Fast reconstruction of atomic-scale STEM-EELS images from sparse sampling

Etienne Monier, Thomas Oberlin, Nathalie Brun, Xiaoyan Li, Marcel Tencé, Nicolas Dobigeon

HAL Id: hal-02949156
https://hal.science/hal-02949156

Submitted on 28 Sep 2020

Fast reconstruction of atomic-scale STEM-EELS images from sparse sampling[*]

Etienne Monier[a,*], Thomas Oberlin[a,b], Nathalie Brun[c], Xiaoyan Li[c], Marcel Tencé[c], Nicolas Dobigeon[a,d]

[a]*University of Toulouse, IRIT/INP-ENSEEIHT, 31071 Toulouse Cedex 7, France*
[b]*University of Toulouse, ISAE-SUPAERO, 31400 Toulouse, France*
[c]*Université Paris-Saclay, CNRS, Laboratoire de Physique des Solides, 91405, Orsay, France*
[d]*Institut Universitaire de France (IUF), France*

Abstract

This paper discusses the reconstruction of partially sampled spectrum-images to accelerate the acquisition in scanning transmission electron microscopy (STEM). The problem of image reconstruction has been widely considered in the literature for many imaging modalities, but only a few attempts handled 3D data such as spectral images acquired by STEM electron energy loss spectroscopy (EELS). Besides, among the methods proposed in the microscopy literature, some are fast but inaccurate while others provide accurate reconstruction but at the price of a high computation burden. Thus none of the proposed reconstruction methods fulfills our expectations in terms of accuracy and computation complexity. In this paper, we propose a fast and accurate reconstruction method suited for atomic-scale EELS. This method is compared to popular solutions such as beta process factor analysis (BPFA) which is used for the first time on STEM-EELS images. Experiments based on real as synthetic data will be conducted.

Keywords: scanning transmission electron microscopy, electron energy loss spectroscopy, atomic-scale images, spectrum-images, partial acquisition, fast reconstruction

1. Introduction

Electron energy loss spectroscopy (EELS) performed in a scanning transmission electron microscope (STEM) has proved to be a powerful tool to analyze chemical components and structures of a sample with a sub-nanometer spatial resolution. A focused electron probe is scanned over the sample and for each probe position an EELS spectrum is acquired, as well as several other signals such as high-angle annular dark field (HAADF). The spectrum-image thus acquired can be used to build not only maps of the spatial distribution of the elements but also maps of edges' fine structures corresponding to local electronic structures.

These important capabilities of modern microscopes are somewhat limited by sample damage, instabilities and poor signal-to-noise ratio (SNR). Indeed, acquiring such EELS data set requires a suitable SNR and typical EELS dwell time (exposure time per location) are in the ms range

$(1 - 100\text{ms})$. These long dwell times proportionally lead to a significant total electron dose received by the sample. This dose increases potential radiation damages to the sample [2]. This is particularly problematic for sensitive materials such as biological samples. Moreover a long acquisition time may increase image distortions caused by time-dependant instabilities of the sample and the microscope. In particular, these instabilities may be substantial at atomic scales. Performing a multi-frame acquisition, followed by a non-rigid alignment step [3], is a promising research domain to improve the spatial resolution and to reduce beam-induced damage. The new generation of direct detection cameras with negligible correlated noise could promote the use of this multi-frame setup with even lower dwell-times. Finally some high resolution acquisitions need to cover large areas (such as in [4] for scanning electron microscopy (SEM)), leading to long acquisition total time, heavy data storage and long processing steps. To increase acquisition speed and/or reduce the full beam exposure, a solution consists in reducing dwell time and subsequently denoising the data as a post-processing operation. Yet, reducing the exposure may be of limited interest since the resulting SNR becomes too low to expect good denoising performance, especially in the case of fine structure analysis.

A recent popular alternative is *sparse (or partial) sampling*. This strategy consists in acquiring the relevant signals only in a small proportion of spatial locations, which

allows for higher dwell time at these positions resulting in the same amount of total electron dose. The resulting acquired image is partially empty and a reconstruction step is required to obtain a fully exploitable image. This paradigm received a renewed interest since the theoretical results of *compressed sensing* (CS) which states that exact recovery of sub-Nyquist rate acquisitions is possible under certain conditions – one of them is that the data should be sparse in an appropriate basis. The CS paradigm states that the data should be projected on n random subspaces with n far below the data size, which is well adapted to electron microscopy tomography [5, 6, 7]. These results raised a lot of interest toward inverse problems which estimate the image based on partial spatial acquisitions which is referred to as *inpainting*. It remains an active research area for STEM [8, 9] and SEM [4], among others.

The two previously described acquisition schemes have pros and cons. Schematically, low dwell-time acquisition usually produces better spatial results while sparse sampling images usually have rich spectral information. Determining which approach is the best is not trivial. To that end, recent works studied and compared these solutions [10, 11] based on experiments conducted on synthetic as well as real images.

Following the second aforementioned acquisition scheme, this paper addresses the problem of reconstructing spatially sub-sampled atomic-scale STEM EELS images. In particular one motivation here aims at reducing computational burden of the inpainting procedure to make its future implementation possible into the acquisition process. The experimenter should be able to visualize the full spectrum-image along the acquisition, which requires both fast computation and a good accuracy. In addition to this online setup, the experimenter should be able to refine the reconstructed spectrum-image afterwards, where very accurate but possibly time-consuming algorithms are allowed. To that end, we propose a new reconstruction method exhibiting a relevant trade-off between accuracy and complexity. We will also show that this proposed technique can serve as a good initialization to accelerate more efficient yet more computationally intensive methods. Moreover, among the compared methods, we propose to apply the popular beta process factor analysis (BPFA), originally dedicated to remote sensing images [12]. Up to our knowledge, it is the first time BPFA is applied to STEM-EELS images, although it was already used in many microscopy works for 2D data restoration such as in SEM [10]. The paper is organized as follows. Section 2 presents an overview of inpainting techniques already used in electron microscopy, with the emphasis placed on 2D and 3D reconstructions. Since no fast and accurate 3D method fulfills all requirements to envisage a fully operational online implementation, Section 3 describes the newly proposed reconstruction method. Section 4

describes the synthetic, semi-real and real data, as well as the experiments conducted to compare the proposed approach to previous works (especially BPFA as a 3D reconstruction method). The experimental results are reported and discussed in Section 5. Section 6 concludes this study.

2. Related works

The focus of this paper deals with reconstructing spatially sub-sampled STEM images. Many works considered this problem with different methods and modalities. Most of them were proposed to process single 2D images while few considered the reconstruction of 3D images. This section discusses these works which are mainly divided into two parts. The first one considers *learning-free* methods which reconstruct the images based on the single acquired dataset. The second one studies *learning-based* methods which capitalize on a learning set to calibrate an operator subsequently used to reconstruct the data.

2.1. Learning-free methods

Among learning-free methods, nearest neighbor (NN) interpolation is a fast and simple solution possibly allowing for dynamic joint acquisition and reconstruction. To avoid piecewise-constant image like reconstruction, preferred solutions interpolate with a weighted mean over a neighborhood. The weights are chosen to be the normalized inverse distance between the interpolated pixel and the neighbor for reconstructing SEM data in [13], energy-dispersive X-ray spectroscopy (EDS) data in [14, 15] and EELS data in [15]. An alternative considered for SEM images in [10] was based on a natural neighbor interpolation, which adjusts the weights based on a Voronoï cell representation [16]. Note that STEM acquires full spectra at particular spatial positions.

Regularized least-square (LS) methods generally offers better results than NN as they additionally constrain the reconstructed image to fulfill a predefined behavior, usually promoted by a well chosen regularization. Akin to the CS paradigm, a classical regularization is the sparsity of the reconstructed image in an appropriate basis, as the ℓ_1-norm regularized LS problem considered in [17] for atomic force microscopy (AFM). This type of regularized LS methods will be referred as ℓ_1-LS in Table 1. In the case of periodic structures (as for atomic-scale images), this basis can be Fourier or discrete cosine transformation (DCT). The authors in [18] proposed an inpainting method for atomic-scale high-angle annular dark-field (HAADF) images based on a thresholded Fourier transform, which constrains the image sparsity in periodic basis. The method in [19] promoted the sparsity of the DCT

representation to reconstruct HAADF images, using the SPGL1 algorithm [20]. In the same way, this regularization can be coupled with a wavelet basis to dynamically reconstruct HAADF data [21]. Another standard regularization is total variation (TV), i.e., the ℓ_1-norm of the image gradient promoting piecewise constant reconstructed image, as considered in [17] for AFM. The block-DCT representation was coupled with TV for reconstructing SEM data in [4]. The ℓ_2-norm of the image gradient is also widely used as a regularization to promote spatial smoothness and is referred to as the Sobolev energy [11]. In the case of multi-band images, spectral regularizations were proposed in addition to the spatial one. In [11], for instance, the 3S method uses the weighted ℓ_2-norm across the EELS spectrum-image bands or simply a nuclear norm (which ensures the low-rank nature of the reconstructed data) in addition to the classical Sobolev energy spatial regularization.

