Multimodal Transfer Functions for Talbot-Lau Grating Interferometry Data

Lucas DA CUNHA MELO[1,2], Bernhard FRÖHLER[2], Johannes WEISSENBÖCK[2],
Johann KASTNER[2], Christoph HEINZL[2]
[1]TU Wien, Vienna, Austria
[2]University of Applied Sciences Upper Austria, Wels, Austria

Contact e-mail: Lucas.daCunhaMelo@fh-wels.at

Abstract. Talbot-Lau grating interferometry X-ray computed tomography (TLGI-XCT) generates three modalities depicting different features in the scanned specimen. The combination of two or even all three modalities is often required for an in-depth analysis of the material. We present a multimodal transfer function widget for this purpose. It enables users to edit the combined transfer functions of three modalities, resulting in a customized fused image. We implement methods to arrange the individual transfer functions in two different intuitive layouts. The weighting for each modality can be easily modified and the resulting fusion image is immediately updated. We provide a combined histogram over all modalities within a triangular visualization which we call trimodal heatmap. Through this interactive analysis technique, the user can quickly gain insights into the analysed material. We show the effectiveness of our technique with the help of real-world TLGI-XCT datasets.

1. Introduction

Due to advances in X-ray computed tomography (XCT) hardware [1, 2], Talbot-Lau grating interferometry (TLGI) has recently become available in laboratory XCT devices. Such TLGI-XCT devices deliver three imaging modalities:

1. The attenuation contrast (AC), resulting from the attenuation of X-rays in the specimen, as also provided in conventional XCT devices.
2. The differential phase contrast (DPC), which contains information about the refraction of X-rays. It is highly sensitive to changes in the material, such as interfaces, delamination or fractures.
3. The dark-field contrast (DFC), which provides information about X-ray scattering. DFC shows high responses for X-ray scattering, such as those from micro-pores.

Specimens are scanned with a TLGI-XCT device in particular when the attenuation contrast from a conventional XCT device provides insufficient information. With TLGI-XCT, it is possible, for example, to obtain additional clues on structures smaller than the voxel size (provided by DFC) or on boundaries between materials with very similar attenuation coefficients (provided by DPC). Therefore, it is necessary to combine the three modalities for a comprehensive analysis.

For most investigation scenarios, it is also vital to explore the transfer function (TF) design space and the blending of the modalities. The user needs to be able to interactively modify the TFs and immediately see the resulting blended image. For a single modality, the

design of TFs is a well-addressed topic, but there are no fully established solutions for multimodal data and proposed techniques offer limited interactivity [9, 10]. Therefore, we propose a technique to interactively define multimodal TFs and blend them based on user-defined weights, specifically for the case where two or three modalities are analysed.

2. Related Work

As TFs for volume rendering are a central topic in the field of scientific visualization, being aware of the relationship with the field of materials science is essential. Heinzl and Stappen [11] identify the intersection of visual computing and materials science in a state-of-the-art report where they show how the two fields profit from each other. One of the acknowledged issues in the analysis of materials is that passive visualizations have reached their limits, emphasizing the need for interactivity.

With the lack of established solutions for interactive multimodal material analysis, it is often the case that the user's only option is side-by-side comparison. Extensions could be considered, such as checkerboard patterns or other multi-image views (e.g., as presented by Malik et al. [12]). However, these techniques applied to materials exploration would not provide as much spatial context as a fully fused view, where different modalities can intersect. Elaborate TFs can enable such views.

Several methods have been proposed to define multidimensional and multimodal TFs. In the next paragraphs, a short introduction to the methods most closely related to our approach is given. For a more comprehensive overview, referring to literature such as the state-of-the-art report on research surrounding TFs for direct volume rendering by Ljung et al. [4] is advised. They classify TFs into six categories, two of which are of particular importance for this work: dimensionality and user interfaces. These two aspects are also closely tied together, as an increased dimensionality poses a challenge in the design of a user interface: dealing with more dimensions means that there are more attributes to manipulate, which is a challenge, considering the 2D limitation of displays (e.g., computer monitors). To help illustrate the problem of dimensionality, consider the simplest type of TF (the 1D TF), which already requires two dimensions: intensity values on the horizontal axis, opacity on the vertical axis.

