Parameter-Space Exploration for Computed Tomography Image Analysis Algorithms

Bernhard Fröhler[1], Christoph Heinzl[1], Johann Kastner[1], Torsten Möller[2]

[1]University of Applied Sciences, Campus Wels Upper Austria, Research Group Computed Tomography, Stelzhamer Straße 23, 4600 Wels, Austria, e-mail: bernhard.froehler@fh-wels.at, christoph.heinzl@fh-wels.at, johann.kastner@fh-wels.at

[2]University of Vienna, Faculty of Computer Science, Research Group for Visualization and Data Analysis, Währinger Straße 29, 1090 Wien, Austria, e-mail: torsten.moeller@univie.ac.at

Abstract

We present a tool which supports computed tomography specialists in exploring the parameter and result space of volume segmentation pipelines. In a preprocessing step, the user samples the parameter space of a number of pipelines to be analysed. Our sampling tool supports arbitrary pipelines that can be parametrized via the command line. The results from sampling the parameter spaces of these pipelines are clustered by image similarity. The user can then explore the space of possible results in an interactive visualization tool. This visualization tool highlights the variation resulting from different parametrizations and pipelines, and allows exploring the influence of each parameter. We evaluate our tool on the multimodal segmentation of a carbon-fiber-reinforced polymer specimen with impact damage, scanned with a Talbot-Lau grating interferometer computed tomography device.

Keywords: image processing, parameter space exploration, visualization, multimodal segmentation

1 Introduction

Image segmentation and feature extraction play a very important role in the investigation of data acquired by computed tomography devices. Oftentimes it is not obvious from the start which pipeline will deliver the most suitable results. In addition, analysis pipelines often consist of multiple steps, and might need to consider multiple imaging modalities. This leads to pipelines with multiple parameters which all require optimization to achieve suitable results. There is therefore a demand for tools that alleviate comparing different pipelines, and analyzing the space of achievable results with the defined pipelines and the influence of each parameter on the result. In this work, we will present a tool that supports such analysis.

Specifically, we explored the analysis of impact damage in a carbon-fiber-reinforced polymer (CFRP) object as a use case. The attenuation contrast of conventional computed tomography can only deliver information on larger crack structures, therefore data from the so-called dark field image from a Talbot-Lau grating interferometer computed tomography (TLGI-CT) device is used as well, which can provide information on small fractures. For such kind of multimodal analysis, no standardized segmentation methods exist yet, thus it was necessary to conduct experiments to find a suitable segmentation pipeline.

The main contribution of this work are methods for the combined analysis of multiple algorithm outcomes in a single tool. Further contributions include the evaluation of this tool with different segmentation pipelines parameterizations, and a pipeline specially optimized for segmenting crack damage in a carbon-fiber-reinforced polymer.

The paper is organized as follows: Section 2 puts our work in relation to similar methods existing in literature. Section 3 describes the dataset that was used for evaluating our tool, while section 4 reports on the segmentation pipelines that we used in our evaluation. Section 5 details the methods for comparing and analyzing segmentation pipelines we implemented in our tool, and section 6 presents our findings in the analysis conducted with that tool. We conclude in section 7 with a summary of the advantages of our tool in comparison to traditional methods.

2 Related Work

A typical approach for segmenting industrial computed tomography is through global thresholding methods. These are however not able to properly segment very complex modalities. Thus, more elaborate segmentation algorithms are required, which come at the cost of more parameters to be tuned. Systematically exploring the parameter space of algorithms has become an important research topic in the visualization community over the last few years. Sedlmair et al. [2] present a recent overview and taxonomy of the approaches in the literature. There is a large body of work on exploring parameter spaces of specific image analysis pipelines, such as the VisTrails tool by Bavoil et al. [3], which provides a spreadsheet-like interface allowing an interactive exploration. For complex pipelines, where the required computation time can range up to several minutes or hours, VisTrails on-demand execution approach is not feasible. Tuner by Torsney-Weir et al. [4], a tool geared towards image segmentation, overcomes this limitation by sampling the parameter space in a preprocessing step. It requires an objective quality measure though, for example through an available ground truth, which we do not require in our approach. Other tools such as Paramorama2 [5] arrange labelled images resulting from sampling with different parameter sets in clusters and support their comparison. While they require using uniform sampling on a fixed Cartesian grid, which does not scale well, we can employ any sampling method. In addition, they cluster by parameter values whereas we cluster by result similarity. In material sciences,

