100 year anniversary
Measurement Science and Technology

PAPER • OPEN ACCESS

3D total variation denoising in X-CT imaging applied to pore extraction in additively manufactured parts

To cite this article: Rob Heylen *et al* 2022 *Meas. Sci. Technol.* **33** 045602

View the article online for updates and enhancements.

3D total variation denoising in X-CT imaging applied to pore extraction in additively manufactured parts

Rob Heylen[1,*] ⓘ**, Aditi Thanki**[2] ⓘ**, Dries Verhees**[1] ⓘ**, Domenico Iuso**[3] ⓘ**, Jan De Beenhouwer**[3] ⓘ**, Jan Sijbers**[3] ⓘ**, Ann Witvrouw**[2] ⓘ**, Han Haitjema**[2] ⓘ **and Abdellatif Bey-Temsamani**[1] ⓘ

[1] Flanders Make vzw., Gaston Geenslaan 8, 3001 Leuven, Belgium
[2] KULeuven MaPS, Celestijnenlaan 300, bus 2420, 3001 Leuven, Belgium
[3] imec-Vision Lab, Universiteitsplein 1, 2610 Antwerpen, Belgium

E-mail: rob.heylen@flandersmake.be

CrossMark

Abstract

X-ray computed tomography (X-CT) plays an important role in non-destructive quality inspection and process evaluation in metal additive manufacturing, as several types of defects such as keyhole and lack of fusion pores can be observed in these 3D images as local changes in material density. Segmentation of these defects often relies on threshold methods applied to the reconstructed attenuation values of the 3D image voxels. However, the segmentation accuracy is affected by unavoidable X-CT reconstruction features such as partial volume effects, voxel noise and imaging artefacts. These effects create false positives, difficulties in threshold value selection and unclear or jagged defect edges. In this paper, we present a new X-CT defect segmentation method based on preprocessing the X-CT image with a 3D total variation denoising method. By comparing the changes in the histogram, threshold selection can be significantly better, and the resulting segmentation is of much higher quality. We derive the optimal algorithm parameter settings and demonstrate robustness for deviating settings. The technique is presented on simulated data sets, compared between low- and high-quality X-CT scans, and evaluated with optical microscopy after destructive tests.

Keywords: denoising, X-CT, additive manufacturing, selective laser melting, segmentation, porosity

(Some figures may appear in colour only in the online journal)

---

* Author to whom any correspondence should be addressed.

1. Introduction

Selective laser melting (SLM) has become a common and relatively mature manufacturing technique for fabrication of complex or near net-shape metallic parts [1, 2]. This process allows a larger design freedom than is possible with conventional manufacturing techniques and it enables the production of functional parts with good material and part properties. Although it offers a great advantage over conventional manufacturing, the process is very sensitive to many factors to form the desired component. This process involves complex multiphysics phenomena, such as absorption, transmission and reflection of laser light [3], material evaporation [4], melt pool dynamics and microstructure evolution [5], and rapid melting and solidification of metallic powder material [6]. The slightest deviation of processing parameters such as laser power and scan speed from their optimum can affect physical phenomena, which highly influences the amount of defects in the SLM parts. Therefore, defects are an inevitable issue in SLM parts [7] and may influence the mechanical and physical properties of parts. Most common defects are keyhole and gas porosities, insufficient fusion, solidification cracking and solid-state cracking.

Different techniques have been used for the evaluation of defects in SLM parts. 2D cross-sectioning of parts has been used for pore characterization of parts, however this method is destructive and non-repeatable [8]. Archimedes' method has also been used for parts with porosity to determine their relative density, but the limitation is that this approach does not provide the size, location and morphological information of defects [9]. X-ray computed tomography (X-CT) has shown the potential for non-destructive assessment of defects in SLM parts [10, 11]. However, multiple factors influence the final X-CT reconstruction quality, and the level at which defects can be resolved depends on the resolution, noise level and artefacts present in the X-CT reconstruction. Imaging artefacts such as cone-beam artefacts, beam hardening and several types of noise influence the quality of the reconstruction and ultimately the porosity segmentation. Image noise in a X-CT projection is caused primarily by physical phenomena affecting the x-ray photons (e.g. scattering and the intrinsic discrete nature of photons), but can also originate from other sources, such as the measurement or device settings (e.g. source voltage, source power, target material, exposure time, averaging and filtering), object properties (e.g. object attenuation coefficient, object geometry, object placement), electronic noise of the detector and the employed image reconstruction algorithm [12].

In this paper, we explore the performance of a 3D image denoising method based on total variation (TV) denoising in improving pore segmentation in parts that are additively manufactured via SLM. TV denoising is an edge preserving denoising technique, as opposed to alternatives like non-local means or sparse domain transformation methods such as wavelet or Fourier transforms, and has already shown promising results in image denoising of regular 2D images and in medical X-CT imaging, e.g. with applications in cardiac, pulmonary and brain imaging [13, 14]. It is demonstrated that this type of denoising can also successfully improve segmentation

results in AM manufacturing, with significant qualitative and quantitative improvements in segmentation metrics. Furthermore, a parameter space exploration is performed for X-CT reconstructions that are typically encountered in SLM applications, leading to a robust set of parameters that can be used to denoise and improve segmentation results. The main goal of this work is to demonstrate the improvements in pore segmentation results that arise from including a recent state of the art TV denoising step in the X-CT segmentation workflow and to provide the details of the employed algorithm and its parameters. Note that many alternative algorithms are available that can perform similarly.

The structure of this paper is as follows: In section 2, the TV denoising algorithm is briefly described, with enough details that allow the interested reader to implement this, and references to recent literature for an in-depth analysis of this algorithm. In section 3, the employed data set is described. Section 4 contains a detailed parameter space exploration to determine ideal settings for the algorithm and to demonstrate its robustness against parameter deviations. Section 5 contains several experiments on real and simulated data sets, followed by the conclusions in section 6.

