y in the range of H such that $l(y) < \infty$;
- 4) $l(x)$ is a proper closed convex function which is lower bounded by a function $f(\|x\|)$ such that $f(x)$ is monotone increasing with

$$\lim_{\alpha \rightarrow \infty} \frac{f(\alpha)}{\alpha} = \infty.$$

Then the following results hold:

- 1) H is a proximal mapping for some proper closed convex function $s(x)$;
- 2) There exists a MAP estimate, \hat{x}_{MAP} , such that

$$p^* = \inf_{x \in \mathbb{R}^N} \{l(x) + \beta s(x)\} = l(\hat{x}_{MAP}) + \beta s(\hat{x}_{MAP});$$

- 3) The plug-and-play algorithm converges in the following sense,

$$\begin{aligned} \lim_{k \rightarrow \infty} \{\hat{x}^{(k)} - \hat{v}^{(k)}\} &= 0; \\ \lim_{k \rightarrow \infty} \{l(\hat{x}^{(k)}) + \beta s(\hat{v}^{(k)})\} &= p^*, \end{aligned}$$

where $\hat{x}^{(k)}$ and $\hat{v}^{(k)}$ denote the result of the k^{th} iteration.

The proof of this theorem, which is presented in Appendix A, depends on a powerful theorem proved by Moreau in 1965 [42]. This theorem states that H is a proximal mapping if and only if it is non-expansive and the sub-gradient of a convex function on \mathbb{R}^N . Intuitively, once we can show that the denoising operator, H , is a proximal mapping, then we know that is effectively implementing an update step from of the ADMM algorithm of equation (8).

The first and second conditions of the theorem ensure that the conditions of Moreau's theorem are met. This is because the doubly stochastic structure of $H(x)$ ensures that H is the gradient of some function ϕ , that ϕ is convex, and that H is non-expansive.

The additional two conditions of Theorem III.1 ensure that the MAP estimate actually exists for the problem. Importantly, this is done without explicit reference to the prior function $s(x)$. More specifically, the third condition ensures that the set of feasible solutions is not empty, and the fourth condition ensures that the MAP cost function takes on its global minimum value, i.e., that the minimum is not achieved toward infinity.

Importantly, in the next section, we will show that real denoising operators can be modified to meet the conditions of this theorem. In particular, the symmetrized non-local means filters investigated by Milanfar [44] are designed to create a symmetric gradient.

IV. NON-LOCAL MEANS DENOISING WITH DOUBLY STOCHASTIC GRADIENT

In order to satisfy the conditions for convergence, the gradient of the denoising operator must be a doubly-stochastic matrix. However, the standard NLM denoising algorithm does not satisfy this condition. So in this section, we introduce a simple modification of the NLM approach, which we refer to as the doubly stochastic gradient NLM (DSG-NLM), that satisfies the required convergence conditions. Interestingly, the symmetrized non-local means filters investigated by Milanfar [44] also achieve a symmetric gradient, but requires the use of a more complex iterative algorithm to symmetrize the operator.

The NLM algorithm is known to produce much higher quality results than traditional local smoothing-based denoising methods [45]. It works by estimating each pixel ² as a weighted mean of all pixels in the image ³. In this section, \tilde{v} will denote a noisy image with voxel values v_s at locations $s \in S$. Generally, S is a discrete lattice, so for 2D images $S = \mathbb{Z}^2$ and for 3D volumes $S = \mathbb{Z}^3$.

Using this notation, the NLM denoising method can be represented as

$$\hat{v}_s = \sum_{r \in \Omega_s} w_{s,r} \tilde{v}_r, \quad (10)$$

where \hat{v}_s is the denoised result, the coefficients $w_{s,r}$ are the NLM weights, and Ω_s is the NLM search window defined by

$$\Omega_s = \{r \in S : \|r - s\|_\infty \leq N_s\}.$$

²All discussion remains valid even if we consider voxels instead of pixels.

³In practice, we only compute the weighted mean of pixels in a search window, instead of the whole image.

Note that the integer N_s controls the size of the NLM search window. In general, larger values of N_s can yield better results but at the cost of higher computational cost.

Using this notation, the plug-and-play denoising operator is given by

$$H(\tilde{v}; \sigma_n) = W\tilde{v},$$

where the matrix

$$W_{s,r} = \begin{cases} w_{s,r} & \text{if } r \in \Omega_s \\ 0 & \text{otherwise} \end{cases}.$$

Now if we fix the weights, then it is clear that

$$\nabla H(v; \sigma_n) = W.$$

So the condition 2 of Theorem III.1, simply requires that W be a doubly stochastic matrix.

In fact, W is guaranteed to be doubly-stochastic if it is symmetric with positive entries and rows (or columns) that sum to 1. The following modified procedure for computing the NLM weights ensures that these properties hold. We start by defining $P_s \in \mathbb{R}^{N_p^2}$ to be a patch of size $N_p \times N_p$ centered at position s . Then we compute the weights through the following 3-step algorithm.

$$w_{s,r} \leftarrow \exp \left\{ \frac{-\|P_r - P_s\|_2^2}{2N_p^2\sigma_n^2} \right\} \quad (11)$$

$$w_{s,r} \leftarrow \frac{w_{s,r}}{\sqrt{(\sum_{r \in \Omega_s} w_{s,r}) (\sum_{s \in \Omega_r} w_{r,s})}} \quad (12)$$

$$w_{s,s} \leftarrow w_{s,s} - \left(\sum_{r \in \Omega_s} w_{s,r} - 1 \right). \quad (13)$$

Notice that all three steps of equations (11), (12), and (13) are symmetric in s and r , so they produce symmetric weights with the property that $w_{s,r} = w_{r,s}$. While equation (12) results in rows and columns that are approximately normalized, and equation (13) guarantees normalization by adjusting the diagonal coefficient of the matrix W . Theoretically, equation (13) could produce a negative coefficient, but in practice this does not occur in real data for two reasons. First, the diagonal coefficient, $w_{s,s}$ is always the largest value generated in step 1 of equation (11) because $\|P_s - P_s\|_2^2 = 0$. Second, the normalization of equation (12) typically makes the subtracted quantity of equation (13) small.

Therefore, this algorithm generates a matrix W which is symmetric with rows and columns that sum to 1, and in all practical cases, non-negative elements. This makes W a doubly stochastic matrix, so it fulfills condition 2 of Theorem III.1 as is required for guaranteed convergence of the plug-and-play algorithm.

V. 3D BRIGHT FIELD EM FORWARD MODEL

In this section, we formulate the explicit form of the inversion operator, $F(x, \sigma_\lambda)$, for the application of 3D bright field EM tomography. For this problem, we adopted both the forward model and optimization algorithms described in [14].

