A two-stage approach to the segmentation of FIB-SEM images of highly porous materials

M. Salzer[a], A. Spettl[a,*], O. Stenzel[a], J.-H. Smått[b], M. Lindén[c], I. Manke[d], V. Schmidt[a]

[a] *Institute of Stochastics, Ulm University*
[b] *Center for Functional Materials, Laboratory of Physical Chemistry, Åbo Akademi University*
[c] *Institute of Inorganic Chemistry, Ulm University*
[d] *Institute of Applied Materials, Helmholtz Zentrum Berlin für Materialien und Energie GmbH*

Abstract

Segmentation of 3D FIB-SEM images is a specific problem for porous materials, where grey intensities are not sufficient to determine the phase represented by a certain voxel. In this paper a new approach is proposed to segmentation of FIB-SEM images, which is especially designed for highly porous materials. It detects and allocates structures based on their last occurrence in z-direction. Afterwards, the obtained segmentation is improved by applying local thresholds. This can be done in an iterative or in a direct way – both methods are described and compared with each other. Furthermore, the final segmentation is compared to a binarisation obtained by adaptive local thresholding.

Keywords: Porous Media, FIB-SEM Tomography, 3D Imaging, Segmentation, Local Thresholding

1. Introduction

Highly porous materials are of current interest within a wide range of applications, including heat flow management related to insulation and heat

*Corresponding author:
Postal address: Institute of Stochastics, Ulm University, 89069 Ulm, Germany
E-mail: aaron.spettl@uni-ulm.de, Phone: +49 731 50-23555, Fax: +49 731 50-23649

exchange [1], drug delivery [2], molecular separation by adsorption or chromatographic separation [3], and catalysis [4] just to name a few. Furthermore, many naturally occurring materials, like bone, rock, and wood, not to mention the extracellular space in the brain, exhibit complex pore architectures. In addition to the total porosity, the distributions of pore shape, pore size, pore connectivity, and effective pore length are important structural characteristics that control the mass transport properties of the materials, which in turn directly relates to the material performance in all above mentioned application examples. For instance, the thermal conduction in porous materials is controlled by the ratio between the mean free path of the gas molecules and the pore diameter [5], but is also highly dependent on the pore connectivity [6]. The diffusion of reactants and products within catalysts [7] or drugs within porous drug carriers is also highly dependent on pore shape and pore connectivity. Depending on the material structure, these parameters may, however, not be easily experimentally accessible; often only averaged global values are obtained. This is especially true for amorphous materials with a wide range of pore sizes, for materials which may contain closed pores, and in cases when larger pores are connected to each other through smaller pores. For example, the pore size of mechanically stable macroporous materials, like macroporous ceramics and metals, is often determined by mercury porosimetry. Here, the intrusion of mercury into the pores is measured as a function of pressure, from which the pore size is derived by applying the Washburn equation, assuming a cylindrical pore shape [8]. This may lead to large errors in cases where the material contains large pores which are connected to the outside through smaller pores, as mercury will not intrude into the larger pores before the pressure is high enough to fill small pores. Also the pores may not have a cylindrical pore shape, which represents another experimental challenge. Furthermore, tortuosity, defined as the ratio between the effective and the projected pore length, is often used in modelling diffusion in porous materials, but also here only a global value is obtained. A global value does not reflect the complexity of the real pore system, which makes it difficult if not impossible to define true structure-property relationships. Different imaging techniques are therefore of increasing interest as an additional means for gaining detailed insight into the 3D structure of porous materials. Depending on the material and the length-scale of interest, confocal microscopy [9], nuclear magnetic resonance imaging (MRI) [10], X-ray computerized tomography [11, 12], 3D transmission electron microscopy (TEM) [13, 14], and focused ion beam scanning electron microscopy

(FIB-SEM) [15, 16] can all be used for generating 3D images of the material. High-quality image analysis, including high-quality binarisation of the image, is the key to accurately compute the above-mentioned parameters, like pore connectivity, pore shape, tortuosity etc.. This, however, is in many cases a non-trivial task.

The exemplary data used in this paper was obtained by focused ion beam (FIB) tomography, which is a common method to analyse nanoscale structures with spatial resolutions from about 8 nm. The technique was developed about 10 years ago [17, 18, 19] and today, it is an established and powerful method for a wide range of applications in materials science; mainly to assess the three-dimensional structure and morphology of the material [20, 21, 22, 23]. FIB-SEM tomography makes use of the imaging ability of a SEM (scanning electron microscope) [24] and the ability of the FIB to remove certain parts of material from the surface of a sample with nanometre accuracy [25]. These two techniques are combined by sequentially acquiring an image with the SEM and then using the FIB to remove the current layer to expose the material behind. By repeating this step many times an image stack is obtained that consists of 2D images each representing a different layer of the sample [26].

To perform statistical microstructure analysis as described above, it is necessary to reconstruct the binary 3D image, where the value of a voxel indicates whether it belongs to the considered phase or not. This reconstruction is usually done by aligning, cropping and combining the 2D images and then segmenting the data with an appropriate segmentation algorithm [27]. However, when one of the phases is transparent, the single 2D images acquired via SEM do not only contain information about the material of the actual layer located directly at the surface, but also about parts that are located in the background [23, 28]. This can lead to relatively high grey values in parts of the reconstructed 3D image representing pore space. On the other hand, when analysing a sample of high porosity, structures of lower thickness appear darker than those with a higher thickness. In the exemplary data used in this paper, some of the cell walls separating different pores are significantly thinner than others. When electrons of the SEM reach such thin structures they penetrate not only the surface but some of them even the whole structure. The electrons surpassing the structure vanish in the empty space behind and therefore cannot be detected by the sensor of the SEM. This leads to significantly lower grey values and an overlap of grey values for thin objects close to the sensor and more distant thick ones. The

overlap prohibits the application of common pre-processing like filtering [29] that is often used to remove negative effects caused by measurements errors.

This complicates the segmentation of the grey scale image into a binary image, where each voxel either belongs to the foreground (solid phase) or to the background (pore space). For example, when using elementary segmentation algorithms like global thresholding, i.e., considering a voxel to be foreground if its grey value is above a certain prechosen threshold, voxels that actually represent empty space are often classified foreground as their grey value is increased significantly by material located behind it. Even more complex binarisation techniques like local thresholding as described in [30] and successfully applied in [31] do not consider the special nature of this issue.

We therefore developed a new approach called *local threshold backpropagation* that compares grey values in z-direction and detects the appearance and disappearance of structures based on this comparison rather than relying on absolute grey values. The algorithm provides good results when it comes to allocating objects in z-direction. However, further filtering is necessary. In particular, a dilation of the foreground adds small missing connections, but it also leads to the expansion of the foreground into the background. In a second step, we correct this negative effect by applying local thresholds based on the surroundings of every voxel.