2. TV denoising

To suppress noise and improve edges in a X-CT reconstruction, we employ a TV denoising method where a target function consisting of a data fidelity term and a TV regularizing term is minimized. A computationally efficient implementation for solving such a minimization problem is based on the primal-dual optimization algorithm of Chambolle and Pock [15], and has recently been presented by several authors in the medical X-CT literature [16–18].

The implementation used here is as following: Consider a 3D voxel cube $y = \{y_{ijk}\}_{ijk}$ with non-negative voxel values $\forall i,j,k : y_{ijk} \in \mathbf{R}^+$ and dimensions $[N_1, N_2, N_3]$: $i \in 1, \ldots, N_1$, $j \in 1, \ldots, N_2$, $k \in 1, \ldots, N_3$. We consider the optimization problem

$$\min_{x \in X} F(x - y) + \lambda G(x) \tag{1}$$

where

- The variable $x \in X$ is the variable to be optimized, which is a 3D voxel data cube with similar size as y
- The function F describes a data fidelity term that aims to keep the denoised solution x close to the original data y. This can be accomplished by choosing the Fröbenius norm for F:

$$F(x - y) = \|x - y\|_F^2 \tag{2}$$

$$= \sum_{ijk} (x_{ijk} - y_{ijk})^2. \tag{3}$$

- The second term $G(x)$ regularized the solution. In our case we choose the TV of x, which is the overall sum of absolute

differences between neighbors. TV penalizes differences between neighbors, and is well fit to denoise several types of noise:

$$G(x) = |x|_{TV} \tag{4}$$

$$= \sum_{(i,j)} |x_i - x_j|. \tag{5}$$

The sum runs over all neighboring voxel pairs (i,j) in the volume x, where we use the six-neighborhood in 3D (top/bottom, left/right and front/back neighbors).

- The constant λ determines the strength of the TV regularizing term.

Several equivalent formulations of the Chambolle–Pock optimization algorithm exist for solving equation (1), with often different names in literature (see [15] for an overview and derivations). The version we employ is described in detail in [19]. First, define the following operators and functions where x is a variable representing a 3D voxel cube and u represents its dual (a 3D vector space, obtained e.g. by taking a numerical gradient of x):

$$K(x) = \nabla(x) \tag{6}$$

$$K^*(u) = -div(u) \tag{7}$$

$$P_{\tau y}(x) = \frac{x + \tau y}{1 + \tau} \tag{8}$$

$$P_g(u) = v \iff \forall ijk : v_{ijk} = \frac{\lambda u_{ijk}}{\max(\overline{u}_{ijk}, \lambda)}. \tag{9}$$

Initialize the variables:

$$x_2 = y \tag{10}$$

$$u_2 = P_g(K(x_2)). \tag{11}$$

The solution is iterated using algorithm 1 until convergence is reached. We refer to [15] and [19] for more details, e.g. on the proper boundary conditions for the numerical gradient and divergence operations. Under certain conditions, this algorithm is shown to quickly converge towards a solution that minimizes equation (1). Convergence can be monitored by evaluating the gap between the primal and dual solutions: when this gap becomes zero, both terms balance each other and convergence has been reached. A detailed parameter analysis is presented in section 4.

We have implemented the above algorithm in Matlab, based upon extending a 2D implementation for image denoising provided by the author of [19]. As the algorithm is an iterative algorithm, its runtime scales linearly with the number of iterations. Note that the algorithm is very well fitted for parallellization, and GPU-accelerated versions should be relatively

---

Algorithm 1: The proposed algorithm.

|	
1	Initialize all variables, choose $\tau > 0$, $\sigma > 0$ and $\rho \in [1,2[$
2	**for** $n \geqslant 0$ **do**
3	$\quad x = P_{\tau y}(x_2 - \tau K^*(u_2))$
4	$\quad u = P_g(u_2 + \sigma K(2x - x_2))$
5	$\quad x_2 = x_2 + \rho(x - x_2)$
6	$\quad u_2 = u_2 + \rho(u - u_2)$

---

easy to implement. A possible concern is the memory requirement: 3D X-CT images are typically large, and the algorithm requires multiple copies of similar size to store the gradient and other intermediate solutions. Possible solutions are the use of disk caching for intermediate solutions or block-wise execution of the algorithm where only subsets of the entire dataset are processed at once.

3. Description of the data set

To assess the performance of the algorithm on improving pore segmentation in additive manufactured parts, a real data set was created under controlled conditions where pores were induced in certain predefined regions by deviating from the optimal system parameters. To this end, a cubic object of size 9 mm × 9 mm × 7.5 mm (length, width, height) was printed in 316 L stainless steel with a 3D Systems ProX DMP320 SLM machine. The bulk of the object was printed with nominal settings with layer thickness 30 μm, laser power 215 W, laser speed 900 mm s^{-1} and interlayer time 8 s. The top layer was remelted, and 14 lines of length 8 mm and with a line distance of 0.5 mm were printed on top of the cube with varying system settings to investigate the effect of the system settings on defect formation (see figure 1).