the PorosityAnalyzer tool [6] enables users to compare and evaluate segmentation pipelines in the context of porosity analysis. In our GEMSe tool [7], we previously investigated the exploration of a single fixed segmentation pipeline for computed tomography datasets as well as data from other sources. In this work, we extended the approach of the GEMSe tool for the comparison and analysis of multiple arbitrary segmentation pipelines.
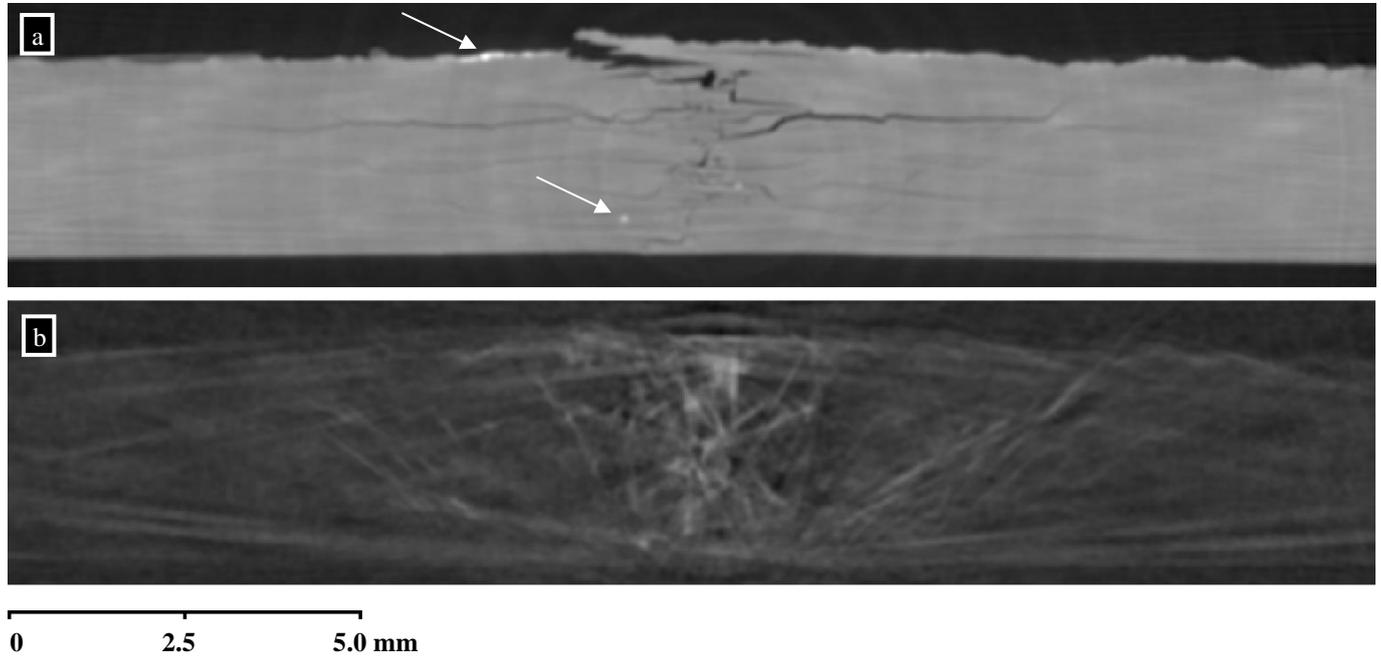
3 Dataset

Figure 1: The CFRP specimen with impact damage used for evaluating our tool. (a) shows the attenuation image, (b) the dark field signal image. It can clearly be seen that (a) shows the horizontal cracks, while (b) shows mostly vertical or 45 degrees slanted fractures. In the attenuation image, some high density inclusions are clearly visible (two are marked with white arrows).

