

Multi-Resolution Data Fusion for Super-Resolution Electron Microscopy

Suhas Sreehari^{1,*}, S. V. Venkatakrishnan², Katherine L. Bouman³,
Jeff P. Simmons⁴, Larry F. Drummy⁴, Charles A. Bouman¹

¹Purdue University, ² Oak Ridge National Lab, ³ MIT, ⁴ Air Force Research Lab

*ssreehar@purdue.edu

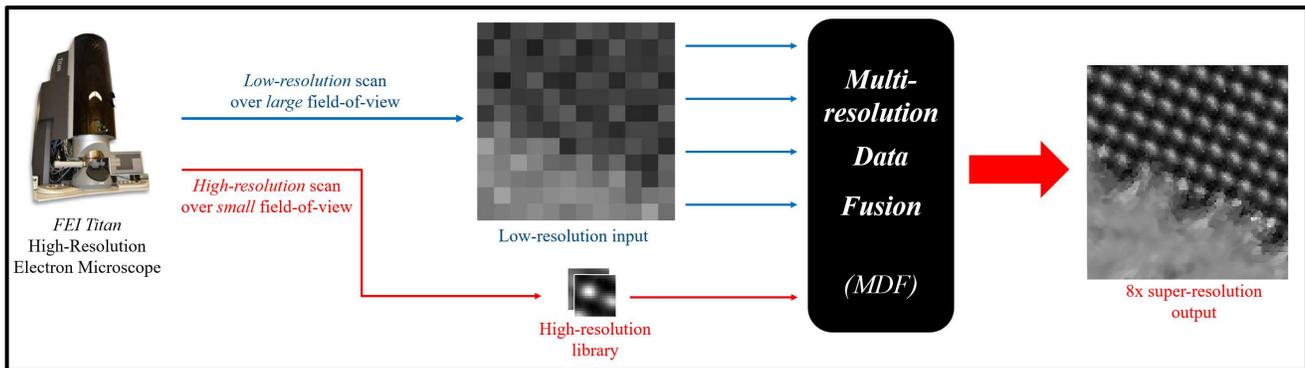


Figure 1: An illustration of the proposed multi-resolution data fusion (MDF) system in which a transmission electron microscope is specially configured to collect both low resolution data over a large field-of-view (FoV) together with a small set of high resolution patches from the same sample. The MDF algorithm uses the high-resolution patches as an image model within the “plug-and-play” framework to synthesize the 8x interpolated super-resolution output image over the full FoV. We demonstrate the MDF system using an aberration-corrected FEI Titan transmission electron microscope and show that it can dramatically speed up image acquisition and reduce dosage as compared to conventional homogeneous raster-scanning microscopy.

Abstract

Perhaps surprisingly, all electron microscopy (EM) data collected to date is less than a cubic millimeter – presenting a huge demand in the materials and biological sciences to image at greater speed and lower dosage, while maintaining resolution. Traditional EM imaging based on homogeneous raster scanning severely limits the volume of high-resolution data that can be collected, and presents a fundamental limitation to understanding physical processes such as material deformation and crack propagation.

We introduce a multi-resolution data fusion (MDF) method for super-resolution computational EM. Our method combines innovative data acquisition with novel algorithmic techniques to dramatically improve the resolution/volume/speed trade-off. The key to our approach is to collect the entire sample at low resolution, while simultaneously collecting a small fraction of data at high resolution. The high-resolution measurements are then used to create a material-specific model that is used within the “plug-and-play” framework to dramatically improve resolution of the low-resolution data. We present results using FEI electron microscope data that demonstrate super-resolution factors of 4x-16x, while substantially maintaining high image quality and reducing dosage.

1. Introduction

Scanning transmission electron microscopes are widely used for characterization of samples at the nano-meter scale. However, raster scanning an electron beam across a large field of view is time consuming and can often damage the sample. In fact, the sum total of all electron microscopy (EM) data collected to date, represents less than a cubic millimeter of material [24]. This fact encapsulates a central challenge of materials and biological sciences to image larger volumes at higher resolution, greater speed, and lower dosage.

Traditional microscopy based on homogeneous raster-order single-resolution scanning severely limits the volume of high-resolution data that can be collected, and represents a fundamental limitation in understanding physical processes such as material deformation, crack propagation, and pyrolysis. For these reasons, there is growing interest in reconstructing/simulating full-resolution images from sparsely-sampled [11, 12] and low-resolution images [8, 18, 29].

In many cases, EM samples contain repeating structures that are similar or identical to each other. This

presents enormous redundancy in conventional data acquisition methods, and consequently an opportunity to design imaging systems with sparse/low-resolution samples.

In our paper, we take a major step in proposing a joint design of the sensor acquisition and algorithmic reconstruction for super-resolution electron microscopy, specifically exploiting the structural regularity in EM samples. Our method allows for dramatic data acquisition speed-ups while still maintaining high resolution and image quality. Specifically, we introduce a novel computational EM imaging method based on multi-resolution data fusion (MDF) that combines innovative data acquisition with algorithmic techniques to dramatically improve the resolution/volume/speed trade-off as compared to traditional EM imaging methods. Measuring fewer samples also leads to lower dosage and less sample damage, which is crucial for imaging modalities like cryo-EM. The key to our approach is the acquisition of sparse measurements at both low and high resolutions from which a full high-resolution image is formed. The speedup in data acquisition results from the fact that we sparsely sample the bulk of the material while densely sampling only a small-but-representative fraction of the material – instead of making dense high-resolution measurements across the whole material. The high-resolution measurements are then used to create a material-specific probabilistic model in the form of a patch-library that is used to dramatically increase resolution.