The paper is organized as follows. First, in Section 2, we describe the material and image data that is used as an exemplary application of our approach. In Section 3, we introduce our method of local threshold backpropagation and show how the obtained binary image is postprocessed. We then present two methods of improving that preliminary segmentation, first by a direct approach (Section 4) and second we use the framework introduced in [27] as an iterative alternative (Section 5). The effect of the parameters and the two improvement methods are discussed in Section 6. Finally, in Section 7, we compare the final segmentation to a binarisation obtained by adaptive local thresholding and discuss the advantages and disadvantages of both techniques.

2. Preliminaries

2.1. Description of Material

As an exemplary material a monolithic foam-like silica material has been chosen, which has been synthesized under acidic conditions using a sol-gel

process [32, 33]. Depending on the kinetics of phase separation and gelation, which in turn can be tuned by changing synthesis temperature for example, different structures can be obtained, including foam-like structures, three-dimensionally interconnected macroporous networks, and isolated particles, as discussed in detail in earlier works [34]. Especially the foam-like structure is difficult to structurally characterize using standard techniques like mercury porosimetry for reasons already stated above. A representative SEM image of this monolithic material is shown in Figure 1, which clearly shows the foam-like structure with large, often close-to-spherical main pores with a wide range of pore sizes, which are connected to each other mainly through smaller pores, often referred to as windows.

2.2. Imaging Technique

A Zeiss 1540EsB CrossBeam® with an ultra-high resolution GEMINI® e-Beam column and a high performance Canion gallium ion column were used for the FIB-SEM tomography measurements. The SEM was operated at an acceleration voltage of 5 kV. To prevent the curtain effect, a Pt protective layer with a thickness of a few 100 nm was applied with the gas injection system (GIS) [25]. A high efficiency annular type in-lens secondary electron (SE) detector was applied. The 3D data consists of a stack of 90 2D images with 1024×768 voxels in the x-y-plane. Each voxel represents a cuboid with 25.29 nm length in x- and y-direction, see Figure 1, and 25 nm in z-direction.

2.3. Data Preprocessing

FIB-SEM images need to be aligned, cropped and combined to a single 3D image. To align the individually acquired 2D images correctly we use a straight forward implementation of minimizing the difference between two images for a manually selected region [35]. More precisely, we consider potential shift vectors and calculate the difference of the shifted 2D slice and a 2D reference image for a certain region. The shift vector which minimizes this difference is chosen as shift for the corresponding 2D slice, i.e. for every 2D slice, a new shift vector is calculated.

As the aligned image still contains small shifts and we strongly rely on correct alignment in Section 3, further filtering is needed. We therefore apply a mean-value filter to each 2D slice of the aligned 3D image I and denote the result by $I_{2\times2}$, a voxel of $I_{2\times2}$ is given by

$$I_{2\times2}(x,y,z) = \frac{1}{4}(I(x,y,z) + I(x+1,y,z) + I(x,y+1,z) + I(x+1,y+1,z))$$

A 2D filter is used because only a small smoothing effect in x-y-direction is desired, which is required to compensate small aligment problems. Note that we use the image $I_{2\times2}$ only in Section 3 and improve the segmentation result in Section 4 and 5 based on the original image I.

3. Local Threshold Backpropagation

A typical approach to image segmentation is global thresholding, i.e., considering a voxel to be foreground if its grey value is above a certain threshold. As the grey value of a given voxel is not a sufficient criterion for segmentation of FIB images, a more advanced framework for segmentation has been introduced in [27]. Although this framework is more flexible, the parameter classes described in [27] still rely on a reasonable global threshold. However, choosing a reasonable global threshold for our data is impossible, due to different grey intensities based on different thicknesses, see Figure 5 iv). We therefore developed a new technique to detect structures based on their relative grey value w.r.t. their neighbours rather than their absolute grey value.

3.1. Local Threshold Backpropagation

Due to the high porosity of the sample and the imaging technique, structures are visible quite some time before they actually appear in the section plane. Therefore, when looking at each 2D image individually it is – even for human vision – difficult to determine the expansion of a structure in z-direction. However, by inspecting a sequence of images it is easy to determine the last time a structure is present due to the sudden decrease in the corresponding grey values, see Figure 2.

This drop of grey values allows us to detect the last occurrence of a structure and to use the last grey value observed to locate the first occurrence. The complete process consists of three steps, see also Figure 3: 1. Detecting the last occurrence of a structure. 2. Estimating a threshold based on the grey value of the last occurrence. 3. Backpropagating the threshold. Note that steps 2 and 3 are only performed if the last occurrence of a structure is detected in the first step. Hence, structures that do not disappear within the observation window are not detected properly. However, this is not a huge issue as it is possible to ignore the last few slices of the image stack.

To process the voxels in the right order we divide the 3D image $I_{2\times2}$ into 1-dimensional images I_{xy} where $I_{xy}(z)$ is given by $I_{2\times2}(x, y, z)$, see Figure 4. We

consider each image I_{xy} individually and process the voxels from $z = 1$ until $z = z_{\max}$. Thus, at the time we process voxel z we can be sure that all voxels $1, \ldots, z-1$ are already processed. We use this fact during backpropagation, when we recognize that the current structure has already been processed partially, see Section 3.1.3.

3.1.1. Detecting Disappearance

We detect the disappearance of structures by comparing the grey value of a voxel to the grey value of its successor. If there is a significant drop in grey value we assume that the current structure disappears in the next slice. To distinguish drops in grey value caused by the disappearing of the current structure from drops caused by measurement errors we choose a global threshold d based on the average measurement errors. More precisely, we will constitute voxel z to be the last voxel representing a certain structure if $I_{xy}(z) > I_{xy}(z+1) + d$ is fulfilled. The condition is met when there is a drop in grey value from voxel z to $z+1$ larger than d, see Figure 3 i). In this case we continue with the following steps 2 and 3. If there is no significant decrease in the grey value we do not continue processing this voxel, though it might be classified later on while processing another voxel.