More specifically, the negative log likelihood function is given by

$$l(x, d, \sigma) = \frac{1}{2} \sum_{k=1}^K \sum_{i=1}^M \beta_{T,\delta} \left((y_{k,i} - A_{k,i,*}x - d_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma} \right) + MK \log(\sigma) + C,$$

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, C is a constant, and $\beta_{T,\delta}(\cdot)$ is the generalized Huber function defined as,

$$\beta_{T,\delta}(x) = \begin{cases} x^2 & \text{if } |x| < T \\ 2\delta T|x| + T^2(1 - \delta) & \text{if } |x| \geq T. \end{cases} \quad (14)$$

The generalized Huber function is used to reject measurements with large errors. This is useful because measurement may vary from the assumed model for many practical reasons. For example, in bright field EM, Bragg scatter can cause highly attenuated measurements that otherwise would cause visible streaks on the reconstruction [48].

To compute the inversion operator F of equation (7), we minimize the cost function below with respect to x , d , and σ .

$$\begin{aligned} c(x, d, \sigma; \tilde{x}, \sigma_\lambda) &= \frac{1}{2} \sum_{k=1}^K \sum_{i=1}^M \beta_{T,\delta} \left((y_{k,i} - A_{k,i,*}x - d_k) \frac{\sqrt{\Lambda_{k,ii}}}{\sigma} \right) \\ &+ MK \log(\sigma) + \frac{\|x - \tilde{x}\|_2^2}{2\sigma_\lambda^2}. \end{aligned} \quad (15)$$

So the inversion operator is computed as

$$F(\tilde{x}; \sigma_\lambda) = \arg \min_{x \geq 0, d, \sigma} c(x, d, \sigma; \tilde{x}, \sigma_\lambda). \quad (16)$$

As in the case of sparse interpolation, we set $c(x, d, \sigma; \tilde{x}, \sigma_\lambda) = +\infty$ for $x < 0$ in order to enforce positivity.

The details of the optimization algorithm required for equation (16) are described in [14]. The optimization algorithm is based on alternating minimization with respect to the three quantities and it uses a majorization based on a surrogate function to handle the minimization of the generalized Huber function [49].

For this complex problem, we note some practical deviations from the theory. First, the negative log likelihood function, $l(x)$, is not convex in this case, so the assumptions of the plug-and-play convergence do not hold. Moreover, with such a non-convex optimization, it is not possible to guarantee convergence to a global minimum, but in practice most optimization algorithms generate very good results. In addition, this cost function also violates condition 4 of Theorem III.1 because it only grows at a linear rate as $\|x\| \rightarrow +\infty$. Again, this condition is used to guarantee that the plug-and-play algorithm does not drift off to a minimum tending to infinity.

However, in practice, we have never observed this to happen with real data sets and useful denoising operators. Finally, the global optimization of equation (16) is approximated by three iterations of alternating minimization with respect to x , d , and σ . Nonetheless, in our experimental results section, we will illustrate our empirical observation that the plug-and-play algorithm consistently converges even with these approximations to the ideal case.

VI. SPARSE INTERPOLATION FORWARD MODEL

In this section, we formulate the explicit form of the inversion operator, $F(x, \sigma_\lambda)$, for the application of sparse interpolation. More specifically, our objective will be to recover and image $x \in \mathbb{R}^N$ from a noisy and sparsely subsampled version denoted by $y \in \mathbb{R}^M$ where $M \ll N$. More formally, the forward model for this problem is given by

$$y = Ax + \epsilon, \quad (17)$$

where $A \in \mathbb{R}^M \times \mathbb{R}^N$ matrix. Each entry $A_{i,j}$ is either 1 or 0 depending on if the j^{th} pixel is taken as the i^{th} measurement. Also, each row of A has exactly one non-zero entry, and each column of A may either be empty or have one non-zero entry. We also define $I(j) = \sum_i A_{i,j}$ so that $I(j) = 1$ when the j^{th} pixel is sampled, and $I(j) = 0$, if it is not. Furthermore, ϵ is an M -dimensional vector of i.i.d. Gaussian random variables with mean zero and variance σ_w^2 .

For such a sparse sampling system, we can write the negative log likelihood function as

$$l(x) = \frac{1}{2\sigma_w^2} \|y - Ax\|_2^2 + C, \quad (18)$$

where C is a constant. In order to enforce positivity, we also modify the negative likelihood function by setting $l(x) = +\infty$ for $x < 0$. We include positivity in $l(x)$ rather than in the denoising operator so that H remains continuously differentiable. Using equation (7), the interpolation inversion operator is given by

$$F(\tilde{x}; \sigma_\lambda) = \operatorname{argmin}_{x \geq 0} \left\{ \frac{1}{2\sigma_w^2} \|y - Ax\|_2^2 + \frac{1}{2\sigma_\lambda^2} \|x - \tilde{x}\|_2^2 \right\}.$$

Due to the simple structure of the matrix A , we can also calculate an explicit pixel-wise expression for F . Moreover, if we let $\sigma_w^2 = 0$, then F reduces to the following form

$$F_i(\tilde{x}; \sigma_\lambda) = \begin{cases} [y_i]_+ & \text{if } I(i) = 1 \\ [\tilde{x}_i]_+ & \text{if } I(i) = 0 \end{cases}. \quad (19)$$

where $[\cdot]_+$ represents zeroing of any negative argument. In this case, the interpolation is forced to take on the measured values at the sample points.

VII. RESULTS

In this section, we present experimental results on both real and simulated data for the applications of bright-field EM tomography and sparse interpolation. For all experiments, we present convergence plots that compare both primal and dual residual convergence resulting from using different priors. The

normalized primal and dual residues [25, p. 18], $r^{(k)}$ and $s^{(k)}$ respectively, at the k -th iteration of the P&P algorithm are given by

$$r^{(k)} = \frac{\|\hat{x}^{(k)} - \hat{v}^{(k)}\|_2}{\|\hat{x}^{(\infty)}\|_2}; \quad (20)$$

$$s^{(k)} = \frac{\|\hat{v}^{(k)} - \hat{v}^{(k-1)}\|_2}{\|u^{(k)}\|_2}, \quad (21)$$

where $\hat{x}^{(k)}$, $\hat{v}^{(k)}$, and $u^{(k)}$ are the values of \hat{x} , \hat{v} , and u respectively 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} .