Another class of reconstruction methods exploit the spatial redundancy in the image, often referred to as *patch-based methods*. They form a very popular and successful class of reconstruction methods which raised a lot of attention in the last decades to solve inverse problems such as denoising, inpainting and deblurring. For example, the exemplar-based inpainting [22] (EBI) reconstructs partially corrupted images by iteratively replacing the image patches by the best matching uncorrupted patch extracted from its neighborhood. To describe spatial redundancy, successful algorithms aims at sparsely representing image patches thanks to atoms of a dictionary jointly learned with the reconstructed data. BPFA is probably the most popular dictionary learning (DL) method in the microscopy community [12]. It was first used for atomic-scale HAADF images in [9] and was used afterwards in many papers for the same kind of data [23, 24]. The authors of BPFA proposed Kruskal-factor analysis (KFA) as a tensor extension of BPFA [25]. KFA was used to reconstruct EELS images based on a multiplexed spectrum-image acquisition [26]. Last, the expected-patch log-likelihood (EPLL) algorithm assumes the patch distribution to be a Gaussian mixture [27]. Yet, its computation time is important and the authors in [15] preferred a simplified but accelerated version called fast EPLL (FE-PLL) to reconstruct SEM images [28]. In addition to the patch-based methods used in the microscopy community, wKSVD [29] and ITKrMM [30, 31] learn the dictionary from incomplete data without assuming particular patch distribution. They remain state-of-the-art methods that will be considered in this paper.

To achieve better performance with reduced acquisition time, several predefined scan patterns were proposed such as regular scan [32], random horizontal lines [24, 17], mixed regular-random scan [18, 32], spiral scans [33, 34, 17] or square-shape scan [17]. These results tend to show that the best performance is achieved by semi-random scan patterns, which introduce randomness and avoid large holes. Last, adaptive sparse scanning enables consequent reconstruction improvement by selecting the pixel to sample based on previously acquired data. In [35], the authors proposed to perform a first low-SNR scan to locate the spatial edges. A second high-SNR scan is then performed on these edges only. Finally, in the low-SNR acquired image, the smooth regions are filtered and the edges are filled with the pixels from the second high-SNR acquisition. An alternative adaptive scanning scheme proposed in [36] consists in iteratively locating possible points of interest to be sampled. Learning-based adaptive sparse scanning are discussed in the next subsection.

2.2. Learning-based methods

Contrary to learning-free methods which recover the full image based only on the partially acquired data, learning-based methods learn an operator based on a possibly large training set. These methods are known to be much more accurate as long as the geometric content of the image to reconstruct is similar to the content of the training data.

For instance, the GOAL algorithm learns a dictionary which maximizes the sparsity of the training dataset representation [37]. The learned dictionary is then used to perform the inpainting task for test data. Similarly, EBI which is originally a learning-free method, can be adapted to benefit from the availability of a training set. To that end, within the conventional EBI framework, instead of extracting the copied patch from the neighborhood, this patch can be chosen from a dictionary learned beforehand on uncorrupted images. This is the strategy followed in [38] to reconstruct 3D SEM data. GOAL and the learning-based counterpart of EBI were used in [10] for 2D SEM images but BPFA seemed to give better results.

Learning-based approaches can also be designed to decide which positions should be sampled in order to minimize the distortion after reconstruction. Indeed, sub-sampled data reconstruction performance highly depends on the sample locations [10]. To improve the reconstruction quality, the supervised learning approach for dynamic sampling (SLADS) learns a function (called *expected reduction in distortion* (ERD)) indicating which location should be sampled to maximally reduce distortion [13]. This learning step is based on a list of descriptors and requires labeled training data, and it was used to dynamically sample SEM images. This method has been also applied to EDS data in [14]. To that end, a convolutional neuronal network (CNN) classifies the test data spectra and the ERD function is computed simultaneously for all labels. The paper [15] modified this approach to allow mixed elements in EELS and EDS. All these approaches needs a fast reconstruction which is achieved thanks to a weighted NN technique.

3

Family	Method	Works 2D	3D	Execution time	Accuracy
NN	NN	-	-	+++	−−
	Weighted neighbor	[16, 10]	-	++	−
LS-regularized	ℓ_1-LS	[17, 19, 21, 4]		+	+
	TV-LS	[17]	-	+	+
	3S	-	[11]	+	+
DL-based methods	BPFA	[9, 10]	*[12]*	−−−	+++
	EBI	[10]	[38]	−	++
	FEPLL	*[28]*,[15]	-	−	++
	wKSVD	-	*[29]*	−−	++
	ITKrMM	*[30]*	*[31]*	−	++
	GOAL	*[37]*,[10]	-	−	++

Table 1: Comparison of the methods proposed in the microscopy literature for reconstructing partially sampled images. Additional references not originated from the microscopy literature are also provided in italics. The execution time and accuracy are qualitatively evaluated based on the results of Section 4.

2.3. Application to EELS and feasibility

The previous subsection focused on the related works in microscopy which are summarized in Table 1. They are rated depending on their computational complexity and accuracy and grouped into three main families: NN, LS-regularized and DL-based methods. The works from the literature related to each method are given and separated depending on their ability to reconstruct 2D mono-band images (e.g., HAADF) or 3D spectrum-images (e.g., EELS).

Among the available methods for EELS reconstruction, NN is fast but gives generally poor reconstruction results while DL-based methods are very efficient but remains computationally expensive, especially when considering 3D patches. Therefore, this literature review shows that none of the available methods can optimally reconstruct a spatially sub-sampled spectrum-image fast enough to be included into an experimental setup for online or mini-batch processing. Note that 3S could satisfy the accuracy and speed requirements but the results of Section 4 will show its regularization is not suited for atomic-scale images. In this work, we will propose to apply BPFA to EELS images, i.e., recovering a dictionary composed of 3D patches, as this was originally designed for.

An alternative for systematically reconstructing a spectrum-image consists in processing *separately* and *in parallel* the 2D images associated with each channel. In this case, note that NN as a 3D reconstruction performs the same as a band-by-band processing. Yet, this is sub-optimal as the reconstruction task is expected to perform better by capitalizing on the information of the whole 3D data. Similar considerations could lead to prefer recon-structing one or several single channel images of interest necessary for element mapping. However such a strategy may be also sub-optimal when no a priori knowledge is available regarding the sample to be imaged.

To conclude, NN and DL-based methods are not suited for on-line reconstruction and only LS-regularized methods combine accuracy and reduced computational cost. As a consequence this paper proposes a method which belongs to the regularized LS family to reconstruct quickly and efficiently an atomic-scale spectrum-image. This method is detailed in the next section and will be compared to existing approaches in Section 4 based on synthetic as well as real data experiments.

3. Proposed method

3.1. Structured sparsity in a well-chosen basis

To get a fast and accurate reconstruction technique, LS-regularized methods exploiting sparsity in an appropriate basis seems to provide a relevant trade-off. In particular, to ensure an acceptable computation time, this basis or dictionary can be chosen beforehand, exploiting some expected properties of the image to be reconstructed. Learning this dictionary during a pre-processing step would be time consuming and the reconstruction results would strongly depend on the training set. As a consequence, an analytic basis such as Fourier, DCT or wavelets will be preferred due to its explicit expression. The choice of the basis will be discussed in Section 5.1 based on experiments. Yet, the periodic structure of atomic-scale sample is expected to favor Fourier or DCT representations.

More precisely, in the case of atomic-scale EELS spectrum-images, each 2D image in a given band is expected to exhibit a periodic pattern which can be described thanks to a sparse representation in an appropriate basis. Conversely, each spectrum measured in any spatial location does not likely exhibit particular periodicity. Thus this sparsity property only holds in the *spatial* direction of the 3D datacube and a band-by-band basis transformation is expected to lead to higher sparsity level (i.e., the proportion of nonzero coefficients) compared to 3D basis transformation. Besides, band-by-band representations are expected to share common characteristics since the spatial structure is likely the same across the channel and only depends on the sample. This *structured sparsity* tends to promote images which non-zero band-by-band representation coefficients are located at the same place. This phenomenon is illustrated in Fig. 1 when considering a DCT representation (the choice of the basis will be discussed in Section 5.1). Note that in the case of an important noise level (such as for the band #1111), some powerful high-frequency coefficients correspond to noise (at the bottom-right corner of the panels). Yet, by jointly analyzing the DCT representations for several channels, the location of the main coefficients can be estimated after removing co-efficients associated with noise.