Our method depends on a Bayesian framework for the fusion of low- and high-resolution data without the need for training. The key step in our solution is to use our library-based non-local means (LB-NLM) denoising filter as a prior model within the P&P framework [19, 23] – to produce visually true textures and edge features in image interpolations, while significantly reducing mean squared error and artifacts such as “jaggies”. The P&P framework is based on the alternating direction method of multipliers (ADMM) [3, 6, 10] and decouples the forward model and the prior terms in the maximum a posteriori cost function. 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 of choice.

We present super-resolution and sparse interpolation results on real microscope images demonstrating data acquisition speedups of 15x to 228x on various datasets, while substantially maintaining image quality.

2. Related Work

Image interpolation and super-resolution have been widely studied to enable high-quality imaging with fewer measurements, leading to faster and cheaper data acquisition.

Though local models have been adapted to solve inverse problems [2, 14, 15], there have been several efforts to solve the sparse image interpolation problem using patch-based models [7, 20, 25]. Dictionary learning [26, 27] and example-based methods [9] have also been proposed for achieving super-resolution from low-resolution measurements. Zhang et al. [28] proposed a steering kernel regression framework to use non-local means to achieve super-resolution. Atkins et al. [1] proposed tree-based resolution synthesis using a regression tree as a piece-wise linear approximation to the conditional mean of the high-resolution image given the low-resolution image. More recently, Timofte et al. [22] discussed the use of libraries of structures and image self-similarity to improve super-resolution, but these methods involved extensive training. Another training-based approach to super-resolution was proposed by Perez-Pellitero et al. [17] where they derive a regression-based manifold mapping between low- and high-resolution images. Data fusion techniques have been studied in panchromatic image sharpening [21].

Apart from these specific solutions for image interpolation and super-resolution, Sreehari et al. [19] proposed a generic Bayesian framework called “plug-and-play” priors (P&P) for incorporating modern denoising algorithms as prior models in a variety of inverse problems such as sparse image interpolation. In this spirit, Brifman et al. [4] have adopted the P&P framework to use sparse-coding and dictionary-learning-based denoisers for achieving super-resolution. In any case, no training-free method that we know of currently exists to fuse sparse/low-resolution data with dense/high-resolution measurements from the same EM sample to form full high-resolution images.

3. Method

We present a new computational imaging method for fast acquisition of high-resolution EM images. Our proposed acquisition method scans the image at low resolution while simultaneously collecting a small-but-representative portion of the sample at a higher resolution. This is central to achieving the data acquisition speed-up. Then, we fuse the low- and high-resolution images to synthesize high resolution over the full field-of-view (FoV). In this section we describe the algorithm used to fuse the low- and high-resolution images. We construct a high-quality patch library containing typical textures and edge features from the high-resolution portion. In order to use such a library as an image model, in Sec. 3.2 we construct a version of the non-local means denoiser that uses patches from the high-resolution library. We then use our library-based non-local means (LB-NLM) denoiser as an image/prior model within the P&P framework (see Sec. 3.1). Using the P&P framework, we super-resolve the low-resolution image acquired over large areas of the sample through Bayesian inversion,

creating a high-resolution image over the full FoV.

3.1. The Plug-and-Play (P&P) Framework

Let $x \in \mathbb{R}^N$ be an unknown image with a prior distribution $p(x)$, and let $y \in \mathbb{R}^M$ be the sparsely subsampled image with conditional distribution $p(y|x)$, which is often referred to as the forward model for the sampling system. Then the MAP estimate of the image x conditioned on the sparse/low-resolution image y is given by

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

where $l(x) = -\log p(y|x)$, $\beta s(x) = -\log p(x)$, and $\beta > 0$ is a parameter used to control the regularization in the MAP reconstruction. In order to use flexible prior models to solve ill-posed inverse problems, it is helpful to separate the forward model inversion from the prior model. This can be effectively achieved through the application of ADMM. The first step in applying ADMM for solving equation (1) is to split the variable x , resulting in an equivalent expression for the MAP estimate given by

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

This constrained optimization problem has an associated scaled-form augmented Lagrangian [3, Section 3.1.1],

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

where σ_λ is the augmented Lagrangian parameter.

The ADMM algorithm consists of iteratively minimizing the augmented Lagrangian with respect to the split variables, x and v , followed by updating the dual variable, u , that drives the solution toward the constraint, $x = v$.

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

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

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

where \hat{v} can be initialized to a baseline reconstruction and u is initialized to zero. By letting $\tilde{x} = v - u$, and $\tilde{v} = x + u$, we can define two operators that help solve Eqs. (4) and (5). 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, 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 Fig. 2.

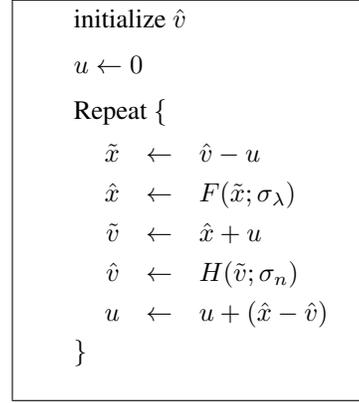


Figure 2: Plug-and-play algorithm for implementation of a general forward model $F(\tilde{x}; \sigma_\lambda)$, and a prior model specified by the denoising operator in $H(\tilde{v}; \sigma_n)$.