A. Bright Field EM Tomography

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 silicon dioxide. We compare four reconstruction methods – filtered backprojection, MBIR with qGGMRF prior [50], plug-and-play reconstructions with 3D NLM and 3D DSG-NLM as prior models. We used qGGMRF, 3D NLM and 3D DSG-NLM as prior models within the plug-and-play framework. Filtered backprojection was used as the initialization for all MBIR-based reconstructions. All the reconstruction results shown below are x - z slices (i.e., slices parallel to the electron beam). The qGGMRF parameters used for all reconstructions are as follows: $q = 1$, $p = 1.2$, and $c = 0.001$. The NLM and DSG-NLM patch size used for all reconstructions is $5 \times 5 \times 5$. In order to meet the conditions of convergence, we stopped adapting the DSG-NLM weights at 20 iterations of the plug-and-play algorithm. The P&P parameters used are given in Table II.

In all the experiments, we observe from Tables III and IV that the DSG-NLM ensures that the plug-and-play algorithm converges fully, while NLM achieves convergence to within a fraction of a percent.

1) *Aluminum spheres (simulated) dataset*: The aluminum spheres simulated dataset contains 47 equally-spaced tilts about the y -axis, spanning $[-70^\circ, +70^\circ]$. The attenuation coefficient of the spheres are assumed to be 7.45×10^{-3} nm. 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.

Fig. 1 shows a 0° tilt projection 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 and DSG-NLM reconstructions have no shadow artifacts, and also have low RMSE values (see Table I). The edges are also sharper in the NLM and DSG-NLM reconstructions.

Table I: RMSE of the reconstructed Al spheres image compared to the ground truth (after 200 P&P iterations)

FBP	qGGMRF	NLM	DSG-NLM
14.608	4.581	2.531	2.529
$\times 10^{-4} \text{ nm}^{-1}$	$\times 10^{-4} \text{ nm}^{-1}$	$\times 10^{-4} \text{ nm}^{-1}$	$\times 10^{-4} \text{ nm}^{-1}$






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



(a) The aluminum spheres phantom (ground truth)



(b) Filtered Backprojection



(c) qGGMRF ($T = 3$; $\delta = 0.5$)



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



(e) 3D DSG-NLM using plug-and-play

Figure 2: Tomographic reconstruction of the simulated aluminum spheres dataset. NLM and DSG-NLM reconstructions are clearer and relatively artifact-free.




Figure 3: Plug-and-play primal and dual residual convergence for tomographic reconstruction of (simulated) aluminum spheres. DSG-NLM achieves complete convergence.

2) *Aluminum spheres (real) dataset*: The aluminum spheres dataset (see Fig. 4) has 67 equally-spaced tilts about the y -axis, spanning $[-65^\circ, +65^\circ]$. Fig. 4 shows a 0° tilt projection of the real aluminum spheres TEM data. Fig. 5 shows three reconstructions along the x - z plane. The NLM-based reconstruction has less smear artifacts than the qGGMRF reconstruction, and more clarity than the filtered backprojection reconstruction. Also, the NLM and DSG-NLM reconstructions have visibly suppressed missing-wedge artifact.





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




(a) Filtered Backprojection



(b) qGGMRF ($T = 3$; $\delta = 0.5$)



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



(d) 3D DSG-NLM using plug-and-play

Figure 5: Tomographic reconstruction of the real aluminum spheres dataset. NLM and DSG-NLM reconstructions are clearer and have less smear and missing-wedge artifacts.

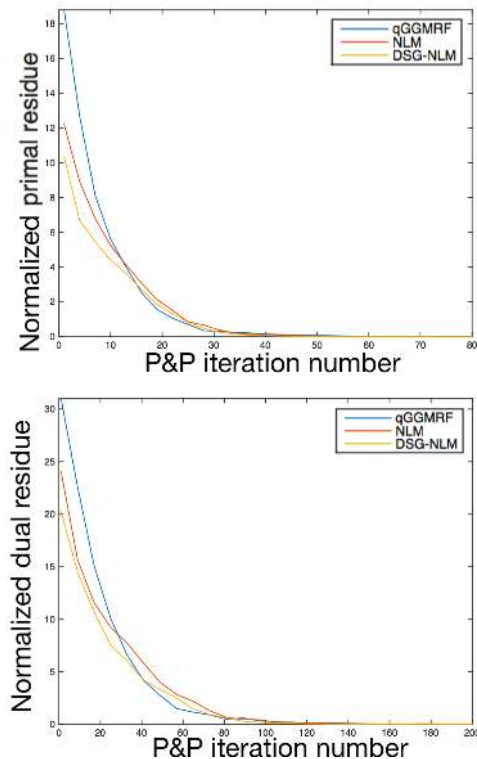


Figure 6: Plug-and-play primal and dual residual convergence for tomographic reconstruction of (real) aluminum spheres. DSG-NLM achieves complete convergence.

3) *Silicon dioxide (real) dataset*: The silicon dioxide dataset (see Fig. 7) has 31 tilts about the y -axis, spanning $[-65^\circ, +65^\circ]$.





Figure 7: Contrast-adjusted version of the 0° tilt of the silicon dioxide (real) dataset.


Fig. 7 shows a 0° tilt projection of the real silicon dioxide TEM data. Fig. 8 shows three reconstructions along the x - z plane. The NLM and DSG-NLM reconstructions have less smear artifacts than the qGGMRF reconstruction, and far more clarity than the filtered backprojection reconstruction.



(a) Filtered Backprojection



(b) qGGMRF ($T = 3$; $\delta = 0.5$)



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



(d) 3D DSG-NLM using plug-and-play

Figure 8: Tomographic reconstruction of the silicon dioxide dataset. NLM reconstruction is clearer and has less smear artifacts. DSG-NLM reconstruction improves upon the NLM result through clear reconstruction of the structure on the left.

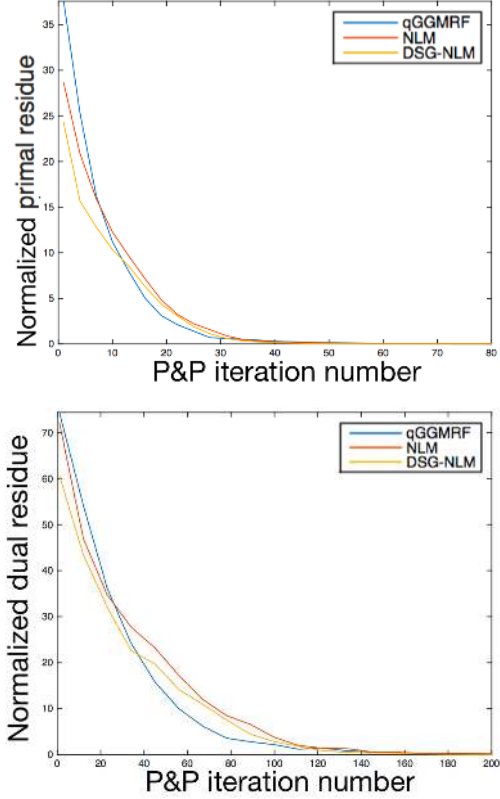


Figure 9: Plug-and-play primal and dual residual convergence for tomographic reconstruction of (real) silicon dioxide. DSG-NLM achieves complete convergence.