3.2. Cosine Least Square reconstruction

Capitalizing on the finding of the previous sections, the proposed method relies on a regularized LS problem whose formulation writes

$$\hat{\mathbf{X}} = \arg\min_{\mathbf{X}} \frac{1}{2}||\mathbf{Y} - \mathbf{X}\Phi||_{\mathrm{F}}^2 + \lambda\mathcal{R}(\mathbf{X}). \qquad (1)$$

where $\hat{\mathbf{X}}$ is the recovered image, Φ is the subsampling operator and $\mathbf{Y}$ is the acquired data. The operator $\mathcal{R}(\cdot)$ is a regularization and λ is a scalar that adjusts the importance of the regularization with respect to the data fidelity term. To enforce the periodic spatial patterns to be similar across the channels, we choose the regularization $\mathcal{R}(\mathbf{X})$ to be

$$\mathcal{R}(\mathbf{X}) = ||\mathbf{X}\Psi||_{2,1} \qquad (2)$$

where Ψ is an orthogonal band-by-band basis transform. The $\ell_{2,1}$-norm is defined as[1]

$$||\mathbf{U}||_{2,1} = \sum_j ||\mathbf{U}_j||_2 = \sum_j \sqrt{\sum_i |\mathbf{U}_{ij}|^2} \qquad (3)$$

where $\mathbf{U}_j$ is the jth column of $\mathbf{U}$, corresponding to the spectrum associated with the jth pixel. Minimizing this norm enforces the structured sparsity as it aims at setting the columns of lowest magnitude to zero while preserving


Figure 1: Illustration of the data structured sparsity. The STEM acquisition is composed of a 2D HAADF image (left panel, 1st row) and the corresponding spectrum-image whose three bands are considered (left panels, 2nd to 4th rows). For each image, the locations of the 2% DCT coefficients of highest magnitude are depicted (right panels). Note that most of these coefficients are located in the same areas of the DCT space for the HAADF as well as corresponding bands, which suggests a certain structured sparsity in the the transform domain.

---

[1] A slightly different definition of the $\ell_{2,1}$-norm is proposed in [39], where the ℓ_2-norm is first computed on the rows of $\mathbf{U}$.

the most powerful columns. In case of a periodic basis (e.g., Fourier or DCT), the regularization (2) promotes similar and sparse frequency representations across all the bands. For further details, see Appendix A. In the sequel of the paper, the resulting method will be referred to as *Cosine Least Square* (CLS).

CLS can be used directly on the partially-acquired spectrum-image. Yet, as a preprocessing step, we propose to perform a principal component analysis (PCA) of the acquired data $\mathbf{Y}$ and to apply CLS on the T first principal components. The formulation (1) remains the same but the measurement matrix $\mathbf{Y}$ and reconstructed image matrix $\mathbf{X}$ are replaced by the PCA spectrum-image $\tilde{\mathbf{Y}}$ and its reconstructed counterpart $\tilde{\mathbf{X}}$. The fully reconstructed image $\mathbf{X}$ can then be derived from $\tilde{\mathbf{X}}$ by the corresponding inverse transformation. While it is not a prerequisite of the proposed method, this preprocessing step has several advantages. First, it implicitly introduces a spectral regularization of the inverse problem by imposing a low-rank structure of the solution. Similar strategies have been widely promoted for various tasks conducted on multiband images, including compressive sensing [40, 41], inpainting [11, 42], fusion [43, 44, 45] or mixture analysis [46, 47]. Second, by reducing the amount of data to process, it allows the computational time to be significantly reduced. This preprocessing may induce some reconstruction artefacts when the estimation of the covariance matrix is not accurate, e.g., in case of a low sampling ratio. Moreover, the number of PCA components to keep (which will be referred to as *PCA threshold* in the following) should be carefully chosen as a too small threshold can make some low-powerful structures disappear [48]. A systematic strategy to adjust this threshold was proposed in [11]. The use of PCA as a pre-processing step is discussed in the supporting document [49].

Besides, the choice of the regularizing parameter λ in (1) is non-trivial in general inverse problems. However, it can be fairly adjusted by exploiting some information about the image to be reconstructed such as the noise level (which can be estimated beforehand) or the sparsity level, i.e., the ratio of non-zero coefficients of the representation. This prior knowledge can be used to assess the quality of the solution $\hat{\mathbf{X}}_\lambda$ obtained by CLS for a given value λ of the regularization parameter. For instance, for a relevant solution, the data fidelity term is expected to be of the order of magnitude of the noise level. Thus a dichotomic search can be conducted to adjust the regularization parameter automatically. In other words, CLS is run for a given value of the parameter and the data-fidelity term is evaluated at convergence. Since the fidelity term increases with λ, if its value is below (resp. above) the noise level, λ should be increased (resp. decreased) and CLS should be run again. For instance, a similar strategy has been successfully followed in [11].

3.3. CLS as a pre-processing step

As it will be shown in Section 4, the CLS algorithm is fast and it efficiently reconstructs spectrum-images based on partially acquired data. Thus we also propose here to apply CLS as a pre-processing step for more advanced algorithms, in particular to initialize DL-based methods. To that end, the CLS-reconstructed data $\hat{\mathbf{X}}$ is decomposed into a dictionary and a sparse representation using the conventional mini-batch DL algorithm proposed in [50] coupled to the orthogonal matching pursuit [51]. These dictionary and code are subsequently used to initialize DL-based reconstruction methods. This initialization is particularly interesting for wKSVD and ITKrMM which solve a nonconvex optimization problem based on a direction alternating scheme. Indeed, for such nonconvex problems, initialization is known to be a crucial issue to ensure convergence to a relevant solution. Alternatively, a similar strategy cannot be easily adopted to initialize BPFA since it implements a Markov chain Monte Carlo algorithm and the dictionary and code distributions depend on hyperparameter distributions. The relevance of CLS as an initialization step will be discussed in Section 5.4.

4. Experiments

4.1. Data

4.1.1. Materials

Two real atomic-scale spectrum-images referred to as R_1 and R_2 have been acquired on a NION UltraSTEM200 at the Laboratoire de Physique des Solides (LPS), Orsay, France. R_1 is a spectrum-image acquired at 100kV on a $La_{1-x}Sr_xMnO_3/Pb(Zr,Ti)O_3$ (LSMO/PZT) heterostructure grown onto $SrTiO_3$ [52]. The spectra were acquired in an energy range corresponding to $Ti-L_{2,3}$, $O-K$, $Mn-L_{2,3}$ and $La-M_{4,5}$ edges. R_2 is a spectrum-image acquired at 100kV on a $NdNiO_3$ thin film grown on a $LaAlO_3$ substrate [53]. The spectra were acquired in an energy range corresponding to $O-K$, $La-M_{4,5}$, $NiL_{2,3}$ and $Nd-M_{4,5}$ edges. To reduce the acquisition noise, the images were acquired with relatively long dwell time. Additional information about these images (such as their sizes and resolutions, dwell times and the PCA threshold used in Section 5.2) are reported in Table 2.

All experiments discussed in this section were conducted on an Intel Xeon CPU E5540 @ 2.53GHz with 8 cores – including hyperthreading – and 50Gb of memory. Note that the high amount of memory is only required for BPFA as a 3D algorithm while other methods can run on a machine with only 13.2Gb of memory.

Image	Sample	Size (x, y, λ) Res. $(\Delta x = \Delta y, \Delta \lambda)$	Dwell time (ms)	Relevant edges		PCA threshold T
				Element	Energy loss (eV)	
R_1	PbZrTiO$_3$ / LaSrMnO$_3$ / SrTiO$_3$	(232, 101, 1530) (0.055nm, 0.27eV)	20	Ti O Mn La	456 532 640 832	9
R_2	NdNiO$_3$/ LaAlO$_3$	(63, 115, 1505) (0.045nm, 0.32eV)	20	O La Ni Nd	532 832 855 978	7
S	cf. R_2	(70, 120, 1435) (0.045nm, 0.32eV)	-	cf. R_2		4

Table 2: Additional information about the R_1, R_2 and S images.

4.1.2. Synthetic and semi-real images

Since R_1 and R_2 are naturally corrupted by noise, computing relevant metrics to assess the performance of the reconstruction methods may be biased by an unknown noise level. To alleviate, noise-free counterparts of R_1 and R_2 were first generated by conducting a PCA and keeping only the first T principal components. The choice of the threshold T is generally a difficult task. In this work, it was set such that the discarded principal components did not carry any spatial information. More details are provided in Appendix B. These denoised images, referred to as $\bar{\mathsf{R}}_1$ and $\bar{\mathsf{R}}_2$, are assumed to be the ground-truth images $\mathbf{X}$ to be reconstructed.

In addition to these two semi-real images, a synthetic spectrum-image was generated based on the R_2 data. To that end, an independent component analysis was conducted on R_2 to extract four characteristic spectra that were filtered afterwards while the associated spectral maps were synthetically produced. These data were subsequently mixed to get the spectrum-image referred to as S (see Appendix B for generation details).

To mimic realistic experimental setups, the three noise-free images $\bar{\mathsf{R}}_1$, $\bar{\mathsf{R}}_2$ and S were subsequently corrupted by an additive white Gaussian noise with signal-to-noise ratio adjusted in agreement with the acquisition process. Finally these pseudo-real images, referred to as R_1^* and R_2^*, and the noisy synthetic image S^*, were uniformly randomly spatially subsampled with a 20% ratio to provide the measurement matrix $\mathbf{Y}$. Note that results obtained for other sampling ratios are reported in the supporting document [49].

4.2. Methods

As discussed in Section 3.2, for all algorithms, a PCA is first conducted on the observed image, keeping the T most relevant principal components. The reconstruction algorithms are then run in this lower-dimensional subspace. The inverse transformation will be applied afterwards to get the reconstructed image $\hat{\mathbf{X}}$. Compared methods are NN, 3S, CLS, ITKrMM, wKSVD and BPFA. In particular, NN is the only one to be applied band-by-band. For all methods, the algorithmic parameters have been adjusted to reach the best performances for each method. In particular, DL-based methods consider 3D patches of size $M \times M \times T$ with $M = 25$ for ITKrMM and wKSVD and $M = 41$ for BPFA.