Kniss et al. [5] propose 2D TFs where the first and second derivatives of the scalar points are used additionally to enable a more precise voxel classification. It allows the user to manipulate regions (such as boxes and pyramids) and assign them a colour. Zhou et al. [6] combine this 2D TF technique with the 1D TF to enable up to three attributes in a 3D TF. Two attributes are selected using the 2D TF by defining regions with a lasso tool, which can be further selected using the third attribute. TFs using parallel coordinates plots for n-dimensional data such as those proposed by Zhao and Kaufman [7] and Guo et al. [8] are attractive due to the increased scalability (limited only by screen size), but at the cost that most relationships between attributes are lost, as each axis in this plot can connect to at most only two others.

Although these methods enable high-dimensional TFs to be interactively manipulated, they do not necessarily lead to satisfactory results for multimodal data, such as the data from TLGI-XCT. The additional attributes used are either extracted from the scalar volume (e.g., gradient magnitude, variance, skewness) or they use inherently multivariate data (e.g., pressure, temperature and speed from a hurricane dataset).

Lawonn et al. [9] offer an overview of multimodal visualization for medical data, where a large volume of publications is cited. However, none of them offer the level of interactivity this work aims at. Based on information theory, Bramon et al. [10] present a
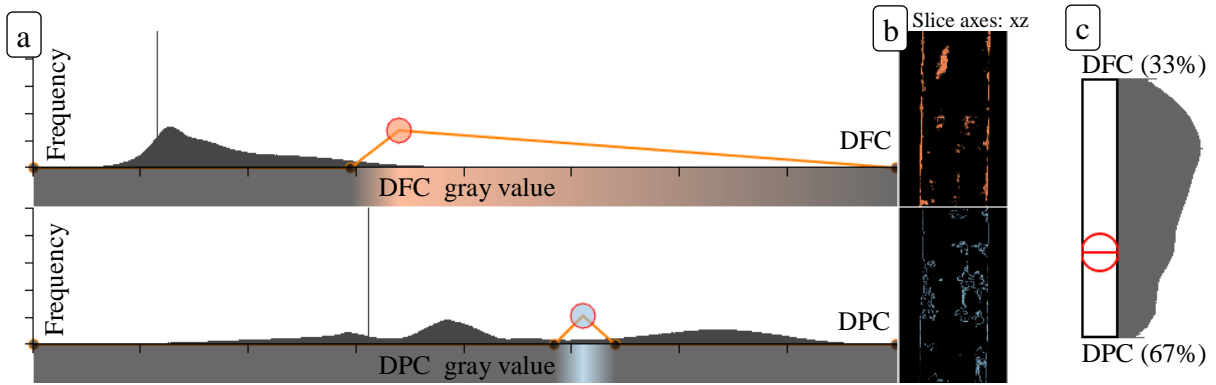
Fig. 1. Bimodal transfer function widget with two modality TFs (a), their respective modality slicers (b) and the weighting widget with co-occurrence histogram (c), which is a visualization of the dominance of grey values in the modalities. For instance: if a voxel from modality *A* has a higher value than the voxel of modality *B* with the same coordinates, this pair of voxels contributes to the histogram in a position closer to modality *A*'s end of the slider.

multimodal TF design as well, but their techniques also defy our interactivity goal, as their method can function entirely without user intervention, optionally allowing the user to make only two choices (choosing a reference image for the colour fusion and proposing a target distribution to decide which features to enhance).

A concept study for the widget proposed in this work has been presented at the 9th international conference on industrial computed tomography, where the developed tool applies the proposed TF in 2D slices only [3]. In this work, we expand that concept by applying the TF to a 3D visualization of the volume, as well as proposing a method to provide the user with visual context of the data.
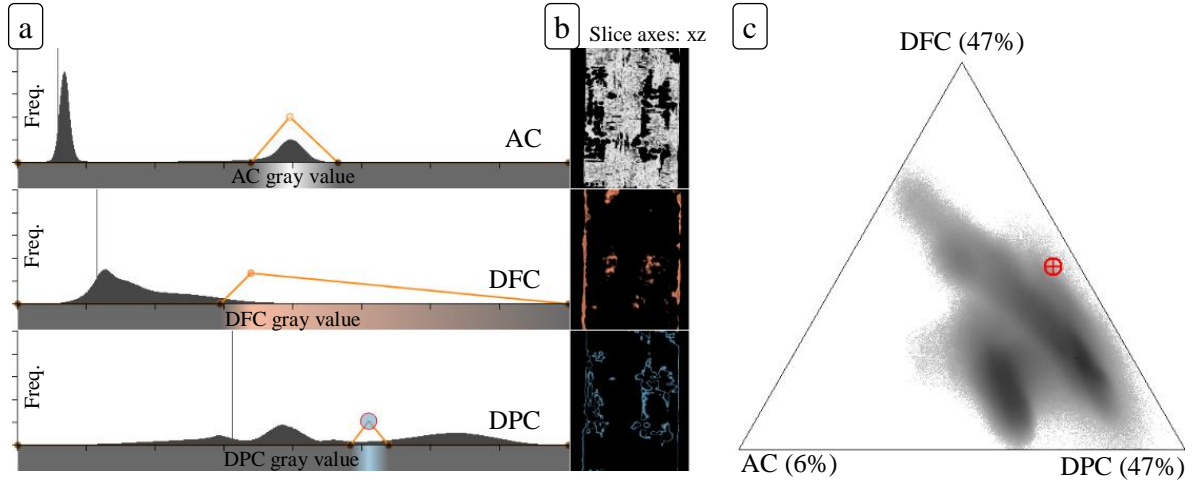
3. Multimodal Transfer Function Design

