ROTATIONALLY-INVARIANT NON-LOCAL MEANS FOR IMAGE DENOISING AND TOMOGRAPHY

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

[a]Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA
[b]Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
[c]Air Force Research Laboratory, Dayton, OH 45433, USA

ABSTRACT

Many samples imaged in structural biology and material science contain many similar particles at random locations and orientations. Model-based iterative reconstruction (MBIR) methods can in principle be used to exploit such redundancies in images through log prior probabilities that accurately account for non-local similarity between the particles. However, determining such a log prior term can be challenging. Several denoising algorithms like non-local means (NLM) successfully capture such non-local redundancies, but the problem is two-fold: NLM is not explicitly formulated as a cost function, and neither can it capture similarity between randomly oriented particles.

In this paper, we propose a rotationally-invariant non-local means (RINLM) algorithm, and describe a method to implement RINLM as a prior model using a novel framework that we call plug-and-play priors. We introduce the idea of patch pre-rotation to make RINLM computationally tractable. Finally, we showcase image denoising and 2D tomography results, using the proposed RINLM algorithm, as we highlight high reconstruction quality, image sharpness, and artifact suppression.

***Index Terms*—** Rotationally-invariant NLM, denoising, tomography, plug-and-play, prior modeling

1. INTRODUCTION

Many important imaging problems in structural biology [1,2], and material science [3] require the reconstruction of many similar particles at random orientations. This presents a huge opportunity in image prior modeling due to the enormous redundancy in the reconstruction of many such similar particles. In fact, cryo electron microscopy (EM) tomography exploits this redundancy but requires very specialized preparation of a thin frozen biological sample that ensures that particles do not overlap [1]. In the more general problem of 3D EM tomography, no solution currently exists which fully exploits the redundancy of similar particles.

Model-based iterative reconstruction (MBIR) [4,5] can in principle be used to capture such redundancies through the choice of a log prior term in the MAP cost function. In practice, formulating such a log prior term that captures similarity between non-local particles can be quite difficult. In previous research, Chen et al. proposed using pixel-wise differences between local patches as a log prior term [6]; Chun et al. used the convex Fair potential [7] based on a convex function of L_2 differences between image patches [8]; and Yang et al. proposed a dynamically updated cost function with properties similar to a non-local means filter [9], resulting in a model that captures the intrinsically non-convex behavior required for modeling distant particles with similar structure. In any case, none of these capture the redundancies of rotationally-invariant particles.

In fact, there has been substantial progress in developing denoising algorithms that capture non-local redundancies. For example, non-local means (NLM) effectively denoises images with non-local redundancies by matching image patches at distant locations [10, 11]. However, it does not account for the rotation of particles in an image. There have been some methods proposed to allow for rotational invariance in the non-local matching process [12–16]. However, these methods are either computationally approximate because they use Fourier descriptors instead of actual patch rotation, or they are computationally expensive because they rotate patches dynamically during the search process. In any case, a fundamental limitation of NLM is that it is a denoising algorithm rather than a prior model. So it is unclear how the methods of NLM can be used to improve tomographic reconstruction quality for images containing similar particles.

In this paper, we present a method for tomographic reconstruction of images containing many similar particles at random positions and orientations. The key to our method is a rotationally invariant non-local means (RINLM) filtering algorithm, together with a novel framework for incorporating RINLM into MAP reconstruction. Our proposed RINLM algorithm allows for the matching of particles or structures at random orientations, and it is computationally efficient because it depends on the pre-rotation of patches to a standard

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.

orientation. In order to use the RINLM denoising algorithm in the MBIR reconstruction framework, we use a method introduced in [17, 18] which we call *plug-and-play priors*. Our experimental results on both simulated and real transmission electron microscopy (TEM) data demonstrate that the RINLM-MBIR reconstruction can substantially improve reconstruction quality as compared to traditional MBIR reconstructions [19]. Moreover, we find that the algorithm has rapid and stable convergence to its final result.

