Pore REconstruction and Segmentation (PORES) method for improved porosity quantification of nanoporous materials

Pore REconstruction and Segmentation (PORES) method for improved porosity quantification of nanoporous materials

G. Van Eyndhoven[a,*], M. Kurttepeli[b], C. J. Van Oers[c], P. Cool[c], S. Bals[b], K.J. Batenburg[a,d,e], J. Sijbers[a]

[a]*iMinds-Vision Lab, University of Antwerp, Universiteitsplein 1, B-2610 Wilrijk, Belgium*
[b]*EMAT, University of Antwerp, Groenenborgerlaan 171, B-2020 Antwerp, Belgium*
[c]*Laboratory of Adsorption and Catalysis, University of Antwerp, Universiteitsplein 1, B-2610 Wilrijk, Belgium*
[d]*Centrum Wiskunde & Informatica, Science Park 123, NL-1090 GB Amsterdam, The Netherlands*
[e]*Mathematical Institute, Universiteit Leiden, Niels Bohrweg 1, NL-2333 CA Leiden, The Netherlands*

Abstract

Electron tomography is currently a versatile tool to investigate the connection between the structure and properties of nanomaterials. However, a quantitative interpretation of electron tomography results is still far from straightforward. Especially accurate quantification of pore-space is hampered by artifacts introduced in all steps of the processing chain, i.e., acquisition, reconstruction, segmentation and quantification. Furthermore, most common approaches require subjective manual user input. In this paper, the PORES algorithm ("*POre REconstruction and Segmentation*") is introduced; it is a tailor-made, integral approach, for the reconstruction, segmentation, and quantification of porous nanomaterials. The PORES processing chain starts by calculating a reconstruction with a nanoporous-specific reconstruction algorithm: the *Simultaneous Update of Pore Pixels by iterative REconstruction and Simple Segmentation* algorithm (SUPPRESS). It classifies the interior region to the pores during reconstruction, while reconstructing the remaining region by reducing the error with respect to the acquired electron microscopy data. The SUPPRESS reconstruction can be directly plugged into the remaining processing chain of the PORES algorithm, resulting in accurate individual pore quantification and full sample pore statistics. The proposed approach was extensively validated on both simulated and experimental data, indicating its ability to generate accurate statistics of nanoporous materials.

Keywords: Electron tomography, reconstruction, segmentation, nanoporous material, pore size distribution

1. Introduction

Although electron tomography provides valuable three-dimensional visualizations of the sample under interest, accurate quantification of pore sizes in nanoporous materials remains a difficult problem, especially if the pores are irregularly shaped.

Quantification of nanoporous materials is important in many applications in the field of sorption/separation or catalysis, in which size selectivity often plays an important role [1, 2]. This makes a reliable and accurate knowledge of the pore size distribution indispensable.

Microporous $(d < 2nm)$ and mesoporous $(2nm < d < 50nm)$ materials are usually characterized by N2-sorption experiments at a temperature of $77K$ [3]. However, the quantification of the pore size distribution based on these measurements is carried out using models that assume a regular pore size, i.e., cylindrical or slit-shaped pores. Unfortunately, no model is available for materials with irregular pores. In this publication, an alternative and reliable approach to determine the pore size in nanoporous materials is proposed.

Transmission electron microscopy (TEM) is an ideal technique to investigate nanoporous materials at a local scale, but conventional TEM is limited to providing two-dimensional (2D) projections of a three-dimensional (3D) microscopy sample [4]. To measure the pore size distribution, a 3D representation of the sample is required, which can be obtained using electron tomography. This technique combines the information of a tilt series of 2D TEM images in a 3D voxel-based reconstruction [5]. The quality of the 3D reconstruction is of critical importance, since it influences further quantification. Computing accurate reconstructions from TEM projection images with classical

*email: geert.vaneyndhoven@uantwerpen.be

analytical algorithms such as weighted back-projection (WBP) [6] or algebraic algorithms like the simultaneous iterative reconstruction technique (SIRT) [7] is a difficult task, mainly because of two issues. First, the limited tilt range of the sample (usually about ±75°) causes elongation of the 3D reconstruction and smearing of the voxel values, often referred to as the "missing wedge" artifact. Secondly, the reconstruction quality also depends on the number of TEM projection images, which is often relatively small to avoid beam damage, especially for sensitive materials.

It has been shown recently that the quality of a 3D reconstruction can be improved by incorporating prior knowledge in the reconstruction process. Assuming that the sample contains just a few a priori known compositions, each occurring in homogeneous regions, the discrete algebraic reconstruction technique (DART) has been able to strongly reduce missing wedge artifacts [8, 9]. The partially discrete algebraic reconstruction technique (PDART) exploits the existence of dense homogeneous particles of which the grey value is known by incorporating this knowledge in the reconstruction algorithm, resulting in more accurate reconstruction quality [10]. Other methods minimize the total variation of the reconstruction, where the sample is assumed to have a sparse gradient, i.e., the number of boundary pixels in the sample is relatively small compared to the total number of pixels [11, 12]. However, the prior knowledge assumptions incorporated in the reconstruction algorithms of the previous examples are not always applicable to nanoporous materials, since the reconstruction may consist of a continuous range of grey values with non-sparsity of the gradient image. In this paper, we propose an approach that exploits a different kind of prior knowledge, which is related uniquely to porous materials: the existence of many local regions of void space.

After the reconstruction step, individual pores can be extracted. To that end, a segmentation step should be applied to separate the pores from the material matrix. Manually or automatically selecting global thresholds can produce satisfactory results if there is a clear separation between the background and the material matrix [13]. However, due to reconstruction artifacts, this separation is not straightforward in practice. In particular for the segmentation of pores that are small compared to the voxel size, this approach is error prone. It can therefore be expected that further analysis of the pores with individual pore statistics such as size, orientation, eccentricity, etc. will be strongly influenced by the results of the two previous steps, i.e., reconstruction and segmentation.

To overcome the limitations discussed above, we present a tailor-made, integral approach, for the reconstruction, segmentation, and quantification of porous nanomaterials: the PORES ("*POre REconstruction and Segmentation*") algorithm. The PORES data processing chain outperforms conventional approaches, since it is optimized for nanoporous structures. The PORES processing chain starts by calculating a porous sample specific reconstruction with the new SUPPRESS ("*Simultaneous Update of Pore Pixels by iterative REconstruction and Simple Segmentation*") algorithm. SUPPRESS reduces artifacts by exploiting prior knowledge about the porous structure of the material, while automatically classifying the interior of the pores. The PORES method continues by applying a watershed algorithm directly to the reconstruction, resulting in accurate segmentation of the pores. This segmentation permits accurate quantification of individual pores, which is employed to generate full sample pore statistics.

The PORES method is described in section 2. In section 3, the method is validated with both simulation and real experiments. The paper is concluded in section 4.

2. Method

This section describes the entire PORES algorithm, which is displayed in the flowchart in Fig. 1. It consists of two parts: the reconstruction algorithm (described in section 2.1 and displayed in the uppermost part of the flowchart in Fig. 1) and the segmentation and quantification (described in section 2.2 and displayed in the bottommost part of the flowchart in Fig. 1).

2.1. Reconstruction

In this section, a novel reconstruction technique is described, which will be referred to as the SUPPRESS algorithm ("*Simultaneous Update of Pore Pixels by iterative REconstruction and Simple Segmentation*"). It exploits a prior that comes naturally for porous materials: the existence of many local regions of void space.