Our current multimodal transfer function designs support the analysis of two or three modalities. For both scenarios, the proposed TF definition is divided into two aspects: the first is the individual TF definition for each input modality, described in Section 3.1; the second is the weighting of each modality (and thus their corresponding TF) in the final visualization, as described in Section 3.2. Currently, all modalities are required to have the same dimensions and resolution.

3.1 Per Modality Transfer Function Definition

For the definition of each modality's individual TF, we employ a 1D TF widget, as shown in Fig. 1(a) and Fig. 2(a). It offers a histogram of binned scalar attribute data to help the user gain context of the modality's intensity distribution. The colour and opacity are set by manipulating control points to create linear piecewise functions. The colour of the TF for each intensity is shown underneath the axes of the chart; the shown colour is however made half-transparent such that the axes are always visible. Where grey is shown in Fig. 1(a), the image therefore shows black. We chose this design due to its popularity, especially outside the visualization research community [4]. In theory, however, any TF design can be used here, e.g., the 2D TF [5] referenced in Section 2.

Next to each of these TF widgets, a volume slicer is displayed, as can be seen in Fig. **1**(b) and Fig. **2**(b). Any change in the individual TFs triggers an instant update of the slicer, showing the result of the defined TF on its corresponding modality. The slicer was chosen due to its relatively cheap rendering, in contrast to a 3D visualization. Displaying either or both is, however, possible. The slice number can be changed, as well as the slice axis (XY,

Fig. 2. Trimodal transfer function widget with three modality TFs in the stacked layout (a), the respective modality slicers (b) and the weighting widget with trimodal heatmap (c).

XZ, YZ). Panning and zooming interactions are also provided. When appropriate TFs are selected, the user can switch to selecting weights such that the blended visualization fits their needs. This process is not strictly sequential; the user can switch back to refining the TFs at any time.

3.2 Weighting

The bi- and trimodal variants of our weighting widgets differ, although both are based on the same principles. For the bimodal case, a slider is provided to set the weights of the modalities, as can be seen in Fig. 1(c). On each end of the slider, a label of one modality is displayed alongside its influence. For a pair (A, B) of modalities, the slider is positioned initially halfway between A and B, resulting in the weights $(0.5, 0.5)$. Moving the slider all the way towards modality A, for instance, results in an increasingly stronger influence of modality A and finally in the weights $(1, 0)$. The slider thus controls an interpolation parameter that is used in the final renderings.

Next to the slider, we propose a histogram that provides information on the co-occurrence of the values of each modality. This histogram is calculated by first normalizing each modality (i.e., linearly mapping its intensity values to the range $[0,1]$) and then performing a voxel-wise normalization across the modalities. In other words, for two modalities A and B and each pair of voxels $a(x) \in A$ and $b(x) \in B$ with equal coordinates $x = (x, y, z)^T$, the resulting voxels are $a'(x) = \frac{a(x)}{a(x)+b(x)}$ and $b'(x) = \frac{b(x)}{a(x)+b(x)}$. Note that $a'(x) = 1 - b'(x)$, which means only the value for a' needs to be computed and stored in a new volume, the value of b' is computed dynamically when required.

The resulting volume is then used to draw the histogram next to the slider. A high frequency towards modality A would mean that there is a large number of voxels where $a'(x) > b'(x)$, whereas a high frequency halfway between A and B means there is a large number of voxels where $a'(x) \approx b'(x)$. In Fig. 1(c), for instance, the higher values towards the DFC modality indicate that a larger weight for the DPC modality may be necessary, hence the chosen weights of 33% for DFC and 67% for DPC.