2. PLUG-AND-PLAY FRAMEWORK

We outline a general framework for adapting any denoising operator as a prior model for our tomographic reconstruction problem. This approach, which we call Plug-and-Play priors, lends the ability of having two separate software modules – one to perform the inversion and the other to denoise; however, both these modules are still incorporated into the MBIR framework allowing for qualitatively optimized experimental reconstructions, and quantitatively accurate reconstructions of the simulated data. The plug-and-play technique is an application of alternating direction method of multipliers (ADMM) [20, 21], in which we split the state variable thus decoupling the forward and prior model terms of the MAP estimation. The key advantage of this framework is that the denoising operator does not need to correspond to an optimization problem in order to function as a prior model.

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

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

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 is a unit less parameter that controls the regularization applied. Increasing the value of β increases the regularization applied, while decreasing β corresponds to lowering the regularization. The variable x (see Eq. (1)) is split into two variables x and v [17], and we reformulate Eq. (1) as,

$$(\hat{x}, \hat{v}) = \arg \min_{\substack{x \geq 0 \\ v = x}} \{ l(y; x) + \beta s(v) \}. \tag{2}$$

The augmented Lagrangian corresponding to Eq. (2) is given by,

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

where u is a scaled dual variable (initialized to zero) and σ_λ is the augmented Lagrangian parameter. In theory, the reconstruction does not depend on σ_λ, but in practice, well-chosen values of σ_λ tend to speed up the ADMM convergence. We discuss how we set the value of σ_λ in Sec. 4.

Each iteration of ADMM algorithm consists of the following steps.

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

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

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

where $\tilde{x} = \hat{v} - u$, $\tilde{v} = \hat{x} + u$, and $\sigma_n = \sqrt{\beta}\sigma_\lambda$ is the assumed noise standard deviation for the denoising operation. Also, $\hat{x}$ and $\hat{v}$ are initialized to the best available guess of the estimate. The inversion operator, $\mathbb{F}$ used in Eq. (4), is given by,

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

The denoising operator, $\mathbb{H}$ used in Eq. (5), is given by,

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

We use the proposed RINLM that we develop in Sec. 3 as the denoising operator $\mathbb{H}$. Although NLM-based methods are not explicitly formulated as solutions to optimization problems, we do observe consistent empirical convergence of $\hat{x}$ and $\hat{v}$ [18] – implying convergence of the ADMM iterations.

In Sec. 4, we showcase tomography results supporting this choice of $\mathbb{H}$.

3. FORMULATION OF RINLM

Let S is the set of all lattice points in the image, $N = |S|$ be the number of lattice points in the image, and $s = (s_1, s_2)^t \in S$ be the s-th lattice point. Also, let y be the tomographic measurement, x be the unknown, and $\hat{x}$ be the estimate of x. Further, let $\tilde{v}_s$ be the gray-level of the s-th pixel of the noisy image, and $\hat{v}_s$ be that of the s-th pixel of the denoised image.

Instead of considering all patches in the image for denoising, we only consider patches corresponding to pixels in a square search window around the pixel. A $(2n+1)^2$ search window around pixel s is defined by,

$$\Omega_s = \{ r \in S : \|r - s\|_\infty \leq n \} \subset S. \tag{9}$$

Next we define a circular window, W, to window out circular patches as we pass through the image.

$$W = \{ s \in \mathbb{Z}^2 : \|s\|_2 \leq R \}, \tag{10}$$

where $R \in \mathbb{N}$ is the radius of the window.

3.1. Patch pre-rotation

To make NLM rotationally-invariant, we need to compare each patch with the rotated versions of patches in the search window. Each such comparison is an L_2 norm, thus making the naive implementation intractable.

Our solution is to pre-rotate all the patches just once – to align all the patches in the image along one matching orientation. We assign each pixel some mass that is proportional to

its gray-level intensity. Then we compute the center of mass of each patch of the image. We match the orientations of the patches by rotating each patch, by an angle, θ_s, such that the center of mass of the rotated patch lies on the positive x-axis of the image, as shown in Fig. 1. The center of mass of the

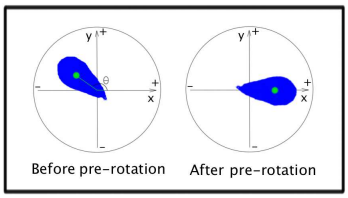

Fig. 1. Axes on the image plane, with the center of masses pointed out before and after rotation. (The green dot is the center of mass of the blue foreground.)