Table II: Plug-and-play parameters for tomographic reconstructions

	Al spheres (simulated)	Al spheres (real)	Silicon dioxide (real)
σ_λ (nm^{-1})	8.66×10^{-4}	8.66×10^{-4}	8.66×10^{-4}
β	3.68	4.77	4.25

Table III: Normalized primal residual convergence error for tomography experiments (after 200 P&P iterations)

Dataset	qGGMRF	NLM	DSG-NLM
Al spheres (simulated)	3.46×10^{-12}	2.12×10^{-3}	2.91×10^{-10}
Al spheres (real)	7.06×10^{-11}	3.66×10^{-4}	6.89×10^{-9}
Silicon dioxide (real)	4.99×10^{-12}	8.12×10^{-3}	4.42×10^{-9}

Table IV: Normalized dual residual convergence error for tomography experiments (after 200 P&P iterations)

Dataset	qGGMRF	NLM	DSG-NLM
Al spheres (simulated)	1.55×10^{-10}	7.22×10^{-3}	8.83×10^{-9}
Al spheres (real)	2.61×10^{-10}	1.12×10^{-3}	3.39×10^{-8}
Silicon dioxide (real)	9.06×10^{-11}	5.49×10^{-2}	5.04×10^{-8}

B. Sparse Interpolation

In this section, we present sparse interpolation results on both simulated and real microscope images. We show that a variety of denoising algorithms like NLM, DSG-NLM, and BM3D can be plugged in as prior models to reconstruct images from sparse samples. In all the sparse interpolation experiments, we stopped adapting the weights of the DSG-NLM after 12 iterations of the plug-and-play algorithm. The P&P parameters used are given in Table V.

Our first dataset is a set of simulated super ellipses that mimic the shapes of several material grains like Ni-Cr-Al alloy [51]. The next dataset is a real microscope image of zinc oxide nano-rods [52]. All the images are scaled to the range $[0, 255]$.

In all experiments, the plug-and-play sparse interpolation results are clearer than Shepard interpolation results. We observe from Table VIII that DSG-NLM typically results in the least RMS interpolation error. The RMSE values are normalized as $\frac{\|x - \hat{x}\|_2}{\|x\|_2}$, where \hat{x} is the interpolated image and x is the ground truth image. Furthermore, we can see from Tables VI and VII that DSG-NLM makes plug-and-play converge fully.

Table V: Plug-and-play parameter, β , for the 10% sampling case.

Image	NLM	DSG-NLM	BM3D
Super ellipses	0.9	0.79	0.55
Zinc oxide nano-rods	0.81	0.74	0.49

Table VI: Normalized primal residual convergence error for the 10% sampling case (after 150 P&P iterations)

Image	NLM	DSG-NLM	BM3D
Super ellipses	1.31×10^{-3}	5.41×10^{-8}	1.20×10^{-3}
Zinc oxide nano-rods	2.02×10^{-3}	3.64×10^{-9}	1.14×10^{-3}

Table VII: Normalized dual residual convergence error for the 10% sampling case (after 150 P&P iterations)

Image	NLM	DSG-NLM	BM3D
Super ellipses	9.10×10^{-3}	3.58×10^{-7}	8.71×10^{-3}
Zinc oxide nano-rods	1.14×10^{-2}	6.33×10^{-8}	3.23×10^{-2}

Table VIII: Interpolation error (after 150 P&P iterations): normalized RMSE of the interpolated image compared to the ground truth