Before elaborating on the proposed reconstruction algorithm, a short introduction to the standard algebraic SIRT algorithm is given, since it is utilized as a subroutine in SUPPRESS. In SIRT, starting from an initial estimate of the sample, the current estimate is refined in each iteration by reducing the distance between the measured TEM projection data and the projection data simulated from the current estimate. It has been proven in [7] that SIRT converges towards a weighted least squares solution of
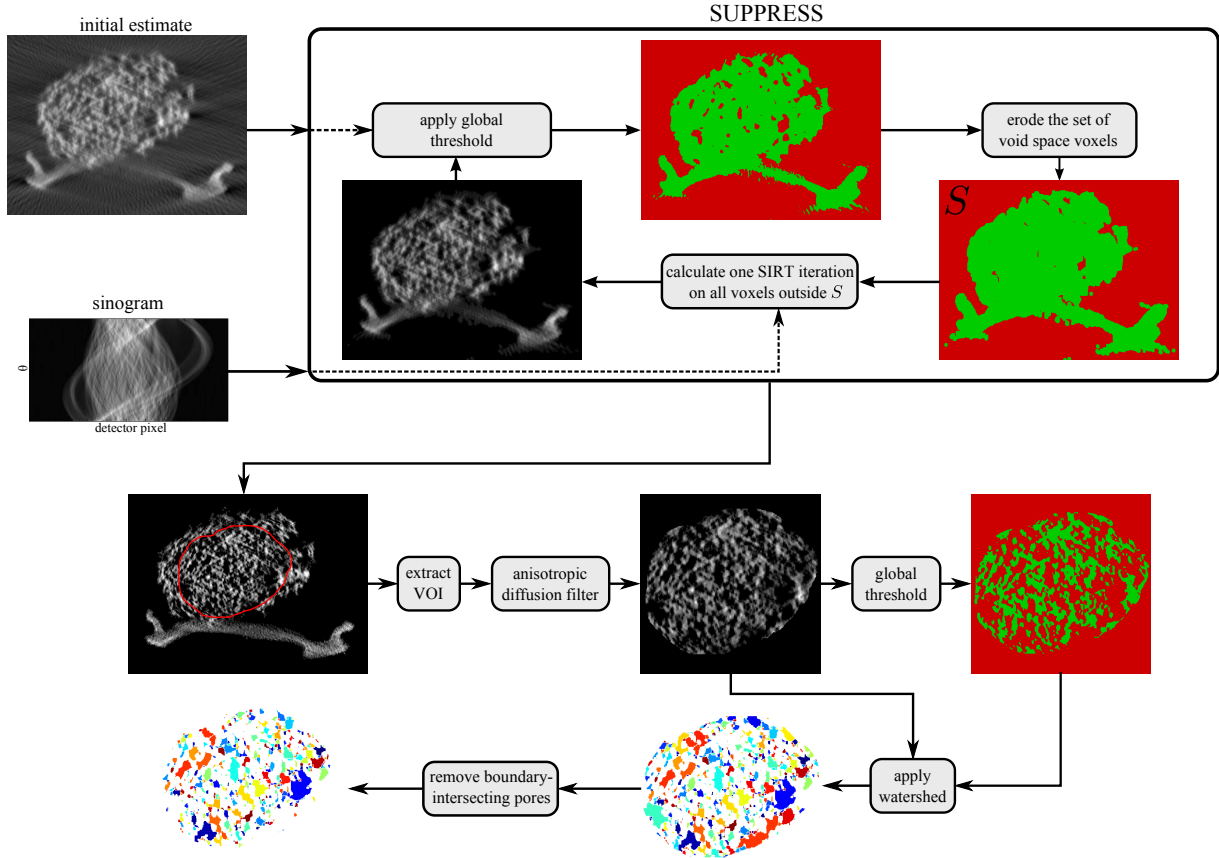
$$Wx = p \quad , \tag{1}$$

Figure 1: Flowchart of the entire PORES algorithm, which consist of the SUPPRESS reconstruction algorithm and the segmentation protocol. Red pixels indicate the estimate for void space, indicated by S in the flowchart. The SIRT update is calculated on the all pixels in the complement of S, i.e., on the green pixels in the illustration.

where $\boldsymbol{x} \in \mathbb{R}^N$ is the discretized object represented on a grid of N voxels, $\boldsymbol{p} \in \mathbb{R}^M$ is the acquired projection data with M the total number of acquired projection values at each angle and each detector pixel and $\boldsymbol{W}$ is the linear operator that models the projection process.

In SUPPRESS, iterative update steps are combined with the prior knowledge that voxels inside void space, i.e., pores and background, should be homogenous and have a lower grey value than the support material. All steps in SUPPRESS are displayed in the uppermost part of the flowchart in Fig. 1. The algorithm starts by generating an initial SIRT reconstruction. Next, a conservative set of void space voxels, S, is estimated with the following two steps. First, the current reconstruction is segmented by a global thresholding operation, i.e., all pixels with a grey value smaller than a specified global threshold τ are selected as possible candidates for the S set. In our approach, τ was chosen as the grey value halfway between the grey value of void space and the

smallest grey value of the support material. Next, this set is eroded, removing voxels at the boundary. This results in the set S, which is a conservative estimate for the void space voxels. The erosion operation reduces the chance that voxels of the material matrix are incorrectly classified as void space voxels (misclassified voxels are found typically on the edge between pore-space and material-space), which is essential for the next step in the algorithm. Based on the assumption that no material is present in the region defined by S, the reconstruction is then continued by applying a SIRT iteration solely to the voxels that belong to the complement of S, while keeping the voxels in S fixed at a grey level of 0 (i.e., no material). This procedure of identifying void space voxels and applying SIRT iterations to the remaining voxels is repeated until a fixed number of iterations is reached.

The key strength of SUPPRESS lies in the fact that the pores (and background) are identified during the re-

3

construction process itself. This results in a more refined estimate for void space at every iteration, which in turn will also results in more accurate reconstruction quality on the voxels not belonging to void space. This is due to the fact that the fixing of all voxels in S to the grey level of 0 actually decreases the number of variables in Eq. (1), while maintaining the same number of equations as the original system, resulting in faster convergence and more accurate reconstruction quality.

The erosion and threshold parameters of the SUPPRESS algorithm control how many voxels are included in the S set. A large erosion size or a small threshold will result in a smaller intermediate void space estimate S, reducing the probability of falsely classifying material-space voxels as void space voxels during iterations, thereby increasing the algorithm's robustness. On the other hand, it is beneficial to include as many voxels into the S set as possible (without classifying material-space voxels as void space voxels), because this decreases the number of variables in Eq. (1) even more, thereby giving faster convergence and more accurate reconstruction quality. Selecting the optimal parameters can be done with automatic procedures such as projection distance minimization schemes [14] or by histogram-based clustering methods such as Otsu's method [15]. This is, however, out of the scope of this paper, and therefore these parameters were selected manually.

Also note that the void space estimate S will not contain all voxels corresponding to the pores; it will typically contain fewer pixels due to the erosion operation. This conservative estimate for the pore-space voxels is, however, not the final segmentation of the pores; it rather serves as an input for the segmentation of the individual pores in the next step which is described in the next section.