In theory, the value of σ_λ does not affect the solution computed by the P&P algorithm; however, in practice a poor choice may slow convergence. Sreehari et al. [19] suggest setting the value of σ_λ based on the amount of variation in a baseline reconstruction¹,

$$\sigma_\lambda^2 \approx \frac{1}{N} \sum_{i=0}^{N-1} \operatorname{var}[x_i|y]. \quad (9)$$

3.2. Library-Based Non-Local Means

Access to high-quality patches can increase the effectiveness of non-local means (NLM) as a prior model. However, in conventional implementations of NLM, we are limited by the degraded patches present in the image. To remove this limitation, in this section, we outline a library-based non-local means (LB-NLM) algorithm.

Our solution is to use an external library of densely-sampled/high-resolution patches that contain typical textures and edge features. We then compute the weighted mean of the center pixels of these patches. Computation of the weighted mean is identical to standard non-local means, only the patches do not come from the image we are denoising. The weight normalization is also identical to standard NLM weight normalization, and is therefore simpler than the symmetric normalization (see [19]).

If we denote the $N_p \times N_p$ patch centered at position s by P_s , and the r -th patch of the library by L_r , we can compute

¹In our experiments, the baseline reconstructions are Shepard or cubic interpolations.

the LB-NLM weights as

$$w_{s,r} \leftarrow \exp \left\{ \frac{-\|P_s - L_r\|_2^2}{2N_p^2\sigma_n^2} \right\} \quad (10)$$

$$w_{s,r} \leftarrow \frac{w_{s,r}}{\sum_{r=1}^{N_l} w_{s,r}}, \quad (11)$$

where N_l is the number of patches in the library. Further, the LB-NLM filtering is given by,

$$\hat{v}_s = \sum_{r=1}^{N_l} w_{s,r} z_r, \quad (12)$$

where \hat{v}_s is the denoised pixel at position s , and z_r is the center pixel of library patch L_r .

We use the LB-NLM denoiser as an implicit prior model within the P&P algorithm. When the patch library contains high-resolution, typical textures and edge features, using LB-NLM as a prior model enables reproduction of visually-accurate textures and edges in the reconstructed image.

Denoising sparse/low-resolution images with the aid of dense/high-resolution patches is, in fact, how we implement the multi-resolution data fusion in every step of the P&P algorithm.

3.3. Super-Resolution Forward Model

In this section, we formulate the explicit form of the inversion operator, $F(x, \sigma_\lambda)$, for the application of super-resolution. Our objective will be to interpolate an image $y \in \mathbb{R}^M$ to its higher resolution version, $x \in \mathbb{R}^N$, where $M \ll N$. The forward model for this problem is given by

$$y = Ax + \epsilon, \quad (13)$$

where $A \in \mathbb{R}^M \times \mathbb{R}^N$ represents the point-spread function (PSF) of the electron microscope. For super-resolution by a factor of L , the PSF could be approximated by averaging the values of L^2 neighborhood pixels of every L^2 -th pixel in x . Furthermore, ϵ is an M -dimensional vector of i.i.d. Gaussian random variables with mean zero and variance σ_w^2 .

We can write the negative log likelihood function as

$$l(x) = \frac{1}{2\sigma_w^2} \|y - Ax\|_2^2 + \frac{M}{2} \log(2\pi\sigma_w^2). \quad (14)$$

In order to enforce positivity, we also modify the negative likelihood function by setting $l(x) = +\infty$ for $x < 0$. Using Eq. (7), the interpolation inversion operator is then given by

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

If we let $\sigma_w^2 = 0$, then the inversion operator for super-resolution by a factor of L reduces to the following simple expression

$$F(\tilde{x}; \sigma_\lambda) = \left[\tilde{x} + A^t \left(y - \frac{1}{L^2} A \tilde{x} \right) \right]_+, \quad (15)$$

where $[\cdot]_+$ enforces positivity.

3.4. 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. Unlike in the super-resolution problem, in sparse interpolation the acquired image resolution is the same as the desired resolution, except that many random pixels are missing from the measurement. More formally, the forward model for this problem is given by

$$y = Ax + \epsilon, \quad (16)$$

where the sparse sampling matrix, $A \in \mathbb{R}^M \times \mathbb{R}^N$. 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. Furthermore, ϵ is an M -dimensional vector of i.i.d. Gaussian random variables with mean zero and variance σ_w^2 . The log likelihood has the exact form as the sparse interpolation case.

Like in the super-resolution case, if we let $\sigma_w^2 = 0$, then the inversion operator, F , reduces to the following explicit expression

$$F(\tilde{x}; \sigma_\lambda) = \left[\tilde{x} + A^t (y - A \tilde{x}) \right]_+. \quad (17)$$

where $[\cdot]_+$ enforces positivity, forcing the interpolation to take on the measured values at the sample points.

4. Results

In this section, we demonstrate the value of our MDF algorithm on three different datasets, all acquired on real electron microscopes. We compare our super-resolution reconstructions against state-of-the-art super-resolution algorithms like the single-image super-resolution (SISR) algorithm [27], and show that fusing multi-resolution data is a powerful framework to form high-quality high-resolution images – at a fraction of the cost, time, and beam dosage compared to conventional methods.

4.1. Datasets

The strength of our experiments lies in the fact that all our test datasets were acquired on real electron microscopes. As opposed to a post-processing technique, our MDF algorithm is a computational electron microscopy method that combines innovative data acquisition with novel reconstruction algorithms.

High Resolution Transmission Electron Microscopy (HR-TEM) of gold nanorods was performed on an aberration-corrected FEI Titan operating at 300 kV [16]. Images were recorded on a 2k by 2k Gatan Ultrascan CCD

at electron optical magnifications ranging from 380 kX to 640kX. The gold atoms image (Fig. 3) is a magnified region of edges of these gold nanorods.

