

An Interactive Visual Comparison Tool for 3D Volume Datasets represented by Nonlinearly Scaled 1D Line Plots through Space-filling Curves

Johannes Weissenböck¹, Bernhard Fröhler¹, Eduard Gröller², Jonathan Sanctorum³, Jan De Beenhouwer³, Jan Sijbers³, Santhosh Ayalur Karunakaran⁴, Helmuth Hoeller⁴, Johann Kastner¹, Christoph Heinzl¹

¹University of Applied Sciences Upper Austria - Wels Campus, Stelzhamerstraße 23 4600 Wels, Austria, e-mail: johannes.weissenboeck@fh-wels.at, bernhard.froehler@fh-wels.at, johann.kastner@fh-wels.at, christoph.heinzl@fh-wels.at

²TU Wien and VRVis Research Center, Vienna, Austria, e-mail: groeller@cg.tuwien.ac.at

³imec-VisionLab, University of Antwerp, Universiteitsplein 1, Antwerp, Belgium, e-mail: jonathan.sanctorum@uantwerpen.be, jan.debeenhouwer@uantwerpen.be, jan.sijbers@uantwerpen.be

⁴FACC Operations GmbH, Fischerstraße 9, 4910 Ried im Innkreis, Austria, e-mail: s.ayalur@facc.com, h.hoeller@facc.com

Abstract

The comparison of many 3D volumes to find subtle differences is tedious, time-consuming and error-prone. Previously we presented Dynamic Volume Lines [1], a novel tool for the interactive visual analysis and comparison of ensembles of 3D volumes, which are linearized using Hilbert spacing-filling curve and represented as 1D line plots. In this paper we further demonstrate the usefulness and capabilities of our method by conducting a detailed visual analysis and evaluation of an artificial specimen from simulated 3D X-Ray Computed Tomography (XCT) and a real-world XCT titanium alloy specimen.

Keywords: Comparative visualization, visual analysis, Hilbert curve, nonlinear scaling, X-ray computed tomography

1 Introduction and Motivation

Typically, the domain experts from the field of reconstruction have to adjust the parameters of their reconstruction pipeline. The procedure of parameter adjustment leads to many volumes, which are slightly different. In order to find suitable reconstruction parameter settings, the domain experts have to compare the different volumes. During this task, regions with a high variation of intensities (e.g., feature edges on interfaces) in the ensemble of volumes are of great interest to them, as the behavior of the reconstruction algorithm can be deduced. Usually, the comparison the different volumes is based on selected 2D slices, which are positioned side by side. Typically, the intensity differences between the various volumetric reconstructions are rather small. In general, this comparison task is error-prone and time-consuming. Therefore, it is difficult for an expert to judge whether an algorithm or parameter set provides better results (e.g., sharper edges) than another one. This problem gets even worse with increasing number of volumes to compare and is further exacerbated by the fact that domain experts not only want to find regions with high variations in individual 2D slices, but also in all of the 3D volumes in the ensemble.

Working closely with domain experts in the field of reconstruction, we analyzed their workflows to design and develop a software tool that helps specialists in the comparison of 3D volume ensembles to find subtle differences. We already presented a method called Dynamic Volume Lines [1], which meets the comparison task. The core idea of it is to abstract 3D volumes as 1D line plots (see Figure 1). Using Hilbert space-filling curves [2] or simple scan line curves (i.e., traversing a volume row by row and slice by slice), all reconstructed volumes are traversed. This ensures that each traversed voxel may be compared directly with one of his neighbor voxels and each voxel position may be compared with the same voxel position in the other datasets of the ensemble. Along the space-filling curve the respective intensities are extracted and plotted in a 1D line plot, i.e., a 1D Hilbert line plot. Inspired by the work of Mindek et al. [3] and Lindow et al. [4], we present a nonlinear scaling for space-filling curves, which is based on the intensity variations in the ensemble of the compared volumes. This nonlinear scaling enables to focus on interesting areas where a lot of intensity changes are happening. This is an effective way of exploiting the available screen width. The nonlinear scaling serves in addition as basis for the interactive visualization techniques introduced. We propose an interactive histogram heatmap of the intensity frequencies as aggregation of larger areas, 1D line plots to focus on the individual reconstructions, and functional box plots [5], to give a statistical view on the local variations in the reconstructions. The nonlinear scaling draws the user's attention to important regions with high variation in the ensemble. The histogram heatmap serves as an overview visualization of the entire ensemble, which gets replaced during zooming into the plot by the detailed 1D line plots. At each level of detail, the scaling widget emphasizes the local ensemble variations by highlighting the nonlinear scaling using a sequential color-coded gradient. The interaction concept of brushing and linking is utilized to highlight voxels in regions with a high ensemble variation also in the 3D spatial context.

In this paper we investigate the usefulness of Dynamic Volume Lines by conducting a detailed visual analysis and evaluation of an artificial specimen from simulated 3D X-Ray Computed Tomography (XCT) and a real-world XCT titanium alloy specimen.

2 Methods

In this section we briefly describe the pre-processing steps and the visualization techniques: histogram heatmap, Hilbert line plots, and scaling widget. More details can be found in the work of Weissenböck et al. [1].