3.1.2. Estimating a Threshold

If z is the last occurrence of a structure we denote it by z_{lo} and use the grey value $I_{xy}(z_{\text{lo}})$ to estimate a reasonable local threshold $T_{\text{first}}(z_{\text{lo}})$. This threshold $T_{\text{first}}(z_{\text{lo}})$ is used to determine the expansion or more precisely the first occurrence of the current structure, see Figure 3 ii). However it is not applied to the voxel z_{lo} because we already classified that voxel as foreground due to the sudden drop in grey value in step 1. To estimate the local threshold $T_{\text{first}}(z_{\text{lo}})$ we apply a linear transformation $x \mapsto \alpha \cdot x + \beta$ to $I_{xy}(z_{\text{lo}})$, where α and β are prechosen global parameters, leading to $T_{\text{first}}(z_{\text{lo}}) = \alpha \cdot I_{xy}(z_{\text{lo}}) + \beta$. How these parameters are chosen strongly depends on the nature of both the data and the analysis to be performed later on. Lower thresholds are more likely to be met at a greater distance and, therefore, structures are likely to be estimated too big but a good connectivity is obtained. On the other hand, higher thresholds may increase accuracy, but we risk missing some structures or connections. For example, when analysing foam, lower thresholds may lead to a better estimate of the connectivity, whereas higher

thresholds may lead to more precise results when analysing the thickness of cell walls. For the exemplary data set, the effect of the parameters will be analysed in Section 6.

3.1.3. Backpropagating the Threshold

The threshold $T_{\text{first}}(z_{\text{lo}})$ derived in the previous step is now compared to the grey values of voxels with smaller z coordinates. Starting with $z-1$ we descend in z-direction and constitute the current voxel to be foreground if its grey value is equal or greater than $T_{\text{first}}(z_{\text{lo}})$. This is continued until a voxel meets one of the following abort criteria.

grey value below threshold: If a voxel has a grey value below the threshold, i.e., $I_{xy}(z) < T_{\text{first}}(z_{\text{lo}})$, we assume to have passed the appearance of the structure (see Figure 3 iii)).

classified as foreground: If a voxel has already been constituted as foreground we assume that we already have detected the appearance of the current structure in a previous step. This is most likely to be the case when a structure disappears not at once but over multiple slices. As the structure gets thinner with each layer being cut off by the FIB the grey values are likely to be lower. By not backpropagating this lower threshold, we automatically use the first threshold to estimate the appearance of the structure. Figure 4 gives a real-world example of this. Note that the last drop detected is used for the disappearance of the structure, even though the first drop determines the local threshold that is backpropagated.

3.2. Postprocessing

Performing the previously described steps already yields a reasonable segmentation that outperforms usual approaches like global thresholding, see Figure 5. However, there are small clusters of misclassified voxels that are removed by the postprocessing described in the following. First, we remove isolated clusters of foreground voxels that do not represent any structure by constituting them as background when they have less than 3 foreground voxels in their $5 \times 5 \times 1$ neighbourhood. After this, we perform a 2D dilation on all x-y-planes and denote the result by B_{backprop}. For the exemplary data set, the structuring element for the dilation is a circle with radius 2. The dilation removes small clusters of background voxels within the foreground

phase. Additionally, it connects otherwise separate foreground voxels when they are close enough to each other, see Figure 6. Such cell walls are not detected properly because they are orthogonal to the x-y-plane and thus do not disappear suddenly but slowly drift away. This leads to a slow decrease in grey values that often stays below the threshold d used in Section 3.1.1 to detect the disappearance of a structure.

Although this kind of postprocessing improves the segmentation significantly, the dilation also leads to an increase of the volume and inaccuracy especially at the edges of both phases. We therefore present two approaches to compensate for this negative effect, see Sections 4 and 5.

4. Improvement by Direct Local Thresholding

In the following we present a method to further improve the segmentation obtained in Section 3 by using a simple local thresholding scheme. These local thresholds are needed to detect background voxels close to an edge that are currently falsely classified as foreground. To classify such voxels, we compare their grey value to the grey values of their neighbouring voxels. If a voxel belongs to the background and is close to the foreground phase it is expected to have a lower grey value than its neighbouring foreground voxels (and also a higher grey value than background voxels not close to the foreground). To increase this effect, we apply a mean-value filter for every x-y-plane, with a large radius of 10 voxels. We denote this smoothed image by M and consider the difference image $I - M$, see Figure 7. For voxels clearly in the foreground or background, the mean-value filter has nearly no effect, thus the grey value in the difference image $I - M$ is approximately zero. For voxels close to the interface between foreground and background, $I - M$ contains negative or positive values: negative values correspond to voxels that have brighter voxels in their neighbourhood, positive values correspond to voxels with darker voxels in the neighbourhood.

We want to detect voxels in the background that are close to the foreground, such that we can remove them from the segmentation B_{backprop} obtained in the previous section, where the dilation expanded the foreground. We apply a global threshold τ to the difference image, the resulting binary image is called B_{local}, and τ is set to a negative value close to zero. Therefore, the dark (background) phase of B_{local} shows the background areas close to the foreground. These voxels are going to be removed from the foreground

of the segmentation B_{backprop} in the next step. We do this by combining the local threshold segmentation B_{local} with the segmentation B_{backprop} by using a voxel-wise minimum operator. This means we consider a voxel to be foreground only if it is constituted as foreground in *both* of the given techniques. This leads to the result image $B_{\text{backprop/local}}$ being defined by $B_{\text{backprop/local}}(x, y, z) = \min(B_{\text{backprop}}(x, y, z), B_{\text{local}}(x, y, z))$.

5. Improvement by Iterative Local Thresholding

In the following we describe and then apply an iterative thresholding algorithm presented in [27] to improve the segmentation B_{backprop} obtained by the local threshold backpropagation.

5.1. Iterative Framework for Automatic Image Segmentation

The approach proposed in [27] focuses on the boundary of the two phases that are to be segmented rather than the phases themselves. The surface given by this boundary is initialized based on a preliminary segmentation (e.g. global thresholding). This initial surface is then evolved iteratively based on a partial differential equation whose parameters are derived from either the original image or *a priori* information. Finally, when the surface has reached a steady state, the corresponding segmentation is extracted and considered as final result. To represent the surface, the approach makes use of the level set methodology, see [36]. The level set function will be initialised based on a preliminary segmentation that is derived by a separate and most likely simpler method like global thresholding. After initialising the surface it is evolved by iteratively solving the partial differential equation $\phi_t + V \cdot \nabla\phi + a = b\kappa$, where ϕ denotes the level set function and κ its mean curvature defined as $\kappa = \phi_{xx} + \phi_{yy} + \phi_{zz}$, a and b are scalar fields controlling the expansion of the surface in normal direction. V is a vector field pulling the surface in a certain direction, e.g. to areas of high intensities of the gradient image. To compute a numerical solution it is necessary to discretize. The spatial component is discretized based on the 3D grid that is already used to represent our image data. The grid we use to discretize time is given by a manually chosen grid granularity Δt. This leads to the update rule $\phi^{n+1} = \phi^n - \Delta t(V \cdot \nabla\phi^n + a - b\kappa^n)$, where ϕ^n denotes the – now discrete – level set function ϕ at time step n, κ^n its curvature. a and b still denote scalar fields and V a vector fields, all being discrete.