For the trimodal case, a triangular weighting widget is provided instead of a slider. The previously mentioned slider can be interpreted as a means of manipulating one interpolation parameter t based on the handle position on the slider axis, which is used to weight two modalities. Analogously, a handle inside the triangle's area is used to manipulate two interpolation parameters α and β, defined as the barycentric coordinates of the handle (with $\gamma = 1 - \alpha - \beta$).
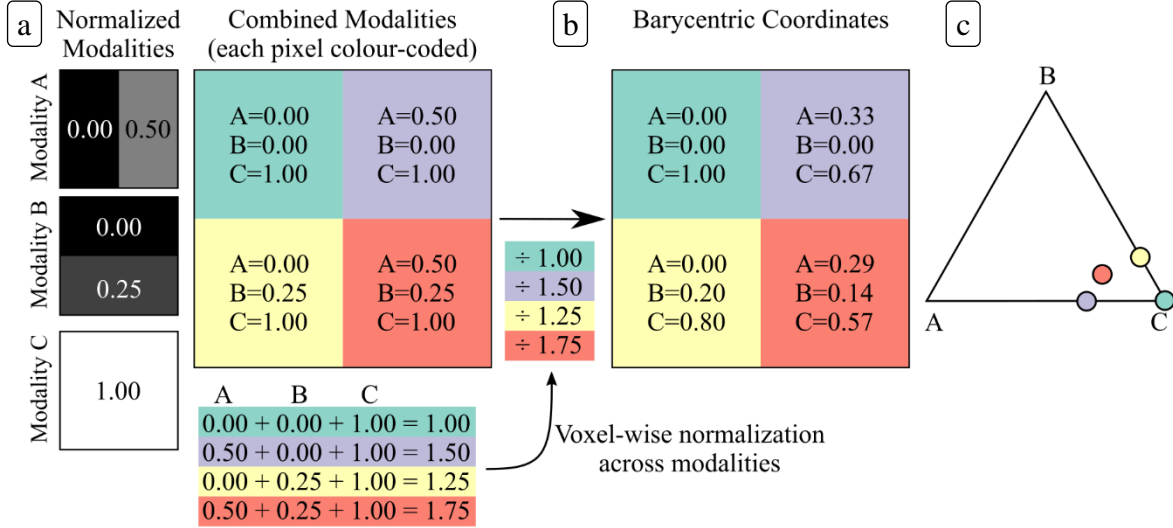
Fig. 3. The trimodal heatmap computation on the example of three 2x2 modality images. First, each modality image is normalized separately to values in the range [0,1] (a). Then for each pixel/voxel a normalization across the modalities takes place (b). The resulting triple is interpreted as barycentric coordinates, the frequency at the corresponding location in the triangle is increased (c). For the trimodal heatmap, each location in the triangle is colour-coded by its frequency, a value of zero is coloured white, while the highest occurring frequency is coloured black; frequencies in between are coloured in an interpolated grey tone.

In the triangle's area, we propose a visualization of the correlation of the data analogous to the one provided for bimodal case, which we call trimodal heatmap, as shown in Fig. 2(c). A graphical representation of how this visualization is computed is shown in Fig. 3. Analogously to the slider's histogram, the three modalities are first normalized, then fused, as shown in Fig. 3(a). Let A, B and C be the modalities. For each voxel triple $a(x) \in A, b(x) \in B$ and $c(x) \in C$ with equal coordinates $x = (x, y, z)^T$, the resulting voxels are $a'(x) = \frac{a(x)}{a(x)+b(x)+c(x)}$, $b'(x) = \frac{b(x)}{a(x)+b(x)+c(x)}$ and $c'(x) = \frac{c(x)}{a(x)+b(x)+c(x)}$, i.e. the values are normalized across modalities such that their sum equals 1, as shown in Fig. 3(b). Again, note that $c(x) = 1 - a(x) - b(x)$, meaning that only two of the volumes need to be calculated and stored, as the third volume can be calculated dynamically based on them.