2.1 Dynamic Volume Lines

The main concept of Dynamic Volume Lines is to linearize 3D volumes along a space-filling curve to be able to compare many volumes through their line plots without occlusion. The resulting line plots are a familiar representation to engineers. A drawback of linearizing volumes is the loss of spatial coherency. Among the many possibilities, we use the Hilbert space-filling curve, which preserves the spatial coherency as much as possible (see Figure 1a and b). In order to linearize the 3D volumes, their three-dimensional regular grids are traversed by the respective space-filling curve. Through the traversal, the grid points can be brought into a one-dimensional linear order. An issue is that with the increasing volume size the line plots get very long. To display, e.g., 200,000 Hilbert indices as points in a 1D Hilbert line plot on a standard PC monitor, it way beyond the horizontal screen resolution, which is currently between 2000 and 3000 pixels. It is not possible to assign each Hilbert index to its own pixel column. Therefore, a nonlinear scaling of the line plots is proposed, so that important regions get the required screen space and unimportant regions are compressed. Figure 1c depicts the calculation of the nonlinear scaling based on the ensemble of volumes to compare. Since domain experts are interested in those regions where the variation of the intensities in the ensemble is high, we first need to compute the maximum local ensemble variation V_h for each index:

$$V_h = \max_{\forall m \in M_h} \text{Intensity}_h(m) - \min_{\forall m \in M_h} \text{Intensity}_h(m)$$

Here, m defines an ensemble member of a local ensemble M_h at a discrete Hilbert index h , Intensity_h defines the intensity at Hilbert index h for member m . Based on V_h , we formulate the local importance-function f_l :

$$f_l(h) = \left(\frac{V_h}{\max V_h} \right)^p$$

To be able to filter for a specific importance-value, we perform a normalization by the maximally occurring local ensemble variation. The nonlinear scaling is amplified by an exponent p , which can be adapted by the user. Setting p to zero means equal importance for all Hilbert indices. Setting it to a value larger than zero increases the importance for Hilbert indices with a high variation in the ensemble. By summing up the local importance-function values, we define the cumulative importance-function f_c , which serves as the nonlinear mapping to compress the distances between the Hilbert indices on the x-axis:

$$f_c(h) = \sum_{i=0}^h f_l(i)$$

The Dynamic Volume Lines tool provides several linked views. The main view consists of two charts, the nonlinearly scaled and the constantly scaled chart, which are vertically stacked. In between is the scaling widget, which highlights the differences between the nonlinearly scaled and the constantly scaled regions (see Figure 1f). The charts provide one of the following: an overview visualization, the histogram heatmap (see Figure 1d) and a detailed visualization, the 1D Hilbert line plots (see Figure 1e, on the left). With brushing and linking, indices in the 1D Hilbert line plots or voxels in a 3-dimensional spatial view can be selected. In addition, the 1D Hilbert line plots can be overlaid with an aggregated view, the functional box plots [5], which provide statistical properties such as the lower and upper whisker, the median, and the interquartile range (see Figure 1e on the right).