Depending on the system settings, these pores will have different pore densities and size distributions, and due to the construction we can relate every observed pore with the system settings that created these. In this work, we focus on a single line, namely the line with ID 4, that shows a high density of keyhole pores. Line ID 4 was printed with 400% nominal energy density, obtained with a nominal laser power (215 W) and 25% laser speed (225 mm s^{-1}). Note that in the nominal bulk material no pores are typically present except for subsurface pores. After detachment from the build plate by wire electrical discharge machining (EDM), the object was CT scanned. The UAntwerp FleXCT scanner [20] was used to acquire and reconstruct the x-ray projections. The CT scans were performed with a peak voltage of 220 kV and 23 W as tube power, using a 1 mm thick copper filter. The x-ray source was placed 63.33 mm distant from the object and 950 mm distant from the detector, while 4283 projections have been acquired. Two CT scans were performed: one with a lower of exposure time per projection (1110 ms), and the other one with a longer exposure (5000 ms). The former scan yields a noisy, low-quality, CT reconstruction compared to the more accurate reconstruction given by the latter scan. The two CT scans
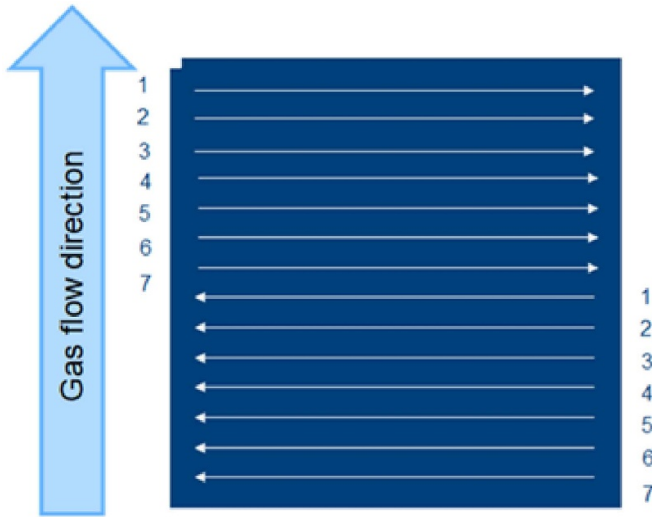
Figure 1. The top layer of the cube with 14 numbered lines. The direction is indicated with an arrow. The cube sides have length 9 mm.
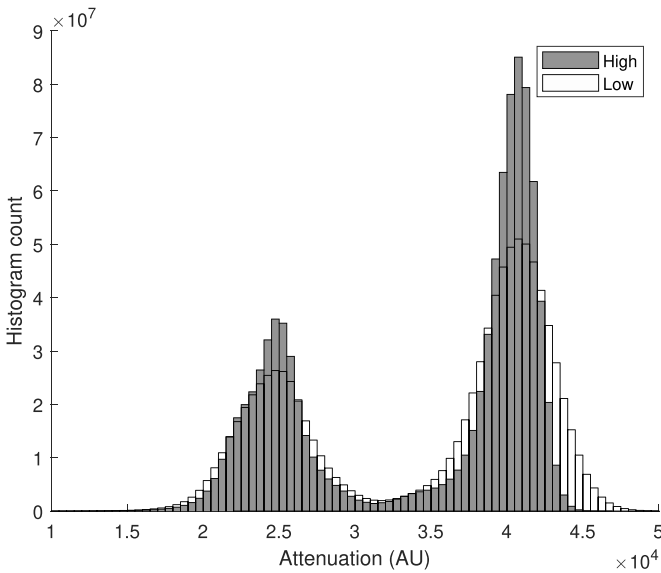
Figure 2. The histograms of the low and high quality X-CT scan reconstructions after histogram matching.

have a resolution of 10 μm, where beam hardening correction was performed to minimize reconstruction artefacts. As both CT data sets were collected successively and with different settings, the obtained reconstructions do not align. To allow comparison, both data sets were geometrically aligned with the CAD model of the object with a software based on a mesh registration algorithm [21], which also aligned both data sets with each other. Furthermore, histogram matching was applied, where the histogram of the high-quality reconstruction was rescaled via a linear transformation so that the two modes of the bimodal distribution representing the air and metal voxels were mapped to each other. The resulting histograms of both voxel distributions are shown in figure 2, and two slices are shown in figure 3.

Furthermore, only a subset of the full data set is considered in each experiment:

- A smaller 3D selection containing many pores is sufficient to demonstrate the technique, and easier to process and visualize.
- As the segmentation will also segment air outside the object, a subset that does not contain this air component is considered for the segmentation quality assessment, either by removing this air component via 3D region selection or by focusing on a selection that contains only bulk metal.

This selection is indicated with a dashed rectangle in figure 3, with a cross-section shown in figure 4 and a 3D visualization of pores segmented by the popular Otsu's thresholding method [22] in figure 5.

4. Parameter space exploration

As the proposed algorithm contains several parameters, a parameter space exploration was carried out first to evaluate the influence of each parameter on the result and to identify a good parameter set. It must be noted that Chambolle–Pock based algorithms typically show a relatively low dependence on several of their algorithmic parameters and will converge to the same final solution for many parameter sets. This was also observed in this case, which illustrated the robustness of the method. The algorithm parameters are the following:

- Relaxation parameter ρ, which should be in the interval $[1, 2[$.
- Step size σ and its dual τ. For the 3D TV regularization algorithm, we choose $\sigma = (16\tau)^{-1}$, as this choice guarantees convergence [15]. Only the step size τ then remains to be chosen.
- The number of iterations N should be chosen large enough to reach convergence.
- The regularization parameter λ weights the strength of the TV term, and plays an important role in setting the denoising level. This parameter will be evaluated further for each experiment.

To assess the influence of these parameters, the low-energy X-CT reconstruction was denoised with different parameter settings. The convergence was assessed by evaluating the primal cost F and dual cost G in equation (1), and the primal-dual gap $F - G$ [15].