2.2. Segmentation and quantification

In the approach we present here, we did not only optimize the 3D reconstruction algorithm, but also the segmentation process. Our segmentation methodology is displayed in the bottommost part of the flowchart in Fig. 1. First, a volume of interest (VOI) is manually indicated. Next, an anisotropic diffusion filter [16, 17] (with the diffusion constant function as it is proposed in [17]) is applied to the reconstruction to reduce noise without compromising the edges. Since the SUPPRESS reconstruction contains little or no missing wedge artifacts, the pore-space can be segmented by global thresholding. Extracting individual pores is a crucial step, since the pore statistics depend strongly on a good inter-pore separation. A well-known approach for the re-

moval of the artificially introduced connectivity in the case of regularly shaped pores, consists of applying the watershed algorithm to a distance transform of the segmented image [18]. Since this method is inadequate for irregularly shaped pores, it needs to be adapted. The SUPPRESS reconstruction provides a conservative estimate of the set of voxels interior to the pores, i.e., the subset of S that does not correspond to the background. To separate individual pores, the watershed algorithm can then be applied directly on the filtered SUPPRESS reconstruction, by flooding regions starting at the grey value corresponding to void space and stop flooding at the global threshold value that was used for the pore-space segmentation. In the last step of the segmentation procedure all pores that coincide with the boundary of the VOI are removed. Once the segmentation has been computed, the equivalent spherical diameter [19] is determined for each individual pore, providing a quantitative measurement of the pore size distribution.

Note that the global threshold for final pore-space segmentation should typically be chosen larger than the threshold value in the SUPPRESS algorithm. The threshold in SUPPRESS should be chosen smaller, to prevent material space pixels to be classified as pore-space pixels during reconstruction.

Also note that the size of the global threshold parameter affects the size of the pore-space. However, because the SUPPRESS algorithm results in an accurate reconstruction with a clear distinction between void space and material-space, the associated pore-space segmentation is less sensitive to changes in the global threshold parameter (in comparison to other reconstruction algorithms such as SIRT or WBP).

3. Experiments and results

In this section, a range of experiments to evaluate our approach and their corresponding results are discussed. First, in section 3.1, the TEM acquisition set-up for an aluminosilicate sample is described. Next, various simulation experiments are reported in section 3.2. In section 3.3, different figures of merit for the validation of our approach are introduced. In section 3.4, the results of all experiments are reported. Finally, the PORES algorithm is applied to the real data in section 3.5.

3.1. Material and acquisition: aluminosilicate sample

The material under study is an amorphous mesoporous aluminosilicate with a wormhole-like pore structure and irregularly shaped pores. The exact experimental conditions and material specifications have
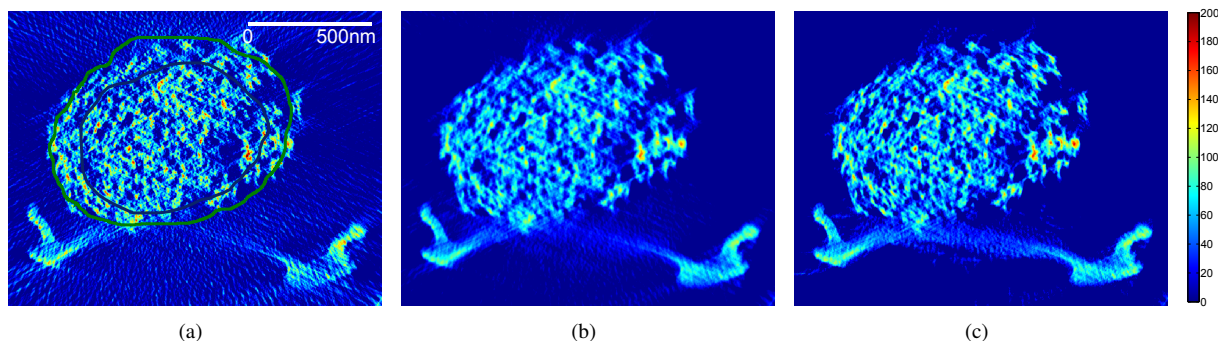
Figure 2: A central slice of reconstructions calculated with different algorithms. (a) WBP with small and large volume of interest (VOI) indicated. (b) SIRT with 300 iterations and a positivity constraint. (c) SUPPRESS with 100 initial SIRT iterations, 200 SUPPRESS iterations and a disk with a two pixel radius as morphological erosion operator.

been elaborated in a previous publication [20]. The TEM samples were prepared by applying drops of ethanol suspension of the powder sample on a carbon coated copper grid. The tomographic imaging was performed using high-angular annular dark-field scanning transmission electron microscopy (HAADF-STEM) [21, 22]. A tilt series was acquired using a FEI Tecnai G2 electron microscope operated at 200 kV in combination with model 2020 Fischione Instruments tomography holder and the FEI XPlore3D software package. The HAADF STEM images were acquired at the detector inner and outer collection angles of 72 mrad and 227 mrad with convergence semi-angle of 10 mrad. Each HAADF-STEM image contains 1024×1024 pixels, which have a 2.06 nm inter-pixel distance. The tilt series was collected over an angular range of $\pm 74°$ with 2° increments and is displayed in Fig. 3. The projection images were aligned with an iterative cross-correlation algorithm together with a manual tilt axis adjustment implemented in FEI Inspect3D software [23]. The STEM image values were shifted in order to have a zero grey value corresponding to void space. To this end, a region where the electron beam clearly encountered void space was manually indicated in every STEM projection image and subsequently the average of the detector pixel values in this region was subtracted from the STEM image values. Reconstructions were calculated on a $1024 \times 1024 \times 1024$ voxel grid of voxel size 2.06 nm^3 in a slice-by-slice manner. The WBP, SIRT and SUPPRESS reconstructions are displayed in Fig. 2.

3.2. Material and acquisition: simulation phantoms

In this section, various simulation phantoms are described, each of which was chosen specifically to validate certain aspects of the SUPPRESS algorithm and the
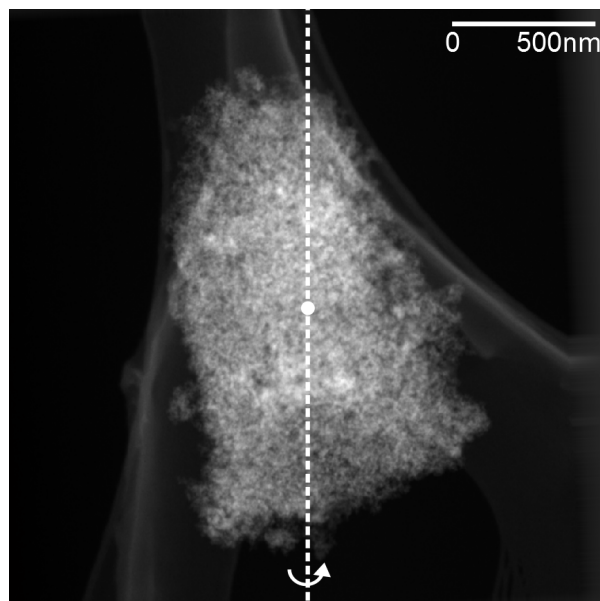

Figure 3: Example HAADF-STEM projection image (at 0° tilt angle) of the aluminosilicate sample. The tilt axis is indicated with the dashed line.

complete PORES approach. A first experiment is a direct validation of the SUPPRESS algorithm by a transmission tomography experiment with various simulation phantoms that contain different pore-space structures (section 3.2.1). In the second experiment, the nanoporous aluminosilicate TEM sample is simulated and the entire PORES processing chain is validated (section 3.2.2). For the final simulation experiment HAADF-STEM data was simulated using the CASINO software package (section 3.2.3), thereby introducing realistic noise into the experiment.