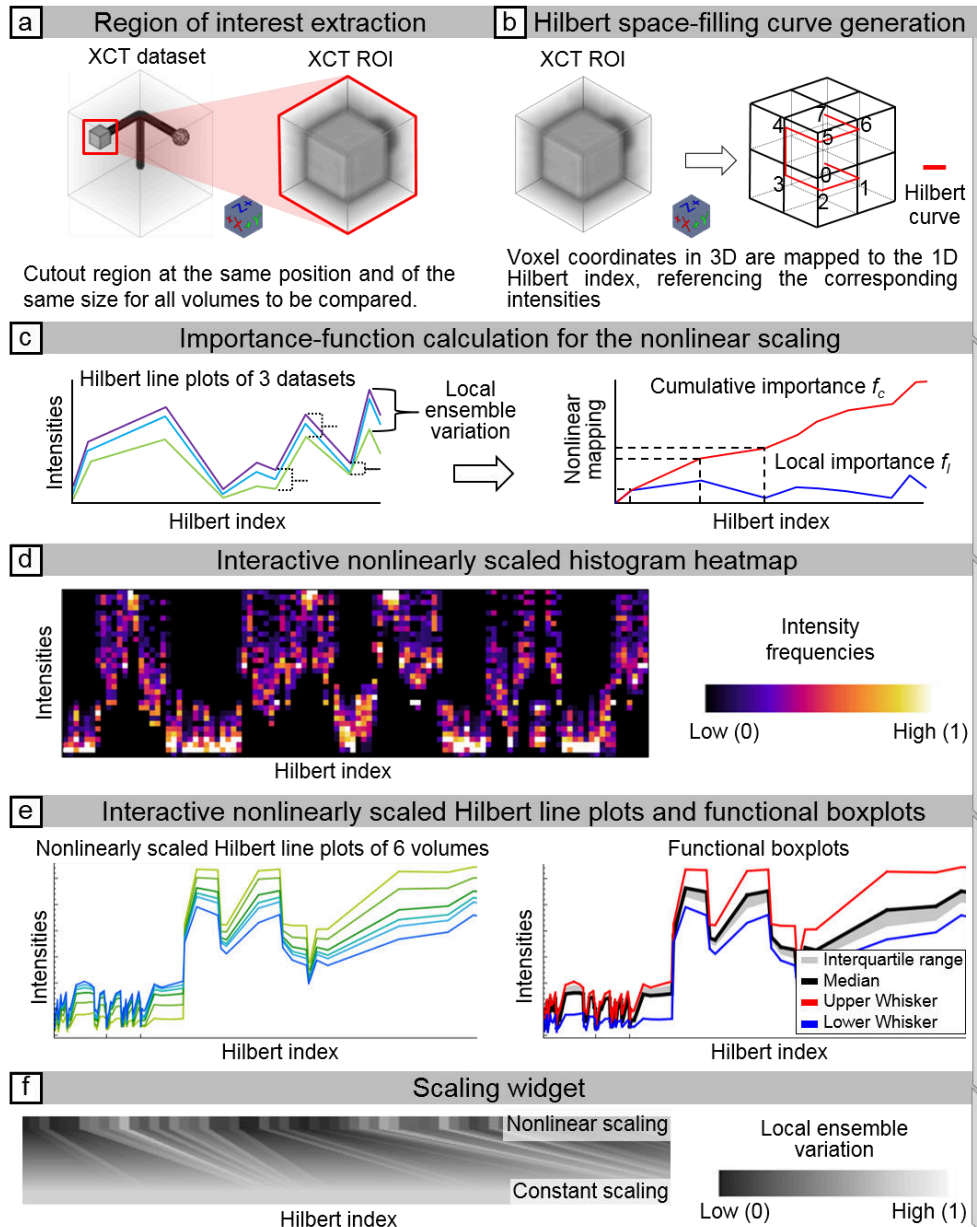


Figure 1: Workflow of the Dynamic Volume Lines: (a) extraction of interesting regions, (b) generation of a space-filling Hilbert curve, (c) nonlinear scaling of the Hilbert line plots, (d) interactive nonlinearly scaled histogram heatmap, (e) interactive nonlinearly scaled 1D Hilbert line plots, (f) and the interactive scaling widget.

2.2 Histogram Heatmap

Since for large volumes even the line plots can be overwhelming. Therefore, we provide an overview visualization as a histogram heatmap to guide the user towards interesting regions in the volume ensemble (see Figure 1d and Figure 2). To be able to fit the entire volume into a chart limited by the screen size, we split the x-axis of the nonlinearly scaled plot into intervals of equal width. For each interval we then compute a single histogram of the intensities from the enclosed Hilbert indices, over all ensemble members. All the single histograms then form the histogram heatmap, which encodes the intensity distribution in the volume ensemble. White regions denote a high concentration of intensities in a single bin. This is the case for example if the variation is low in that region. In contrast, a high variation is indicated by a broader distribution in violet, red, and orange. If intensities are below a user-defined threshold, they are considered to come from background areas. These background areas are assigned a very low importance-value and are therefore highly compressed by the nonlinear scaling.

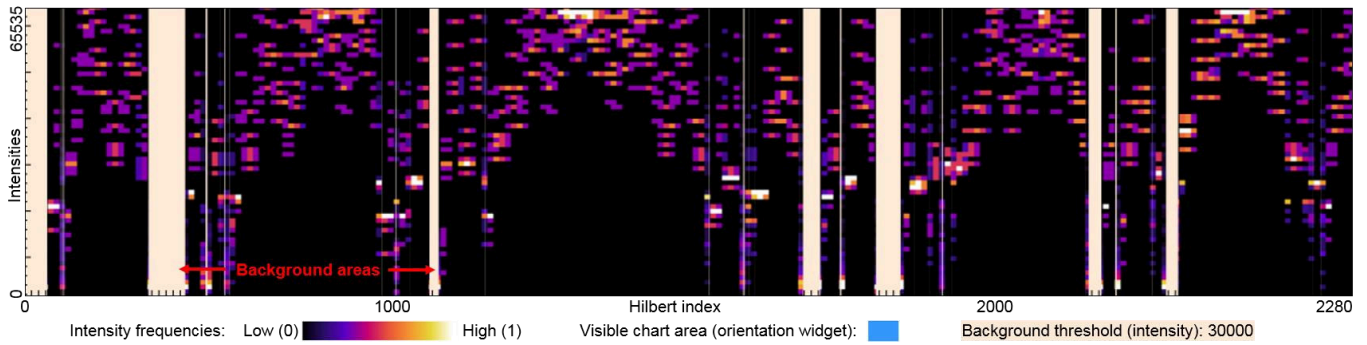


Figure 2: Histogram heatmap overview visualization of the nonlinearly scaled chart. White areas indicate a high frequency of intensities, i.e., if all ensemble members agree on a small range of intensities, where the variation in this region is low. A high variation is indicated by a broader distribution in violet, red, and yellow. Light orange background areas hide the uninteresting intensities below 30000.