The resulting values are then interpreted as barycentric coordinates and visualized in a trimodal heatmap on the triangular widget, as Fig. 3(c) illustrates. High frequencies towards modality A_1 would mean that there is a large number of voxels where $a(x) > b(x)$ and $a(x) > c(x)$, whereas high frequencies near the triangle's barycentre means there is a large number of voxels with $a(x) \approx b(x) \approx c(x)$. High frequencies near the edge between, for example, modalities A_1 and A_2 would mean that there is a large number of voxels with $c(x) < a(x)$ and $c(x) < b(x)$.

3.3 Trimodal Transfer Function Layouts

The layout presented in Fig. 2, which we call the "stacked" layout (as TFs are positioned on top of each other), occupies an ample horizontal space on the screen. As a more compact alternative, we propose the "triangle" layout, shown in Fig. 4. In this layout, the TF widgets are positioned on the edges of a triangle, and the weighting triangle is rotated so that its corners (each corresponding to a modality) touch their respective TF widget. That results in an intuitive manipulation of the handle on the weighting triangle: the closer the handle is placed to a TF widget (i.e., a vertex of the triangle), the more influence its modality will have in the final visualizations.

This layout leaves available space on the corners of the triangle created by the TF widgets. We use this space to display the modality slicers, previously positioned to the right
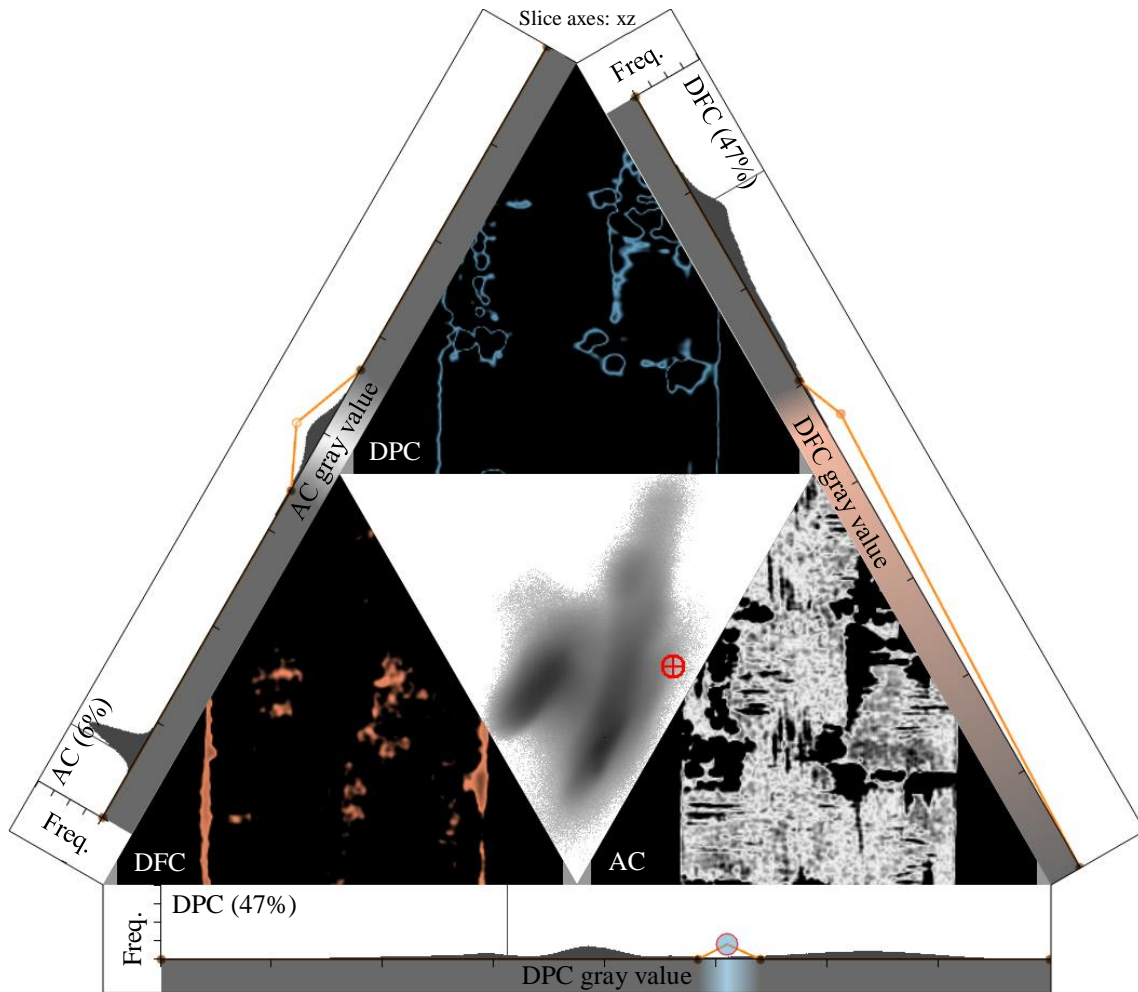
Fig. 4. Trimodal transfer function widget in its triangular layout.

of each TF widget. The slicers are placed on the triangle region on the opposite side of their corresponding TF widget.