The typical evolution of the primal and dual cost as a function of the number of iterations of the algorithm is shown in figure 6. The inset shows the primal-dual gap, which is the absolute difference between primal and dual cost. The erratic behavior of this gap after iteration 200 is due to numerical noise and indicates that convergence has been reached down to the level of the employed numerical accuracy. The parameter values are indicated in the caption. This graph indicates that iterating for more than 200 iterations is not useful, as the two terms in the optimization equation (1) balance each other
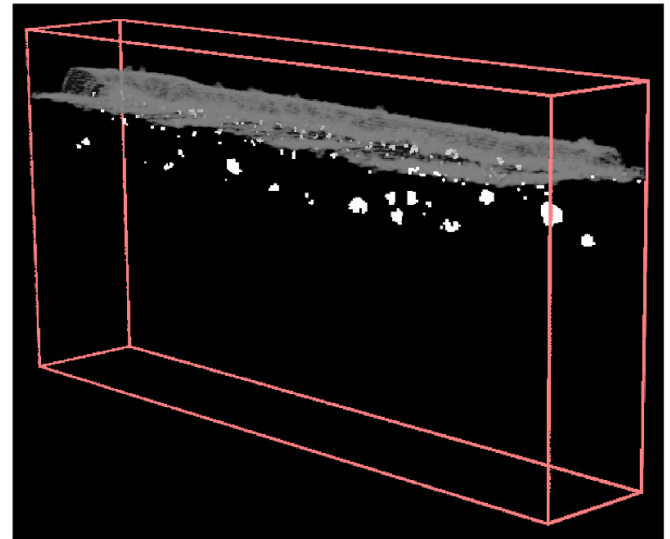
(a) Low quality
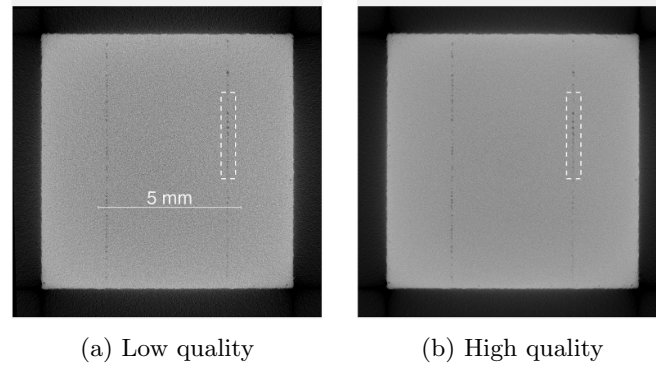
(b) High quality

Figure 3. A X-CT slice of the specimen parallel to the build platform for the low (a) and the high-quality (b) aligned and histogram-matched reconstructions. Voxels are cubic with size $(10\ \mu m)^3$. Two distinct lines of pores are clearly visible, corresponding to line ID 4 with 400% nominal energy density, obtained with nominal laser power and 25% of the nominal laser speed. The total data set has a size of $1081 \times 1081 \times 809$ voxels. The top view of the selection used below is indicated with the dashed rectangle.

(a) Low quality

(b) High quality
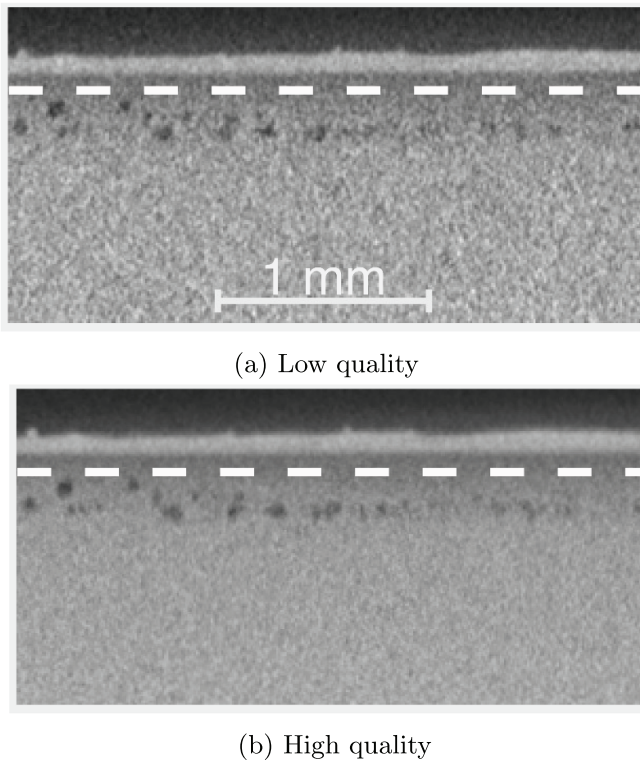
Figure 4. A selection of a vertical slice through the line with pores on the right for the low (a) and high (b) quality X-CT reconstructions. Note the increased noise level in the low-quality reconstruction. Only the bulk material below the dashed line was employed for quality assessment to avoid air and border segmentation issues.


Figure 5. A 3D representation of the segmented pores in this selection, obtained via Otsu's threshold applied to the high-quality reconstruction and removal of the air component outside the object. Pores are indicated in white, the border of the selection in red, and the top surface of the object in transparent grey. Voxels are cubic with size $(10\ \mu m)^3$. The red selection has size 3.0 mm $\times$ 1.5 mm $\times$ 0.50 mm.

up to the level of numerical noise. Note that the decrease in the primal-dual gap is exponential, so good results are already reached well before this point.

The dependence on the algorithmic parameter τ is shown in figure 7. The algorithm converges for all values of τ, ranging over several orders of magnitude, but the convergence speed is affected by the choice of τ. The fastest initial decrease is observed for $\sigma = \tau = 1/4$, a value also often suggested as a

good starting point in literature, and convergence is reached after only 50 iterations.

Dependence on the algorithmic relaxation parameter ρ is shown in figure 8, and is surprisingly low. The choice of this parameter in its allowable range $\rho \in [1, 2[$ has but a minor effect on the convergence speed. Fastest convergence is observed for high values. From these graphs, we can conclude that the algorithm is robust for a wide range of parameter settings, and that the results after convergence are very similar. We therefore fix these algorithmic settings to the optimal set observed in these tests: $\sigma = \tau = 1/4$, $\rho = 1.99$ and $N_i t = 100$.