2.3 1D Hilbert Line Plots and Scaling Widget

The nonlinearly scaled chart (see Figure 3a) and the constantly scaled chart (see Figure 3c) are arranged on top of each other. If the user zooms into the histogram heatmap, the 1D Hilbert line plots are automatically activated as soon as the range of currently visible Hilbert indices fits on the screen without aggregation. Each volume is represented by a 1D Hilbert line plot. The intensities are displayed on the y-axis, the Hilbert indices on the x-axis. The current mouse position is indicated by a vertical orange line in both diagrams and in the scaling widget (see Figure 3b). An additional tooltip provides the Hilbert index and the intensity at the current mouse position. The user can perform multiple selections in the charts based on the indices or in the 3D spatial view based on voxels. In addition, in the diagrams Hilbert indices can be automatically selected by specifying a user-defined importance range (degree of variation in the ensemble). The selections of voxels or indices are displayed individually for each dataset in a separate 3D rendering view. Thus, different or similar areas in the ensemble of volumes can be easily identified.

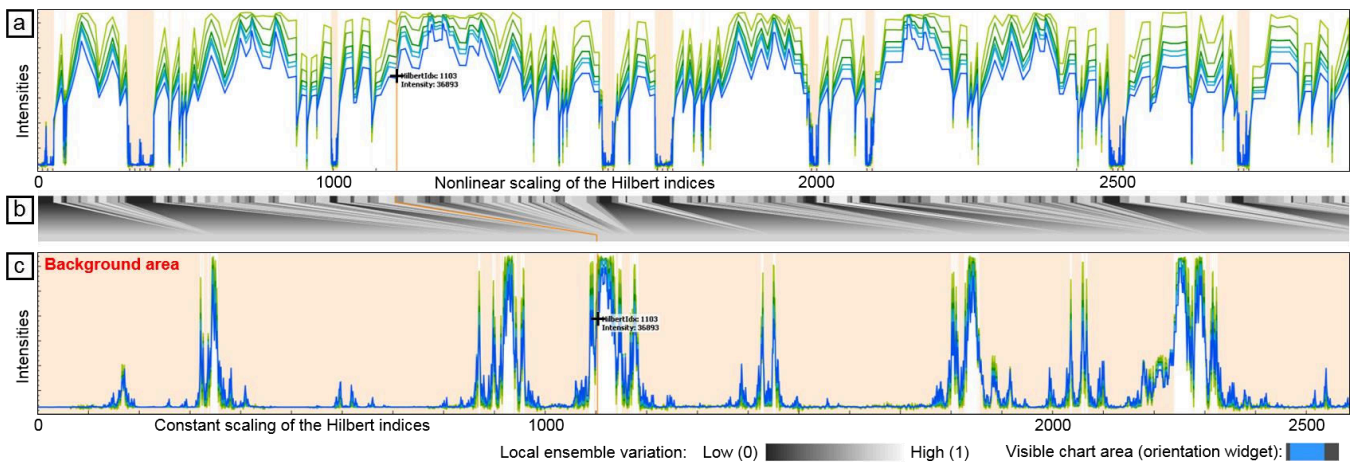


Figure 3: (a) The nonlinearly scaled chart shows the 1D Hilbert line plots differently colored for each volume. The intensities are plotted on the y-axis over the Hilbert indices on the x-axis. (b) Depicts the scaling widget, which emphasizes the nonlinear scaling (compression/expansion). (c) Shows the constantly scaled chart.

3 Results

In this section, an ensemble of six artificial XCT datasets generated by CT simulation and another ensemble consisting of three real-world XCT datasets of a titanium alloy specimen are analyzed with respect to ensemble differences using Dynamic Volume Lines.

3.1 Artificial Dataset from Simulated XCT

The first ensemble consists of six 3D datasets with different smoothing levels of an artificial specimen from simulated XCT [6]. With this artificially generated sample it is possible to demonstrate the basic functionality of Dynamic Volume Lines under well-defined conditions.

The specimen is built up by three cylindrical bars perpendicular to each other, with a sphere and a cube attached to the ends of two bars. The voxel datatype is 16 bit, meaning that intensities in the dataset range from 0 to 65535. The dimension of each dataset is $128 \times 128 \times 128$ voxels (see Figure 4a). The isometric voxel size is $0.64 \mu\text{m}$. The artificial projection images are generated

by calculating penetration lengths of primary monochromatic X-rays through the specimen. Attenuations are calculated by applying Lambert-Beer's law. The X-ray scatter and the blurring effects were disabled in the simulation. The virtual projection images are processed using the conventional reconstruction algorithm by Feldkamp et al. [7]. A Gaussian filter with increasing smoothing effect was applied to generate five additional volumes. No smoothing was applied to the first dataset. The individual variances of the Gaussian smoothing in the range [0.2, 1.0] are increased by a step of 0.2. Through preprocessing, a user-defined region of interest (ROI) is selected and applied to all ensemble volumes (see Figure 4b). The resolution of the ROIs is $16 \times 16 \times 16$ voxels. Figure 4c depicts an yz-slice of the ROI without smoothing. Figure 4d–h show yz-slices of the different results of the Gaussian smoothing filter.




Figure 4: (a) 3D reconstruction dataset of the artificial specimen from simulated XCT, (b) with an ROI of a sphere. (c) Shows the yz-slice in the center of the dataset without Gaussian smoothing. (d–h) Depict the yz-slice in the center of the five datasets with an increasing variance parameter of the Gaussian smoothing between 0.2 and 1.0.