5.2. Application to Iterative Local Thresholding

The dilation performed in Section 3.2 led to an expansion of the foreground into the background. Thus the surface given by the phase boundary is located in the background. The key idea is to iteratively expand this surface until it reaches relatively high grey values which are expected to belong to the foreground phase.

Let M again denote the mean-value filtered version of the original image I and now consider the difference $M - I$. This difference image is expected to have negative values for foreground voxels close to an edge and positive values for background voxels also located close to an edge. We define the scalar field a by $a = \max\{0, M - I - \tau\}$, where τ is chosen as described in Section 4. This keeps the surface expanding (while in the background) until it reaches areas with grey intensities higher than their surroundings. Setting all negative values of $M - I - \tau$ to zero guarantees that the volume enclosed by the surface is not shrinking anywhere, even if the grey intensities of the underlying voxels are lower than those of their – possibly in a different slice located – surroundings.

All other parameters V and b are set to zero as we do not need them here. This simplifies the update rule to $\phi^{n+1} = \phi^n - \Delta t a$. To initialise ϕ we use the segmentation B_{backprop} obtained by local threshold backpropagation and denote the final result by $B_{\text{backprop/iterative}}$, see Figure 8.

6. Discussion of Parameters and Methods

6.1. Backpropagation Parameters

To get an impression of the influence of the backpropagation parameters, we computed the result of local threshold backpropagation for various parameter combinations of α and β. Due to a missing reference binarisation we can only compare different sets of parameters to each other. Therefore, we computed the porosity and the spherical contact distribution function for a certain set of parameters. The spherical contact distribution function is widely used in spatial statistics to characterize random sets. It is defined as the cumulative distribution function $H(r)$, $r > 0$, of the minimum distance from an arbitrary location in one phase to the other. For example, if H_B is the spherical contact distribution function of the foreground to the background of binary image B, then $H_B(x)$ represents the volume fraction of the foreground that has a distance to the background smaller than or equal to

x. Figure 9 i) shows that both parameters α and β have a significant influence on the result obtained by the local threshold backpropagation. Even small changes of 10% already have a significant impact on the total porosity. Furthermore, Figure 9 ii) shows the spherical contact distribution function for various values of α with all other parameters remaining constant. The early increase represents the high amount of thin structures (cell walls) that is confirmed by visual inspection. For $\alpha = 1.0$, about 85% of all voxels have a distance less than or equal to 5 voxels to the background phase. The value increases to 96% when α is increased to 1.1. This represents the higher amount of thin structures that is natural to appear when higher thresholds are used. For lower values of α (0.9, 0.8) the reverse effect occurs. A similar behaviour can be seen for β.

6.2. Detecting Disappearance Parameters

The backpropagation – controlled by the previously discussed parameters α, β – is only initiated after a drop in grey value larger than d. Hence, one would expect this parameter to be the most influential one. However, even for a relatively wide range of values of d (10 to 60) the porosity stays in a relatively small range (between 0.315 and 0.4). For high values of d the porosity is decreasing. This is expected, because higher thresholds make it less likely for a structure to be detected. Surprisingly, for lower values the same effect can be seen. One possible explanation is that for structures that disappear not at once, but over multiple slices, the first disappearance is detected earlier with lower values for d. At this earlier stage the structure still has a relatively high grey value and therefore the threshold T_{first} is higher. A higher value of T_{first} then leads to thinner structures as already seen in the previous section.

6.3. Direct vs. Iterative Local Thresholding

When comparing the iterative approach to the direct one there are almost no additional background voxels. This is to be expected because the iterative method evolves the surface based on the same local thresholds. However, 0.34% of all voxels are classified foreground with the iterative approach but not with the direct one. On the one hand the majority of those voxels are part of very small clusters. On the other hand there are also more serious differences, i.e., clusters classified as foreground by the iterative method that are classified background when using the direct approach. Visual inspection shows that most of these clusters represent actual foreground, see Figure 10.

Most of the voxels classified differently have a grey value lower than the grey values of voxels in their neighbourhood. This leads to them being constituted as background voxels by the direct method although visual inspection shows that they are actually representing solid space. The iterative approach which is based on a surface does not remove those voxels from the foreground because they are not reached by the surface.

We conclude that the iterative approach provides better results for these regions while performing as good as the direct one for others. However this comes to the cost of higher complexity in implementation and computation, e.g., one has to be careful about numerical stability when using the iterative approach. In contrast most of the operations needed for the direct approach (e.g. mean-value filtering, thresholding) can be performed by standard image processing software. Additionally, applying the direct approach takes less memory and by far less computation time. This is especially a matter during manual – and therefore labour-consuming – parameter tuning. Considering the fact that significant differences only occur at larger areas of connected foreground voxels, the direct approach might be the method of choice for data without the corresponding property.

7. Comparison to Adaptive Local Thresholding

Because the correct segmentation is unknown, it is not possible to directly quantify the quality of the proposed segmentation technique. A global thresholding is clearly not adequate (compare Figure 1), therefore we adapted a local thresholding method proposed in [37]. The method is based on Niblack's algorithm [38], which uses the local mean and standard deviation of voxels to determine local thresholds. We use the efficient implementation given in [39], adapted to 3D. This second segmentation allows us to discuss the advantages and disadvantages of both techniques.

7.1. Local Adaptive Thresholding

For a given grayscale image I, the aim is to compute a suitable threshold $T(x, y, z)$ for every voxel (x, y, z), the resulting binary image is then given by

$$B(x,y,z) = \begin{cases} 255, & \text{if } I(x,y,z) \geq T(x,y,z), \\ 0, & \text{otherwise.} \end{cases}$$

Using the method proposed in [37], the local threshold is computed considering the mean and standard deviation of the greyscale intensities in a

symmetric window centred around the voxel. The local thresholds are given by

$$T(x,y,z) = m(x,y,z) \cdot \left[1 + k \cdot \left(\frac{s(x,y,z)}{R} - 1\right)\right]$$

where $m(x,y,z)$ denotes the mean and $s(x,y,z)$ the standard deviation of grey values in the window centred around voxel (x,y,z). The window size and $k > 0$ are parameters of the algorithm. The parameter k controls the effect of the standard deviation on the threshold. R is the range of the standard deviation, i.e., for our 8-bit greyscale image it is 128. Therefore, a very high contrast results in the threshold being approximately the mean value. For a low contrast (i.e., small standard deviation) the threshold is lower than the mean value.