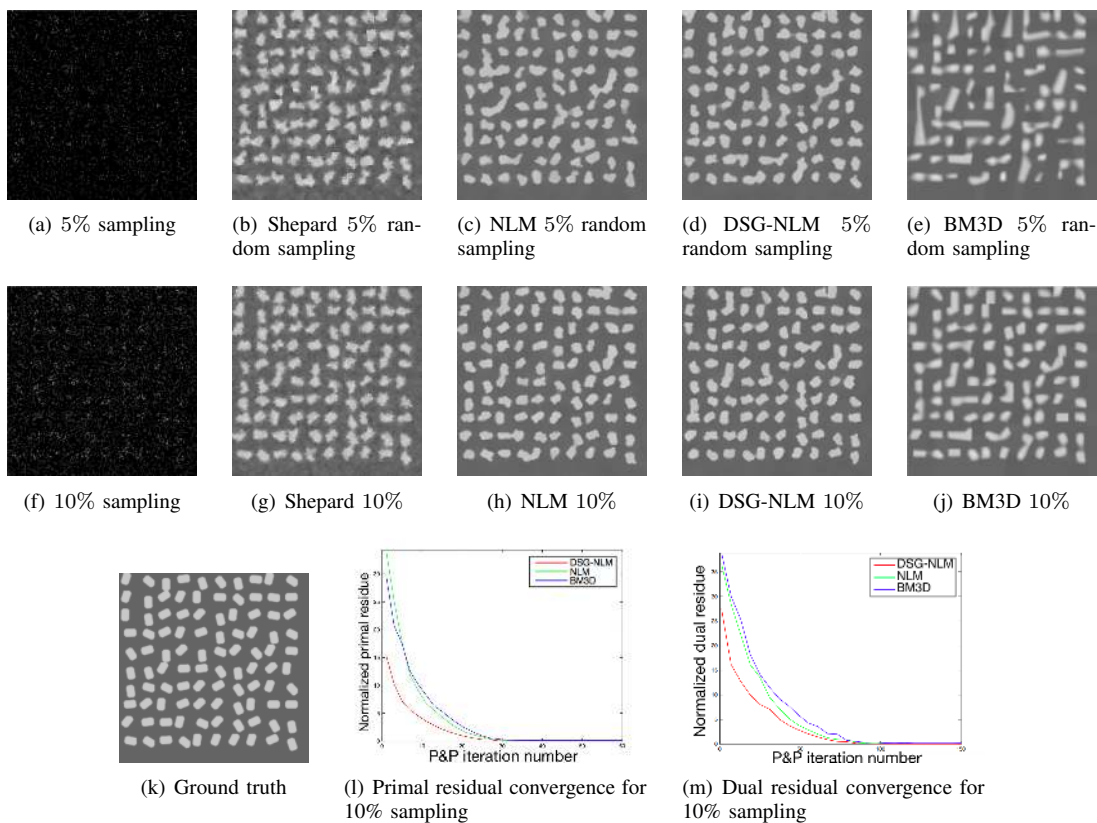
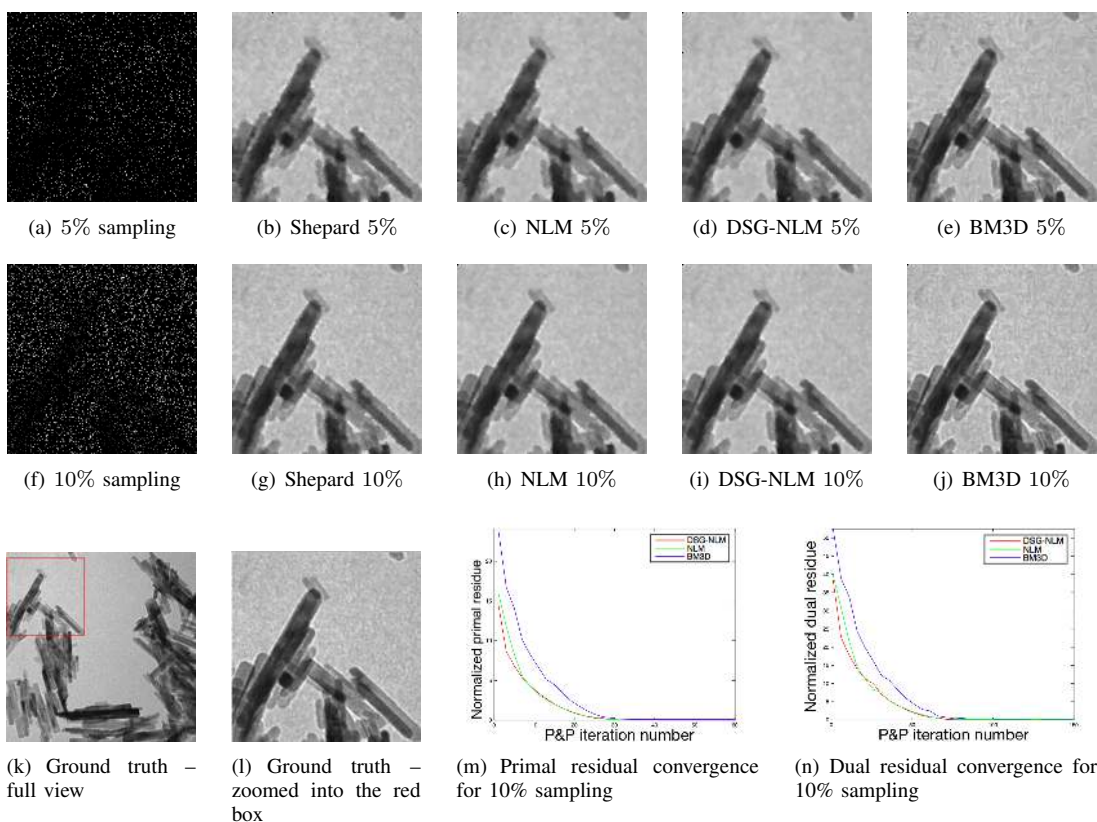
Image	Method	5% random sampling	10% random sampling
Super ellipses	Shepard	10.61%	8.99%
	NLM	8.51%	7.12%
	DSG-NLM	8.33%	6.98%
	BM3D	9.75%	7.46%
Zinc oxide nano-rods	Shepard	6.01%	5.49%
	NLM	4.35%	3.67%
	DSG-NLM	4.18%	3.39%
	BM3D	4.72%	3.80%

VIII. CONCLUSIONS

Microscope images of material and biological samples contain several repeating structures at distant locations. High quality reconstruction of these samples is possible by exploiting non-local repetitive structures. Though model-based iterative reconstruction (MBIR) could in principle exploit these repetitions, practically choosing the appropriate log probability term is very challenging. To solve this problem, we presented the “plug-and-play” (P&P) framework which is based on ADMM. ADMM is a popular method to decouple the log likelihood and the log prior probability terms in the MBIR cost function. Plug-and-play takes ADMM one step further by replacing the optimization step related to the prior model by a denoising operation. This approach has two major advantages: First, it allows the use of a variety of modern denoising operators as implicit prior models; and second, it allows for more modular implementation of software systems for the solution of complex inverse problems.

We next presented and proved theoretical conditions for convergence of the plug-and-play algorithm which depend on the gradient of the denoising operator being a doubly stochastic matrix. We also re-designed the non-local means (NLM) denoising algorithm to have a doubly stochastic gradient, thereby ensuring plug-and-play convergence.

In order to demonstrate the value of our method, we applied the plug-and-play algorithm to two important problems: bright field electron tomography and sparse image interpolation. The results indicate that the plug-and-play algorithm when used with the NLM and DSG-NLM priors were able to reduce artifacts, improve clarity, and reduce RMSE (for the simulated dataset) as compared to the filtered back-projection and qGGMRF reconstructions. Then we performed sparse interpolation on simulated and real microscope images with as little as 5% of the pixels sampled – using three denoising operators: NLM, doubly-stochastic gradient NLM (DSG-NLM), and BM3D. We then compared the results against Shepard’s interpolation as the baseline. In all experiments, DSG-NLM resulted in the least RMSE and also complete convergence of the plug-and-play algorithm, as predicted by theory.

Figure 10: Interpolation of a 256×256 grayscale image of a set of super ellipses.Figure 11: Interpolation of a 414×414 grayscale image of zinc oxide nano-rods.

APPENDIX A

PROOF OF PLUG AND PLAY CONVERGENCE THEOREM

This appendix provides a proof of Theorem III.1. We start by defining a proximal mapping as any function $H : \mathbb{R}^N \rightarrow \mathbb{R}^N$ which can be expressed in the form

$$H(x) = \arg \min_{v \in \mathbb{R}^N} \left\{ \frac{\|x - v\|^2}{2} + s(v) \right\}, \quad (22)$$

where $s : \mathbb{R}^N \rightarrow \mathbb{R} \cup \{+\infty\}$ is a proper closed convex function on \mathbb{R}^N . With this definition, we can formally state the theorem proved by Moreau in 1965 [42] which gives necessary and sufficient conditions for when H is a proximal mapping.

Theorem A.1. (Moreau 1965 [42]) *A function $H : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is a proximal mapping if and only if*

- (1) H is non-expansive and,
- (2) H is the sub-gradient of a convex function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$.