The parameter λ weighs the TV term and indirectly determines the magnitude of the TV noise correction. For $\lambda = 0$ there is no correction and the output of the algorithm equals the
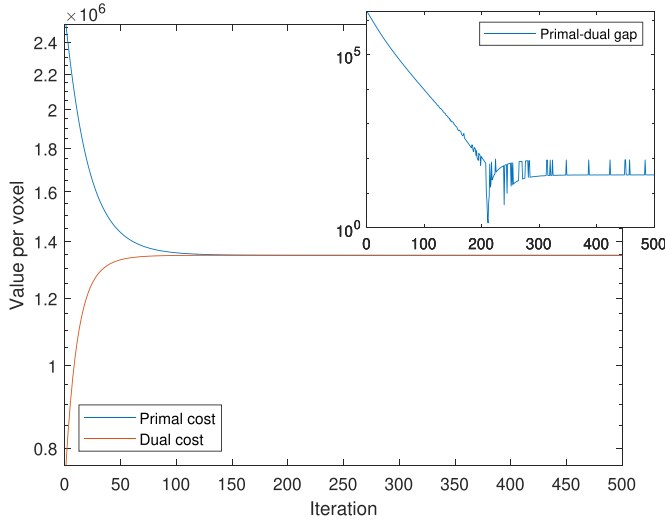
Figure 6. Typical evolution of the primal and dual cost as a function of the number of iterations, and the primal-dual gap (inset). The parameters were $\lambda = 1000$, $\tau = 0.01$, $\rho = 1.99$.
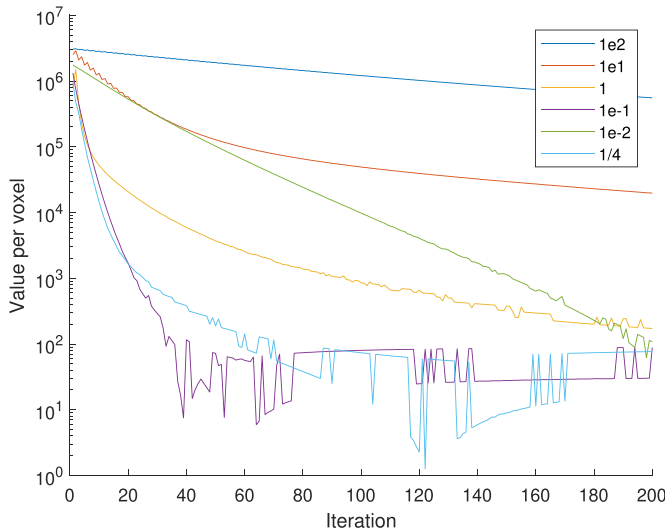
Figure 7. The evolution of the primal-dual gap for different values of τ, indicated in the legend. The other parameters were $\lambda = 1000$, $\rho = 1.99$.

input. For very large values of λ the TV term becomes dominant, which results in one or several large uniform values across the reconstruction (the so-called cartoon-like effect of strong TV regularization, extended to 3D). Note that this parameter is scale-dependent: rescaling the input data set requires appropriate rescaling of the λ parameter to obtain qualitatively similar results. The scale of the employed data set can be seen from its histogram in figure 2.

The effect of using different values of λ is illustrated in figure 9, where the original data set and several denoised versions are put next to each other for comparison. The effect of thresholding with a fixed threshold $T = 33664$ (in the employed unit-less `uint16` X-CT attenuation values) obtained via Otsu's method applied on the original data set is also shown.
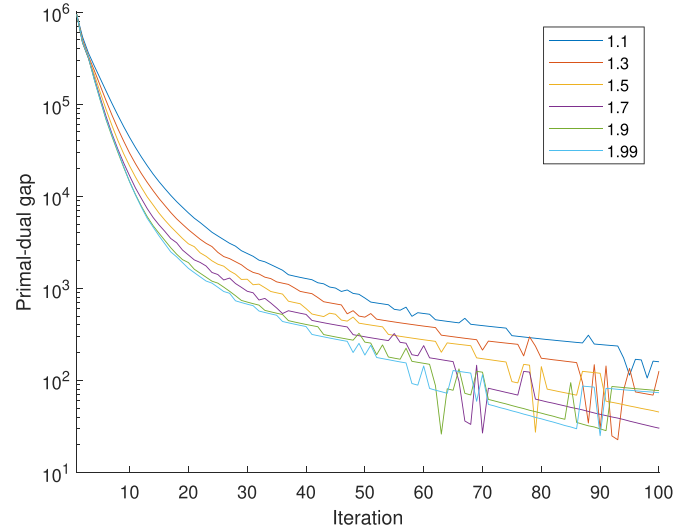
Figure 8. The evolution of the primal-dual gap for several values of the relaxation parameter ρ, indicated in the legend.

5. Experiments

### 5.1. Overview of experiments and quality metrics

We design three different experiments to assess the proposed technique's performance in improving X-CT pore segmentation in additive manufactured parts:

- We use the high-quality X-CT reconstruction to create a pore segmentation that is considered as ground truth, and compare the pore segmentation on the low-quality X-CT with this ground truth.
- We select a small number of pores from the high-quality X-CT reconstruction, use these to construct a simulated low-quality reconstruction, and compare the results with the known segmentation.
- We assess the segmentation against ground truth obtained via destructive testing of the object.