5

3.2.1. First set of simulation phantoms

For validating the SUPPRESS algorithm for a series of different pore-space structures, 100 phantoms similar to the 5 phantoms shown in Fig. 4 were generated. Since these are 2D phantoms, mimicking a slice of a
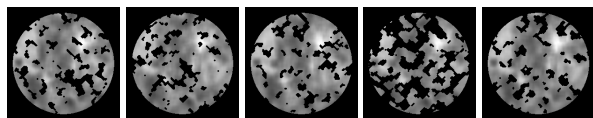

Figure 4: Five examples of the 100 simulation phantoms yielding different pore concentrations.

3D object, their corresponding projections are 1D. The 100 phantoms were generated by introducing random pores in a fixed material phantom of cylindrical shape. The random pores were created by performing a global thresholding operation on different instances of 2D Perlin noise [24]. In total, 70 equiangular projections between ±72° were simulated with a strip kernel [25] and a higher resolution version of the phantom, i.e., on a 256×256 isotropic pixel grid. Poisson distributed noise was applied to the simulated projection data, i.e., each individual noise-free projection value was replaced by a value sampled from the Poisson distribution defined by an expectation value (and variance) equal to the noise-free projection value. Reconstructions were calculated on a 128×128 isotropic pixel grid and with a linear projection model [25].

3.2.2. Second simulation phantom

Direct validation of the SUPPRESS reconstruction algorithm and the subsequent segmentation and quantification of the pore size distribution on the real TEM data of the aluminosilicate is difficult, since no underlying accurate reference image is available. Therefore, a simulation phantom similar to the aluminosilicate and corresponding simulated projection data was created as follows. First, from the HAADF-STEM series of the aluminosilicate, a SUPPRESS reconstruction of one of the more central slices was calculated on a 1024×1024 pixel grid. The SUPPRESS reconstruction parameters were 100 initial SIRT iterations, 200 SUPPRESS iterations and a disk with a two pixel radius as morphological erosion operator. This reconstruction is displayed in Fig. 2(c). Starting from this reconstruction, void space surrounding the sample was manually indicated and assigned a zero grey value. Subsequently, pores were segmented using the watershed algorithm as described above. The resulting pore-space pixels were also set to zero. The resulting 1024×1024 reference image is

displayed in Fig. 7(a). Based on the reference image, artificial projection data was generated along the same 75 projection angles as the real tilt series. Poisson distributed noise was applied to the projection data. With this approach, a reference image is available, and hence an elaborate validation can be performed.

3.2.3. Third simulation phantom

To validate the SUPPRESS algorithm under more realistic noise conditions, HAADF-STEM projection data was simulated with the CASINO Monte Carlo simulation software [26, 27] over an angular range of ±90° with 2° increments. The created sample is displayed in Fig. 5(a) and consists of pores of ellipsoid shape with different lengths for the semi-principal axes. The mate-
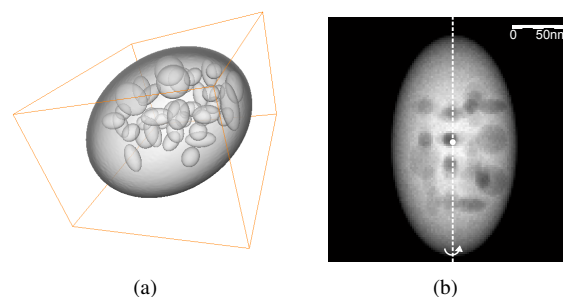

(a)   (b)

Figure 5: (a) The third simulation phantom. The outer contours are defined by an ellipsoid with semi-principal-axis lengths of 78 nm, 108 nm and 61.8 nm. (b) An example HAADF-STEM projection image (at 0° tilt angle) that was generated from the third simulation phantom with the CASINO software. The tilt axis is indicated by the dashed line.

rial matrix was set to contain weight fractions of 0.6% Al, 46.0% Si, 0.3% Na and 53.1% O. Reasonable values for these weight fractions were determined by applying an electron probe micro-analyzer (EPMA) to the aluminosilicate sample (section 3.1). The experiment was set up with a 200 keV-microscope with a high-angular annular dark-field detector that collects electrons scattered between 72 and 227 mrad. Furthermore, a beam semi-angle of 10 mrad was assumed and 60000 electrons were simulated per detector pixel. Each simulated HAADF-STEM projection image was acquired by probing the porous sample with a pixel size of 2.06 nm² in a 116×116 grid. A simulated HAADF-STEM projection image is displayed in Fig. 5(b). Reconstructions were calculated in a slice-by-slice fashion on a $116 \times 116 \times 116$ isotropic voxel grid.

3.3. Figures of merit

For validation, we used two different measures. A first figure of merit is the relative root mean square error

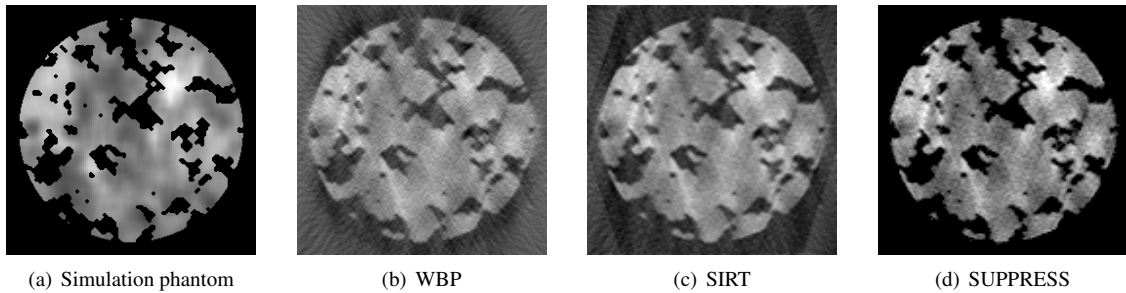| (a) Simulation phantom | (b) WBP | (c) SIRT | (d) SUPPRESS |

Figure 6: A simulation phantom and several different reconstructions. The SUPPRESS reconstruction has less artifacts than the WBP and the SIRT reconstruction.

(RRMSE). It is defined as

$$RRMSE = \sqrt{\frac{\sum_{i=1}^{N}(\hat{x}(i) - x(i))^2}{\sum_{i=1}^{N}(x(i))^2}} \quad , \qquad (2)$$

where $\hat{x} \in \mathbb{R}^N$ denotes the calculated reconstruction and $x \in \mathbb{R}^N$ denotes the phantom used to generate the data.

Finally, to validate the quality of the pore-space segmentation, the Number of Misclassified Pixels (NMP) was calculated, which is defined as the number of pixels that were falsely classified (with respect to the reference image) as either pore or material, reported in percentage of the total number of image pixels.

These validation measures were calculated on the pixel or voxel grid of the phantom, hence, whenever necessary, the reconstructions were upsampled to the resolution of the phantom.

3.4. Results of the simulation experiments

3.4.1. First set of simulation phantoms

The first set of simulation phantoms (of which 5 examples are displayed in Fig. 4) was reconstructed with WBP, SIRT and SUPPRESS. The SIRT algorithm was applied with 200 iterations and a positivity constraint. The SUPPRESS algorithm performed 100 initial SIRT iterations followed by 100 iterations of the main loop. The erosion operator was chosen to be a disk of 1 pixel radius. One particular instance of the simulation phantom and a WBP, SIRT and SUPPRESS reconstruction are shown in Fig. 6(a-d), respectively. Fig. 6 allows for a

first visual assessment of the reconstruction results. The SUPPRESS has more detail and has little influence of the limited angular range over which the projection data was acquired, whereas the WBP and SIRT reconstruction clearly suffer from missing wedge artifacts, which would hamper further analysis. In total, the experiment was repeated 100 times, each time with a different phantom instance. The average results over all experiments are summarized in Table 1. Since the RRMSE assesses the image quality directly and NMP assesses the segmentation directly, it can be concluded from Table 1 that SUPPRESS performs better than SIRT and WBP.