To apply this method to our data, we use a window of size $w_x \times w_y \times w_z$. It is useful to be able to choose the size for all three dimensions separately due to the different illumination in y-direction and the small depth in z-direction. Empirical tests showed that a size of $101 \times 101 \times 11$ voxels yields good results. For our data, the best value for the parameter k is slightly negative, i.e., $k = -0.2$, because this has the effect that in the centre of large pores less structure is (wrongly) detected. Note that the choice of parameters is a trade-off between the width or even presence of detected cell walls and wrongly detected structures located especially in large pores. Figure 11 shows a slice (for fixed z), where the local thresholds seem to work well, although it is clearly visible that the structures are in most cases too thick.

7.2. Comparison to Local Threshold Backpropagation

As mentioned above, the local adaptive thresholding proposed in [37] works surprisingly well when looking at a planar section in x-y-direction. It is hard to choose a window size because the pore sizes are in a very large range, but the local standard deviation makes it possible to reduce the misclassification of voxels located in the interior of pores larger than the window. Nonetheless, small cell walls are often lost, see Figure 12, which shows a planar section in x-z-direction. The local threshold backpropagation uses the special nature of FIB-SEM images to detect cell walls regardless of their low contrast, which is not possible for usual local thresholding techniques (or, only slightly visible structures in the background of pores would be detected, too). Furthermore, using the local threshold backpropagation, the width of the detected cell walls can be controlled directly, see Section 6. Note also

that the postprocessing and the direct/iterative improvement only refine the result of the local threshold backpropagation, i.e., they have no large effect themselves.

8. Summary and Conclusion

A new approach to automatic segmentation of FIB-SEM images has been proposed. The approach is based on the detection of disappearing structures and subsequent threshold backpropagation, where the choice of parameters has influence on which features of the 3D image are represented best. The binarisation obtained for an exemplary data set was compared to a binarisation obtained by the local thresholding method proposed in [37].

In a first step, the stack of 2D SEM images has been preprocessed, where especially the correct alignment of the 2D images is important. Because of the high porosity of our monolithic foam-like silica material, every slice also shows structures located behind the current layer. The alignment of the slices allows us to track changes of grey values in z-direction, where a drop in grey values above a certain limit d corresponds to the disappearance of a structure. Applying a linear function $x \mapsto \alpha x + \beta$ to the grey value of its last appearance yields a local threshold, which can be used to detect the first appearance of the same structure. Then, we used a dilation to close small gaps between the detected foreground voxels. This enhances the segmentation but also causes an expansion of foreground into background. This is corrected with either a direct or iterative improvement method, where local thresholds are used to reconstruct the surface. With both methods good results are obtained, yet they have different advantages. In addition to comparing the two improvement methods, we also analyse the effect of the choice of parameters. The parameter d is relatively insensitive, but α and β should be chosen according to the features that should be represented best by the segmentation.

The advantage of the algorithm proposed in the present paper is that it takes the special nature of FIB-SEM images of porous materials into account, i.e., visible structures although they are located in a different layer. In particular, when choosing a global threshold to preserve all existing cell walls, a huge amount of the cell wall that belongs to another slice is almost completely classified as foreground as well. On the other hand, if the global threshold is chosen to partially remove that cell wall, most of the other cell walls are destroyed. Existing adaptive thresholding techniques can only par-

tially overcome this problem, because small structures in the foreground are often very similar to structures located in the background, visible through large pores. Furthermore, it is very hard to control the thickness of cell walls without destroying connectivity.

Acknowledgements. The authors would like to thank Holger Kropf for performing FIB-SEM tomography measurements.

References

[1] Clyne TW, Golosnoy IO, Tan JC, Markaki AE. Porous materials for thermal management under extreme conditions. Phil Trans R Soc A 2006;64:125–46.

[2] Khanafer K, Vafai K. The role of porous media in biomedical engineering as related to magnetic resonance imaging and drug delivery. Heat Mass Transfer 2006;42:939–53.

[3] Dullien FAL. Porous Media, Fluid Transport, and Pore Structure. 2nd ed. New York: Academic Press Inc.; 1992.

[4] Kortunov P, Vasenkov S, Kärger J, Fé Elía M, Perez M, Stöcker M, et al. Diffusion in fluid catalytic cracking catalysts on various displacement scales and its role in catalytic performance. Chem Mater 2005;17:2466–74.

[5] Loeb LB. The Kinetic Theory of Gases. New York: McGraw-Hill; 1934.

[6] Lu TJ, Stone TA, Ashby MF. Heat transfer in open-cell metal foams. Acta Mater 1998;46:3619–35.

[7] Christensen CH, Norskov FK. A molecular view of heterogeneous catalysis. J of Chem Phys 2008;128:182503.

[8] Rouquerol F, Rouquerol J, Sing K. Adsorption by Powders and Porous Solids: Principles, Methodology and Applications. London: Academic Press; 1999.

[9] Fredrich JT. 3D imaging of porous media using laser scanning confocal microscowith application to microscale transport processes. Phys Chem Earth Pt A 1999;24:551–61.

[10] Baldwin CA, Sederman AJ, Mantle MD, Alexander P, Gladden LFJ. Determination and characterisation of the structure of a pore space from 3D volume images, J Colloid Interface Sci 1996;181:79–92.

[11] Manke I, Markötter H, Tötzke C, Kardjilov N, Grothausmann R, Dawson M, et al. Investigation of energy-relevant materials with synchrotron X-rays and neutrons. Adv Eng Mater 2011;13:712–29.

[12] Coker DA, Torquato S, Dunsmuir JH. Morphology and physical properties of Fontainebleau sandstone via a tomographic analysis, J Geophys Res B Solid Earth 1996;101:17497–506.

[13] Frank J. Electron Tomography, Three-Dimensional Imaging with the Transmission Electron Microscope. 1st ed. New York: Plenum Press; 1992.

[14] Yao Y, Czymmek KJ, Pazhianur R, Lenhoff AM. Three-dimensional pore structure of chromatographic adsorbents from electron tomography. Langmuir 2006;22:11148–57.

[15] Karwacki L, de Winter DAM, Aramburo LR, Lebbink MN, Post JA, et al. Architecture-dependent distribution of mesopores in steamed zeolite crystals as visualized by FIB-SEM tomography. Angew Chem Int 2011;501294–8.

[16] Holzer L, Cantoni M. Review of FIB-tomography. In: Utke I, Moshkalev SA, Russel Ph, editors. Nanofabrication Using Focused Ion and Electron Beams: Principles and Applications, Oxford: Oxford University Press; 2011, in press.