Scanning Electron Microscopy (SEM) of a surface crack in the shell of the mollusk, *Hinea brasiliana* [5] (Fig. 4) was done on a FEI XL30 at 5 kV accelerating voltage.

For the sparse interpolation experiment, we use a TEM image of iridovirus assemblies [13] (Fig. 5), which were cast onto amorphous carbon support films and imaged using a Phillips CM200 Transmission Electron Microscope.

Note that the *Hinea brasiliana* image was acquired natively on the FEI Titan microscope at two different resolutions under similar imaging conditions. The structure in Fig. 4(b) and (c) were acquired at 4x resolution compared to the structure in Fig. 4(d) and (h). The gold atoms image, on the other hand, was acquired at a single resolution (Fig. 3(b)). We artificially reduced the resolution by 4x, 8x, and 16x (see Figs. 3(c), 3(g), and 3(k), respectively).

4.2. Experiments

We apply our MDF algorithm to achieve 4x, 8x, and 16x resolution synthesis results on the atomic-resolution TEM image of gold atoms (Fig. 3), 4x resolution synthesis on the SEM image of *Hinea brasiliana* (Fig. 4), and 5% sparse image interpolation on the TEM image of iridovirus assemblies (Fig. 5). We show that a variety of denoising algorithms such as doubly-stochastic gradient non-local means (DSG-NLM) [19] and library-based non-local means (LB-NLM) can be plugged in as prior models within the P&P framework to enable data acquisition speedups in EM.

In the super-resolution experiments, we compared our MDF algorithm against a state-of-the-art single-image super-resolution (SISR) algorithm from [27]². Since the high-resolution library image is not available before scanning, it is imperative to train the patch library online. The SISR algorithm is designed to run on a pre-trained dictionary, but for quality comparisons, we still used the SISR algorithm by training the image library online, which takes several hours for every dataset, making our training-free MDF algorithm much faster.

To implement the LB-NLM, we first selected non-overlapping 40×40 square regions called “super-patches” uniformly randomly from the high-resolution library images (see Figs. 3(a), 4(a), and 5(a)). We then extracted patches from these super-patches, and stacked them together to create the high-resolution patch library. In general, the high-resolution pixels used in MDF accounted for less than 10% of all acquired pixels.

We did not add noise to the measurements, and therefore set $\sigma_w = 0$. As stated earlier, the converged results of

²In the sparse interpolation experiment, we could not compare against the SISR algorithm since it is not designed for sparse interpolation problems.

P&P do not depend on the choice of σ_λ^2 ; however, using the described procedure, we set $\sigma_\lambda^2 = 64$ for sparse interpolation, $\sigma_\lambda^2 = 55$ for super-resolution of *hinea brasiliana*, and $\sigma_\lambda^2 = 72$ for super-resolution of gold atoms. The only free parameter for P&P reconstruction is then the unit-less parameter β , which we selected to minimize the mean squared error of the reconstruction compared to the ground truth (see Table 4). For all experiments, we present the P&P residual convergence resulting from using different priors (see Tables 3). The normalized residue [3, p. 18], $r^{(k)}$ is given by

$$r^{(k)} = \frac{\|\hat{x}^{(k)} - \hat{v}^{(k)}\|_2}{\|\hat{x}^{(\infty)}\|_2} \quad (18)$$

where $\hat{x}^{(k)}$ and $\hat{v}^{(k)}$ are the values of \hat{x} and \hat{v} respectively after the k -th iteration of the P&P algorithm, respectively, and $\hat{x}^{(\infty)}$ the final value of the reconstruction, \hat{x} .

We initialized the P&P super-resolution algorithm with cubic-interpolation and sparse-interpolation with Shepard-interpolation. Further, to implement 8x and 16x super-resolution, we used a hierarchical strategy for enhancing resolution. For example, in the 16x case, we super-resolved the low-resolution image to 4x, then the 4x image to 8x, and then the 8x image to 16x.

RMS Error of the Super-Resolution Experiments			
Dataset / resolution synthesized	Cubic	SISR	MDF
<i>Hinea brasiliana</i> / 4x	10.33%	6.14%	4.06%
Gold atoms / 4x	16.12%	11.73%	7.92%
Gold atoms / 8x	39.36%	32.58%	11.13%
Gold atoms / 16x	45.71%	38.11%	15.08%

Table 1: RMSE between the super-resolution interpolated image and the high-resolution ground truth.

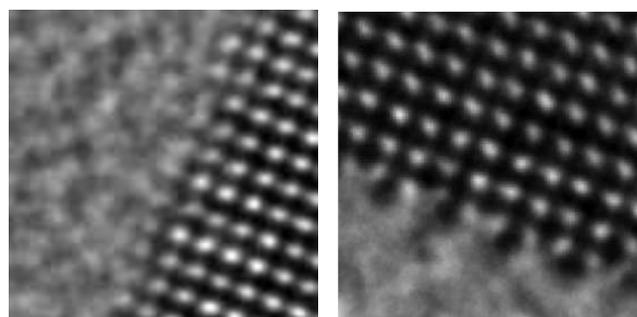
RMS Error of the Sparse Interpolation Experiment			
Experiment	Shepard	DSG-NLM	MDF
Sparse interpolation	13.10%	10.64%	9.77%

Table 2: RMSE between the interpolated image and the ground truth.