The binary segmentation S_x is obtained by thresholding the 3D volume x with the threshold T: $S_x = x < T$. This way, voxels belonging to pores are considered logical true. The metric employed to assess the correspondence between two different segmentations S_x and S_y is the Dice coefficient $D(S_x, S_y)$, also known by several alternative names such as F1 score or Sørensen-Dice index:

$$D(S_x, S_y) = \frac{2|S_x \cap S_y|}{|S_x| + |S_y|} \tag{12}$$

where $|.|$ is the cardinality operator of a set, returning the number of elements in this set, and the union indicates the overlap between two sets. The Dice coefficient is often used to assess binary segmentation problems in 2D and 3D. It ranges from zero to one, reaching one for identical segmentations. Note that several alternative segmentation metrics exist, e.g. the Jaccard index or intersection-over-union [23]. These other metrics lead to similar conclusions, and preference was given
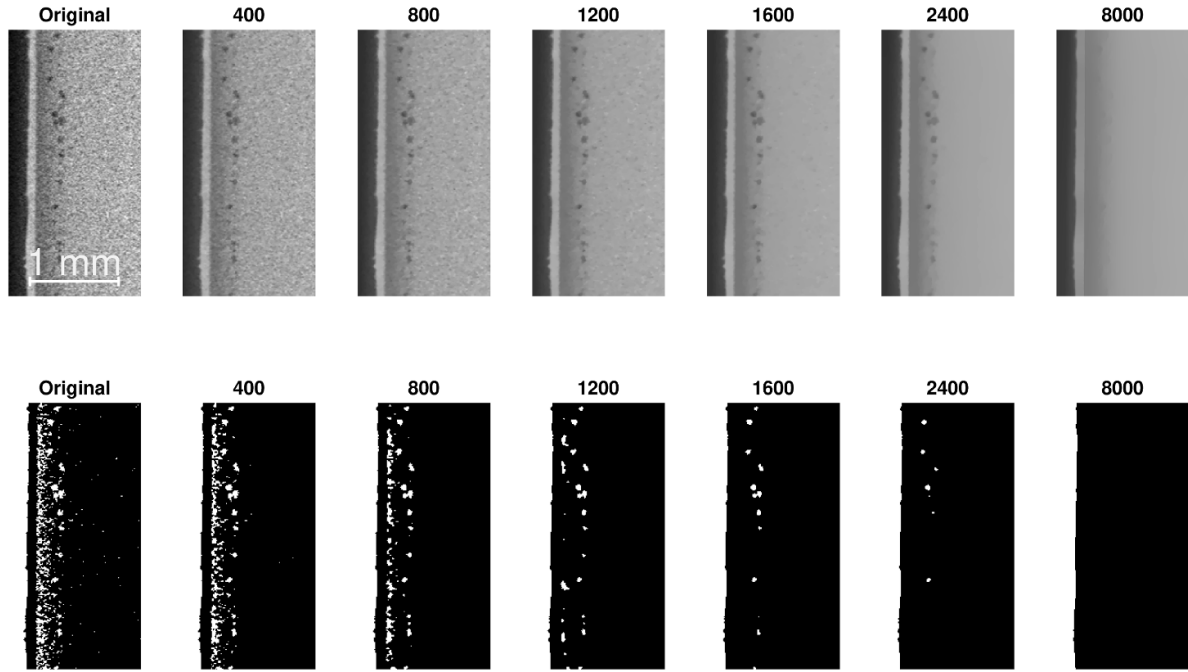
Figure 9. A slice of the X-CT data set and denoised versions with increasing λ, indicated above each figure (top row). The segmentations obtained with the same threshold calculated by Otsu's method on the original data set (bottom row).

to using only a single well-known segmentation metric to improve readability.

5.2. Comparison with high-quality X-CT

In the first experiment the pore ground truth was extracted from the high-quality X-CT reconstruction and the lower-quality X-CT data set was employed to assess the proposed technique (a similar technique was used in [24]). The first difficulty in this approach is to properly segment the high-quality X-CT reconstruction to extract the ground truth, as also this data set contains noise, although significantly less than the low-quality reconstruction. Possible segmentation levels are shown in figure 10, which shows that also here a balance must be found between selecting undesired areas such as border artefacts with lower attenuation values and missing too much of the true pore volumes. For this experiment, the threshold between the yellow and green ranges was selected, corresponding to the unit-less attenuation value 32 125.

Once this ground truth was generated, the low-quality X-CT reconstruction was denoised with the ideal settings found via the parameter space exploration, thresholded with several fixed thresholds, and the Dice coefficient versus the ground truth segmentation was calculated. This quality metric is plotted in figure 11 as a function of the denoising level λ and the applied threshold.

A contour plot of the Dice coefficient as function of the denoising level λ and the applied threshold value is shown in figure 12. Note that the plateau where the Dice coefficient is maximal is relatively broad and flat, as there are many parameter pairs that yield a similar performance and the method is relatively robust against parameter settings that deviate from
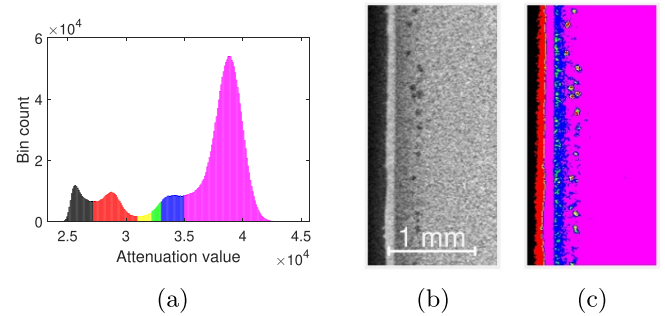

Figure 10. Segmentation of the high-quality reconstruction with different thresholds: the histogram of the distribution with different color-coded value ranges (a), a slice through the original data set (b) and the segmented image with the same color codes as the histogram (c).

the optimum. Furthermore, the performance metric is well-behaved and does not seem to show sharp peaks, discontinuities or local minima, which allows the application of gradient-based or steepest descent maximization strategies. In figure 13 the denoised result and its histogram and threshold are shown, along with the original data set for comparison, for $\lambda = 1000$ and a threshold $T = 32\,500$.