Our test specimen is a CFRP object with impact damage, provided by the company FACC. The fracture damage was introduced using an energy of 20 Joule. The dataset was acquired using a SkyScan 1294 Talbot-Lau grating interferometer computed tomography device [8]. The scan was performed at a tube voltage of 35 kV and 920 µA tube current, using a 0,25 mm Al filter. The voxel size is (22.8 µm)³. 1200 projection images were taken, each at an acquisition time of 650 ms, using 7 times averaging. Slices of the reconstructed image can be seen in Figure 1. We use data from two modalities. In Figure 1a, the attenuation image is shown, which delivers information on the density of the material at a location, as known from standard computed tomography. The second used modality, displayed in Figure 1b, is the so-called dark field image, which measures the amount of radiation scattered at small angles. It can resolve high-frequency structures with a size even below the voxel resolution, and delivers a high signal (indicated by white color in Figure 1b) at location with such structures. The main goal is to characterize the damage that occurred because of the impact. The focus here is on a qualitative, visual evaluation, as for a quantitative inspection there is so far no reliable reference data available.

4 Segmentation Pipelines

As segmentation pipelines, we explored the usage of several algorithms. First experiments were conducted with a pipeline around a graph-based segmentation algorithm, the Random Walker, which is described in section 4.1. These experiments showed that a specialized pipeline is required for this kind of dataset, which is described in section 4.2.

4.1 Graph-Based Segmentation Pipeline (SVM+ERW)

Graph-based algorithms are successfully used in a number of contexts, such as medical imaging [11] or the geosciences [12]. Thus we applied a pipeline based on Support Vector Machines (SVM) [9] and the Extended Random Walker (ERW) [10], below referred to as SVM+ERW. The SVM can be thought of as a multidimensional thresholding, and in our pipeline is used in its probabilistic version to deliver the input prior model for the ERW. The fusion of image modalities here happens implicitly in both the SVM and the ERW. SVM considers all modalities as separate feature values. In the ERW they are considered in a weighted fashion for the distance metric in the optimization graph. The pipeline comes with three general parameters for each SVM and ERW, and three additional parameters per modality for ERW, so with our two modalities, we have a total of twelve parameters, which are listed in Table 1. The pipeline is depicted in Figure 2a, a typical result is shown on the right side. The visual inspection of results of this pipeline unfortunately showed this method to be lacking for this specific use case. For one, as

the dark field image shows a high signal value not only at the location of damage, but also at the border of the specimen, we can see in Figure 2a on the right, that on the border of the specimen there is a layer segmented as damage (the violet region marked with an arrow). Also, cracks were typically oversegmented.

The first problem, the segmentation of a layer at the border of the specimen as damage, can be alleviated by an additional surrounding filtering step, consisting of a global thresholding performed on the attenuation image, followed by a morphological closing to remove all internal air-filled structures. This results in a mask for the whole specimen, which after the other segmentation steps is used to mask out regions outside of the specimen, and to mark them accordingly. The second problem however, the oversegmentation of cracks, seems to be a systematic problem with the algorithms used in this pipeline. So while these experiments gave us a good start for result comparisons, further experiments with algorithms more specific to this problem were required.
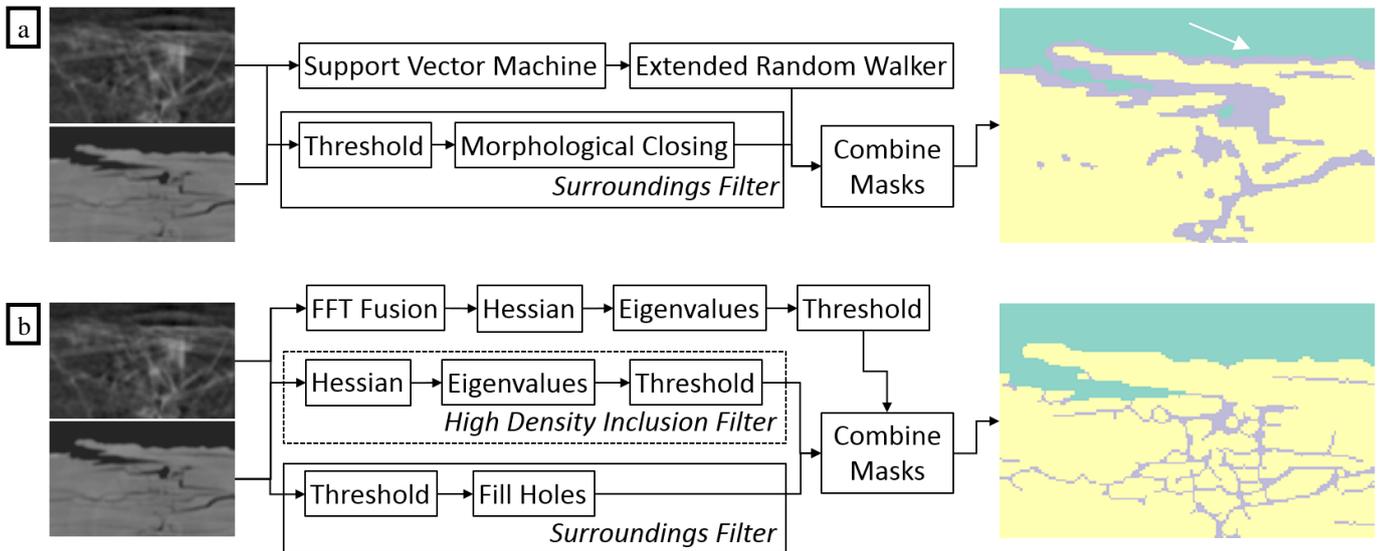