The ITKrMM and wKSVD implementations used in these experiments are the Matlab codes provided by Prof. K. Schnass[2]. The implementation of BPFA is the Matlab code provided by Dr. Z. Xing[3]. The other methods have been implemented by the authors of this paper and are available in a Python library called `inpystem`[4]. The codes to reproduce the experiments described in this paper are also available online[5].

4.3. Metrics

To evaluate the reconstruction quality, several quantitative measures will be used to compare the ground truth $\mathbf{X}$

and reconstructed $\hat{\mathbf{X}}$ images. First, the normalized mean-squared error (NMSE) is chosen as an error measure and is computed according to

$$\text{NMSE}(\hat{\mathbf{X}}, \mathbf{X}) = \frac{||\hat{\mathbf{X}} - \mathbf{X}||_{\text{F}}^2}{||\mathbf{X}||_{\text{F}}^2}. \tag{4}$$

The smaller NMSE, the better. Then, this error measure is turned out as a performance measure by considering its negative-logarithm, defining the signal-to-noise ratio (SNR)

$$\text{SNR}(\hat{\mathbf{X}}, \mathbf{X}) = -10 \log_{10}\left(\text{NMSE}(\hat{\mathbf{X}}, \mathbf{X})\right). \tag{5}$$

The higher SNR, the better. Additionally, we also consider the average spectral angle distance (aSAD) defined as [54, 55]

$$\text{aSAD}(\hat{\mathbf{X}}, \mathbf{X}) = \frac{1}{P}\sum_{j=1}^{P} \text{acos}\left(\frac{\langle \hat{\mathbf{X}}_j, \mathbf{X}_j \rangle}{||\hat{\mathbf{X}}_j||_2 \times ||\mathbf{X}_j||_2}\right), \tag{6}$$

where P is the number of pixels. The aSAD, which is a measure of spectral distorsion, is independent of scaling and should be close to zero. Finally, the structural similarity index (SSIM) averaged over all the bands is considered as a criterion to assess spatial reconstruction [56, 57]. A value close to 1 (resp. to 0) indicates that the spatial structures are similar (resp. different). The closer to 1, the better.

5. Results

5.1. Appropriate basis for sparsity

As explained in the previous section, the data sparsity level highly depends on the image and on the basis chosen for its representation. In the case of atomic-scale images, previous works considered wavelets [34], Fourier [18] or DCT [19, 4] bases. However, no systematic comparison of these transforms has been conducted in these works. This section proposes to fill this gap. Intuitively, Fourier transform or DCT are expected to provide best results as they are known to lead to highly sparse representations when applied to periodic structures. To confirm this intuition, we propose to monitor the magnitude of the data representation coefficients which are expected to decrease as fast as possible when sorted in decreasing order. A significant decreasing would mean that less coefficients are required to accurately represent the data. Thus, the reconstruction error is evaluated for each image when represented with the following transforms: Fourier basis, DCT basis, Daubechies and symlet wavelets (with 3, 10 and 20 vanishing moments). The reconstruction error is evaluated as a function of ratio $r \in (0,1)$ of nonzero representation coefficients, i.e., keeping only the $r\%$ representation

coefficients of highest magnitude. The inverse basis transformation provides an estimate $\hat{\mathbf{X}}$ of $\mathbf{X}$ which is degraded as r decreases. The reconstruction error correspond to the normalized mean square error between $\mathbf{X}$ and $\hat{\mathbf{X}}$. To summarize, the most suitable transform for a given image should be the one whose reconstruction error drops fastest with r as this shows that less coefficients are necessary to describe the image.

This procedure has been applied to R_1, R_2 and S after a PCA. The results are depicted in Fig. 2. They show that DCT is always better than Fourier as the measurements are real-valued data, which introduces a redundancy of factor 2 for each axis in the 2D Fourier space whereas DCT does not suffer from this redundancy. Moreover, our intuition is confirmed since DCT and Fourier perform better for the three images compared to all the wavelet transforms. Finally, in the case of atomic-scale images considered in this work, these experiments show that DCT offers the sparsest transform and it will be chosen as the orthogonal band-by-band transform Ψ in (2) of the proposed CLS method.

5.2. Results on synthetic and semi-real images

In this section, the reconstruction methods were applied to the two semi-real images R_1^* and R_2^* and to the synthetic image S^*.

The metric values as well as the execution times are reported in Table 3. These results show that there is a clear performance gap, both in terms of quality and time, between NN and DL-based techniques. This is particularly true for BPFA which exhibits a prohibitive computation time, yet often leading to the best image reconstruction. The LS-regularized methods 3S and CLS seem to fill this gap, both in terms of accuracy and speed. CLS performs particularly well as it gives a SNR close to the best methods with a very small computational time. 3S gives lower performance results (even if better that NN) since its regularization is not appropriate to accurately describe the periodic structure of atomic-scale images. Note also that BPFA gives globally the best SNR and aSAD, except for R_1^* for which CLS gives a better aSAD, while NN gives the worst aSAD.

The reconstruction of a non-sampled pixel located in Fig. 3 is also depicted in Fig. 4. In this figure, reference data refers to the noise-free image $\bar{\mathsf{R}}_2$. Equivalent plots for a sampled pixel are omitted here since they do not bring any meaningful insight: the reconstructed spectra are close to be distinguished. These plots show that the NN-reconstructed spectrum is significantly shifted with respect to the reference while BPFA and CLS are close to the reference spectrum. Error maps are reported in the supporting document [49].
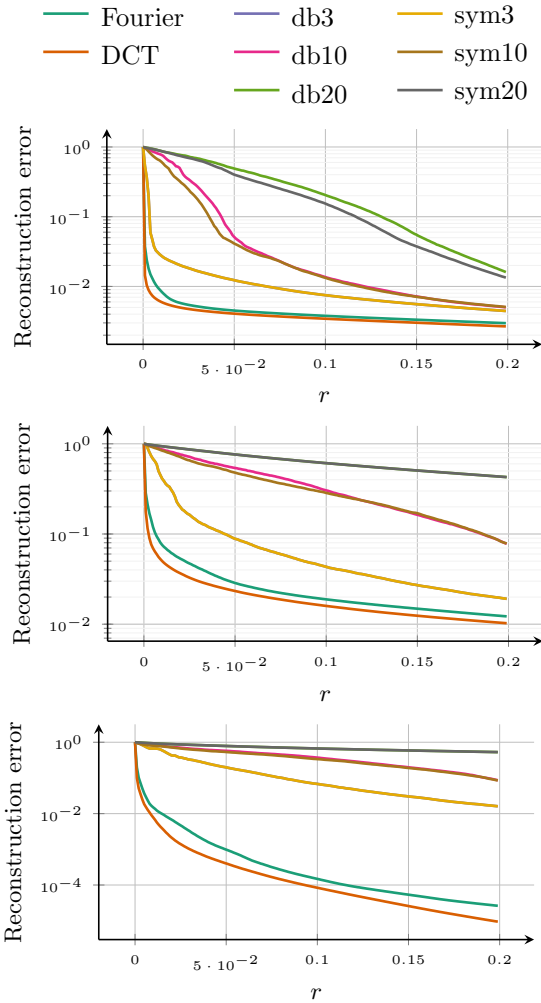
8

Figure 2: Reconstruction error in term of NMSE for several bases when representing R_1 (top), R_2 (middle) and S (bottom). The faster the curve decreases, the better as it means the image needs less representation coefficients to be accurately represented. The DCT basis gives the best results for all images.

(a) R_1^*

Method	SNR	aSAD $(100\times)$	SSIM	Time(s)
NN	30.81	1.384	0.680	**0.423**
3S	32.30	1.080	0.643	15.2
CLS	36.18	**1.061**	0.912	18.4
ITKrMM	36.52	1.097	0.923	6.85e+04
wKSVD	36.93	1.091	0.931	7.97e+04
BPFA	**37.02**	1.089	**0.933**	1.36e+05

(b) R_2^*

Method	SNR	aSAD $(100\times)$	SSIM	Time(s)
NN	28.71	1.815	0.635	**7.94e-02**
3S	30.17	1.496	0.621	3.38
CLS	33.15	1.233	0.790	3.09
ITKrMM	33.67	1.253	0.819	9.55e+03
wKSVD	34.52	1.163	0.841	2.59e+04
BPFA	**35.01**	**1.106**	**0.852**	6.18e+04

(c) S^*

| Method	SNR	aSAD $(100\times)$	SSIM	Time(s)
NN	21.32	1.462	0.735	**6.82e-02**
3S	22.12	1.174	0.710	3.33
CLS	42.14	0.224	0.997	1.48
ITKrMM	44.16	0.338	0.998	9.61e+03
wKSVD	45.59	0.277	0.999	1.57e+04
BPFA	**52.70**	**0.150**	**1.000**	4.06e+04

Table 3: Reconstruction performance in terms of SNR, aSAD and SSIM for semi-real R_1^* and R_2^* and synthetic S^* images. The execution time is also given to be considered jointly with accuracy. There is a clear performance gap in terms of quality and execution time between NN and the DL-based methods . LS-regularized methods fill this gap, especially CLS which performs well in comparison to 3S.