4. Evaluation

The material scientists we worked with expressed a preference over the stacked layout, as, with the triangle layout, the interpretation of histograms becomes more difficult due to their differing rotations. In the implemented tool, the user can dynamically switch between the two layouts.

We evaluate our technique on two datasets; both scanned with a Bruker SkyScan 1294. The device delivers pre-registered datasets of the three imaging modalities AC, DPC and DFC. In Section 4.1, an open cell polyurethane foam is explored, followed by a carbon fibre reinforced polymer in Section 4.2.

4.1 Case Study: Foam data

We used our technique to explore a 6 mm x 6 mm x 6 mm open-cell polyurethane (PU) foam. This specimen was measured in our TLGI-XCT device at a tube voltage of 35kV with 1mA current, acquiring 1800 projection images using an Aluminium filter of thickness 0.25 mm with 4 phase steps. The measurement time was approximately 19 hours. On the detector side,
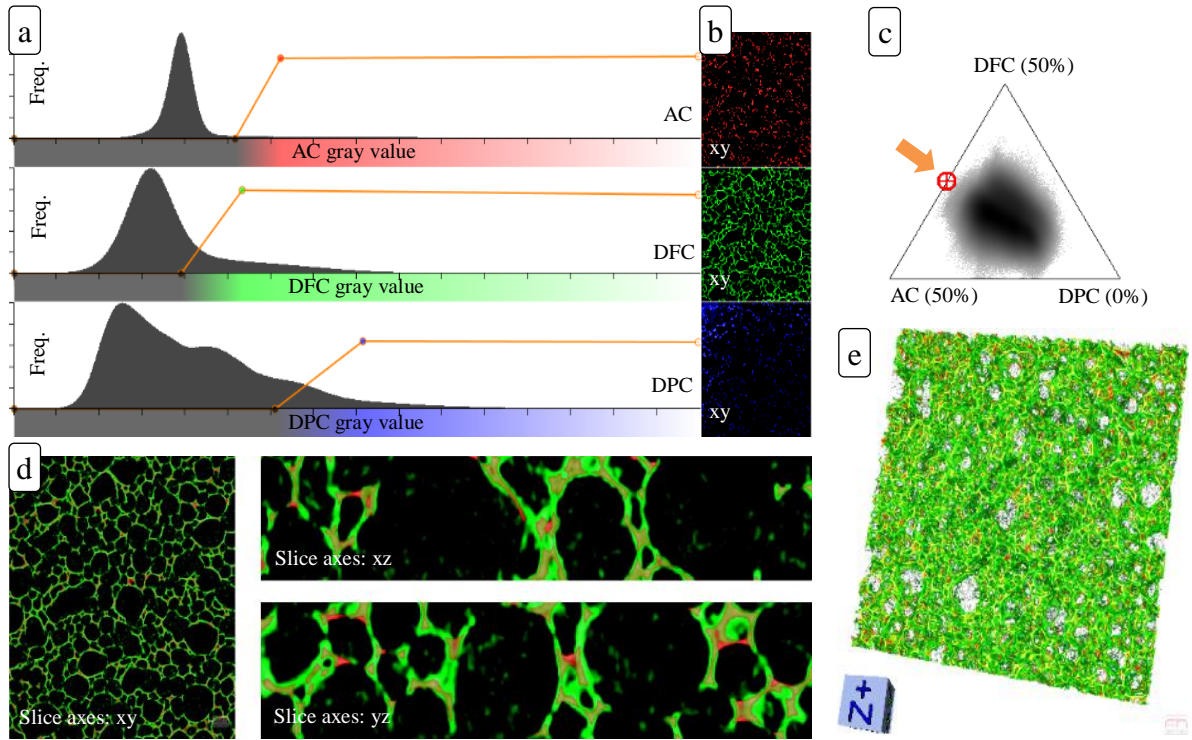
Fig. 5. A polyurethane foam dataset is analysed with the proposed technique. The individual TFs for each modality (a) are shown next to a slice image of each modality (b). Blending is controlled by the control point (marked with an orange arrow) in a triangular widget (c) which also contains the trimodal heatmap. Blended images are shown in axis-aligned slicer images (d) and in a 3D view (e). The slices with axes *xz* and *yz* were zoomed-in to make details better visible.

five-fold averaging and 2x2 binning at an integration time of 1440 ms was used, resulting in an effective voxel size of 11.4 μm.