In fact, if there exists a function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$ such that $\forall x \in \mathbb{R}^N$

$$H(x) = \nabla \phi(x),$$

then we say that $H(x)$ is a conservative function or vector field. The concept of conservative functions is widely used on electromagnetics, for example. The next well known theorem (see for example [43, Theorem 2.6, p. 527]). gives necessary and sufficient conditions for a continuously differentiable function to be conservative on \mathbb{R}^N .

Theorem A.2. *Let $H : \mathbb{R}^N \rightarrow \mathbb{R}^N$ be a continuously differentiable function. Then $H(x)$ is conservative if and only if $\forall x \in \mathbb{R}^N$, $\nabla H(x) = [\nabla H(x)]^t$.*

In general, the sum of two proper closed convex functions, $h = f + g$, is not necessarily proper. This is because the intersection of the two sets $A = \{x \in \mathbb{R}^N : f(x) \leq \infty\}$ and $B = \{x \in \mathbb{R}^N : g(x) \leq \infty\}$ might be empty. Therefore, the following lemma will be needed in order to handle the addition of proper closed convex functions.

Lemma A.3. *Let f and g both be proper closed convex functions and let $h = f + g$ be proper. Then h is proper, closed, and convex.*

Proof: A proper convex function is closed if and only if it is lower semi-continuous. So therefore, both f and g must be lower semi-continuous. This implies that h is also lower semi-continuous. Since h is formed by the sum of two convex function, it must be convex. Putting this together, h is proper, convex, and lower-semi-continuous, and therefore it must be closed. Therefore, h is a proper, closed, and convex function on \mathbb{R}^N .

Using these results, we next provide a proof of Theorem III.1.

Proof: Without loss of generality, we will assume $\beta = 1$ and $\sigma_n^2 = 1$ in order to simplify the notation of the proof.

We start by showing result 1 of the theorem, that H is a proximal mapping for some proper, closed, and convex function $s(x)$. To do this, we use Moreau's result stated above in Theorem A.1. In order to meet the conditions of Moreau's theorem, we first show that H is the sub-gradient of a convex

function $\phi : \mathbb{R}^N \rightarrow \mathbb{R}$. Since $\nabla H(x)$ is assumed to be a doubly stochastic matrix, we know that $\nabla H(x) = [\nabla H(x)]^t$. Then by Theorem A.2 above, we know that $H(x)$ is conservative and there must exist a function ϕ so that

$$H(x) = \nabla \phi(x).$$

Furthermore, since $\nabla H(x)$ is a doubly stochastic matrix, it must have real eigenvalues in the range $(0, 1]$. Since the eigenvalues are positive, ϕ must be convex. Furthermore, since the eigenvalues are ≤ 1 , H must also be non-expansive. So therefore, we know that H is a proximal mapping of some proper, closed, and convex function $s(x)$. More specifically, we know that there exists a proper, closed, and convex function, $s(x)$, on \mathbb{R}^N such that H can be expressed as

$$H(x) = \arg \min_{v \in \mathbb{R}^N} \left\{ \frac{\|x - v\|^2}{2} + s(v) \right\}. \quad (23)$$

We next show result 2 of the theorem, that a MAP estimate exists. This is equivalent to saying that the function $h(x) = l(x) + s(x)$ takes on its global minimum value for some $x = \hat{x}_{MAP}$.

First define the function $h(x) = l(x) + s(x)$. By condition 3 of Theorem III.1 there exists an x and y such that $y = H(x)$ and $l(y) < \infty$. Since, $y = H(x)$ we also know that $s(y) < \infty$. Therefore, $h(y) < \infty$ and h is proper. By Lemma A.3, h must also be proper, closed, and convex.

Now to show that $h(x)$ takes on its global minimum, we need only show that there exists an threshold $\alpha \in \mathbb{R}$ such that the sublevel set of h is a non-empty compact set, that is

$$A_\alpha = \{x \in \mathbb{R}^N : h(x) \leq \alpha\}$$

is a non-empty compact subset of \mathbb{R}^N . Since h is a closed function, A_α must be a closed set. Therefore, it is only necessary to show that A_α is nonempty and bounded.

Define

$$p^* = \inf_{x \in \mathbb{R}^N} h(x).$$

Then since $h(x)$ is proper, closed, and convex, we know that $\infty > p^* > -\infty$. Select any $\alpha > p^*$. So clearly, A_α is nonempty.

Next we show that A_α is bounded. Since $s(x)$ is a proper closed convex function, we know that it must have an affine lower bound, i.e., there exist a finite row vector b and constant c so that for all $x \in \mathbb{R}^N$

$$s(x) \geq bx + c.$$

By condition 4 of Theorem III.1, it is always possible to choose $r > 1$ so that

$$\frac{f(r)}{r} > \|b\| + \|c\| + \alpha.$$

In this case, it is easy to show that for all $\|x\| > r$, we have that

$$\begin{aligned} h(x) &= l(x) + s(x) \\ &\geq f(r) - \{\|b\|r + |c|\} \\ &\geq r\{\|b\| + \|c\| + \alpha\} - \{\|b\|r + |c|\} \\ &\geq \alpha. \end{aligned}$$

So therefore, we know that $\forall x \in A_\alpha$, $\|x\| < r$, and that A_α is a nonempty bounded and therefore compact set. Consequently, h must take on its global minimum value for some value \hat{x}_{MAP} in the compact set A_α .

Finally, we show result 3 of the theorem, that the plug-and-play algorithm converges. Since the plug-and-play algorithm is just an application of the ADMM algorithm, we can use standard ADMM convergence theorems. We use the standard theorem as stated in [25, p. 16]. This depends on two assumptions. The first assumption is that $l(x)$ and $s(x)$ must be a proper, closed, and convex functions, which we have already shown. The second assumption is that the standard (un-augmented) Lagrangian must have a saddle point. The standard Lagrangian for this problem is given by,

$$L(x, v; \lambda) = l(x) + s(v) + \lambda^t(x - v), \quad (24)$$

and the associated dual function is denoted by

$$g(\lambda) = \inf_{x, v \in \mathbb{R}^N} L(x, v; \lambda).$$

We say that $x^* \in \mathbb{R}^N$, $v^* \in \mathbb{R}^N$, $\lambda^* \in \mathbb{R}^K$ are a saddle point if

$$L(x, v; \lambda^*) \geq L(x^*, v^*; \lambda^*) \geq L(x^*, v^*; \lambda).$$

Now we have already proved that a solution to our optimization problem exists and is given by $x^* = v^* = \hat{x}_{MAP}$. So we know that the primal problem has a solution given by

$$\begin{aligned} p^* &= \inf_{\substack{x, v \in \mathbb{R}^N \\ v=x}} \{l(x) + s(v)\} \\ &= l(x^*) + s(v^*). \end{aligned}$$