Convergence Error of P&P		
Experiment	DSG-NLM	MDF
Sparse interpolation	$4.66 \times 10^{-8}\%$	$7.57 \times 10^{-9}\%$
Super-resolution (<i>hinea brasiliana</i>)	–	$3.12 \times 10^{-7}\%$
Super-resolution (gold atoms)	–	$4.54 \times 10^{-6}\%$

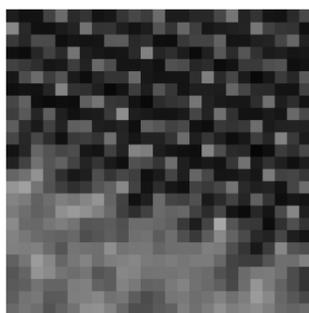
Table 3: Convergence error of the P&P algorithm

Our MDF reconstructions result in 15x to 228x data acquisition speedup across our datasets. The speedup factors

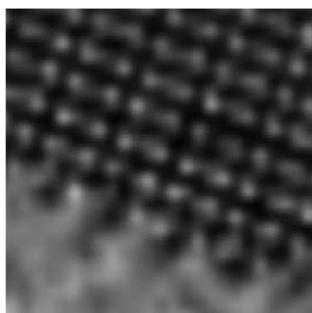


(a) A portion of the library image

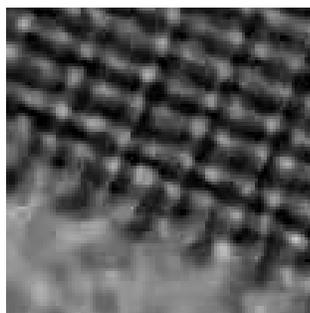
(b) Ground truth



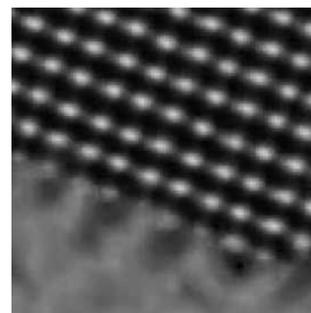
(c) Low-resolution input



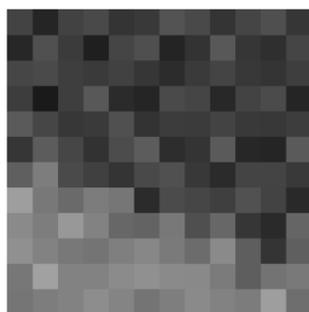
(d) 4x cubic interpolation



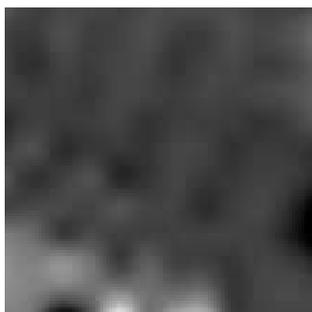
(e) 4x SISR



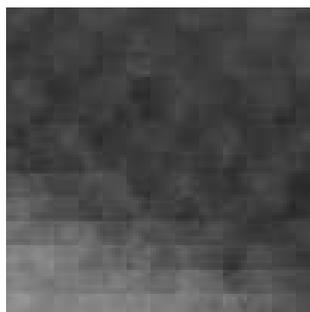
(f) 4x MDF (proposed algorithm)



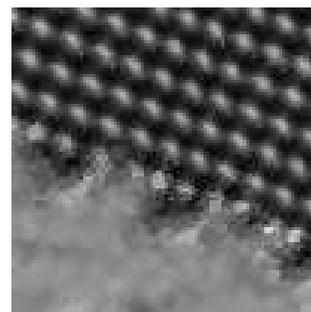
(g) Low-resolution input



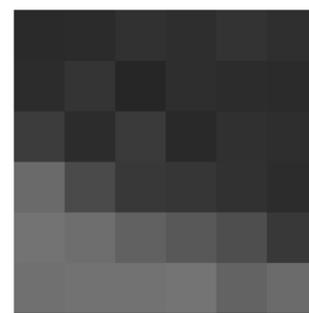
(h) 8x cubic interpolation



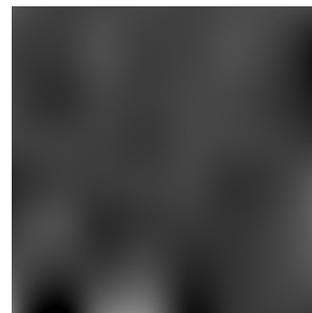
(i) 8x SISR



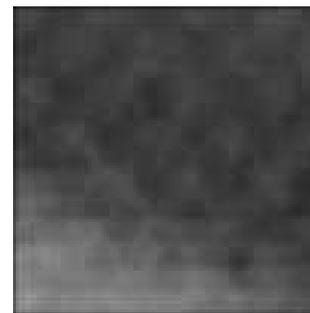
(j) 8x MDF (proposed algorithm)



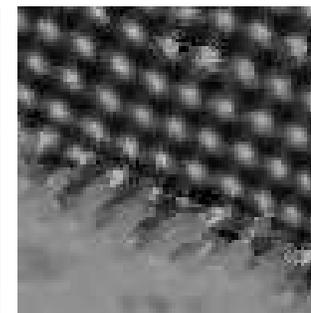
(k) Low-resolution input



(l) 16x cubic interpolation



(m) 16x SISR



(n) 16x MDF (proposed algorithm)

Figure 3: 4x, 8x, and 16x super-resolution of an atomic-resolution TEM image of gold atoms on a carbon substrate – zoomed in to show details. Subfigures (c), (g), and (k) were obtained by artificially reducing the resolution of the ground truth image (b), by 4x, 8x, and 16x, respectively. These low-resolution images have been rendered by pixel replication to match the size of the ground truth image. Our MDF algorithm resolves the structure of the gold atoms in all of the cases (see (f), (j), and (n)), even when the other methods struggle (see (h), (i), (l), and (m)).