(a) Band #2 of $\bar{R}_2$
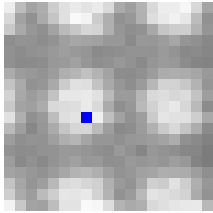
(b) Band #2 of $\bar{R}_2$ (zoom)

(c) Mask (zoom)

Figure 3: Location of the sampled spectrum represented in Figure 4. Image 3(a) shows the 2nd principal component of the semi-real image $\bar{R}_2$ and locates the sampled pixel (blue dot) whose spectrum is considered in Fig. 4. Fig. 3(b) and 3(c) shows zooms on the image region-of-interest and on the sampling mask. The sampling mask white (resp. black) pixels stand for sampled (resp. non-sampled) pixels.

As a consequence, CLS appears as a relevant trade-off between accuracy and complexity since it gives good reconstruction with small computational time. This method could be interesting as an experimental tool whereas BPFA could be used as a post-processing refinement method. Combining both methods will be discussed in Section 5.4.

5.3. Results on a real image

Some illustrative results are also provided for R_2. More precisely, the real spectrum-image R_2 is spatially subsampled with a ratio of 20% and then reconstructed as in the previous subsection.

Visual representations of the reconstructed spectrum-images around some interesting edges are provided in Fig. 5 and the reconstruction of a non-sampled spectrum is shown in Fig. 6. In these figures, note that reference data refers to the real, possibly noisy, image R_2.

Similarly to the previous findings, DL-based methods visually give excellent reconstructed maps when compared to other methods. Results regarding reconstructed spectra show that the NN-reconstructed spectrum for the non-sampled pixel exhibits a significant shift while other algorithms provide smaller biased results.

(a) Non-sampled pixel

(b) Non-sampled pixel (zoom)

Figure 4: Spectra reconstruction results for R_2^* for a non-sampled pixels located in Fig. 3. Reference spectra corresponds to the noise-free image $\bar{R}_2$. The zoom represents the region-of-interest highlighted as shaded a blue area. NN and 3S spectra are significantly shifted compared to the reference whereas the spectra recovered by CLS and DL-based methods are close to the reference spectrum.
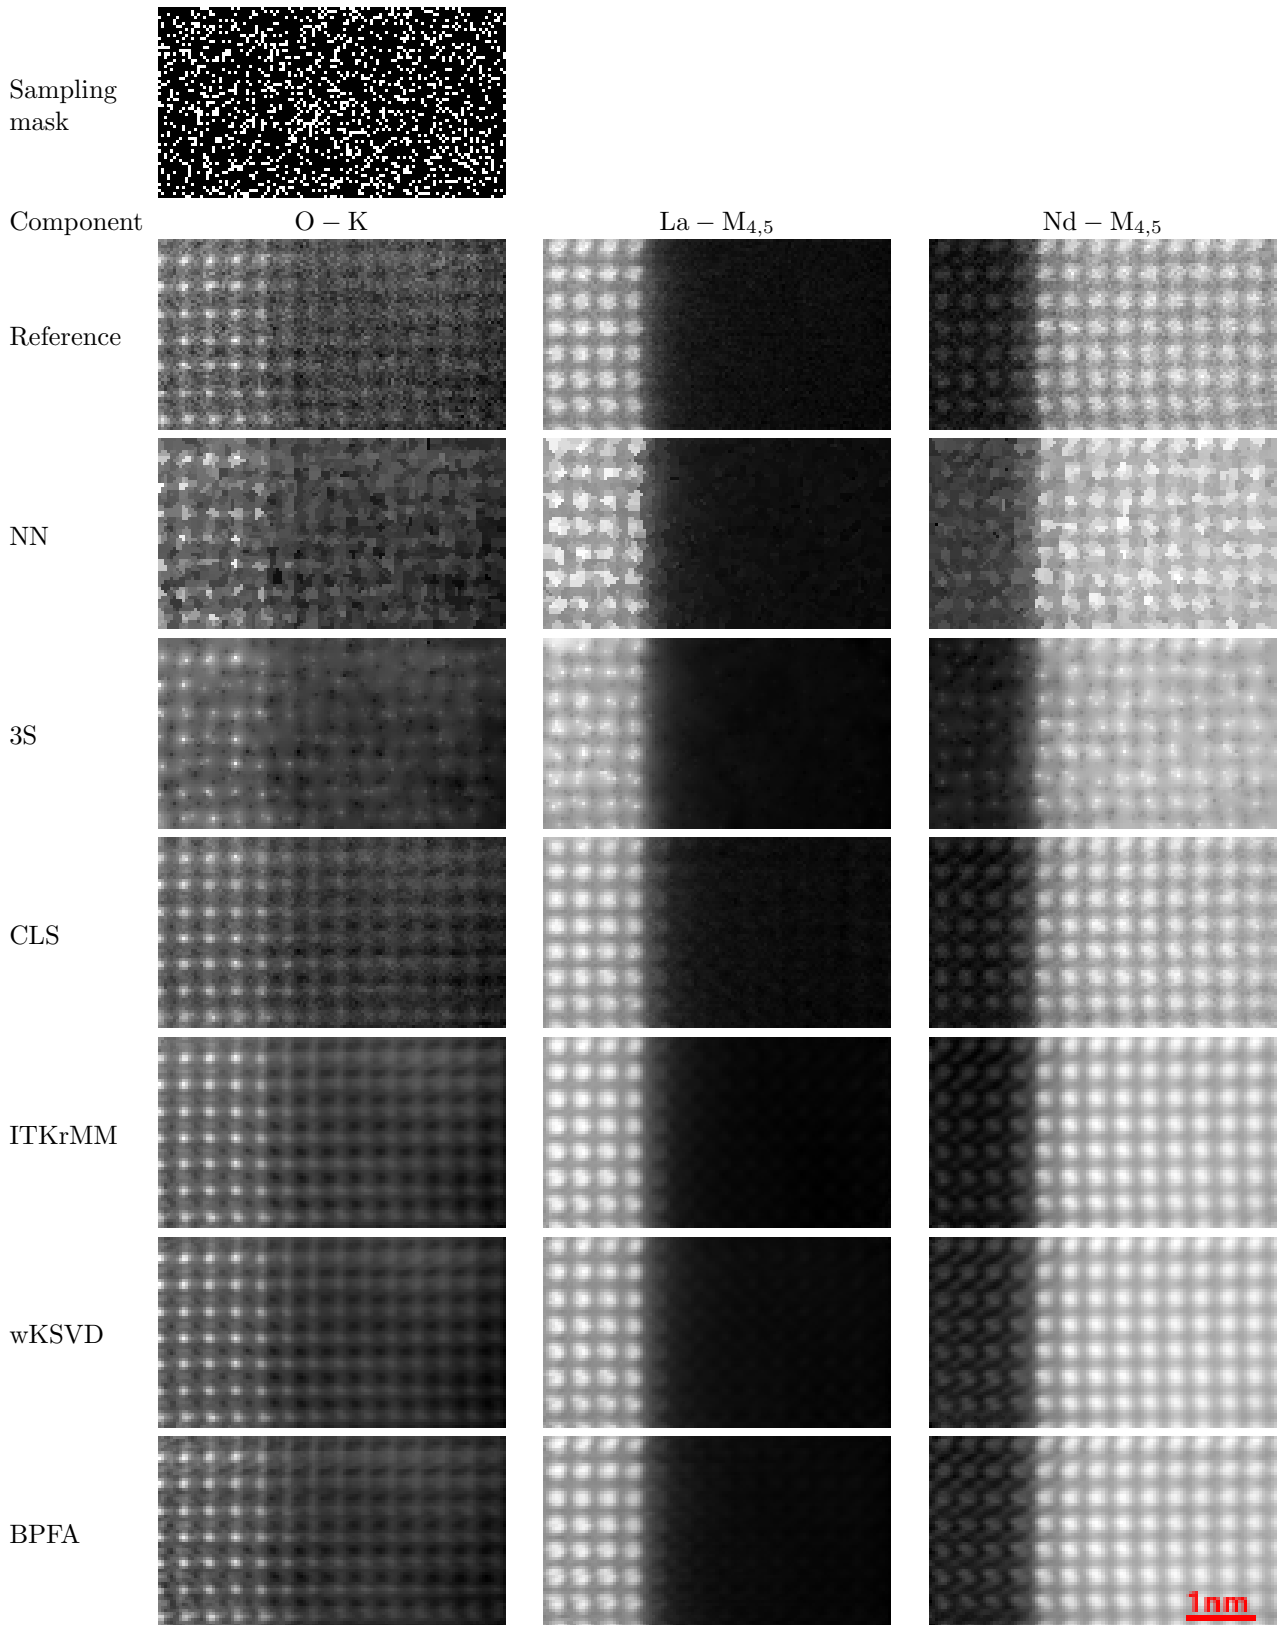
10

Figure 5: Reconstruction results for R_2. The images show the sum of 5 bands around 3 particular edges (O − K, La − $M_{4,5}$ and Nd − $M_{4,5}$). The reference corresponds to the real, possibly noisy, image R_2. The sampling mask is also provided in the first row where white (resp. black) pixels stand for sampled (resp. non-sampled) pixels. These results confirm the performance gap between NN whose images are not smooth enough and DL-based methods which are close to the reference with an additional denoising effect. CLS performs clearly better than NN and 3S and its results are close to DL-based methods.