The ROI for this analysis covers the end of the bar with the attached sphere, as shown in Figure 4a. The goal of the analysis for these datasets is to identify, above all, interesting regions that correspond to areas where most of the changes in the ensemble take place.

A filter is set for regions with high local fluctuations, for this an importance-value range between 0.5 and 1 is used. The 3D spatial view displays the respective regions in each member, as shown in Figure 5b. The voxels shown there clearly indicate that the areas with the biggest changes are in the shells of the spheres. The selection is also highlighted in the nonlinearly scaled Hilbert line plot and in the scaling widget in Figure 5c and d. The scaling widget visualizes the importance through gray-scale coding) and by the trapezoidal shapes, which are much wider at the top than at the bottom. The background threshold is set to 30000. Intensities below this threshold correspond to air and are not of interest for this analysis. The parameter p to influence the nonlinear scaling is set to 2. This setting ensures a good balance for this ensemble of datasets, emphasizing regions with high variation and leaving background regions and regions with a small variation to be just visible. In addition, the Hilbert line plots show the behavior of the Gaussian smoothing filter. The blue line, corresponding to the dataset with the strongest smoothing effect, has the lowest contrast compared to all other datasets (see Figure 5c).

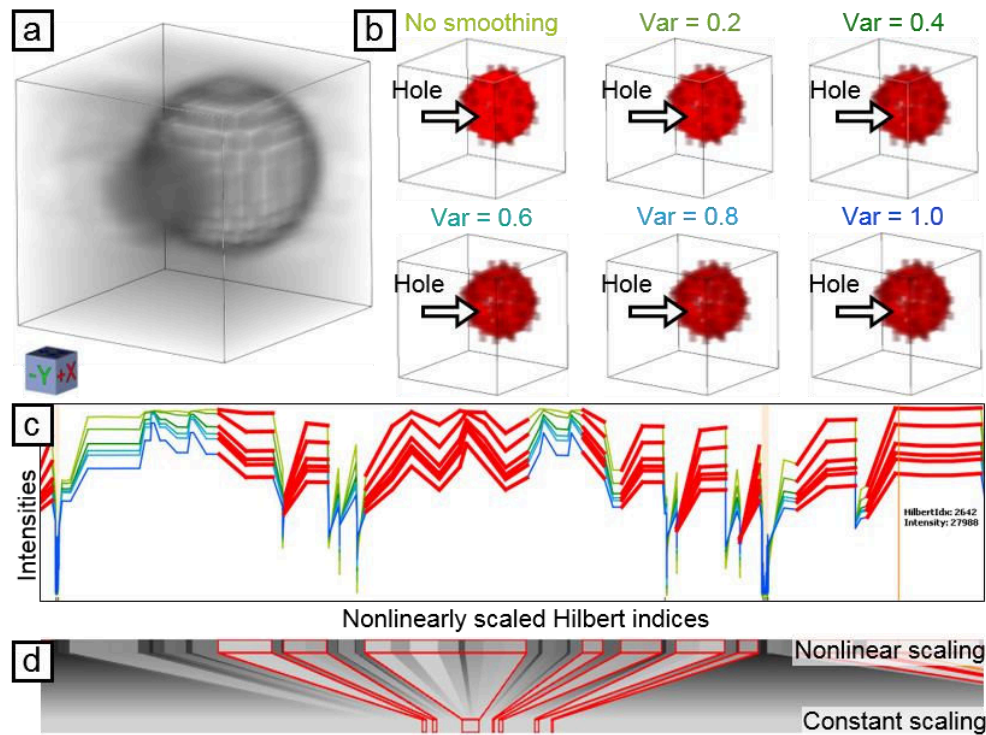


Figure 5: (a) 3D rendering of the dataset without smoothing. (b) 3D volume renderings of the selected voxels for all six ensemble members with the biggest changes in the shells of the spheres. (c) Corresponding Hilbert line plots with selected regions of high cumulative importance-values. (d) Scaling widget with the selection highlighted in red.

3.2 Real-world XCT Dataset

The second ensemble consists of three 3D reconstruction datasets of a Ti6Al4V specimen, which was additively manufactured by SLM in the Center for Smart Engineering (CSM) at the University of Applied Sciences Upper Austria (Campus Wels). A Concept Laser M2 Cusing device (Concept Laser GmbH, Lichtenfels Germany) at a laser power of 100 watt and 800 mm/s speed produced the specimen by moving in a chessboard pattern at a track width of 0.16 mm and a layer thickness of 30 μm . The Ti6Al4V powder has a grain size between 20 and 63 μm . No heat treatment was performed. The specimen was scanned with a Nanotom 180NF (GE phoenix | X-ray) sub-micro computed tomography system at an isometric voxel size of 5 μm , 1600 projections, 700 ms exposure time, 100 kV tube voltage, 210 μA tube current, 5 W target power, and a 0.2 mm copper pre-filter. A beam hardening correction was applied to all three 3D reconstructions.