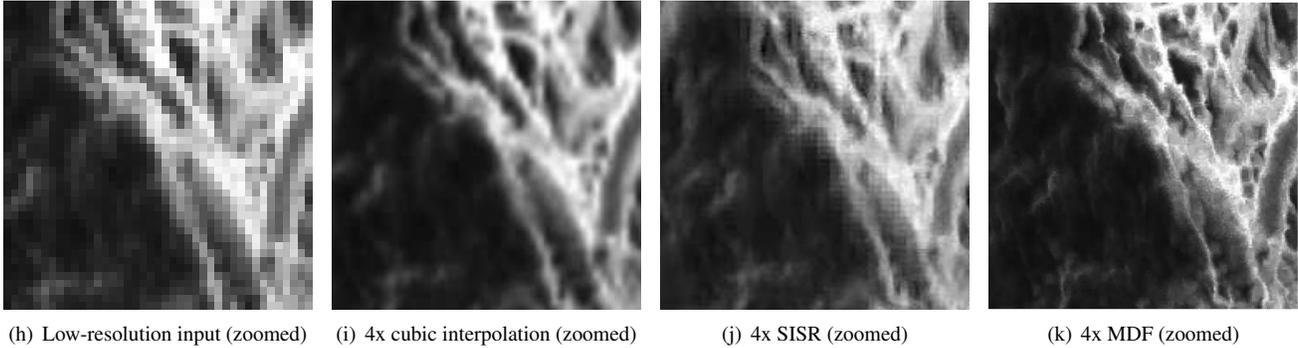
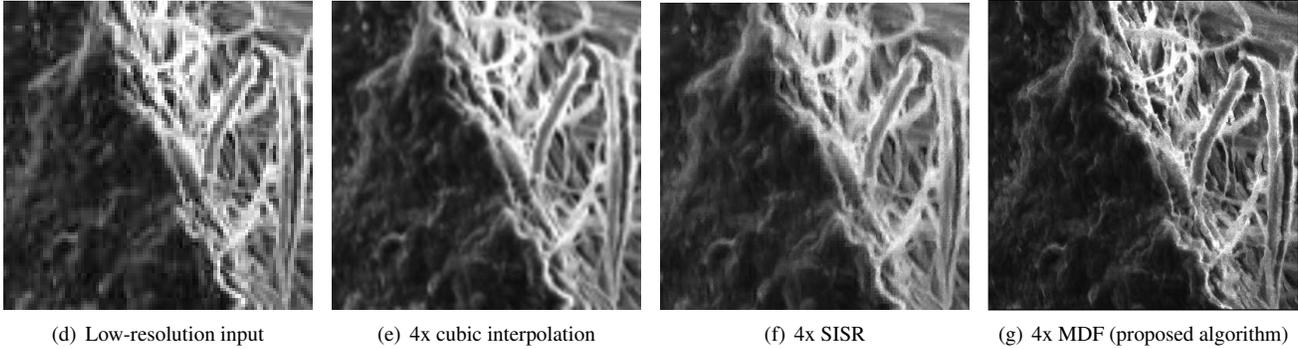
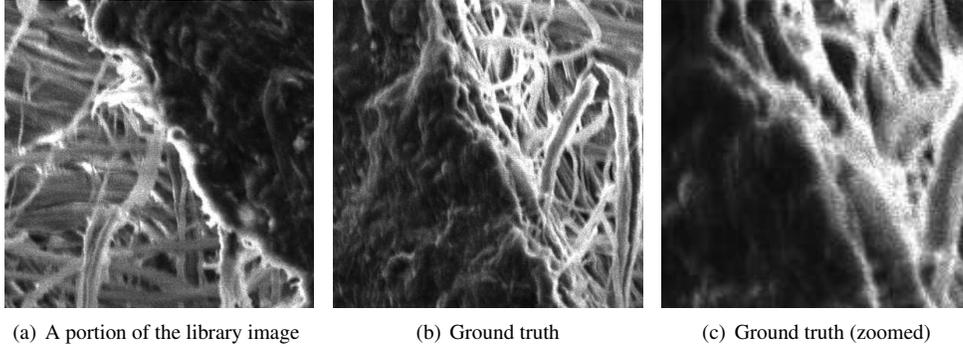


Figure 4: 4x super-resolution of an SEM image of surface crack in the shell of the marine mollusk *Hinea brasiliiana* – zoomed in to show details. The SEM image in (b) was acquired at 4× resolution compared to the image in (d), but under similar imaging conditions. Our MDF results in (g) and (k) show the intricate textures and clear edge features that are visually consistent with the ground truth (b). Subfigures (i) and (j) are zoomed in from (e) and (f) respectively. In (i), we see that the cubic interpolation leads to an overly smoothed interpolation, while we see in (j) that the SISR algorithm produces block artifacts.

Regularization Parameter, β		
Experiment	DSG-NLM	MDF
Sparse interpolation	1.05	1.15
Super-resolution (<i>hinea brasiliiana</i>)	–	1.40
Super-resolution (gold atoms)	–	1.20

Table 4: Regularization parameter, β

are computed as the ratio of the number of pixels in the reconstructed image to the number of pixels measured. In addition to specifying the speedup factors, we specify the

size of our high-resolution library relative to the size of the low-resolution image acquired through the quantity,

$$\rho = \frac{\# \text{ of high-resolution pixels}}{\# \text{ of total pixels}}. \quad (19)$$

Table 5 gives the value of ρ and the speedup factors for each of our experiments.