11
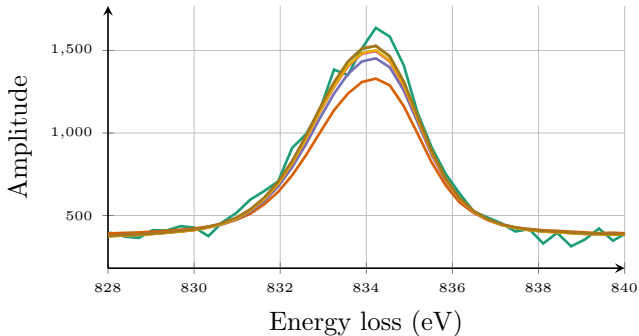
(a) Non-sampled pixel

(b) Non-sampled pixel (zoom)

Figure 6: Spectra reconstruction results for R_2 for a non-sampled pixel with the same locations as in Fig. 3. Reference spectra corresponds to the real image R_2. The zoom represent the region-of-interest highlighted as a shaded blue area. As for the synthetic results, the NN spectrum is significantly shifted compared to the reference spectrum while the results of CLS and DL-based methods (especially BPFA) are close to the reference. Again, CLS appears to be a relevant trade-off.

5.4. CLS as an initialization

Previous subsections compared the reconstruction method performances on synthetic, semi-real and real images. They illustrated the interest of CLS as a fast and efficient method for reconstructing subsampled atomic-scale spectrum-images. Besides, DL-based reconstruction methods are shown to provide higher performances at the price of a higher computational cost. However, coupled with CLS, these methods may be used as post-experiment refinement. In this section, we will confirm this interest based on an additional experiment conducted on the semi-real image R_2^*. Similarly to the previous paragraph, this image has been subsampled and then reconstructed with wKSVD while considering two distinct initializations. The first one is purely random, as initially implemented by the algorithm. The second kind of initialization consists in initializing wKSVD with a dictionary and a corresponding sparse code derived from the CLS-based reconstructed image. Then, the reconstruction quality is monitored as a function of the iterations of the wKSVD algorithm by plotting the reconstruction SNR. Results are depicted in Fig. 7 for both initializations.

The results show that CLS-based initialization leads to less iterations to get a target SNR compared to random initialization. It confirms the interest of CLS to be used as an initialization to accelerate more elaborated yet more computationally intensive techniques. Note that this strategy can be adopted for gradient-descent-based algorithms such as wKSVD or ITKrMM, but is not suited for Markov chain Monte-Carlo algorithms such as BPFA.


Figure 7: SNR as a function of wKSVD iterations for random and CLS initializations. Colored filling corresponds to standard deviation interval computed from 10 runs. The results show that CLS-based initialization needs less iterations to get a target SNR compared to random initialization. This promotes CLS as an initialization for DL-based reconstruction methods.

6. Conclusion

In this paper, we introduced a new reconstruction method referred to as CLS for STEM-EELS imaging that

is both fast and accurate. Experiments on synthetic and real data showed it was much faster than dictionary learning-based reconstruction methods and more accurate than NN methods. The combination of these two advantages is highly interesting to envisage its practical implementation into an online experimental setup.

Moreover, time-consuming but more accurate dictionary learning-based methods were also used as a post-processing of the results provided by CLS. These methods performed reconstruction as a 3D task and gave excellent reconstruction quality. However, the price to pay was a high computational time, which could be reduced by using the reconstructed spectrum-image recovered by CLS as an initialization.

Conducting online microscopy acquisition and reconstruction is an active research task as it speeds up acquisition procedures and identification of components. Performing reconstruction with CLS as an online algorithm and coupling it with adaptive sampling is an interesting perspective towards dynamic STEM-EELS imaging. Such an acquisition protocol would allow sample drifts to be handled, which was not the case in the experiments reported in this paper.

Appendix

A. CLS implementation

This appendix provides some details regarding the implementation of CLS introduced in Section 3.2. Let $\mathbf{X} \in \mathbb{R}^{B \times P}$ denote the unknown spectrum-image to be reconstructed where B is the number of bands and P is the number of pixels. The subsampling acquisition process writes

$$\mathbf{Y} = \mathbf{X}\Phi + \mathbf{E} \qquad (7)$$

where $\mathbf{Y} \in \mathbb{R}^{B \times N}$ is the observation matrix, N is the number of sampled pixels, $\Phi \in \mathbb{R}^{P \times N}$ is the sampling operator and $\mathbf{E}$ is a residual term associated with error modeling and measurement noise. The elements of $\mathbf{E}$ are assumed to be independent and identically distributed according to

a zero-mean Gaussian noise. Note that the sampling operator Φ is the concatenation of N columns extracted from the identity $P \times P$-identity matrix.

Reconstructing the full spectrum-image $\mathbf{X}$ from $\mathbf{Y}$ can be formulated as the following optimization problem

$$\hat{\mathbf{X}} = \arg\min_{\mathbf{X}} \underbrace{\frac{1}{2}||\mathbf{Y} - \mathbf{X}\Phi||_F^2}_{f(\mathbf{X})} + \underbrace{\lambda||\mathbf{X}\Psi||_{2,1}}_{g(\mathbf{X})} \qquad (8)$$

where Ψ is the band-by-band DCT transform operator. This optimization problem can be easily solved with iterative algorithms such as the fast iterative shrinkage-thresholding algorithm (FISTA) [58]. This algorithm splits the function to be minimized into two terms

- f which is a smooth convex function of the type $\mathcal{C}^{1,1}$, i.e., continuously differentiable with Lipschitz continuous gradient $L(f)$,

- g which is a continuous convex and possibly nonsmooth function.

Based on this decomposition, the main steps of the algorithm are described in Algo. 1, which requires to evaluate three elements, namely, the gradient ∇f, an upper bound of the Lipschitz constant $L > L(f)$ and the proximal operator $\mathrm{prox}_{g/L}(\mathbf{X}) = \arg\min_{\mathbf{U}}\{g(\mathbf{U})/L + \frac{1}{2}||\mathbf{U} - \mathbf{X}||^2\}$. The gradient function ∇f can be easily derived as

$$\nabla f(\mathbf{X}) = (\mathbf{X}\Phi - \mathbf{Y})\Phi^T. \qquad (9)$$

The ∇f function Lipschitz constant can be computed following

$$||\nabla f(\mathbf{X}_1) - \nabla f(\mathbf{X}_2)|| = ||\mathbf{X}_1 - \mathbf{X}_2|| \cdot ||\Phi\Phi^T|| \qquad (10)$$

with $L(f) = ||\Phi\Phi^T|| = 1$. Last, the proximal operator is column-separable and is computed for each column $\mathbf{X}_j$ of $\mathbf{X}$ (with $j = 1, \ldots, P$) [59]

$$[\mathrm{prox}_{g/L}(\mathbf{X})]_j = \begin{cases} 0 & \text{if } ||\mathbf{X}_j||_2 < \frac{\lambda}{L} \\ \left(1 - \frac{\lambda/L}{||\mathbf{X}_j||_2}\right)\mathbf{X}_j & \text{otherwise} \end{cases} \qquad (11)$$

As mentioned in Section 3.2, the reconstruction method can be applied to the full measurement matrix $\mathbf{Y}$ or to its representation $\tilde{\mathbf{Y}}$ in a lower-dimensional subspace identified by PCA. The threshold T to be set can be chosen as described in the following section.