[17] Dunn DN, Hull R. Reconstruction of three-dimensional chemistry and geometry using focused ion beam microscopy. Appl Phys Lett 1999;75:3414–6.

[18] Inkson BJ, Steer T, Mobus G, Wagner T. Subsurface nanoindentation deformation of Cu-Al multilayers mapped in 3D by focused ion beam microscopy. J Microsc 2001;201:256–69.

[19] Holzer L, Indutnyi F, Gasser P, Munch B, Wegmann M. Three-dimensional analysis of porous BaTiO3 ceramics using FIB nanotomography, J Microsc 2004;216:84–95.

[20] Velichko A, Mucklich F. Quantitative 3D characterisation of graphite morphology in cast iron – correlation between processing, microstructure and properties. Int J Mater Res 2009;100:1031–7.

[21] Kato M, Ito T, Aoyama Y, Sawa K, Kaneko T, Kawase N, et al. Three-dimensional structural analysis of a block copolymer by scanning electronmicroscocombined with a focused ion beam. J Polym Sci Pt B – Polym Phys 2007;5:677–83.

[22] De Winter DAM, Schneijdenberg C, Lebbink MN, Lich B, Verkleij AJ, Drury MR, et al. Tomography of insulating biological and geological materials using focused ion beam (FIB) sectioning and low-kV BSE imaging. J Microsc 2009;233:372–83.

[23] Zils S, Timpel M, Arlt T, Wolz A, Manke I, Roth C. 3D visualization of PEMFC electrode structures using FIB nanotomography. Fuel Cells 2010;10:966–72.

[24] Goldstein JI, Newbury DE, Echlin P, Joy DC, Lyman CE, Lifshin E, et al. Scanning Electron Microscopy and X-ray Microanalysis. 3rd ed. New York: Kluwer Academic/Plenum Publishers; 2003.

[25] Giannuzzi LA, Stevie FA. Introduction to Focused Ion Beams: Instrumentation, Theory, Techniques and Practice. New York: Springer; 2005.

[26] Park KH, Kishimoto H, Kohyama A. 3D analysis of cracking behaviour under indentation in ion-irradiated β-SiC. J Electron Microsc 2004;53:511–3.

[27] Jørgensen PS, Hansen KV, Larsen R, Bowen JR. A framework for automatic segmentation in three dimensions of microstrutural tomography data. Ultramicroscopy 2010;110:216–28.

[28] Schulenburg H, Schwanitz B, Krbanjevic J, Linse N, Scherer GG, Wokaun A, et al. 3D Imaging of catalyst support corrosion in polymer electrolyte fuel cells. J Phys Chem C 2011;115:14236–43.

[29] Efford N. Digital Image Processing. New York: Addison Wesley; 2000.

[30] Blayvas I, Bruckstein A, Kimmel R. Efficient computation of adaptive threshold surfaces for image binarization. Pattern Recognit 2006;39:89–101.

[31] Thiedmann R, Hassfeld H, Stenzel O, Koster LJA, Oosterhout SD, van Bavel SS, et al. A multiscale approach to the representation of 3D images, with application to polymer solar cells. Image Anal Stereol 2011;30:19–30.

[32] Smått JH, Schunk S, Lindén M. Versatile double-templating synthesis route to silica monoliths exhibiting a multimodal hierarchical porosity. Chem Mater 2003;15:2354–61.

[33] Amatani T, Nakanishi K, Hirao K, Kodaira T. Monolithic periodic mesoporous silica with well-defined macropores, Chem Mater 2005;17:2114–9.

[34] Nakanishi K. Pore structure control of silica gels based on phase separation. J Porous Mater 1997;4:67–112.

[35] Sarjakoski T, Lammi J. Least squares matching by search, Proceedings of the XVIII isprsCongress Vienna Austria, XXXI, 1996, p. 724–28.

[36] Osher S, Fedkiw R. Level Set Methods and Dynamic Implicit Surfaces. New York: Springer; 2003.

[37] Sauvola J, Pietikainen M. Adaptive document image binarization, Pattern Recognit 2000;33:225–36.

[38] Niblack W. An Introduction to Image Processing, Englewood Cliffs: Prentice-Hall; 1986.

[39] Shafait F, Keysers D, Breuel TM. Efficient Implementation of Local Adaptive Thresholding Techniques Using Integral Images, Proceedings of Document Recognition and Retrieval XV San Jose CA USA, 2008, p. 681510.