5.3. Comparison with X-CT simulations

While the previous experiment illustrates the improvements in the segmentation obtained by TV denoising when compared with a high-quality X-CT scan, there are some drawbacks to this type of experiment:
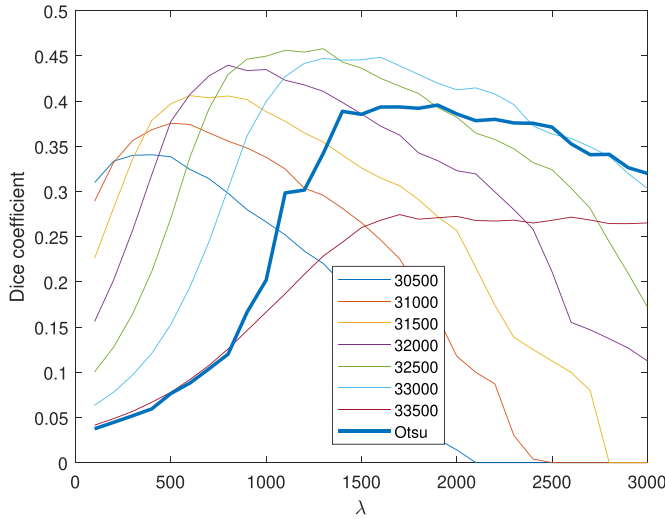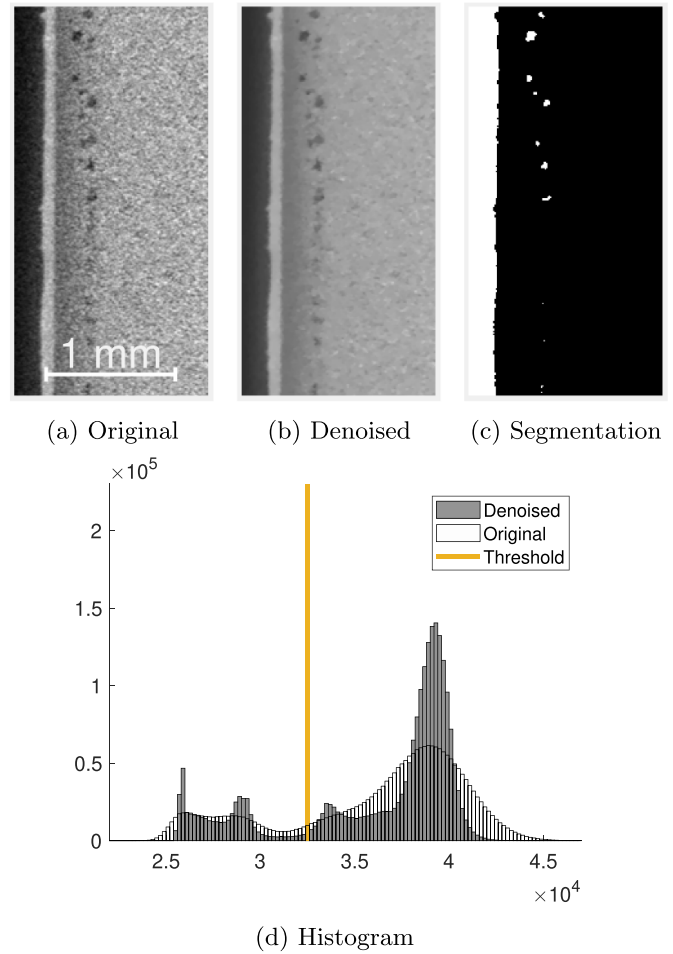
Figure 11. The Dice coefficient as a function of the denoising level λ, for several fixed thresholds and Otsu's threshold indicated in the legend.
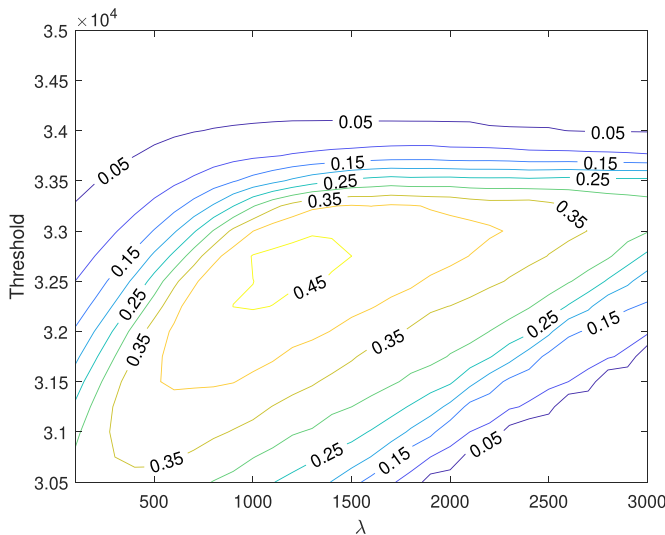
Figure 12. The contour plot of the Dice coefficient as a function of the denoising level λ and the fixed threshold.

(a) Original   (b) Denoised   (c) Segmentation

(d) Histogram

Figure 13. A 2D slice of the original (a) and denoised (b) data set, the segmentation result obtained by thresholding the denoised image (c), and the histogram of the voxel values before and after denoising with the employed threshold indicated (d). The employed parameter values yield the highest Dice coefficient.

- Although a lot of care was put into the alignment of the low- and high-quality reconstructions, these alignments cannot be guaranteed to be perfect. Furthermore, the performed alignment can only be spatially accurate down to discrete steps equal to the voxel size, and further subvoxel alignment might be required.
- Geometrical transformations of a X-CT data cube (translation/rotation) for alignment require a voxel interpolation method such as trilinear interpolation or nearest-neighbor, which will introduce interpolation errors. In the case of trilinear interpolation, noise statistics will be affected also.
- The high-quality X-CT reconstruction is used for ground truth generation, but this reconstruction also suffers from noise, partial volume effects and X-CT artefacts, although to

a lesser extent than the low-quality reconstruction. Furthermore, a proper threshold must also be selected. This introduces errors and uncertainty in the ground truth.