Figure 2: The two main pipelines for impact damage segmentation that were analysed. Pipeline (a) is based on Support Vector Machines (SVM) and the Extended Random Walker (ERW), and includes a Surroundings Filter, while pipeline (b) uses an image fusion based on filtering high frequency parts from both modalities via a Fast Fourier Transform (FFT); then the Hessian is calculated, followed by the calculation of its eigenvalues, on which thresholding is performed. Also here a surroundings filter and optionally a high density inclusions filter were used.

Pipeline	Algorithm	Parameter	Description
SVM+ERW	SVM	C	Soft classification penalty
		γ_{svm}	Gaussian RBF kernel width parameter
		n	Number of channels to consider
	ERW	β	Normalization constant for Gaussian neighborhood
		γ_{erw}	Weight of prior model versus neighborhood
		m	Maximum number of iterations in linear solver
	ERW Modality-specific	p_i	Number of PCA components considered for modality i
		w_i	Weight of modality i in neighborhood information
		d_i	Index of distance metric used for modality i
FFTF+HEV		t_1	Threshold on Hessian eigenvalues for small structures
		t_2	Threshold on Hessian eigenvalues for large structures
		t_s	Threshold for attenuation to be used as surroundings mask
+HDI		t_h	Threshold on attenuation for high density inclusion filter

Table 1: Parameters for each of the two considered pipeline types

4.2 FFT Fusion + Hessian Eigenvalues Pipeline (FFTF+HEV[+HDI])

Several specialized segmentation pipelines have been proposed for damage detection in computed tomography images. For an overview we refer interested readers to the work by Ehrig et al. [13]. As a compromise between speed and quality, we choose a pipeline built around Hessian Eigenvalues, preceded by an Image Fusion. We start by applying a Fast Fourier Transform (FFT) to our two modalities. Subsequently, the inverse-FFT-projected high frequency parts of the dark field image are subtracted from the attenuation image, since we want the image to be dark in regions of a damage. Similarly, the high frequency parts of the attenuation image are added to the result, to emphasize damage. On the result of the fusion, the Hessian matrix is calculated. On the eigenvalues of this matrix, we segment, using a global thresholding, with two different thresholds account for larger and smaller crack structures. In addition, morphological thinning is performed, and structures below a certain size are removed, to

clean the result. Currently, all these steps are performed on a slice-by-slice basis and then merged again to a 3D volume. In the rest of the paper, this pipeline will be referred to as FFTF+HEV.