Now the pair (x^*, v^*) is a strictly feasible solution to the constrained optimization problem because x^* and v^* meet the constraint and they both fall within the open set \mathbb{R}^N . This means Slater's conditions hold, and by Slater's theorem, strong duality must also hold for some λ^* [53], [54]. More specifically, we know that there must exist a $\lambda^* \in \mathbb{R}^K$ such that

$$p^* = g(\lambda^*).$$

Using this result, we have that

$$\begin{aligned} L(x^*, v^*; \lambda^*) &= l(x^*) + s(v^*) + [\lambda^*]^t(x^* - v^*) \\ &= l(x^*) + s(v^*) \\ &= p^* = g(\lambda^*) \\ &\leq L(x, v; \lambda^*). \end{aligned}$$

So we have that $L(x, v; \lambda^*) \geq L(x^*, v^*; \lambda^*)$. Furthermore since $x^* = v^*$, we know that $L(x^*, v^*; \lambda^*) = L(x^*, v^*; \lambda)$ for all λ . So putting together these two results, we have that $L(x, v; \lambda^*) \geq L(x^*, v^*; \lambda^*) \geq L(x^*, v^*; \lambda)$, thus proving the existence of a saddle point of the un-augmented Lagrangian, $L(x, v; \lambda)$.

Adapting the theorem of [25, p. 16], we then have the stated convergence results of equation (10).

$$\lim_{k \rightarrow \infty} \{x^{(k)} - v^{(k)}\} = 0; \quad (25)$$

$$\lim_{k \rightarrow \infty} \{l(x^{(k)}) + s(v^{(k)})\} = p^*. \quad (26)$$

ACKNOWLEDGMENT

The authors thank Gregory Buzzard, professor and head of the Mathematics department at Purdue University, for many useful discussions regarding the conditions of convergence of the plug-and-play algorithm. They would also like to thank Marc DeGraef, professor of material science at Carnegie Mellon University, for providing simulated aluminum spheres tomography datasets.

REFERENCES

- [1] K. Grünewald, P. Desai, D. C. Winkler, J. B. Heymann, D. M. Belnap, W. Baumeister, and A. C. Steven, "Three-dimensional structure of herpes simplex virus from cryo-electron tomography," *Science*, vol. 302, no. 5649, pp. 1396–1398, 2003.
- [2] M. Bárcena and A. J. Koster, "Electron tomography in life science," *Seminars in Cell & Developmental Biology*, vol. 20, no. 8, pp. 920 – 930, 2009.
- [3] P. A. Midgley and R. E. Dunin-Borkowski, "Electron tomography and holography in materials science." *Nature Materials*, vol. 8, no. 4, pp. 271–280, 2009.
- [4] H. S. Anderson, J. Ilic-Helms, B. Rohrer, J. Wheeler, and K. Larson, "Sparse imaging for fast electron microscopy," pp. 86 570C–86 570C–12, 2013. [Online]. Available: <http://dx.doi.org/10.1117/12.2008313>
- [5] A. Stevens, H. Yang, L. Carin, I. Arslan, and N. D. Browning, "The potential for Bayesian compressive sensing to significantly reduce electron dose in high-resolution STEM images," vol. 63, no. 1, pp. 41–51, 2014.
- [6] A. Mohammad-Djafari, "Joint estimation of parameters and hyperparameters in a bayesian approach of solving inverse problems," in *Image Processing, 1996. Proceedings., International Conference on*, vol. 1, Sep. 1996, pp. 473 –476 vol.2.
- [7] C. Bouman and K. Sauer, "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. on Image Processing*, vol. 2, no. 3, pp. 296 –310, Jul. 1993.
- [8] C. A. Bouman, *Model Based Image And Signal Processing*, 2013.
- [9] S. Z. Sullivan, R. D. Muir, J. A. Newman, M. S. Carlsen, S. Sreehari, C. Doerge, N. J. Begue, R. M. Everly, C. A. Bouman, and G. J. Simpson, "High frame-rate multichannel beam-scanning microscopy based on lissajous trajectories," *Opt. Express*, vol. 22, no. 20, pp. 24 224–24 234, Oct 2014. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-22-20-24224>
- [10] S. Oh, A. Milstein, C. Bouman, and K. Webb, "A general framework for nonlinear multigrid inversion," *IEEE Trans. on Image Processing*, vol. 14, no. 1, pp. 125 –140, Jan. 2005.
- [11] E. Mumcuoglu, R. Leahy, S. Cherry, and Z. Zhou, "Fast gradient-based methods for Bayesian reconstruction of transmission and emission PET images," *IEEE Trans. on Medical Imaging*, vol. 13, no. 4, pp. 687 –701, Dec. 1994.
- [12] D. B. Husarik, D. Marin, E. Samei, S. Richard, B. Chen, T. A. Jaffe, M. R. Bashir, and R. C. Nelson, "Radiation dose reduction in abdominal computed tomography during the late hepatic arterial phase using a model-based iterative reconstruction algorithm: how low can we go?" *Investigative radiology*, vol. 47, no. 8, pp. 468–474, 2012.
- [13] S. Venkatakrisnan, L. Drummy, M. A. Jackson, M. De Graef, J. Simmons, and C. Bouman, "A model based iterative reconstruction algorithm for high angle annular dark field-scanning transmission electron microscope (HAADF-STEM) tomography," *IEEE Transactions on Image Processing*, vol. 22, no. 11, pp. 4532–4544, 2013.
- [14] S. Venkatakrisnan, L. Drummy, M. Jackson, M. D. Graef, J. Simmons, and C. Bouman, "Model-based iterative reconstruction for bright-field electron tomography," *IEEE Transactions on Computational Imaging*, vol. 1, no. 1, pp. 1–15, 2014.
- [15] A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling & Simulation*, vol. 4, no. 2, pp. 490–530, 2005.
- [16] A. Wong and J. Orchard, "A nonlocal-means approach to exemplar-based inpainting," in *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*. IEEE, 2008, pp. 2600–2603.
- [17] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Non-local sparse models for image restoration," in *Computer Vision, 2009 IEEE 12th International Conference on*. IEEE, 2009, pp. 2272–2279.