3.4.2. Second simulation phantom

From the simulated projections of the aluminosilicate simulation phantom, WBP, SIRT and SUPPRESS reconstructions were computed. The SUPPRESS reconstruction was calculated with the same parameters as described in section 3.2.2 and the SIRT reconstruction was calculated using a positivity constraint and 300 iterations. The phantom and the SIRT and SUPPRESS reconstruction are displayed in Fig. 7(a), 7(b) and 7(c), respectively. It is clear that the SIRT reconstruction suffers from missing wedge artifacts. This is especially visible on the carbon grid, indicated by the green arrow on top of the SIRT reconstruction in Fig. 7(b), which is smeared out in the vertical direction. Furthermore, the SIRT reconstruction has captured less details in comparison to the SUPPRESS reconstruction, which is clearly illustrated by observing the difference images with respect to the reference image in Fig. 8. The experiment

	WBP	SIRT	SUPPRESS
NMP	10.75% ± 0.67%	6.36% ± 0.91%	4.12% ± 0.88%
RRMSE	0.394 ± 0.010	0.303 ± 0.019	0.245 ± 0.019

Table 1: Validation measures for experiment with the first simulation phantom, reported as *mean±std*.
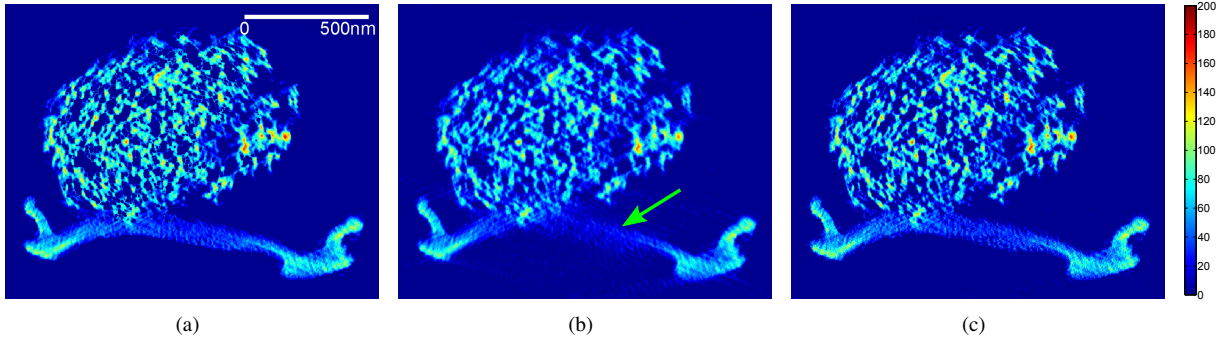
Figure 7: Visual comparison between (a) the reference image, (b) the SIRT reconstruction and (c) the SUPPRESS reconstruction. The missing wedge artifacts, indicated on the SIRT reconstruction by the green arrow, are less pronounced in the SUPPRESS reconstruction.
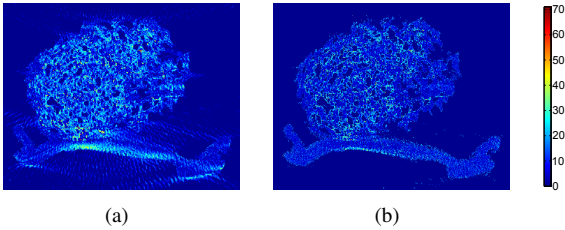

Figure 8: Absolute difference images between the reference image (displayed in Fig. 7(a)) and (a) the SIRT reconstruction and (b) the SUPPRESS reconstruction (displayed in Fig. 7(b) and Fig. 7(c), respectively). Because of missing wedge artifacts, more erroneous pixels are observed in the SIRT reconstruction.

as described above was repeated 100 times (every time with new instances of the Poisson distributed noise) and the calculated statistics were averaged over all experiments. The results are summarized in Table 2. These quantitative results confirm the visual comparison that was made in Fig. 8: The SUPPRESS reconstruction results in the lowest NMP, indicating its ability to accurately capture pore-space, and the lowest RRMSE, indicating that the reconstruction with the highest quality is generated by the SUPPRESS algorithm. Also, the full PORES procedure was applied and hence individual pores were extracted with the procedure as described in section 2.2. Again, after repeating the experiment 100 times, an average histogram of equivalent circular diameter (being the 2D analog of the equivalent spherical

diameter) was composed, shown in Fig. 9. For this histogram, all pore-sizes were considered, even pores corresponding to one single pixel. In practice, these measurements should not be considered, since they are inaccurate because the size of the feature is comparable to the pixel size. In this experiment, however, we were able to compare all pore sizes because of the availability of a ground truth reference image. The histogram shows that the estimation of small pores based on a regular SIRT reconstruction performs significantly worse than based on the SUPPRESS reconstruction.

3.4.3. Third simulation phantom

Tilt series with missing wedge were generated based on the tilt series that was simulated over the full angular range of $\pm90°$ with $2°$ increments. Subsets were taken from the full angular range dataset, representing the angular ranges $\pm(90 - \omega)°$ with $2°$ increments, where ω represents the size of the missing wedge. The SUPPRESS reconstructions were calculated with 100 initial SIRT iterations, 200 SUPPRESS iterations and a four pixel radius disk as morphological erosion operator. SIRT reconstructions were calculated using a positivity constraint and 300 iterations. Calculating the RRMSE and NMP as a function of ω results in Fig. 10.

From the RRMSE plot in Fig. 10, one can notice that for a small missing wedge (i.e., $\omega < 4°$) the reconstruction quality of SIRT is slightly better than for the SUPPRESS reconstruction. This can be contributed

	WBP	SIRT	SUPPRESS
NMP	$1.75\% \pm 0.02\%$	$1.45\% \pm 0.03\%$	$1.37\% \pm 0.03\%$
RRMSE	1.0900 ± 0.0016	0.2900 ± 0.0002	0.2280 ± 0.0002