During our experiments with the FFTF+HEV pipeline and discussing the results with our material experts, we noticed some artefacts from high density inclusions. These are regions with much higher attenuation values than the rest of the polymer matrix, which were also resulting in high eigenvalues in the Hessian matrix, and thus were segmented as defects. We therefore added additional calculation steps to the pipeline in order to reduce these artefacts. We did this by performing a thresholding on the input data for regions which most likely only contained matrix material; we masked out all regions of the image where the intensity was below a certain threshold, and left all other intensity values as they were. We then performed the Hesse transform on this masked image, and did the thresholding on the eigenvalues with the same threshold as used for the whole image. This delivers a binary mask for the crack-like artefacts coming from high-density inclusions, which then can be used to filter them out from the final segmentation result. The pipeline including this high-density inclusion filtering step will be referred to as FFTF+HEV+HDI in the following. The FFTF-HEV pipeline depicted in Figure 2b includes the HDI filtering as optional part.

5 Evaluation Interface

We have implemented a sampling interface for arbitrary volume segmentation pipelines. It uses a description of the required parameters and their ranges, and determines the values to be sampled following the selected sampling strategy (currently either random, on a Cartesian grid or using Latin Hypercubes sampling), and invokes the image processing pipeline with the sampled parameter sets. From every resulting labeled image we also calculate derived output such as the number of connected components. For probabilistic segmentation algorithms, such as the Random Walker, which per pixel do not only deliver a single label, but the probability of belonging to each label, it also calculates the average uncertainty.
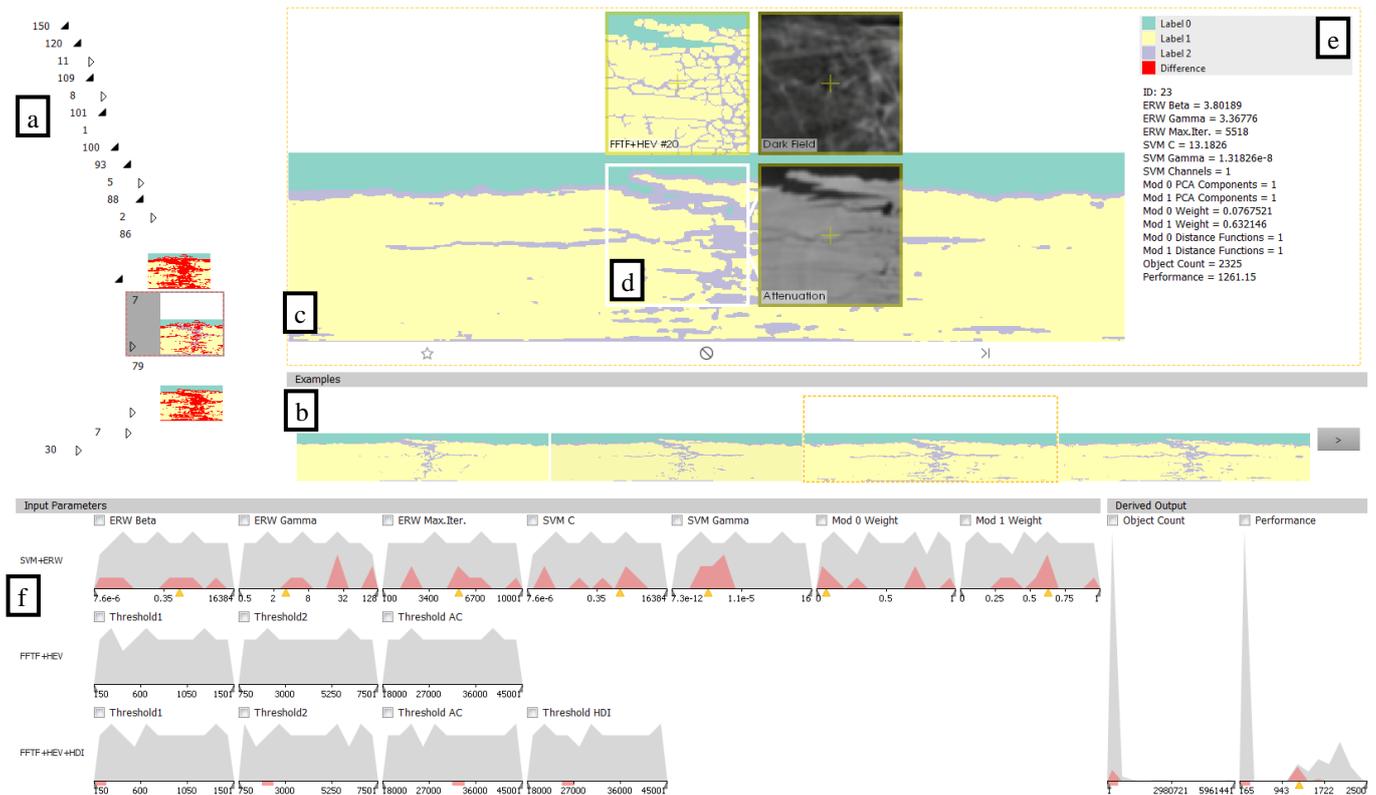