- [18] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *Image Processing, IEEE Transactions on*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [19] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE*. IEEE, 2013, pp. 945–948.
- [20] Y. Chen, J. Ma, Q. Feng, L. Luo, P. Shi, and W. Chen, "Nonlocal prior bayesian tomographic reconstruction," *Journal of Mathematical Imaging and Vision*, vol. 30, no. 2, pp. 133–146, 2008.
- [21] S. Chun, Y. Dewaraja, and J. Fessler, "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.*, 2013, pp. 62–5.
- [22] R. Glowinski and A. Marroco, "Sur l'approximation, par elements finis d'ordre un, et la resolution, par penalisation-dualite d'une classe de problemes de dirichlet non lineaires," *ESAIM: Mathematical Modelling and Numerical Analysis-Modélisation Mathématique et Analyse Numérique*, vol. 9, no. R2, pp. 41–76, 1975.
- [23] D. Gabay and B. Mercier, "A dual algorithm for the solution of nonlinear variational problems via finite element approximation," *Computers & Mathematics with Applications*, vol. 2, no. 1, pp. 17–40, 1976.
- [24] J. Eckstein and D. P. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, no. 1-3, pp. 293–318, 1992.
- [25] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [26] R. C. Fair, "On the robust estimation of econometric models," in *Annals of Economic and Social Measurement, Volume 3, number 4*. NBER, 1974, pp. 117–128.
- [27] Z. Yang and M. Jacob, "Nonlocal regularization of inverse problems: a unified variational framework," *Image Processing, IEEE Transactions on*, vol. 22, no. 8, pp. 3192–3203, 2013.
- [28] G. Gilboa and S. Osher, "Nonlocal operators with applications to image processing," *Multiscale Modeling & Simulation*, vol. 7, no. 3, pp. 1005–1028, 2008.
- [29] G. Peyré, S. Bougleux, and L. Cohen, "Non-local regularization of inverse problems," in *Computer Vision—ECCV 2008*. Springer, 2008, pp. 57–68.
- [30] C. Guillemot and O. Le Meur, "Image inpainting : Overview and recent advances," *Signal Processing Magazine, IEEE*, vol. 31, no. 1, pp. 127–144, Jan 2014.
- [31] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, ser. SIGGRAPH '00. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 417–424. [Online]. Available: <http://dx.doi.org/10.1145/344779.344972>
- [32] M. Elad, J.-L. Starck, P. Querre, and D. Donoho, "Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA)," *Applied and Computational Harmonic Analysis*, vol. 19, no. 3, pp. 340 – 358, 2005, computational Harmonic Analysis - Part 1. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1063520305000655>
- [33] W. Dong, L. Zhang, R. Lukac, and G. Shi, "Sparse representation based image interpolation with nonlocal autoregressive modeling," *IEEE Transactions on Image Processing*, vol. 22, no. 4, pp. 1382–1394, 2013.
- [34] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, Mar. 2010. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1756006.1756008>
- [35] J. Yang, X. Liao, X. Yuan, P. Lull, D. Brady, G. Sapiro, and L. Carin, "Compressive sensing by learning a gaussian mixture model from measurements," *Image Processing, IEEE Transactions on*, vol. 24, no. 1, pp. 106–119, Jan 2015.
- [36] F. Li and T. Zeng, "A universal variational framework for sparsity-based image inpainting," *Image Processing, IEEE Transactions on*, vol. 23, no. 10, pp. 4242–4254, 2014.
- [37] A. Danielyan, V. Katkovnik, and K. Egiazarian, "BM3D frames and variational image deblurring," *Image Processing, IEEE Transactions on*, vol. 21, no. 4, pp. 1715–1728, 2012.
- [38] P. Milanfar, "A tour of modern image filtering: New insights and methods, both practical and theoretical," *Signal Processing Magazine, IEEE*, vol. 30, no. 1, pp. 106–128, 2013.
- [39] P. Chatterjee and P. Milanfar, "Patch-based near-optimal image denoising," *Image Processing, IEEE Transactions on*, vol. 21, no. 4, pp. 1635–1649, 2012.
- [40] H. Talebi and P. Milanfar, "Global image denoising," *Image Processing, IEEE Transactions on*, vol. 23, no. 2, pp. 755–768, Feb 2014.
- [41] S. Sreehari, S. Venkatakrishnan, L. F. Drummy, J. P. Simmons, and C. A. Bouman, "Advanced prior modeling for 3D bright field electron tomography," in *IS&T/SPIE Electronic Imaging*. International Society for Optics and Photonics, 2015, pp. 940 108–940 108.
- [42] J.-J. Moreau, "Proximité et dualité dans un espace hilbertien," *Bulletin de la Société mathématique de France*, vol. 93, pp. 273–299, 1965.
- [43] R. E. Williamson, R. H. Crowell, and H. F. Trotter, *Calculus of vector functions*. Prentice-Hall, 1972.
- [44] P. Milanfar, "Symmetrizing smoothing filters," *SIAM Journal on Imaging Sciences*, vol. 6, no. 1, pp. 263–284, 2013.
- [45] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 2. IEEE, 2005, pp. 60–65.
- [46] B. Wahlberg, S. Boyd, M. Annergren, and Y. Wang, "An admm algorithm for a class of total variation regularized estimation problems," *arXiv preprint arXiv:1203.1828*, 2012.
- [47] E. Ghadimi, A. Teixeira, I. Shames, and M. Johansson, "Optimal parameter selection for the alternating direction method of multipliers (admm): quadratic problems," *Automatic Control, IEEE Transactions on*, vol. 60, no. 3, pp. 644–658, 2015.
- [48] M. De Graef, *Introduction to conventional transmission electron microscopy*. Cambridge University Press, 2003.
- [49] R. Stevenson and E. Delp, "Fitting curves with discontinuities," in *Proc. of the first international workshop on robust computer vision*, 1990, pp. 127–136.
- [50] J.-B. Thibault, K. D. Sauer, C. A. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multislice helical ct," *Medical physics*, vol. 34, no. 11, pp. 4526–4544, 2007.
- [51] H. Zhao and M. L. Comer, "A unified Markov random field/Markov point process image model and its applications to computational materials," in *MURI report*, 2014.
- [52] H. Koerner, J. Kelley, J. George, L. Drummy, P. Mirau, N. S. Bell, J. W. Hsu, and R. A. Vaia, "ZnO nanorod-thermoplastic polyurethane nanocomposites: Morphology and shape memory performance," *Macromolecules*, vol. 42, no. 22, pp. 8933–8942, 2009.
- [53] M. Slater, "Lagrange multipliers revisited," Cowles Foundation for Research in Economics, Yale University, Tech. Rep. 80, November 1950.
- [54] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge UK: Cambridge University Press, 2009.