Table 2: Validation measures for experiment with the second simulation phantom, reported as *mean±std*.
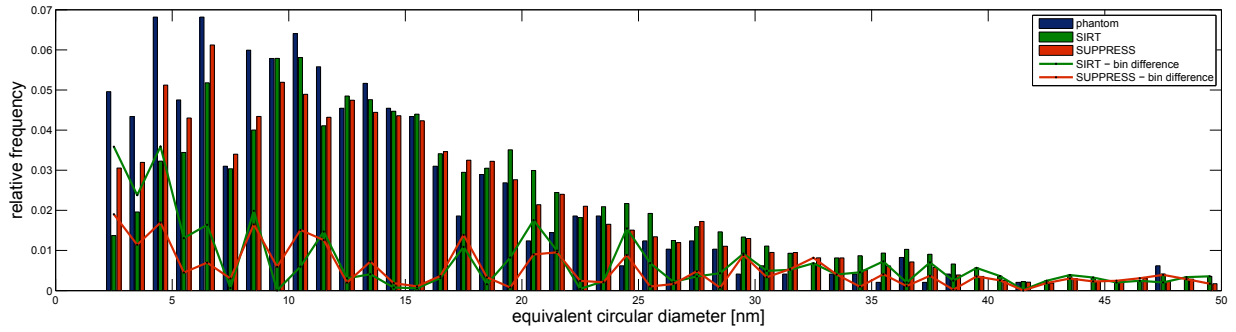
Figure 9: Result of PORES procedure in terms of a histogram of equivalent circular diameters. In the figure, SUPPRESS refers to the regular PORES procedure where the SUPPRESS reconstruction is used as input for the segmentation processing steps and SIRT refers to the PORES procedure where the SIRT reconstruction was used as input for the segmentation processing steps. The green and red curves indicate the absolute difference in relative frequency per bin, for the SIRT reconstruction and SUPPRESS reconstruction, respectively.
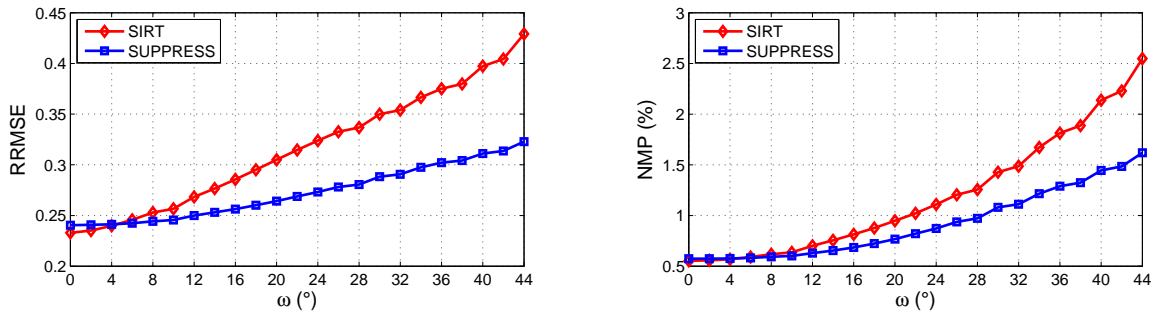

Figure 10: RRMSE (left) and NMP (right) as a function of missing wedge size ω for the third simulation phantom. For this experiment, the projections were simulated with the Monte Carlo method from the CASINO software package.
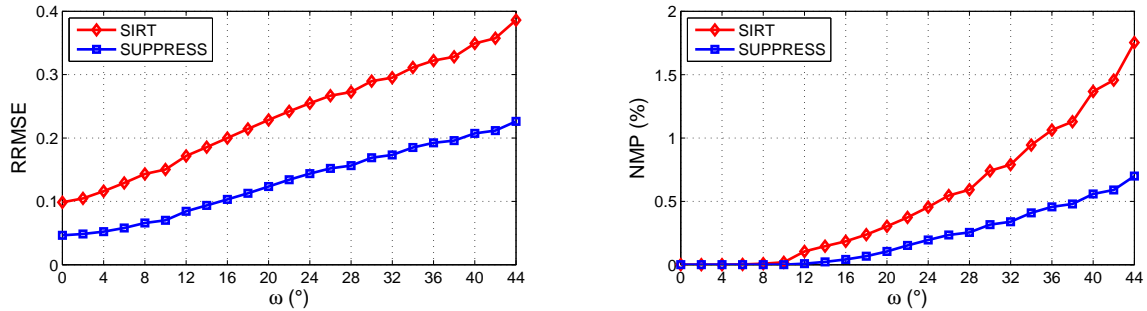

Figure 11: RRMSE (left) and NMP (right) as a function of missing wedge size ω for the third simulation phantom. For this experiment, the projections were simulated with a simple linear model and without noise.

to the fact that in SIRT the noise in the projection images is redistributed over the entire reconstruction domain, whereas the SUPPRESS reconstruction has to distribute it over the smaller set of voxels outside void space, which can result in a slightly larger RRMSE value in comparison to a SIRT reconstruction for small missing wedge values. Indeed, if the same experiment is repeated with projections generated from the phantom with a simple linear model and without noise (see Fig. 11), the SUPPRESS algorithm no longer suffers from this problem and clearly outperforms SIRT for every missing wedge size. Although, for the experiment with the realistic Monte Carlo simulated projections, the reconstruction quality of SIRT is slightly better in

comparison to SUPPRESS for a small missing wedge (which is quantified by the RRMSE plot in Fig. 10), the segmentation quality (which is quantified by the NMP in Fig. 10) is the same. If the missing wedge increases (i.e., $\omega > 4°$) both reconstruction quality (RRMSE) and pore-space segmentation (NMP) is better for the SUPPRESS reconstruction than for SIRT reconstruction. With this increasing missing wedge size, the advantage of SUPPRESS in terms of reconstructing from projection data containing a missing wedge becomes increasingly apparent and the noise effects are no longer dominating.

As an example, the result of the pore-space segmentation of SIRT and SUPPRESS reconstructions for a mis-

sing wedge of 28° is displayed in Fig. 12. To generate Fig. 12, the phantom (see Fig. 5(a)) was first voxelized onto the same $116 \times 116 \times 116$ voxel grid as the reconstructions. From the resulting phantom representation, the pore-space was compared to the segmented pore-space based on the SIRT and SUPPRESS reconstructions. Voxels that were misclassified in both segmentations (a total amount of 0.85% of all voxels) are not visualized, since they do not indicate the difference between both segmentations. From Fig. 12, it becomes clear that the difference in misclassified voxels is caused primarily by the missing wedge artifacts, since the horizontal smearing clearly has more influence on the SIRT-based segmentation than on the SUPPRESS-based seg-

(a) The uniquely misclassified voxels (0.13% of the total number of voxels) of the SIRT-based segmentation interior to the sample's edge (i.e., near the pores).

(b) The uniquely misclassified voxels (0.08% of the total number of voxels) of the SUPPRESS-based segmentation interior to the sample's edge (i.e., near the pores).

(c) Uniquely misclassified voxels of (a) and (b) in one visualization. The phantom of Fig. 5(a) is superimposed to clearly indicate where exactly the misclassifications are located.

(d) The uniquely misclassified voxels (0.28% of the total number of voxels) of the SIRT-based segmentation near the sample's edge and in the background.

(e) The uniquely misclassified voxels (0.04% of the total number of voxels) of the SUPPRESS-based segmentation near the sample's edge and in the background.