Figure 3: Our result exploration interface. Results are clustered by similarity in the *cluster tree view* (a), results from the selected cluster are shown in the *example view* (b). The *detail view* (c) shows a single selected result with a magic lens for comparisons (d), to the right of the detail view is a label legend and information about the parameters of the currently selected image/cluster (e). The *histogram view* (f) displays the distribution of the parameters and derived output of all results.

Our visualization tool, depicted in Figure 3, then clusters the results by similarity, and displays the result in the *cluster tree view*, as can be seen in Figure 3a. In this tree view, we show the number of results contained in the cluster, as well as a representative image for each cluster. The representative image encodes locations with differences inside the cluster with a marker color (labelled "Difference" in the legend in Figure 3e). Currently unimportant clusters can be collapsed, which means that their representative image is not shown anymore, as can be seen in Figure 3a, where the top twelve cluster are collapsed. In the middle of the interface, selected results from the currently selected cluster are shown in the *example view* (Figure 3b). In the center, a larger version of the currently selected cluster representative image or the selected example image is shown in the *detail view* (c). A magic lens (Figure 3d) in the detail view enables comparison with the input image or other results. The information box to the left of the detail view (Figure 3e) shows a label legend, as well as the actual parameter values used for calculating a

particular result. On the bottom, we display the distribution of input parameters and of the derived output in the *histogram view* (f). All views are linked – when the user zooms into a region of interest in one view, the same interaction is applied to all other image views. When selecting a cluster of results, the histograms are updated to show details about this selection. By navigating through the cluster hierarchy, users can determine the influence of each parameter on the result.

The sampling as well as the visualization interface are implemented in C++ using the Qt framework for the user interface and VTK [14] and ITK [15] libraries for visualization and image processing. Its source code is available as part of the open_iA framework on github [16].

6 Results

The main goal in our use case was to characterize the damage that occurred because of the impact in the CFRP specimen. Therefore, we needed to segment cracks and fractures. Larger fracture structures show up in the attenuation image, while smaller cracks only appear in the dark field image. Thus, we needed to combine both modalities to get all damage information. From considering the dark field image, an additional challenge arises in differentiating between actual damage and surrounding air, as the dark field image delivers high signal not only in fracture regions, but also at the border of the specimen. Additionally, the attenuation image shows several high density inclusions that, if left unhandled, produce artefacts in the final segmentation. We thus employed and improved the pipelines described in section 4 on this dataset, and used the tool described in section 5 to compare them and their parametrizations. Figure 3 shows our tool as it is used in the comparison of segmentation masks resulting from three different pipelines: SVM+ERW, FFTF+HEV and FFTF+HEV+HDI. Through navigation in the cluster hierarchy and through examining the representative image, example images from each cluster, as well as the histogram, we could quickly navigate to suitable results as well as the correlated parameters and thus pipelines that produced them.