Three ROIs from the 3D reconstruction datasets were selected at the interface between material (Ti6Al4V) and air (see Figure 6). The ROI datasets have a size of 64 \times 64 \times 64 voxels and the datatype of the intensities is 16 bit. No smoothing filter was applied to the original ROI dataset shown in Figure 6a. A median smoothing filter with a kernel size of 3 \times 3 \times 3 voxels was applied to the dataset depicted in Figure 6b. A non-local means smoothing filter of the Avizo software (Thermo Fisher Scientific Inc.) with the parameters: search window 30, local neighborhood 4, and similarity value 0.5, was applied to the dataset shown in Figure 6c. The non-local means filter is edge-preserving and smoothing.




Figure 6: Three 3D reconstructed ROI datasets of a Ti6Al4V specimen, depicting the interface between material and air (top: 3D rendering view, bottom: yz-slice view in the center). (a) Original ROI dataset without smoothing filter. (b) ROI dataset with a 3x3x3 median smoothing filter. (c) ROI dataset with a non-local means smoothing filter.

The goal of the analysis for these datasets, is to examine the results of the two different smoothing filters, median and non-local means, with Dynamic Volume Lines and to show the differences in their mode of operation.

The histogram heatmap overview visualization in Figure 7a shows that the datasets are split up in a material part and in an air part by the Hilbert curve traversal. The background threshold is set to zero as all intensities will be analyzed. If the user zooms into the histogram heatmap, e.g., into the regions of material, the detailed Hilbert line plots are visualized. Figure 7b shows the Hilbert line plot of the non-local means dataset. The red spikes in the region of material represent lower intensities. By selecting these red spikes in the Hilbert line plot, the affected voxels are highlighted in red in the 3D rendering view. Figure 7c depicts the highlighted red voxels, which correspond to pore positions, e.g., in the original ROI dataset. With Dynamic Volume Lines, for example, pores can be selected without also selecting the surrounding air.

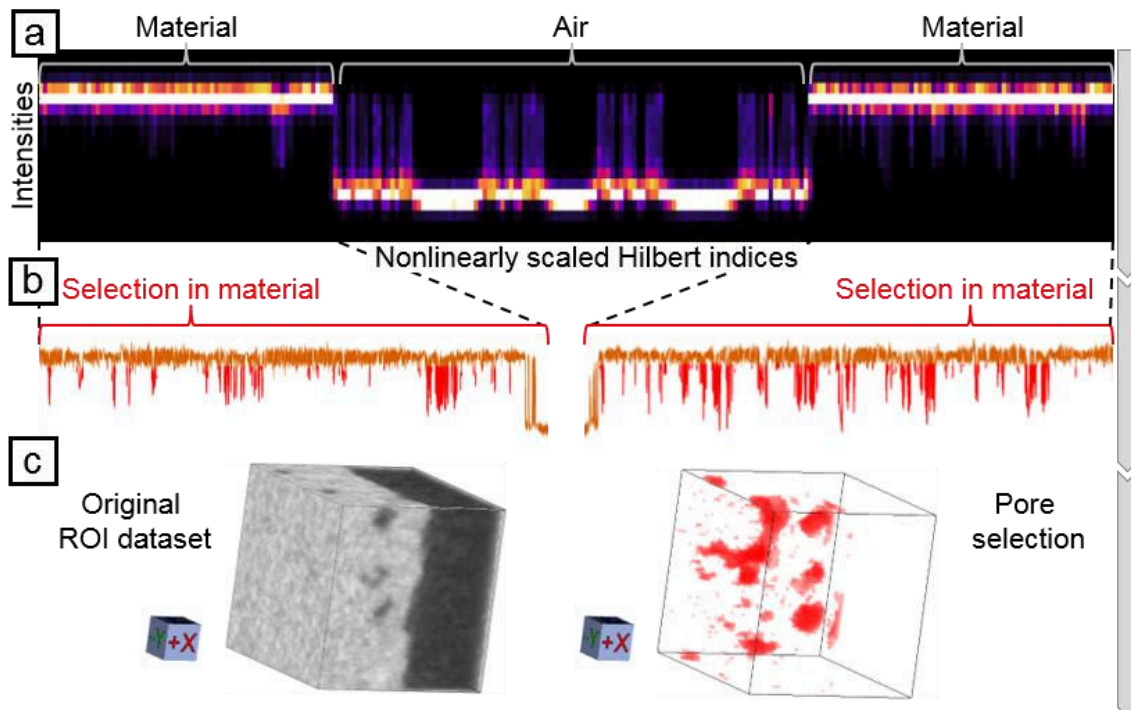


Figure 7: (a) Histogram heatmap overview visualization of the three real-world datasets. (b) Hilbert line plot of the non-local means dataset in the region of material with selected spikes in red. (c) The selected spikes correspond to pores.

The detailed visualization of the Hilbert line plots in Figure 8a shows the three real-world datasets (original, median, and non-local means). The parameter p to influence the nonlinear scaling is set to 5. Regions with a high variation within the ensemble are expanded in the Hilbert line plots and are emphasized in light gray in the scaling widget. Detail X corresponds to low

intensities (air) in the Hilbert line plots, Detail Y corresponds to high intensities (material). In both views (see Figure 8a), Detail X and Detail Y, it can be seen that the purple line, which corresponds to the original ROI dataset, has much higher intensity fluctuations (on the y-axis) than the other datasets. This is due to the higher noise content in the intensities in the nonsmoothed original ROI dataset. In comparison, the median and non-local means smoothed datasets have much lower intensity fluctuations (see Figure 8b).