Importantly, the MDF reconstructions are clearer than Shepard and cubic interpolations, despite the high speedup factors achieved. Also, MDF results in the least RMSE (see Tables 1 and 2), while substantially reducing jaggies and producing better textures and features compared to the SISR algorithm, cubic and Shepard interpolations.

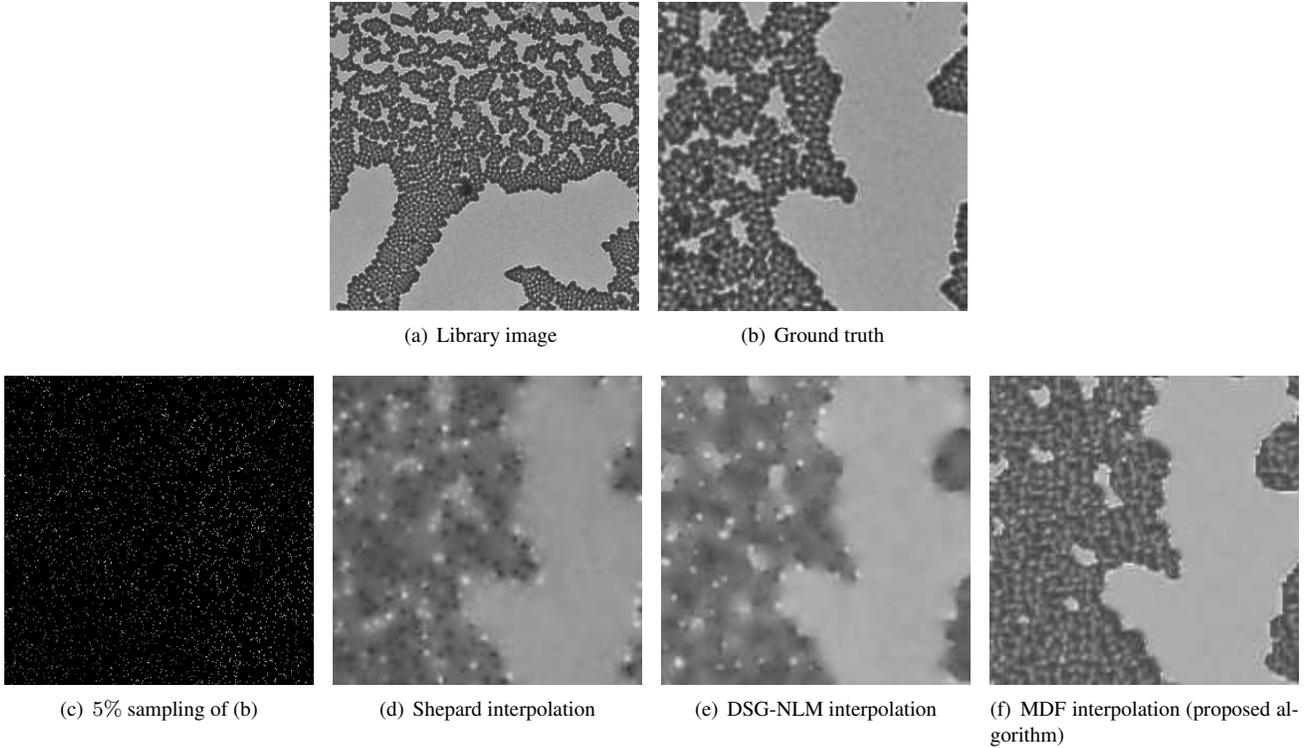


Figure 5: Interpolation of a TEM image of iridovirus assemblies on a carbon substrate; (c) shows the 5% randomly sampled image that serves as the input to all the interpolation methods; (f) shows the texture synthesized by our MDF algorithm, which is visually more accurate compared to that generated by DSG-NLM in (e) and Shepard’s interpolation in (d). In addition, the MDF interpolation results in the least mean squared error of all the methods.

Data Acquisition Speedup		
Experiment	ρ	Speedup factor
Super-resolution (hinea brasiliana / 4x)	6.11%	15x
Super-resolution (gold atoms / 4x)	0.76%	15.87x
Super-resolution (gold atoms / 8x)	2.96%	62.1x
Super-resolution (gold atoms / 16x)	10.88%	228.14x
Sparse interpolation (5% sampling)	8.74%	18.25x

Table 5: Data acquisition speedup achieved by our MDF algorithm. The ratio of the high-resolution pixels measured to the low-resolution pixels measured is given by ρ .

5. Conclusion

In this paper, we introduced a novel computational electron microscopy imaging method based on multi-resolution data fusion (MDF) that combines innovative data acquisition with novel algorithms to dramatically improve the resolution/speed trade-off as compared to traditional EM imaging methods. High-resolution measurements taken over a small field-of-view are used to create a material-specific

probabilistic model in the form of a patch-library that is used to dramatically increase resolution. We introduced a Bayesian framework for the fusion of high and low-resolution data without the need for training. We outlined a library-based non-local means (LB-NLM) algorithm that uses typical high-resolution patches, and we used it as a prior model within the *plug-and-play priors* framework to produce visually true textures and edge features in the interpolations, while significantly reducing mean squared error. Finally, we presented 4x, 8x, and 16x super-resolution and 5% sparse image interpolation results demonstrating data acquisition speedups of up to 228x on various datasets, while substantially maintaining image quality. We demonstrated the power of our method on real electron microscope images and compared our results against the state-of-the-art single-image super-resolution (SISR) algorithm. In our experiments, we found our MDF algorithm to consistently produce clearer, sharper image reconstructions with visibly truer textures and lower mean squared error compared to the SISR, cubic, and Shepard interpolations.