Figure 4 shows a selection of the best results for each algorithm. These were identified through exploration with our tool, among a collection of a total 150 images; for each pipeline, 50 samples were created. The sampling itself took much longer, especially the SVM+ERW pipeline, it took approximately 30 minutes per segmentation run. For FFTF+HEV and FFTF+HEV+HDI the time for a single segmentation run was approximately three minutes. These calculations could however be performed over night and fully automated. The actual exploration of the final results only took approximately 10 minutes.
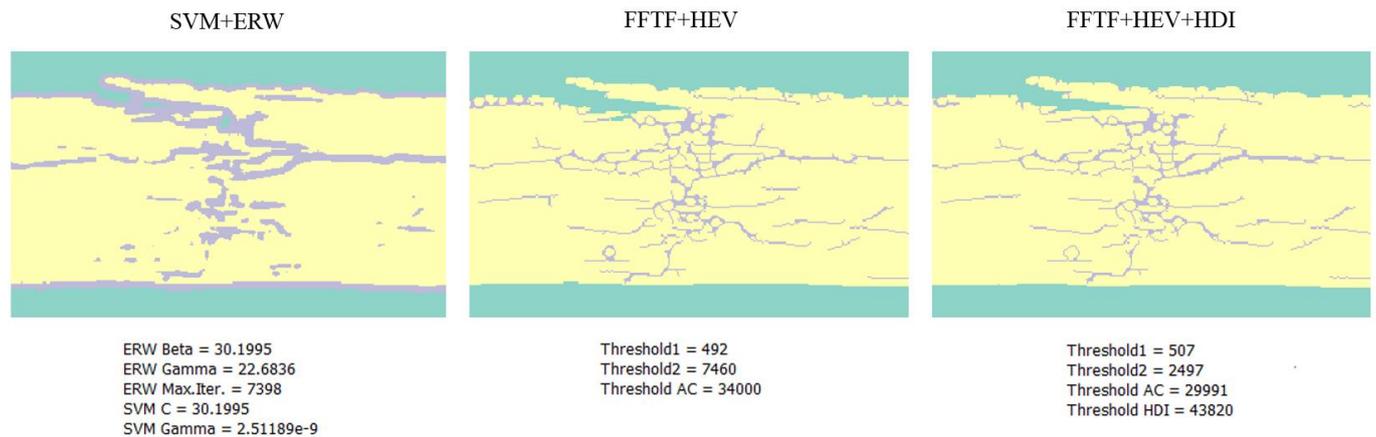

SVM+ERW	FFTF+HEV	FFTF+HEV+HDI
ERW Beta = 30.1995 ERW Gamma = 22.6836 ERW Max.Iter. = 7398 SVM C = 30.1995 SVM Gamma = 2.51189e-9	Threshold1 = 492 Threshold2 = 7460 Threshold AC = 34000	Threshold1 = 507 Threshold2 = 2497 Threshold AC = 29991 Threshold HDI = 43820

Figure 4: Most suitable results for each of the three considered segmentation pipelines.

7 Conclusion and Outlook

We have shown that with our tool, CT experts can easily compare results of different segmentation pipelines, as well as different parametrizations of the same pipeline. Our tool also allows to find the most suitable results for a given pipeline and identify the parameters or parameter ranges that produce these results. It alleviates the need to memorize results from previous runs, by allowing a direct side-by-side comparison.

One limitation of our tool is that it is mainly targeted at situations where several different pipelines are known already. For an evolving pipeline, our tool provides indirect support only, in that one can sample an arbitrary number of times and include multiple samplings in one comparison. However, our tool facilitates the comparison not only between fundamentally different segmentation pipelines, but also between different development stages of the same pipeline.

Future additions could take a deeper look into how to improve and speed up the assembly of new pipelines, or the modification of existing ones. To find suitable methods and a first approximate parameter interval to analyze more closely, we had to do manual experiments on regions of interest. We intend to extend our methods to better support this exploratory workflow. We further intend to perform tests in the direction of acquiring reference data, to be able to perform quantitative analysis on our impact damage CFRP test specimen.

Acknowledgements

References

[1] S. Arens, G. Domik, A survey of transfer functions suitable for volume rendering, In IEEE/ EG Symposium on Volume Graphics (2010), Westermann R., Kindlmann G., (Eds.), The Eurographics Association, doi: 10.2312/VG/VG10/077-083.

[2] M. Sedlmair, C. Heinzl, S. Bruckner, H. Piringer, T. Möller, Visual parameter space analysis: A conceptual framework. IEEE Transactions on Visualization and Computer Graphics 20, 12 (Dec. 2014), 2161–2170, doi: 10.1109/TVCG.2014.2346321.