patch P_s around pixel s is given by,

$$m_s = (m_{s,1}, m_{s,2})^t \in \mathbb{R}^2 = \frac{\sum_{r \in W} r\tilde{v}_{s+r}}{\sum_{r \in W} \tilde{v}_{s+r}}. \tag{11}$$

We compute the angle the center of mass makes with the positive x-axis as below.

$$\theta_s = \arctan\{m_{s,1}, m_{s,2}\}. \tag{12}$$

We modify the computation of θ_s to make it robust, especially when m_s gets arbitrarily close to the center of the patch.

$$\tilde{\theta}_s = \frac{\|m_s\|\theta_s}{\|m_s\| + (R/\rho)}, \tag{13}$$

where R is the radius of the patch, and $\rho \gg 1$ is a scaling factor for the guard term. (We use $\rho = 10$ in our implementation.) Further, let $A_s = \begin{pmatrix} cos(\tilde{\theta}_s) & sin(\tilde{\theta}_s) \\ -sin(\tilde{\theta}_s) & cos(\tilde{\theta}_s) \end{pmatrix}$ be the rotation matrix. Then, we compute the gray-levels of the pixels of the rotated patches as below.

$$P_{s,j} = \tilde{v}_{s+j'}, \tag{14}$$

where $P_{s,j}$ is the gray-level of the j-th pixel of patch P_s, and $j' = A_s^{-1}j$. When j' does not lie on the discrete lattice grid, we interpolate j' using bilinear interpolation.

3.2. RINLM algorithm

In this section, we present a brief formulation of rotationally-invariant NLM. We propose to use the pre-rotated patches, $\{P_s\}_{s \in S}$ to compute the similarity between the reference and candidate patches. Pixel r exerts a weight $w_{s,r}$ on pixel s, given by,

$$w_{s,r} = \exp\left\{\frac{-\|P_r - P_s\|_2^2}{\sigma_n^2}\right\}. \tag{15}$$

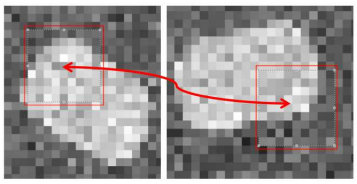

Fig. 2. Illustration of RINLM's ability to match rotated patches: Left and right shapes are rotated versions of each other. Weight assigned by NLM to the patches marked in red = 4.9997e-04; weight assigned by RINLM = 1.63e-02; RINLM treats these patches as being similar, despite the rotation.

After computing the set of all weights, $\{w_{s,r}\}_{s \in S, r \in \Omega_s}$, we normalize the weights as below.

$$\tilde{w}_{s,r} = \frac{w_{s,r}}{\sum_{r \in \Omega_s} w_{s,r}}. \tag{16}$$

The denoised image, $\hat{v}$, is computed as follows.

$$\hat{v}_s = \sum_{r \in \Omega_s} \tilde{w}_{s,r}\tilde{v}_r. \tag{17}$$

The key benefits of RINLM are two-fold: (1) RINLM is significantly faster than naive implementation due to our pre-rotation step, and (2) RINLM offers high quality of reconstruction due to higher weights assigned to similar, but rotated, patches in the noisy image.

4. RESULTS

In this section, we compare the performance of the proposed algorithm (RINLM) with BM3D, qGGMRF [22], K-SVD, and NLM. For this comparison, we specifically look at image denoising and 2D tomographic reconstruction.

4.1. Image denoising

We perform the denoising on randomly oriented snowflakes simulated to induce symmetry and fine structure. The background intensity is 56, the foreground intensity is 200, and the phantom measures 256 pixels $\times$ 256 pixels. The qGGMRF parameters used were $p = 1.05, q = 0.5$, and $c = 0.001$. The NLM patch size we used was 5×5, and the RINLM patch radius was 3. The value of σ_n used for both NLM and RINLM was 22. RINLM-denoised images have clear, sharp foreground, together with background that is almost artifact-free. Table 1 shows that RINLM also consistently results in the lowest RMSE values.