In summary, our MDF computational EM method enables high-resolution EM data to be collected at a fraction of the time as conventional methods, while maintaining high image quality – offering an *integrated imaging* solution to a grand challenge problem in electron microscopy.

References

- [1] C. B. Atkins, C. A. Bouman, and J. P. Allebach. Tree-based resolution synthesis. In *PICS*, pages 405–410, 1999.
- [2] C. Bouman and K. Sauer. A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Trans. on Image Processing*, 2(3):296–310, July 1993.
- [3] 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*, 3(1):1–122, 2011.
- [4] A. Brifman, Y. Romano, and M. Elad. Turning a denoiser into a super-resolver using plug and play priors. In *Image Processing (ICIP), 2016 IEEE International Conference on*, pages 1404–1408. IEEE, 2016.
- [5] D. D. Deheyn and N. G. Wilson. Bioluminescent signals spatially amplified by wavelength-specific diffusion through the shell of a marine snail. *Proceedings of the Royal Society of London B: Biological Sciences*, page rspb20102203, 2010.
- [6] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.
- [7] 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*, 19(3):340–358, 2005. Computational Harmonic Analysis - Part 1.
- [8] S. Farsiu, M. D. Robinson, M. Elad, and P. Milanfar. Fast and robust multiframe super resolution. *Image processing, IEEE Transactions on*, 13(10):1327–1344, 2004.
- [9] W. T. Freeman, T. R. Jones, and E. C. Pasztor. Example-based super-resolution. *Computer Graphics and Applications, IEEE*, 22(2):56–65, 2002.
- [10] D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.
- [11] G. Godaliyadda, D. H. Ye, M. D. Uchic, M. A. Groeber, G. T. Buzzard, and C. A. Bouman. A framework for dynamic image sampling based on supervised learning (slads). *arXiv preprint arXiv:1703.04653*, 2017.
- [12] C. Guillemot and O. Le Meur. Image inpainting : Overview and recent advances. *Signal Processing Magazine, IEEE*, 31(1):127–144, Jan 2014.
- [13] S. B. Juhl, E. P. Chan, Y.-H. Ha, M. Maldovan, J. Brunton, V. Ward, T. Dokland, J. Kalmakoff, B. Farmer, E. L. Thomas, et al. Assembly of wiseana iridovirus: viruses for colloidal photonic crystals. *Advanced Functional Materials*, 16(8):1086–1094, 2006.
- [14] S. S. Kubatur and M. L. Comer. Simulation of grain growth in polycrystalline materials: A signal processing perspective. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 853–857. IEEE, 2015.
- [15] S. S. Kubatur and M. L. Comer. Rare event simulation for markov random fields with application to grain growth in crystals. In *Image Processing (ICIP), 2016 IEEE International Conference on*, pages 3748–3752. IEEE, 2016.
- [16] K. Park, L. F. Drummy, R. C. Wadams, H. Koerner, D. Nepal, L. Fabris, and R. A. Vaia. Growth mechanism of gold nanorods. *Chemistry of Materials*, 25(4):555–563, 2013.
- [17] E. Pérez-Pellitero, J. Salvador, J. Ruiz-Hidalgo, and B. Rosenhahn. Psycho: Manifold span reduction for super resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1837–1845, 2016.
- [18] E. Shechtman, Y. Caspi, and M. Irani. Space-time super-resolution. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(4):531–545, 2005.
- [19] S. Sreehari, S. Venkatakrishnan, B. Wohlberg, G. T. Buzzard, L. F. Drummy, J. P. Simmons, and C. A. Bouman. Plug-and-play priors for bright field electron tomography and sparse interpolation. *Computational Imaging, IEEE Transactions on*, 2016.
- [20] A. Stevens, H. Yang, L. Carin, I. Arslan, and N. D. Brown-ing. The potential for Bayesian compressive sensing to significantly reduce electron dose in high-resolution STEM images. 63(1):41–51, 2014.
- [21] C. Thomas, T. Ranchin, L. Wald, and J. Chanussot. Synthesis of multispectral images to high spatial resolution: A critical review of fusion methods based on remote sensing physics. *IEEE Transactions on Geoscience and Remote Sensing*, 46(5):1301–1312, 2008.
- [22] R. Timofte, R. Rothe, and L. Van Gool. Seven ways to improve example-based single image super resolution. *arXiv preprint arXiv:1511.02228*, 2015.
- [23] 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*, pages 945–948. IEEE, 2013.
- [24] D. B. Williams and C. B. Carter. The transmission electron microscope. In *Transmission electron microscopy*, pages 3–17. Springer, 1996.
- [25] J. Yang, X. Liao, X. Yuan, P. Llull, D. Brady, G. Sapiro, and L. Carin. Compressive sensing by learning a gaussian mixture model from measurements. *Image Processing, IEEE Transactions on*, 24(1):106–119, Jan 2015.
- [26] J. Yang, J. Wright, T. Huang, and Y. Ma. Image super-resolution as sparse representation of raw image patches. In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.
- [27] R. Zeyde, M. Elad, and M. Protter. On single image scale-up using sparse-representations. In *International conference on curves and surfaces*, pages 711–730. Springer, 2010.
- [28] K. Zhang, X. Gao, D. Tao, and X. Li. Single image super-resolution with non-local means and steering kernel regression. *Image Processing, IEEE Transactions on*, 21(11):4544–4556, 2012.
- [29] N. Zhao, Q. Wei, A. Basarab, N. Dobigeon, D. Kouamé, and J.-Y. Tourneret. Fast single image super-resolution using a new analytical solution for $l_2 - l_2$ problems. *IEEE Transactions on Image Processing*, 25(8):3683–3697, 2016.