[3] L. Bavoil, S. P. Callahan, P. J. Crossno, J. Freire, C. E. Scheidegger, C. T. Silva, and H. T. Vo, VisTrails: Enabling interactive multiple-view visualizations, IEEE Conference on Visualization, 2005, pp. 135–142, doi: 10.1109/VISUAL.2005.1532788.

[4] T. Torsney-Weir, A. Saad, T. Möller, B. Weber, H.-C. Hege, J.-M. Verbavatz, and S. Bergner, Tuner: Principled Parameter Finding for Image Segmentation Algorithms Using Visual Response Surface Exploration, IEEE Transactions on Visualization and Computer Graphics, Vol. 17, No. 12, pp. 1892–1901, Nov. 2011, doi: 10.1109/TVCG.2011.248.

[5] A. J. Pretorius, Y. Zhou, and R. A. Ruddle, Visual parameter optimisation for biomedical image processing, BMC Bioinformatics, Vol. 16, No. Suppl 11, p. S9, 2015, doi: 10.1186/1471-2105-16-S11-S9.

[6] J. Weissenböck, A. Amirkhanov, E. Gröller, J. Kastner, C. Heinzl, PorosityAnalyzer: Visual Analysis and Evaluation of Segmentation Pipelines to Determine the Porosity in Fiber-Reinforced Polymers, IEEE Conference on Visual Analytics Science and Technology (IEEE VAST), 2016.

[7] B. Fröhler, T. Möller, C. Heinzl, GEMSe: Visualization-Guided Exploration of Multi-channel Segmentation Algorithms, Computer Graphics Forum (Proceedings Eurographics Conference on Visualization), Vol. 35, No. 3, Groningen, Netherlands, June 2016, pp. 191-200, doi: 10.1111/cgf.12895.

[8] S. Senck, D. Salaberger, C. Gusenbauer, B. Plank, B. Plank, G. Rao and J. Kastner, Talbot-Lau grating interferometer XCT for the quantitative characterization of damage in polymers after impact and static tensile testing, 19th World Conference on Non-Destructive Testing (WCNDT2016), June 2016, Munich, Germany, 2016.

[9] Chih-Chung Chang and Chih-Jen Lin, LIBSVM: a library for support vector machines, ACM Transactions on Intelligent Systems and Technology, 2:27:1--27:27, 2011, Available online at http://www.csie.ntu.edu.tw/~cjlin/libsvm, Last accessed on 15-Dec-2016.

[10] L. Grady, Multilabel random walker image segmentation using prior models, Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 1, pp. 763–770, June 2005, doi: 10.1109/CVPR.2005.239.

[11] L. Grady and G. Funka-Lea, Multi-Label Image Segmentation for Medical Applications Based on Graph-Theoretic Electrical Potentials, in Proceedings of the 8th ECCV04, Workshop on Computer Vision Approaches to Medical Image Analysis and Mathematical Methods in Biomedical Image Analysis, pp. 230-245, 2004.

[12] X. Kang, S. Li, L. Fang, M. Li, and J. A. Benediktsson, Extended Random Walker-Based Classification of Hyperspectral Images, IEEE Transactions on Geoscience and Remote Sensing, Vol. 53, No. 1, pp. 144–153, 2015, doi: 10.1109/TGRS.2014.2319373

[13] K. Ehrig, J. Goebbels, D. Meinel, O. Paetsch, S. Prohaska, V. Zobel, Comparison of Crack Detection Methods for Analyzing Damage Processes in Concrete with Computed Tomography, International Symposium on Digital Industrial Radiology and Computed Tomography (DIR 2011), 20-22 June 2011, Berlin, Germany.

[14] W. Schroeder, K. Martin, B. Lorensen, The Visualization Toolkit, 4th ed., Kitware, 2006.

[15] T. S. Yoo, M. J. Ackermann, W. E. Lorensen, W. Schroeder, V. Chalana, S. Aylward, D. Metaxas, R. Whitaker, Engineering and Algorithm Design for an Image Processing API: A Technical Report on ITK - the Insight Toolkit. Studies in Health Technology and Informatics (2002), 586–592.

[16] open_iA, Available online at https://3dct.github.io/open_iA/, Last accessed on 15-Dec-2016.