(f) Uniquely misclassified voxels of (d) and (e) in one visualization. The phantom of Fig. 5(a) is superimposed to clearly indicate where exactly the misclassifications are located.
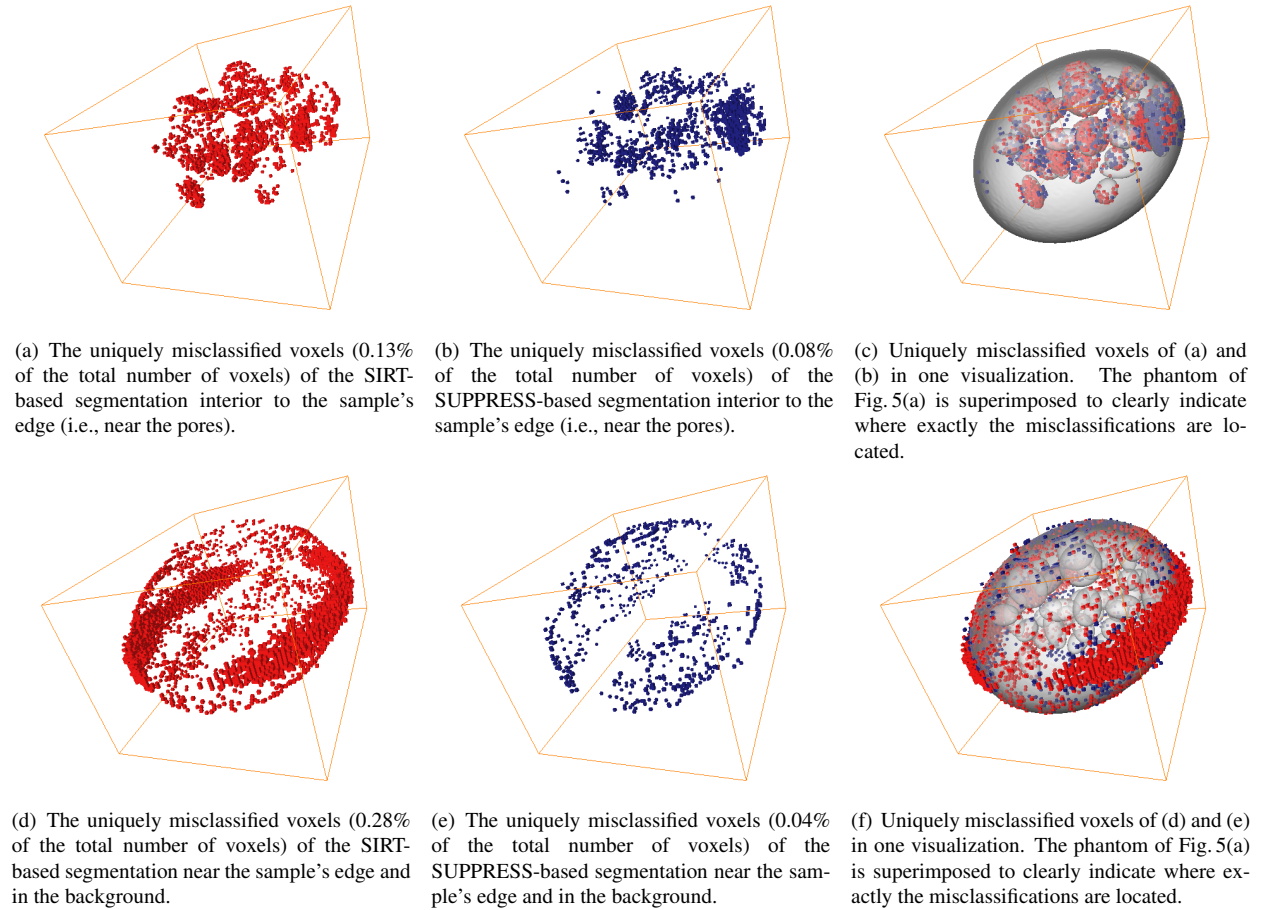
Figure 12: Visualization of misclassified voxels of both the SIRT-based segmentation (red) and the SUPPRESS-based segmentation (blue) for the third simulation phantom with projection data containing a missing wedge of $\omega = 28°$. Voxels that were misclassified by both pore-space segmentations (i.e., based on the SIRT and SUPPRESS reconstructions), representing a total amount of 0.85% of all 116^3 voxels, are not visualized, since they do not indicate the difference between the two methods. Red voxels refer to voxels that were misclassified uniquely by the segmentation based on the SIRT reconstruction while blue voxels refer to voxels that were uniquely misclassified by the SUPPRESS reconstruction. From the figures, it is obvious that the missing wedge (resulting in the horizontal smearing) has far more influence on the SIRT-based segmentation than on the SUPPRESS-based segmentation.

mentation.

3.5. Results for the aluminosilicate sample

The entire PORES 3D processing chain of reconstruction, segmentation and quantification was also applied to the full HAADF-STEM tilt series, i.e., all slices were reconstructed and the resulting 3D reconstruction was used for further processing. The SIRT and SUPPRESS reconstruction were calculated with the same parameters as for the second simulation phantom (section 3.2.2 and 3.4.2). A visual comparison for the WBP, SIRT and SUPPRESS reconstructions can be made in Fig. 2 (a)-(c). In this figure, it is noticeable that the carbon grid (on which the aluminosilicate sample was mounted) is smeared out in the vertical direction due to the missing wedge for the WBP and SIRT reconstruction, which is no longer the case in the SUPPRESS reconstruction. Also, the SUPPRESS reconstruction appears sharper in comparison to the SIRT and WBP reconstruction. To assess the robustness of the entire PORES processing chain, the equivalent spherical diameters in both a larger VOI and a smaller VOI were calculated. The shape of the VOI is indicated on top of the WBP reconstruction of Fig. 2(a). Equivalent spherical diameters corresponding to a single voxel volume (i.e., an equivalent spherical diameter of 2.56 nm) were discarded, because these measurements are inaccurate in practice. The resulting histograms are displayed in Fig. 13. It is obvious that the histograms for the small VOI and the large VOI are highly similar. We can conclude that the quantification based on the segmentation of the SUPPRESS reconstruction is robust. As an illustration, the histogram obtained with a global thresholding operation applied on a basic SIRT reconstruction (a method that is employed often in practice) is displayed in Fig. 14. The histograms in this figure were generated by calculating relative frequencies on bins placed

Figure 14: Comparison of histograms obtained with different methods: the red dashed curve is the histogram based on an estimate of pore-space that was calculated by applying a global thresholding operation to a conventional SIRT reconstruction, the blue curve is the histogram obtained with a nitrogen sorption experiment and the black dotted curve was calculated with the pores method.

around integer pore diameters and fitting a smooth curve through it. Pore diameters corresponding to a single voxel (i.e., pore diameters of 2.56 nm) were removed from the results. It follows from all previous validation experiments (section 3) that the histogram obtained from a globally thresholded SIRT reconstruction is less accurate than the histogram obtained with the PORES algorithm. Furthermore, a comparison to the pore size distribution characterized by a nitrogen sorption experiment can also be done in Fig. 14. This nitrogen sorption measurement has been performed on a Quantachrome Quadrasorb SI unit, after degassing the sample under high vacuum conditions for a duration of 16 h at 473 K. Subsequently, the pore size distribution has been determined by applying the Barret-Joyner-Halenda (BJH) method on the desorption branch of the nitrogen sorption isotherm. It is clear that the PORES method is in
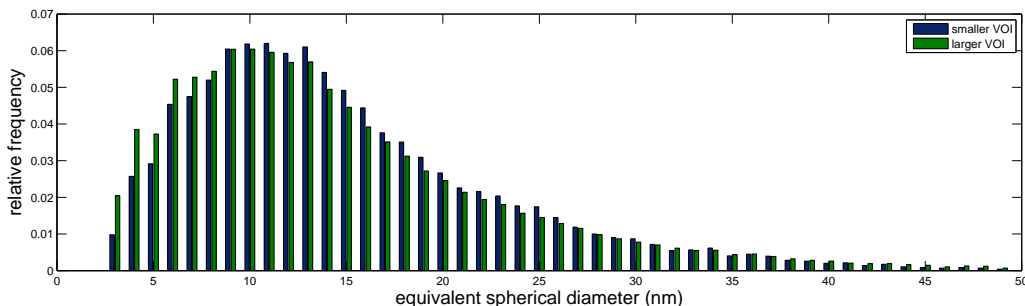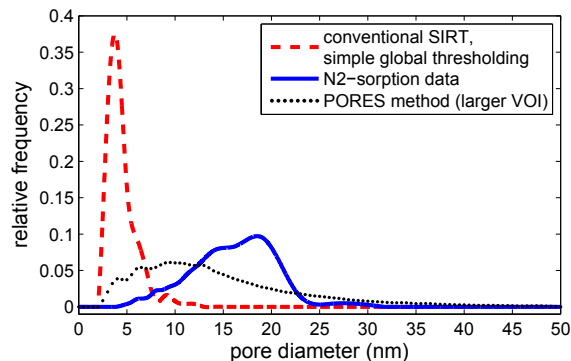