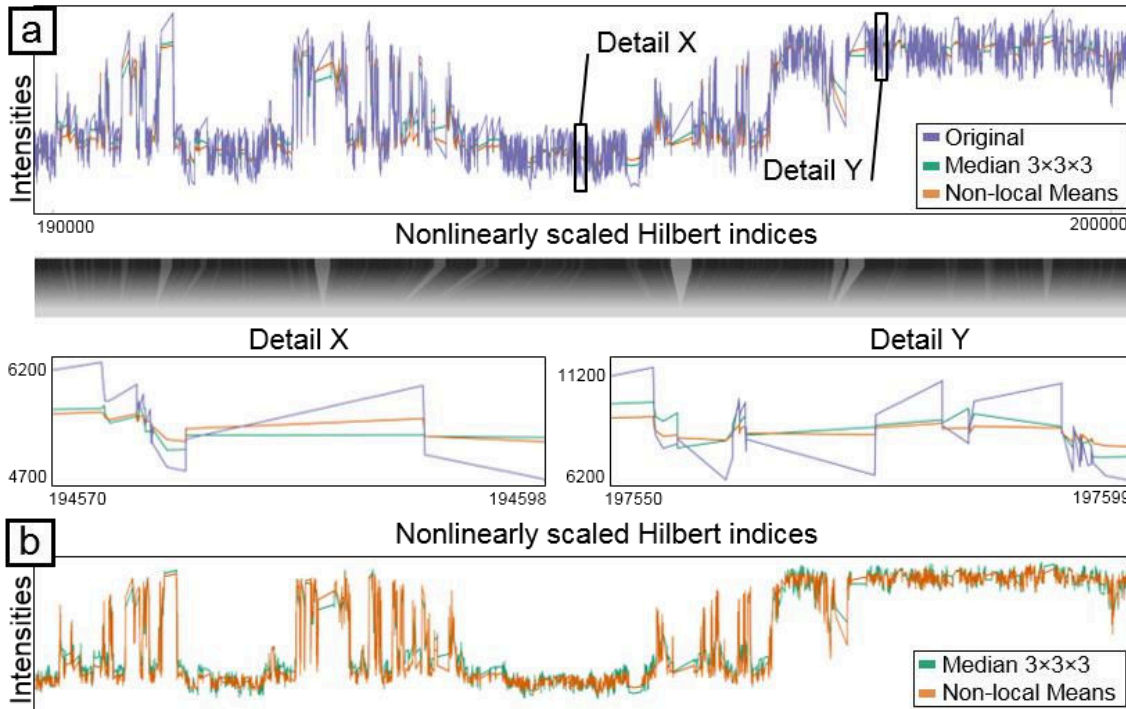


Figure 8: (a) Detailed visualization of the Hilbert line plots of the three real-world datasets and the scaling widget. Detail X corresponds to low intensities (air) in the Hilbert line plots, Detail Y to high intensities (material). (b) The median and non-local means datasets have much lower intensity fluctuations than the original ROI dataset.

In order to show how the two smoothing filters work, Dynamic Volume Lines are used to separately compare the smoothed datasets to the original ROI dataset and to analyze their Hilbert line plots. The background threshold is set to zero as all intensities will be analyzed. The parameter p to influence the nonlinear scaling is set to 1.6.

In order to find regions with relatively high intensity fluctuations (relatively high differences) between the non-local means smoothed dataset and the original ROI dataset, we have selected the cumulative importance-values between 0.3 and 1. This identifies the largest differences between the two datasets. Figure 9c depicts the red marked regions of the Hilbert line plots, which show a high variation. The spatial context of the voxels selected in the Hilbert line plots is depicted in Figure 9b, which depicts the largest differences between the non-local means smoothed dataset and the original ROI dataset. In the center of Figure 9b, there is a vertical white area reflecting the interface between material and air (see Figure 9a). This white area is due to the edge-preserving characteristic of the non-local means smoothing filter, which makes almost no changes to the intensities here. To the left and right in the immediate vicinity of the vertical white area, the data is smoothed relatively strongly. In the outer regions of the dataset, the air region (right) is smoothed less than the material region (left). This is because of the smaller variances of the intensities in the area of air compared to the area of material.

In order to find regions with relatively high intensity fluctuations between the median $3 \times 3 \times 3$ smoothed dataset and the original dataset, we have selected the cumulative importance-values between 0.3 and 1. This identifies the largest differences between the two datasets. Figure 9f indicates in the red regions of the Hilbert line plots that exhibit a high variation. The spatial context of the voxels selected in the Hilbert line plots is depicted in Figure 9e, which depicts the largest differences between the median $3 \times 3 \times 3$ smoothed dataset and the original ROI dataset. Since the median filter iterates over the entire dataset, whereby the value of the filtered voxel is set to the median of the voxels of the filter window, outliers in the intensities (within a filter window) are eliminated. Figure 9e shows a higher number of red voxels in the left half (material) than in the right half (air) (see Figure 9d). This is because of the smaller variances of the intensities in the area of air compared to the area of material.