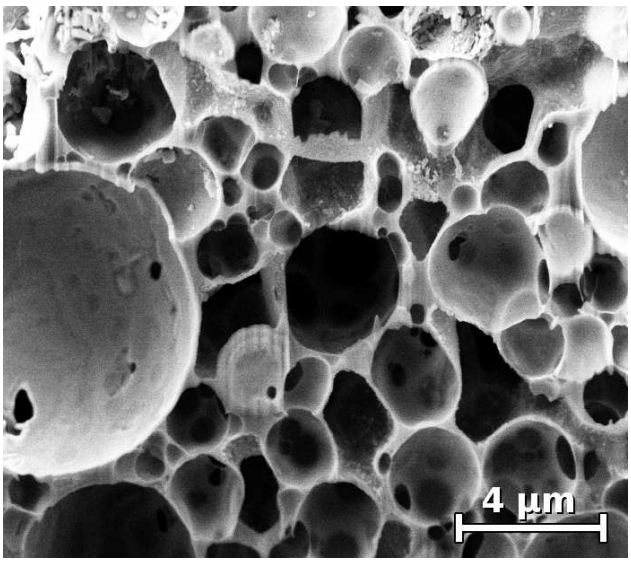

Figure 1: Example of a 2D slice acquired by SEM

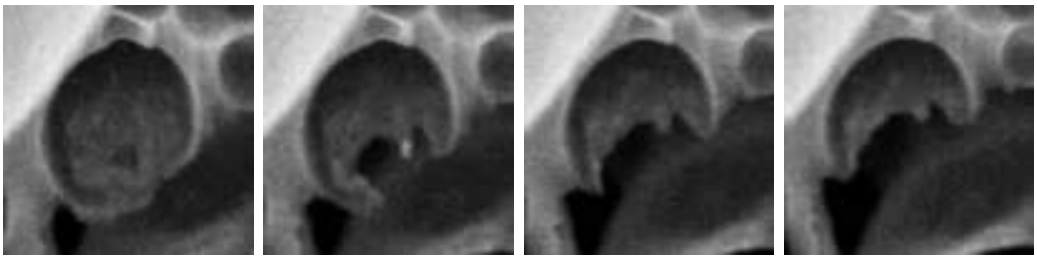

Figure 2: Sequence of images showing a small pore while being cut off by the FIB

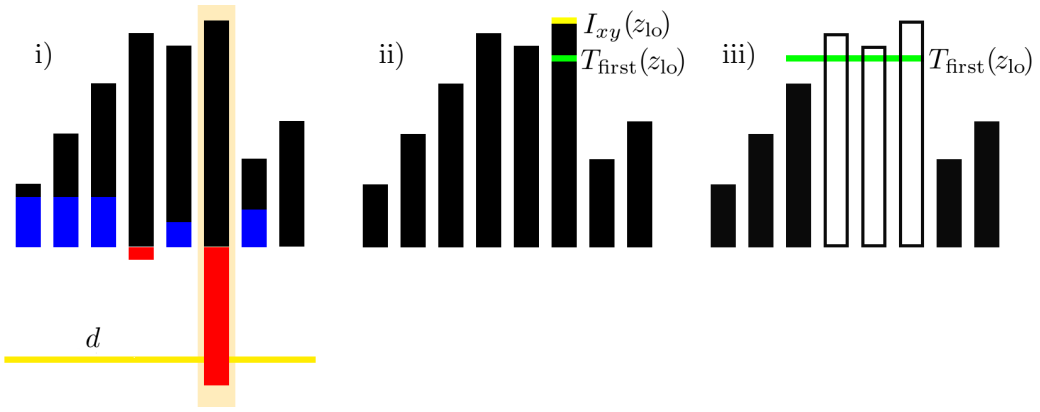

Figure 3: Schematic example of local threshold backpropagation: i) detecting last occurrence with blue and red bars representing the difference between the current and the following grey value ii) deriving a lower threshold to compensate for measurement errors iii) backpropagation that stops at the first voxel with a grey value below the threshold

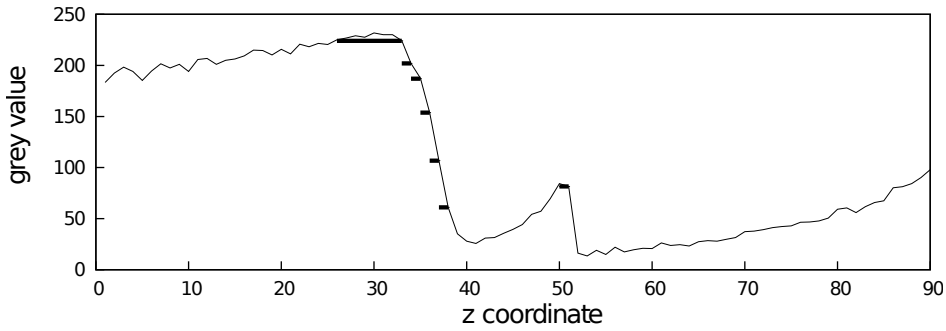

Figure 4: Example of a 1D image I_{xy}; the bold horizontal lines indicate the backpropagation

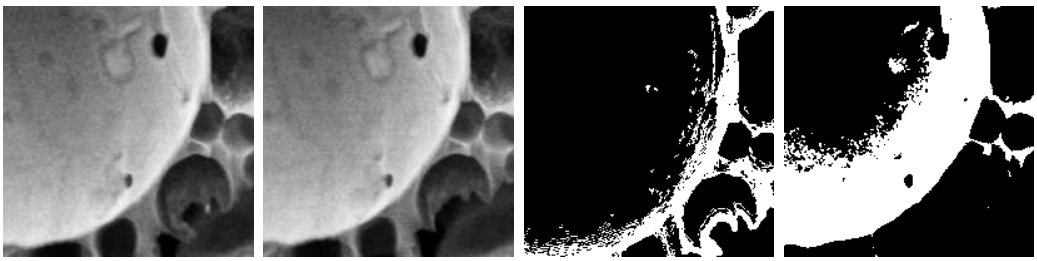

Figure 5: i) original image, ii) its successor as a reference, iii) the backpropagation result (before postprocessing), iv) a simple binarisation generated with global thresholding. Note that the threshold in iv) is not high enough to yield a reasonable segmentation of the bigger pore, however most of the smaller pores are already removed.

Figure 6: i) original image of cell wall, ii) result of backpropagation before postprocessing, iii) result after postprocessing (dilation), iv) for comparison: result after improvement by iterative local thresholding, which is presented in Section 5

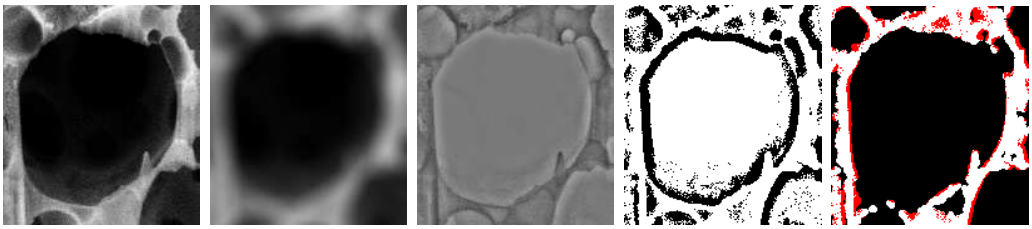

Figure 7: i) original image I, ii) mean-value filtered image M, iii) difference image $I - M$, iv) thresholded image B_{local}, where black indicates voxels that are to be removed from the foreground phase, v) shows the final result, where red voxels have been classified as foreground by local threshold backpropagation, but were removed from the foreground by local thresholding

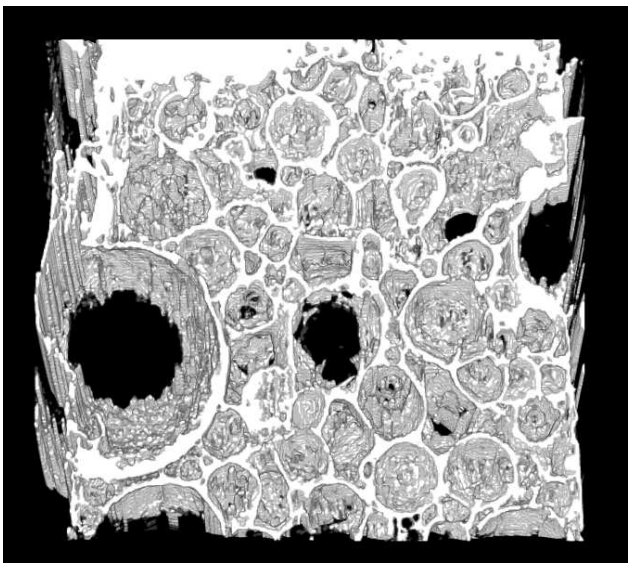

Figure 8: The final result of local threshold backpropagation followed by iterative improvement, denoted by $B_{\text{backprop/iterative}}$.