Figure 13: Histogram of equivalent spherical diameter for a larger VOI and a smaller VOI obtained from with the PORES method (based on the SUPPRESS reconstruction). The histograms have a large similarity, indicating the robustness of the quantification method.

better agreement with the experimental nitrogen sorption data than the histogram obtained by a global thresholding operation applied on a SIRT reconstruction, although still a distinct discrepancy is detected between the PORES and the N2 sorption method. This can be appointed to the complicated unordered structure of the porous sample. Indeed, the pore size distribution based on the nitrogen sorption measurement has been calculated by the BJH model, assuming that the porous structure only contains cylindrical pores, which is not in full accordance with the actual situation. Unfortunately, no better alternative is possible, since no calculation model has yet been developed to determine the pore size distribution based on nitrogen sorption measurements of an unordered aluminosilicate structure.

4. Conclusions

In conclusion, the PORES algorithm was proposed; it is an integral approach for the reconstruction, segmentation and quantification of nanoporous materials. As the proposed processing chain is tailored specifically for nanoporous materials, accurate quantification becomes possible. The first step, i.e., the SUPPRESS reconstruction, significantly reduced missing wedge artifacts in the reconstruction by the incorporation of prior knowledge in the reconstruction algorithm. Individual pores were reliably extracted, allowing for quantification by calculating individual pore statistics. The SUPPRESS and PORES algorithm were extensively validated with different experiments, varying in sample properties and the way in which data was simulated. In contrast to the standard N2-sorption method for determination of the pore size distribution, the PORES method does not assume cylindrical or slit-shaped pores. Furthermore, besides providing an overall pore size distribution (the result of the N2-sorption experiment), our method also allows for quantification of individual pores. Beside pore size, any other quantification is also possible, e.g., eccentricity, orientation, perimeter, etc. Furthermore, the information about interconnectivity between nanopores can also be extracted, which is important to improve mass transport and catalytic effectiveness in nanomaterials.

Acknowledgements

References

[1] F. Schuth, K. S. Sing, J. Weitkamp, Handbook of Porous Solids, volume 1, Wiley, 2002.

[2] G. Ertl, H. Knoezinger, F. Schuth, J. Weitkamp, Handbook of Heterogeneous Catalysis, volume 1, SIAM, 2008.

[3] S. Lowell, J. E. Shields, M. A. Thomas, M. Thommes, Characterization of porous solids and powders: surface area, pore size and density, volume 16, Kluwer Academic Publishers, 2004.

[4] L. Reimer, H. Kohl, Transmission electron microscopy: physics of image formation, volume 36, Springer, 2008.

[5] M. Weyland, P. A. Midgley, Electron tomography, Mater. Today 7 (2004) 32–40.

[6] L. Feldkamp, L. Davis, J. Kress, Practical cone-beam algorithm, J. Opt. Soc. Am. A 1 (1984) 612–619.

[7] J. Gregor, T. Benson, Computational analysis and improvement of SIRT, IEEE Trans. Med. Imag. 27 (2008) 918–24.

[8] K. J. Batenburg, J. Sijbers, DART: a practical reconstruction algorithm for discrete tomography, IEEE Trans. Imag. Proc. 20 (2011) 2542–53.

[9] K. J. Batenburg, S. Bals, J. Sijbers, C. Kübel, P. A. Midgley, J. C. Hernandez, U. Kaiser, E. R. Encina, E. A. Coronado, G. Van Tendeloo, 3D imaging of nanomaterials by discrete tomography., Ultramicroscopy 109 (2009) 730–40.

[10] T. Roelandts, K. J. Batenburg, E. Biermans, C. Kübel, S. Bals, J. Sijbers, Accurate segmentation of dense nanoparticles by partially discrete electron tomography, Ultramicroscopy 114 (2012) 96–105.

[11] B. Goris, W. Van den Broek, K. J. Batenburg, H. Heidari Mezerji, S. Bals, Electron tomography based on a total variation minimization reconstruction technique, Ultramicroscopy 113 (2012) 120–130.

[12] R. Leary, Z. Saghi, P. A. Midgley, D. J. Holland, Compressed sensing electron tomography., Ultramicroscopy 131 (2013) 70–91.

[13] E. Biermans, L. Molina, K. Batenburg, S. Bals, G. Van Tendeloo, Measuring porosity at the nanoscale by quantitative electron tomography, Nano letters 10 (2010) 5014–5019.

[14] W. Van Aarle, K. Batenburg, J. Sijbers, Optimal threshold selection for segmentation of dense homogeneous objects in tomographic reconstructions, IEEE Trans. Med. Imag. 30 (2011) 980–989.

[15] N. Otsu, A threshold selection method from gray-level histograms, IEEE Trans. Syst., Man, Cybern. 9 (1979) 62–66.

[16] G. Gerig, O. Kubler, R. Kikinis, F. a. Jolesz, Nonlinear anisotropic filtering of MRI data., IEEE Trans. Med. Imag. 11 (1992) 221–32.

[17] P. Perona, J. Malik, Scale-space and edge detection using anisotropic diffusion, IEEE Trans. Pattern Anal. Machine Intell. 12 (1990) 629–639.

[18] F. Meyer, Topographic distance and watershed lines, Signal processing 38 (1994) 113–125.

[19] B. Jennings, K. Parslow, Particle size measurement: the equivalent spherical diameter, Proc. R. Soc. Lond. A 419 (1988) 137–149.

[20] C. J. Van Oers, M. Kurttepeli, M. Mertens, S. Bals, V. Meynen, P. Cool, Zeolite β nanoparticles based bimodal structures: Mechanism and tuning of the porosity and zeolitic properties, Micropor. Mesopor. Mat. 185 (2014) 204–212.

[21] P. Midgley, M. Weyland, 3d electron microscopy in the physical sciences: the development of z-contrast and EFTEM tomography, Ultramicroscopy 96 (2003) 413 – 431.

[22] S. Pennycook, Z-contrast transmission electron microscopy: direct atomic imaging of materials, Annu. Rev. Mater. Sci. 22 (1992) 171–195.

[23] R. H. M. Schoenmakers, R. A. Perquin, T. F. Fliervoet, W. Voorhout, High resolution, high throughput electron tomography reconstruction, Microsc. Microanal. 11 (2005) 312–313.

[24] K. Perlin, Improving noise, in: ACM T. Graphic., volume 21, ACM, 2002, pp. 681–682.

[25] A. C. Kak, M. Slaney, Principles of Computerized Tomographic Imaging, SIAM, 2001.

[26] H. Demers, N. Poirier-Demers, D. Drouin, N. de Jonge, Simulating stem imaging of nanoparticles in micrometers-thick substrates, Microsc. Microanal. 16 (2010) 795.

[27] H. Demers, N. Poirier-Demers, A. R. Couture, D. Joly, M. Guilmain, N. de Jonge, D. Drouin, Three-dimensional electron microscopy simulation with the CASINO Monte Carlo software, Scanning 33 (2011) 135–46.