4.2. 2D tomography

We use a real data set that captures the solidification of Al-Cu alloy. The 2D data set has 145 equi-spaced projection angles in $[0, \pi]$. The reconstructed image measures 256 pixels $\times$

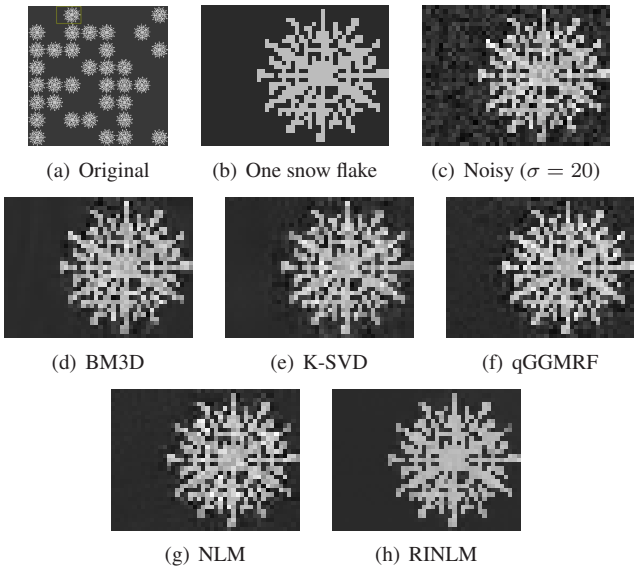

(a) Original (b) One snow flake (c) Noisy ($\sigma = 20$)

(d) BM3D (e) K-SVD (f) qGGMRF

(g) NLM (h) RINLM

Fig. 3. Denoising experiment: (a-c) Simulated data; (d) BM3D produces background artifacts and foreground noise; (e) K-SVD has artifacts beyond visible portion, and high foreground peripheral noise; (f) qGGMRF produces an overall noisy image; (g) NLM foreground noise is comparable to BM3D, and has a grainy background; (h) RINLM renders an almost-speckle-free foreground, while the background is free of artifacts and almost all noise.

.

Table 1. RMSE of various denoising algorithms

Noise σ	BM3D	K-SVD	gGGMRF	NLM	RINLM
5	3.44	3.31	5.39	3.72	**2.64**
10	6.89	6.91	8.77	7.04	**5.95**
20	12.36	12.87	14	12.58	**8.23**
25	14.92	15.28	18.49	14.76	**11.99**

256 pixels, and has been plotted in the range $[0, 19.8 \times 10^{-3}]$ nm^{-1}. The qGGMRF parameters used were $p = 1.1, q = 0.5$, and $c = 0.001$. The NLM patches measured 5×5, while the RINLM patches had a radius of 5. The augmented Lagrangian parameter, σ_λ was chosen to be equal to the amount of variation in the filtered backprojection reconstruction, $\hat{x}^{(0)}$, which we used as the initialization in the plug-and-play algorithm, i.e., $\sigma_\lambda \approx$ std. dev. $(\hat{x}^{(0)})$. The value of σ_λ was chosen to be 8.66×10^{-3} nm^{-1}, while the value of β was set to 4.2, 1.96, 3.74, and 3.25 for K-SVD, BM3D, NLM, and RINLM respectively.

In Fig. (4-f), we show the convergence rate for various prior models. The y-axis denotes 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, $\hat{x}$