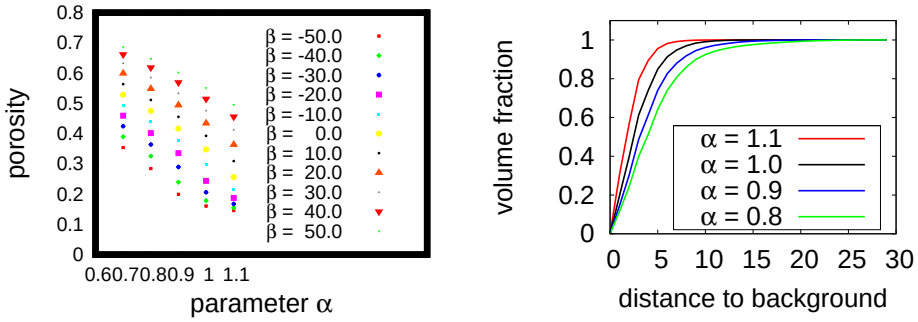

Figure 9: i) porosity for different combinations of α and β, ii) empirical spherical contact distribution function for different values of α. The value d is kept constant.

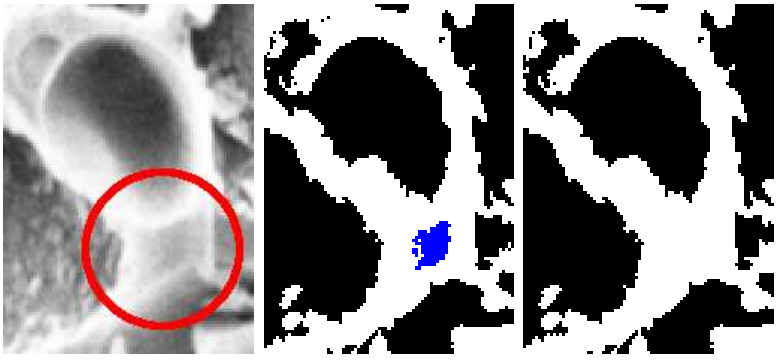

Figure 10: Example of misclassification of the surface below the big pore, marked by the red circle in i), ii) shows the result of the direct approach with blue voxels being highlighted background voxels, and iii) shows the result of the iterative approach with correct classification of that region.

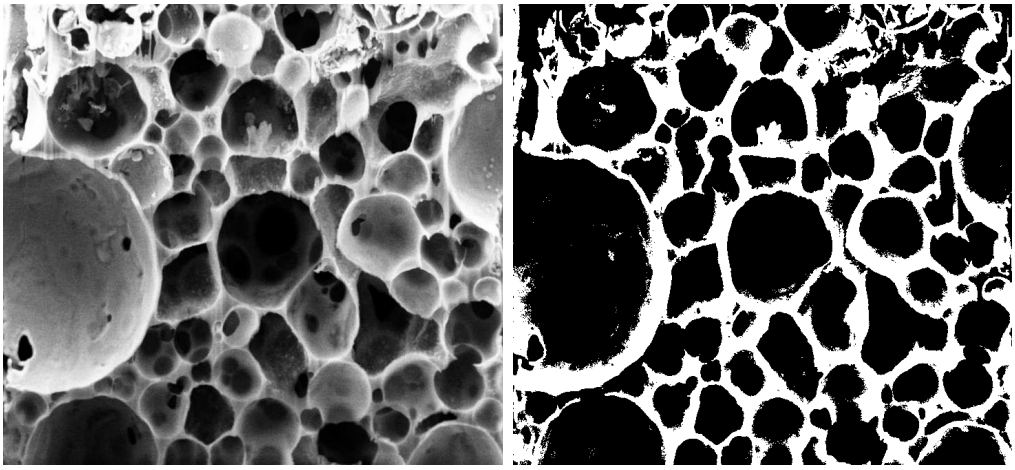

Figure 11: i) 2D slice of original image ii) binarisation by adaptive local thresholding

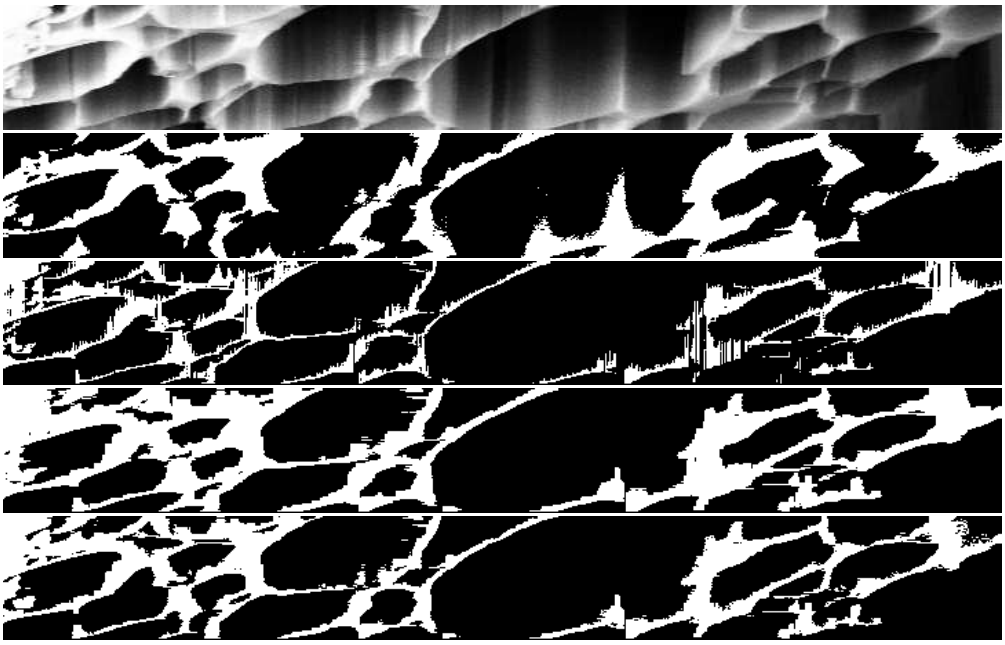

Figure 12: Cross-section in x-z-direction: i) original image after alignment, ii) binarisation by adaptive local thresholding, iii) backpropagation without postprocessing, iv) backpropagation with postprocessing, i.e., dilation, v) backpropagation with iterative improvement