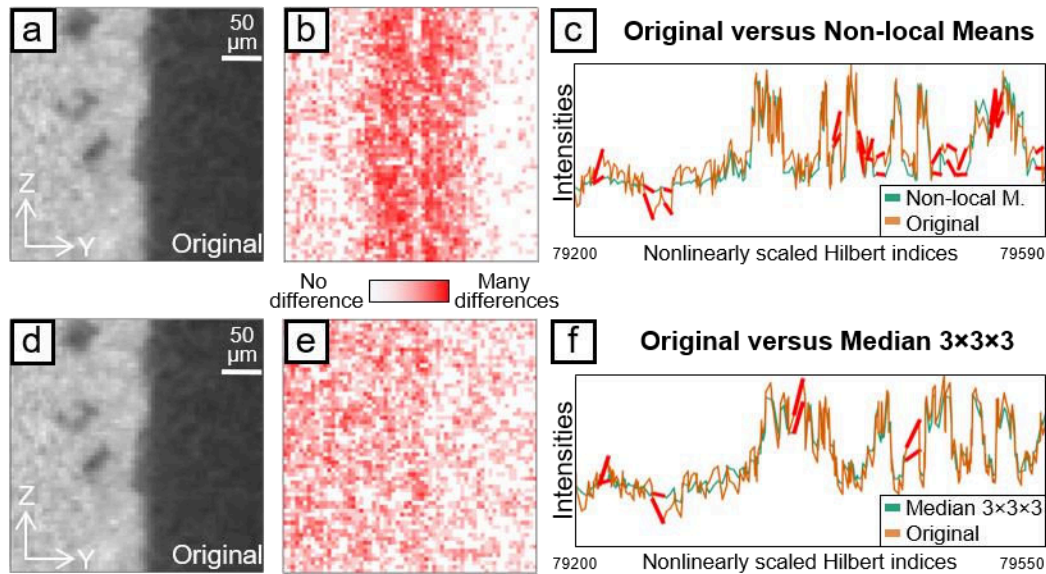


Figure 9: Comparison of the different modes of operation between the non-local means and median smoothing filter. (a) The original ROI dataset. (b) Shows the non-local means smoothing filter result of the largest differences between the smoothed dataset and the original ROI dataset. (c) The Hilbert line plots with high variations in red. (d) The original ROI dataset. (e) The median smoothing filter result of the largest differences between the smoothed dataset and the original ROI dataset. (f) The Hilbert line plots with high variations in red.

4 Conclusion

In this paper we showed that different visualization techniques like the histogram heatmap, the nonlinear scaling of the 1D Hilbert plots, and the scaling widget of the Dynamic Volume Lines approach support the domain experts in performing comparisons of many 3D volumes and allowed them to identify, e.g., small intensity differences in volume ensembles. We showed that the results of different parameterizations of a smoothing filter can be analyzed and evaluated. Furthermore, we visually demonstrated the mode of operation of the non-local means and median smoothing filters. By using the interaction and selection techniques in the 1D Hilbert line plots, we were able to segment features such as pores in a simple way.

Acknowledgements

The research has received funding from the FFG Bridge Early Stage project no. 851249: Advanced multimodal data analysis and visualization of composites based on grating interferometer micro-CT data (ADAM), the Research Foundation Flanders (FWO) and the Austrian Science Fund (FWF) under the grant numbers G0F9117N and I3261-N36 respectively: Quantitative X-ray tomography of advanced polymer composites. This work was partly written in collaboration with the VRVis Competence Center. VRVis is funded by BMVIT, BMWFW, Styria, SFG, and the Vienna Business Agency in the scope of COMET - Competence Centers for Excellent Technologies (854174), which is managed by FFG.

References

- [1] J. Weissenböck, B. Fröhler, E. Gröller, J. Kastner, C. Heinzl, Dynamic Volume Lines: Visual Comparison of 3D Volumes through Space-filling Curves, *IEEE Trans. Vis. Comput. Graph.* 25(1) (2019) 1040–1049.
- [2] C.H. Hamilton, A. Rau-Chaplin, Compact Hilbert Indices for Multi-Dimensional Data, *First International Conference on Complex, Intelligent and Software Intensive Systems*, pp. 139–146, 2007.
- [3] P. Mindek, G. Mistelbauer, E. Gröller, S. Bruckner, Data-sensitive Visual Navigation, *Computers & Graphics* 67(C) (2017) 77–85.
- [4] N. Lindow, D. Baum, H.C. Hege, Perceptually Linear Parameter Variations, *Comput. Graph. Forum*, 31(2pt4) (2012) 535–544.
- [5] Y. Sun, M.G. Genton, Functional Boxplots, *J. Comput. Graph. Stat.* 20(2) (2011) 316–334.
- [6] M. Reiter, M. Erler, C. Kuhn, C. Gusenbauer, J. Kastner, SimCT: a simulation tool for X-ray imaging. In *Proceedings of the 6th Conference on Industrial Computed Tomography*, 2016.
- [7] L.A. Feldkamp, L.C. Davis, J.W. Kress, Practical cone-beam algorithm, *J. Opt. Soc. Am. A* 1(6) (1984) 612–619.