The task, in this case, was to detect the foam cell walls. To some degree, the cell walls are visible in all modalities. In the AC image, however, visualized in red in Fig. 5, they are visible only partially. As can be seen in the enlarged slicer images showing the blended images Fig. 5(e), the DFC image, shown in green, shows much more signal on the cell wall structures. The DPC has an inhomogeneous distribution of grey values, as can be noticed through the visible wall structures in blue in the upper left corner of the DPC image in Fig. 5 (b). Therefore, we chose to assign a weight of 50% to both AC and DFC for the blended visualization. We see that the DFC modality can resolve the cell wall structures much better in comparison to the AC modality. This comparison through superposition helped the material experts gain insight into how well the foam structures can be recognized in each modality for this scan. The fusion by weighting made the early detection of the low relevance of the DPC modality possible, and its influence was decreased accordingly.

4.2 Case Study: CFRP data

The second analysed specimen is a carbon fibre reinforced polymer (CFRP) laminate. In this case, the specimen was approximately 1.4 cm x 1.8 cm x 0.3 cm. It was measured in our TLGI-XCT device at a tube voltage of 35kV with 1mA current. 1499 projection images were taken using an Aluminium filter of thickness 0.25 mm with 4 phase steps. Here, we employed six-fold averaging and 4x4 binning at an integration time of 650 ms, resulting in an effective voxel size of 22.8 μm. A typical task in the analysis of CFRP materials is identifying fibres and pores, which is set as the goal of this case study. For that, we use the proposed bimodal widget, as the DPC and DFC modalities contain the sought information. The DFC image contains high responses for fibre structures and the DPC image, for the hulls of pores.
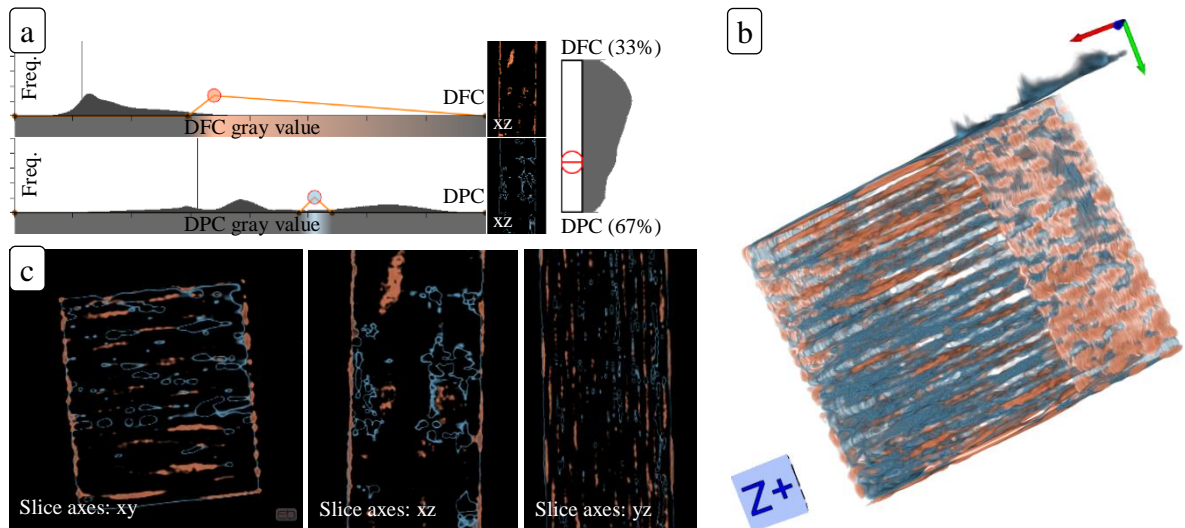
Fig. 6. A carbon fibre reinforced polymer is analysed with the proposed tool. The individual TFs of each modality are positioned next to their respective slicer and the interpolation slider (a). These components are configured exactly as in Fig. 1. The blended results are shown in two types of visualizations: a 3D view (b) and axis-aligned slicer images (c).