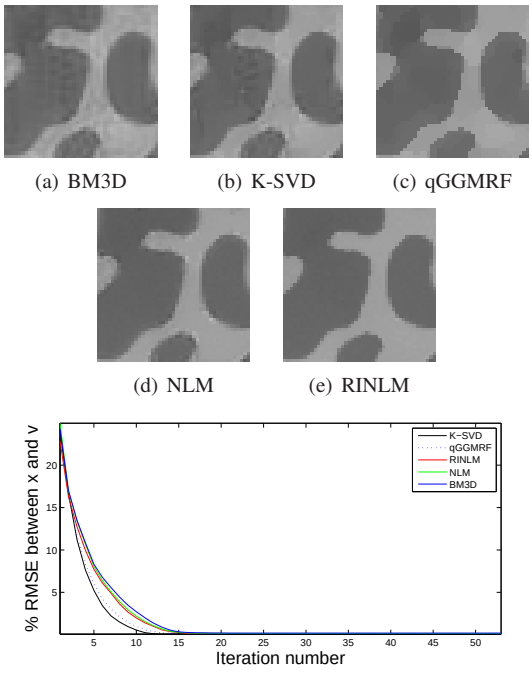

(a) BM3D (b) K-SVD (c) qGGMRF

(d) NLM (e) RINLM

(f) Convergence of various prior models

Fig. 4. 2D tomographic reconstruction: (a) BM3D produces smear-like artifacts in the background and noisy foreground regions; (b) K-SVD results in artifacts in the background, foreground speckles and blotches; (c) qGGMRF produces background blotches and dark foreground merge blots; (d) NLM reconstruction shows foreground peripheral noises with fine-grained background artifacts; (e) RINLM noticeably suppresses contours within the foreground regions, while the background is almost free from blotchiness and any visible contours; (f) Consistent plug-and-play convergence.

The foreground merger regions in RINLM are smooth and lack darker pigmentation. Also, RINLM results in sharper edges, lower noise, and minimizes various types of artifacts.

5. CONCLUSIONS

In this paper, we presented a computationally-efficient rotation-invariant NLM-based denoising algorithm that produces high quality tomographic reconstructions. Denoising results demonstrate visibly lower artifacts, alongside significantly lower RMSE values than competing methods like BM3D, K-SVD, NLM, and qGGMRF. With the plug-and-play framework, we incorporated RINLM as a prior model to achieve high-quality tomographic reconstructions, with least contour artifacts and pigmentation compared to BM3D, K-SVD, NLM, and qGGMRF as prior models.

6. ACKNOWLEDGMENT

We thank Julie Fife and Peter Voorhees of Northwestern University for the Al-Cu coarsening tomography dataset.

7. 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] 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.

[5] 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

[6] 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.

[7] 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.

[8] 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.

[9] 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.

[10] 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.

[11] P. Coupé, P. Yger, S. Prima, P. Hellier, C. Kervrann, and C. Barillot, "An optimized blockwise nonlocal means denoising filter for 3-d magnetic resonance images," *Medical Imaging, IEEE Transactions on*, vol. 27, no. 4, pp. 425–441, 2008.

[12] S. Zimmer, S. Didas, and J. Weickert, "A rotationally invariant block matching strategy improving image denoising with non-local means," in *Proc. 2008 International Workshop on Local and Non-Local Approximation in Image Processing*, 2008, pp. 135–142.

[13] S. K. Alexander, E. R. Vrscay, and S. Tsurumi, "A simple, general model for the affine self-similarity of images," in *Image Analysis and Recognition*. Springer, 2008, pp. 192–203.

[14] Y. Zhuo, J. Liu, J. Ren, and Z. Guo, "Nonlocal based super resolution with rotation invariance and search window relocation," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 853–856.

[15] A. Maximo, R. Patro, A. Varshney, and R. Farias, "A robust and rotationally invariant local surface descriptor with applications to non-local mesh processing," *Graphical Models*, vol. 73, no. 5, pp. 231–242, 2011.

[16] Y. Lou, P. Favaro, S. Soatto, and A. Bertozzi, "Nonlocal similarity image filtering," in *Image Analysis and Processing–ICIAP 2009*. Springer, 2009, pp. 62–71.

[17] 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.

[18] S. Sreehari, S. V. Venkatakrishnan, L. F. Drummy, J. P. Simmons, and C. A. Bouman, "Advanced prior modeling for 3D bright field electron tomography," *To appear in Proc. SPIE 9401, Computational Imaging XIII (February 8, 2015)*, 2015.

[19] S. V. Venkatakrishnan, L. F. Drummy, M. De Graef, J. P. Simmons, and C. A. Bouman, "Model based iterative reconstruction for bright field electron tomography," *To appear in Computational Imaging, IEEE Transactions on*, 2014.

[20] 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.

[21] 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.

[22] 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.