B. Generation of the synthetic and semi-real spectrum-images

Generating the synthetic spectrum-image S. To generate the synthetic spectrum-image S, we propose to follow

Algorithm 1: FISTA with constant step size [58]
1 Input : $L > L_f$ an upper bound of L_f and the observation $\mathbf{Y}$
2 Initialisation : Set $\mathbf{Z}^{(1)} = \mathbf{X}^{(0)} \in \mathbb{R}^p$, $\theta^{(1)} = 1$, $i = 1$ **while** *stopping rule not satisfied* **do**
3 $\quad \mathbf{X}^{(i)} = \text{prox}_{g/L}\left(\mathbf{Z}^{(i)} - \frac{1}{L}(\mathbf{Z}^{(i)}\Phi - \mathbf{Y})\Phi^T\right)$
4 $\quad \theta^{(i+1)} = \frac{1}{2}\left(1 + \sqrt{1 + 4(\theta^{(i)})^2}\right)$
5 $\quad \mathbf{Z}^{(i+1)} = \mathbf{X}^{(i)} + \left(\frac{\theta^{(i)}-1}{\theta^{(i+1)}}\right)\left(\mathbf{X}^{(i)} - \mathbf{X}^{(i-1)}\right)$
6 $\quad i \leftarrow i + 1$

the strategy adopted in [11]. It consists in linearly mixing $N_c = 4$ components according to realistic proportion maps. The spectral components were obtained by smoothing N_c signatures identified by an independent component analysis conducted on the real-spectrum image R_2. They are gathered in the $B \times N_c$ matrix $\mathbf{M}$ and depicted in Fig. 8. The proportion maps were generated synthetically by superimposing periodic structures onto a smoothly varying background similar to the content of the real spectrum-image R_2. These maps, gathered in the $N_c \times P$ matrix $\mathbf{A}$, are represented in Fig. 9. Finally, the synthetic image S can be generated as $\mathbf{MA}$.
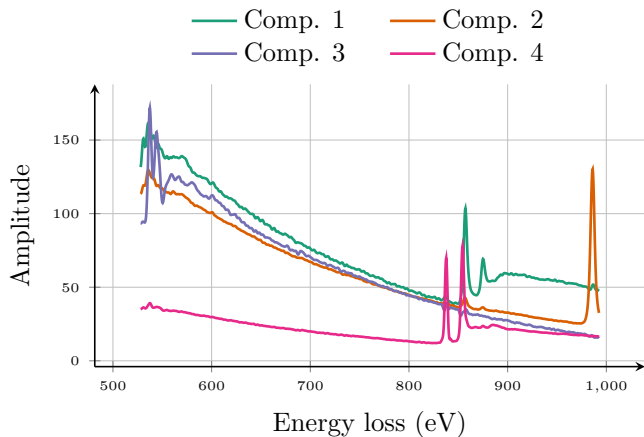

Figure 8: Spectral components used for generating the synthetic image S.

Generating the semi-real spectrum-images R_1^* *and* R_2^*. The reference noise-free spectrum-images $\bar{\mathsf{R}}_1^*$ and $\bar{\mathsf{R}}_2^*$ used in Section 5.2 are generated by keeping the first T principal components identified by PCA conducted on R_1 and R_2, respectively. Generally, determining the PCA threshold T is not a trivial. In this work, we propose to choose T such as the $B-T$ remaining components do not contain spatial information. To quantify the presence or absence of spatial information in a principal components, we monitor their whiteness based on the metrics proposed in [60, Chap. 3]


(a) Comp. 1   (b) Comp. 2

(c) Comp. 3   (d) Comp. 4

Figure 9: Proportion maps associated with the spectral components represented in Fig. 8 used for generating the synthetic image S.

as

$$||r||_2^* = \sqrt{\sum_{\tau \neq 0} r(\tau)^2} \qquad (12)$$

where $r(\tau)$ is the 2D-autocorrelation function. The higher value, the more information contained by the image. To illustrate, this criterion is represented in Figure 10 as a function of the index of the principal components for the two spectrum-images R_1 and R_2. The PCA threshold T is chosen as the maximal index sufficient to get to a stationary curve behavior. These values are reported in Table 2. The semi-real spectrum-images R_1^* and R_2^* are finally obtained by corrupting the reference noise-free spectrum-images $\bar{\mathsf{R}}_1$ and $\bar{\mathsf{R}}_2$ with additive white Gaussian noises whose variances have been adjusted to reach SNR values in agreement with those of the corresponding real spectrum-images R_1 and R_2, respectively.
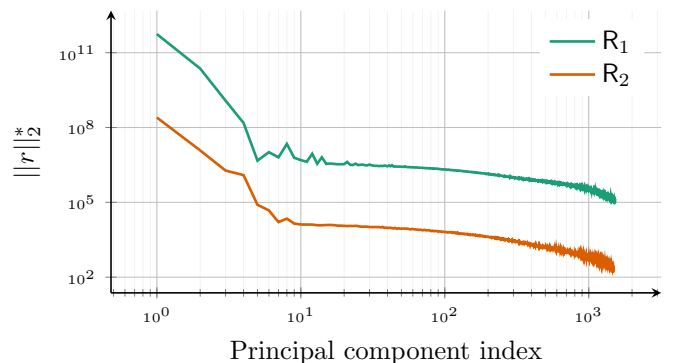

Figure 10: Whiteness criterion $||r||_2^*$ as a function of the principal component index for the real image R_1 and R_2. The first powerful principal components exhibit more spatial content than the last ones. The PCA threshold is chosen as the maximal index sufficient to reach a stationary curve behavior.

References

[1] E. Monier, T. Oberlin, N. Brun, N. Dobigeon, Reconstruction of partially sampled STEM-EELS images with atomic resolution, in: Proc. Workshop on Signal Processing with Adaptative Sparse Structured Representations (SPARS), Toulouse, France, 2019.

[2] R. Egerton, P. Li, M. Malac, Radiation damage in the TEM and SEM, Micron 35 (6) (2004) 399 – 409.

[3] A. Zobelli, S. Y. Woo, A. Tararan, L. H. Tizei, N. Brun, X. Li, O. Stéphan, M. Kociak, M. Tencé, Spatial and spectral dynamics in stem hyperspectral imaging using random scan patterns, Ultramicroscopy (2019) 112912.

[4] H. S. Anderson, J. Ilic-Helms, B. Rohrer, J. Wheeler, K. Larson, Sparse imaging for fast electron microscopy, in: Computational Imaging XI, Vol. 8657, Int. Soc. for Opt. and Photonics, 2013, p. 86570C.

[5] P. Binev, W. Dahmen, R. DeVore, P. Lamby, D. Savu, R. Sharpley, Compressed sensing and electron microscopy, in: Modeling Nanoscale Imaging in Electron Microscopy, Springer, 2012, pp. 73–126.

[6] M. Jacob, L. El Gueddari, G. Navarro, M.-C. Cyrille, P. Bayle-Guillemaud, P. Ciuciu, Z. Saghi, Statistical machine learning and compressed sensing approaches for analytical electron tomography – Application to phase change materials, Microsc. Microanal. 25 (S2) (2019) 156–157.

[7] M. Jacob, T. Sanders, N. Bernier, A. Grenier, R. B. Pinheiro, F. Mazen, P. Bayle-Guillemaud, Z. Saghi, Multivariate analysis and compressed sensing methods for spectroscopic electron tomography of semiconductor devices, Microsc. Microanal. 24 (S1) (2018) 500–501.

[8] A. Béché, B. Goris, B. Freitag, J. Verbeeck, Compressed sensing for beam sensitive materials imaging in scanning transmission electron microscopy, in: Proc. Eur. Microsc. Congr., Wiley, 2016, pp. 340–341.

[9] A. Stevens, H. Yang, L. Carin, I. Arslan, N. D. Browning, The potential for Bayesian compressive sensing to significantly reduce electron dose in high-resolution STEM images, Microscopy 63 (1) (2013) 41–51.

[10] P. Trampert, F. Bourghorbel, P. Potocek, M. Peemen, C. Schlinkmann, T. Dahmen, P. Slusallek, How should a fixed budget of dwell time be spent in scanning electron microscopy to optimize image quality?, Ultramicroscopy 191 (2018) 11 – 17.

[11] E. Monier, T. Oberlin, N. Brun, M. Tencé, M. de Frutos, N. Dobigeon, Reconstruction of partially sampled multiband images – Application to STEM-EELS imaging, IEEE Trans. Comput. Imag. 4 (4) (2018) 585 – 598.

[12] Z. Xing, M. Zhou, A. Castrodad, G. Sapiro, L. Carin, Dictionary learning for noisy and incomplete hyperspectral images, SIAM J. Imag. Sci. 5 (1) (2012) 33–56.

[13] G. M. D. P. Godaliyadda, D. H. Ye, M. D. Uchic, M. A. Groeber, G. T. Buzzard, C. A. Bouman, A framework for dynamic image sampling based on supervised learning, IEEE Trans. Comput. Imag. 4 (1) (2018) 1–16.

[14] Y. Zhang, G. M. D. Godaliyadda, N. Ferrier, E. B. Gulsoy, C. A. Bouman, C. Phatak, Reduced electron exposure for energy-dispersive spectroscopy using dynamic sampling, Ultramicroscopy 184 (2018) 90 – 97.

[15] K. A. Hujsak, E. W. Roth, W. Kellogg, Y. Li, V. P. Dravid, High speed/low dose analytical electron microscopy with dynamic sampling, Micron 108 (2018) 31–40.

[16] R. Sibson, A brief description of natural neighbor interpolation, in: Interpolating Multivariate Data, V. Barnett, 1981, Ch. 2, pp. 21–36.

[17] G. Han, B. Lin, Optimal sampling and reconstruction of undersampled atomic force microscope images using compressive sensing, Ultramicroscopy 189 (2018) 85 – 94.

[18] A. Stevens, L. Luzi, H. Yang, L. Kovarik, B. L. Mehdi, A. Liyu, M. E. Gehm, N. D. Browning, A sub-sampled approach to extremely low-dose STEM, Appl. Phys. Lett. 112 (4) (2018) 043104.

[19] A. Béché, B. Goris, B. Freitag, J. Verbeeck, Development of a fast electromagnetic beam blanker for compressed sensing in scanning transmission electron microscopy, Appl. Phys. Lett. 108 (9) (2016) 093103.