The colours are defined in the TF in Fig. 6(a) and each of the resulting visualizations in Fig. 6(b-c) displays the two features in a combined view. These results accelerate and facilitate the comparison of structures from different modalities, otherwise only possible with a side-by-side comparison.

5. Discussion and Future Work

Material experts working with our technique voiced very positively about the interactivity of our approach. Especially the separate 2D slicers of each modality together with the simultaneous blended visualization can tremendously improve the analysis through their quick feedback.

However, the experts remarked that determining a suitable TF for single modalities remains a difficult task. Therefore, in future work, we plan to incorporate additional guidance for defining these TFs, as well as methods of providing context to the multimodal data, additionally to the trimodal heatmap.

The first developed widget [3] was also limited to exactly three modalities. The experts expressed a wish to extend this to other numbers of modalities, which we have begun to address with the bimodal widget. We plan to address the exploration of a larger number of modalities in future work.

6. Conclusion

We propose a technique for the interactive analysis of TFs for multimodal datasets. For each modality, an individual TF can be specified, and its result is visualized immediately. The modalities can be weighted individually to arrive at a blended image that merges the required information from all modalities.

The two presented cases demonstrate that the introduction of a weighting component (in the presented widgets: a slider and triangle) is sufficient to accelerate and facilitate the exploration of the multiple modalities of Talbot-Lau grating interferometry X-ray computed tomography datasets. Overall, the experts noted that our tool is much easier to use for the

task of defining and manipulating multiple transfer functions and the blending of multimodal datasets in comparison to more generalized visualization applications.

Acknowledgements

References

[1] F. Pfeiffer, T. Weitkamp, O. Bunk, C. David, Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources, Nature Physics, 2 (2006) 258–261.

[2] F. Pfeiffer, O. Bunk, C. David, M. Bech, G. L. Duc, A. Bravin, P. Cloetens, High-resolution brain tumor visualization using three-dimensional X-ray phase contrast tomography, Physics in Medicine and Biology, 52 (23) (2007) 6923–6930.

[3] B. Fröhler, L. da Cunha Melo, J. Weissenböck, J. Kastner, T. Möller, H.-C. Hege, E. Gröller, J. Sanctorum, J. De Beenhouwer, J. Sijbers, C. Heinzl, Tools for the Analysis of Datasets from X-Ray Computed Tomography based on Talbot-Lau Grating Interferometry, Proceedings Industrial Computed Tomography Conference (2019), Padua, Italy.

[4] P. Ljung, J. Krüger, M. E. Gröller, M. Hadwiger, C. Hansen, A. Ynnerman, State of the Art in Transfer Functions for Direct Volume Rendering, Computer Graphics Forum, 35 (3) (2016) 669–691.

[5] J. M. Kniss, G. L. Kindlmann, C. D. Hansen, Multi-Dimensional Transfer Functions for Interactive Volume Rendering, IEEE Transactions on Visualization and Computer Graphics, 8 (2002) 270–285.

[6] L. Zhou, M. Schott, C. D. Hansen, Transfer function combinations, Computers & Graphics, 36 (2012) 596–606.

[7] X. Zhao, A. E. Kaufman, Multi-dimensional Reduction and Transfer Function Design using Parallel Coordinates, Volume Graphics, (2010) 69–76.

[8] H. Guo, H. Xiao, X. Yuan, Scalable Multivariate Volume Visualization and Analysis Based on Dimension Projection and Parallel Coordinates, IEEE Transactions on Visualization and Computer Graphics, 18 (9) (2012) 1397–1410.

[9] K. Lawonn, N. N. Smit, K. Bühler, B. Preim, A Survey on Multimodal Medical Data Visualization, Computer Graphics Forum, 37 (1) (2018) 413–438.

[10] R. Bramon, M. Ruiz, A. Bardera, I. Boada, M. Feixas, M. Sbert, Information Theory-Based Automatic Multimodal Transfer Function Design, IEEE Journal of Biomedical and Health Informatics, 17 (2013) 870–880.

[11] C. Heinzl, S. Stappen, STAR: Visual Computing in Materials Science, Computer Graphics Forum, 36 (3) (2017) 647-666.

[12] M. M. Malik, C. Heinzl, M. E. Gröller, Comparative Visualization for Parameter Studies of Dataset Series, IEEE Transactions on Visualization and Computer Graphics, 16 (5) (2010), 829-840.