[20] E. van den Berg, M. P. Friedlander, Probing the Pareto frontier for basis pursuit solutions, SIAM J. Sci. Comput. 31 (2) (2008) 890–912.

[21] X. Li, O. Dyck, S. V. Kalinin, S. Jesse, Compressed sensing of scanning transmission electron microscopy (STEM) on non-rectangular scans, arXiv preprint arXiv:1805.04957 (2018).

[22] A. Criminisi, P. Perez, K. Toyama, Region filling and object removal by exemplar-based image inpainting, IEEE Trans. Image Process. 13 (9) (2004) 1200–1212.

[23] D. Mucke-Herzberg, P. Abellan, M. C. Sarahan, I. S. Godfrey, Z. Saghi, R. K. Leary, A. Stevens, J. Ma, G. Kutyniok, F. Azough, et al., Practical implementation of compressive sensing for high resolution STEM, Microsc. Microanal. 22 (S3) (2016) 558–559.

[24] L. Kovarik, A. Stevens, A. Liyu, N. D. Browning, Implementing an accurate and rapid sparse sampling approach for low-dose atomic resolution STEM imaging, Appl. Phys. Lett. 109 (16) (2016) 164102.

[25] A. Stevens, Y. Pu, Y. Sun, G. Spell, L. Carin, Tensor-dictionary learning with deep Kruskal-factor analysis, arXiv preprint arXiv:1612.02842 (2017).

[26] A. Stevens, L. Kovarik, H. Yang, Y. Pu, L. Carin, N. D. Browning, Compressive STEM-EELS, Microsc. Microanal. 22 (S3) (2016) 560–561.

[27] D. Zoran, Y. Weiss, From learning models of natural image patches to whole image restoration, in: Proc. IEEE Int. Conf. Computer Vision (ICCV), 2011, pp. 479–486.

[28] S. Parameswaran, C. Deledalle, L. Denis, T. Q. Nguyen, Accelerating GMM-Based patch priors for image restoration: Three ingredients for a 100× speed-up, IEEE Trans. Image Process. 28 (2) (2019) 687–698.

[29] J. Mairal, M. Elad, G. Sapiro, Sparse representation for color image restoration, IEEE Trans. Image Process. 17 (1) (2008) 53–69.

[30] V. Naumova, K. Schnass, Fast dictionary learning from incomplete data, EURASIP J. Adv. Signal Process. 2018 (1) (2018) 12.

[31] V. Naumova, K. Schnass, Dictionary learning from incomplete data for efficient image restoration, in: Proc. European Signal Process. Conf. (EUSIPCO), 2017, pp. 1425–1429.

[32] E. Monier, T. Oberlin, N. Brun, M. D. Frutos, M. Tencé, N. Dobigeon, Reconstruction of partially sampled EELS images, in: Proc. IEEE GRSS Workshop Hyperspectral Image SIgnal Process.: Evolution in Remote Sens. (WHISPERS), 2018, pp. 1–5.

[33] X. Sang, A. R. Lupini, R. R. Unocic, M. Chi, A. Y. Borisevich, S. V. Kalinin, E. Endeve, R. K. Archibald, S. Jesse, Dynamic scan control in STEM: Spiral scans, Adv. Struct. and Chem. Imag. 2 (1) (2017) 6.

[34] X. Li, O. Dyck, S. V. Kalinin, S. Jesse, Compressive sensing on diverse STEM scans: real-time feedback, low-dose and dynamic range, arXiv preprint arXiv:1805.04957 (2018).

[35] T. Dahmen, M. Engstler, C. Pauly, P. Trampert, N. De Jonge, F. Mücklich, P. Slusallek, Feature adaptive sampling for scanning electron microscopy, Scientific Rep. 6 (2016) 25350.

[36] T. Dahmen, P. Trampert, An adaptive sparse sampling scheme for scanning electron microscopy using Delauney triangulation, Microsc. Microanal. 25 (S2) (2019) 154–155.

[37] S. Hawe, M. Kleinsteuber, K. Diepold, Analysis operator learning and its application to image reconstruction, IEEE Trans. Image Process. 22 (6) (2013) 2138–2150.

[38] P. Trampert, S. Schlabach, T. Dahmen, P. Slusallek, Exemplar-based inpainting based on dictionary learning for sparse scanning electron microscopy, Microsc. Microanal. 24 (S1) (2018) 700–701.

[39] M. Kowalski, Sparse regression using mixed norms, Appl. Comput. Harmonic Anal. 27 (3) (2009) 303 – 324.

[40] Q. Zhang, R. Plemmons, D. Kittle, D. Brady, S. Prasad, Joint

segmentation and reconstruction of hyperspectral data with compressed measurements, Appl. Opt. 50 (22) (2011) 4417–4435.

[41] G. Martín, J. M. Bioucas-Dias, A. Plaza, Hyca: A new technique for hyperspectral compressive sensing, IEEE Trans. Geosci. Remote Sens. 53 (5) (2014) 2819–2831.

[42] L. Zhuang, J. M. Bioucas-Dias, Fast hyperspectral image denoising and inpainting based on low-rank and sparse representations, IEEE J. Sel. Topics Appl. Earth Observations Remote Sens. 11 (3) (2018) 730–742.

[43] E. Wycoff, T.-H. Chan, K. Jia, W.-K. Ma, , Y. Ma, Nonnegative sparse promoting algorithm for high resolution hyperspectral imaging, in: Proc. IEEE Int. Conf. Acoust., Speech and Signal Process. (ICASSP), Vancouver, Canada, 2013.

[44] M. Simões, J. Bioucas Dias, L. Almeida, J. Chanussot, A convex formulation for hyperspectral image superresolution via subspace-based regularization, IEEE Trans. Geosci. Remote Sens. 6 (53) (2015) 3373–3388.

[45] Q. Wei, N. Dobigeon, J.-Y. Tourneret, Fast fusion of multiband images based on solving a Sylvester equation, IEEE Trans. Image Processing 24 (11) (2015) 4109–4121.

[46] N. Dobigeon, S. Moussaoui, M. Coulon, J.-Y. Tourneret, A. O. Hero, Joint Bayesian endmember extraction and linear unmixing for hyperspectral imagery, IEEE Trans. Signal Processing 57 (11) (2009) 4355–4368.

[47] N. Dobigeon, N. Brun, Spectral mixture analysis of EELS spectrum-images, Ultramicroscopy 120 (2012) 25–34.

[48] N. Mevenkamp, B. Berkels, M. Duchamp, Denoising electron-energy loss data using non-local means filters, Microsc. Microanal. 23 (S1) (2017) 106–107.

[49] E. Monier, T. Oberlin, N. Brun, M. Tencé, X. Li, N. Dobigeon, Fast reconstruction of atomic-scale STEM-EELS images from sparse sampling – Complementary results, Tech. rep., University of Toulouse, IRIT/INP-ENSEEIHT, France (Feb. 2020). URL http://dobigeon.perso.enseeiht.fr/papers/Monier_TechReport_2020.pdf

[50] J. Mairal, F. Bach, J. Ponce, G. Sapiro, Online dictionary learning for sparse coding, in: Proc. Int. Conf. Machine Learning (ICML), 2009, pp. 689–696.

[51] Y. C. Pati, R. Rezaiifar, P. S. Krishnaprasad, Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition, in: Proc. IEEE Asilomar Conf. Signals, Systems, Computers, IEEE, 1993, pp. 40–44.

[52] X. Li, A. Gloter, D. Preziosi, Charge, strain and polarization profiling in Ferroelectric/Ferromagnetic epitaxial heterostructures, American Cancer Society, 2016, pp. 885–886.

[53] D. Preziosi, L. Lopez-Mir, X. Li, T. Cornelissen, J. H. Lee, F. Trier, K. Bouzehouane, S. Valencia, A. Gloter, A. Barthélémy, et al., Direct mapping of phase separation across the metal–insulator transition of NdNiO3, Nano Lett. 18 (4) (2018) 2226–2232.

[54] N. Keshava, Distance metrics and band selection in hyperspectral processing with applications to material identification and spectral libraries, IEEE Trans. Geosci. Remote Sens. 42 (7) (2004) 1552–1565.

[55] Y. Sohn, N. S. Rebello, Supervised and unsupervised spectral angle classifiers, Photogrammetric Eng. and Rem. Sens. 68 (12) (2002) 1271–1282.

[56] Z. Wang, A. C. Bovik, Mean squared error: Love it or leave it? a new look at signal fidelity measures, IEEE Trans. Signal Process. 26 (1) (2009) 98–117.

[57] Zhou Wang, A. C. Bovik, H. R. Sheikh, E. P. Simoncelli, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Signal Process. 13 (4) (2004) 600–612.

[58] A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, SIAM J. Imag. Sci. 2 (1) (2009) 183–202.

[59] R. Jenatton, J. Mairal, G. Obozinski, F. Bach, Proximal methods for hierarchical sparse coding, J. Mach. Learning Research 12 (July) (2011) 2297–2334.

[60] P. Riot, Residual whiteness for image denoising, Ph.d. thesis, Télécom ParisTech (Feb. 2018). URL https://pastel.archives-ouvertes.fr/tel-01828883