To resolve these issues, a second experiment was designed where a simulated X-CT data set was created, with a known pore segmentation ground truth. This was accomplished by the following steps:

(a) A rectangular 3D subset of the high-quality X-CT data set that contains pores of different sizes was manually selected. Care was taken that no pores are observed within two voxels of the subset borders.

(b) A pore segmentation was created with a manually selected threshold, a connected-component analysis was performed on the segmented voxels, and all connected components with less than eight voxels were removed. The minimal pore size is thus eight voxels. This results in the segmentation shown in figure 14.
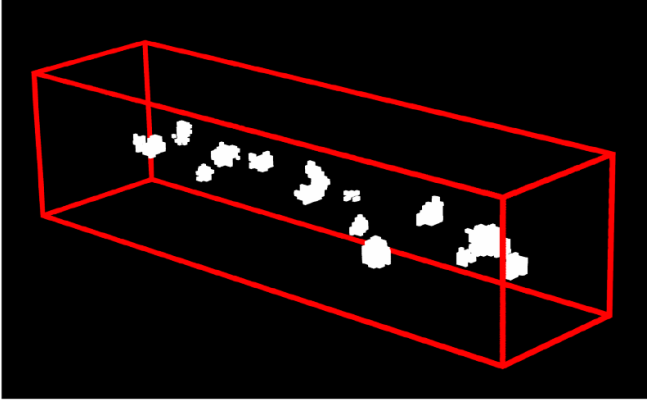
Figure 14. The ground truth pore segmentation used as input for the simulated X-CT volume. The volume indicated in red has $161 \times 41 \times 41$ voxels.
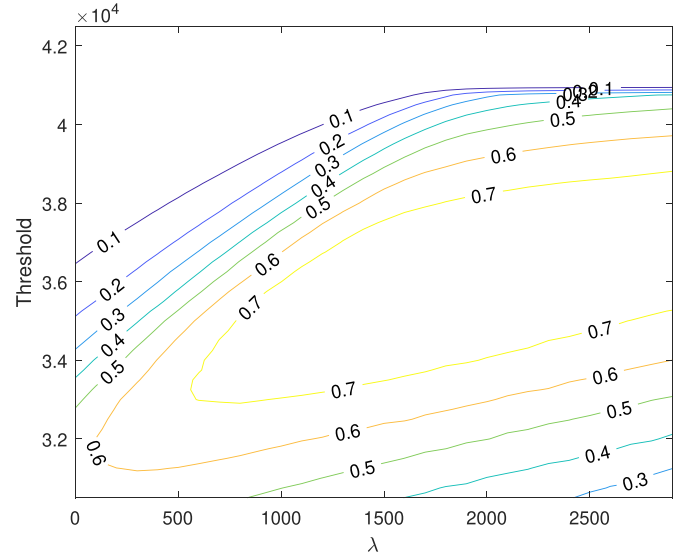

(a) Segmentation
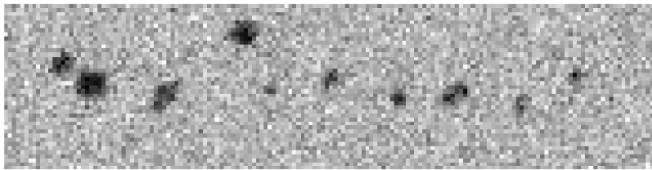
(b) Simulated CT

Figure 15. The central slice through the ground truth pore segmentation (a) and the corresponding simulated X-CT slice (b). Each image has 161×41 pixels.

(c) X-CT values were assigned to each voxel in this segmentation corresponding with the modes observed in figure 2 (4.1×10^4 for metal voxels and 2.3×10^4 for pore/air voxels).

(d) A partial volume effect was simulated by convolving the volume with a $5 \times 5 \times 5$ truncated and normalized Gaussian kernel, and Gaussian noise with a standard deviation $\sigma = 3000$ was added. The noise level was derived from the noise observed in the histogram in figure 2. See figure 15 for an example slice through the simulated volume.

This result can well be compared to the literature [23, 25] where an in-depth study on the generation of simulated X-CT data for pore detection in metal parts was made. The data set created in this way has a similar distribution of voxel values as the real data sets, has known ground truth based on real pores, and is visually indiscernible from a real X-CT data set.

A similar evaluation as in previous subsection was performed on this data set, where it was denoised and segmented for different values of λ and the threshold, and the

Figure 16. The contour plot of the Dice coefficient as a function of λ and threshold.

resulting Dice coefficients with the ground truth segmentation were calculated. The contour plot of this analysis can be found in figure 16. This plot shows that also in this experiment, a large plateau of significantly improved values can be found compared to the original segmentations (given by the values at $\lambda = 0$), which illustrates the improvements in segmentation that this technique can deliver.

The optimal segmentations for the original data set and the denoised data set, in terms of maximal Dice coefficient, are shown in figure 17 as concatenated 3D visualizations. The typical segmentation issues in a regular X-CT data set are clearly present here:

- Voxel noise shows up as individually segmented voxels. This is not present in the denoised segmentation.
- The partial volume effect combined with voxel noise causes ragged edges in the segmentation, which improves significantly in the denoised version.
- Small protruding parts of a pore are difficult to segment in either method, although this problem occurs more often without denoising.

5.4. Destructive testing

After the X-CT scans of the object were acquired and validated, the object was prepared for metallography by mounting in resin, and mechanically cut along one of the tracks that contains a high pore density via EDM wire cutting. This cut plane was polished to reveal a 2D cross section of pores, which was imaged at high resolution with optical microscopy. A careful manual alignment of the X-CT data with metallography cross-section was made.

The cross-section image is shown in figure 18, along with the original X-CT slice that corresponds with this cut, and the X-CT slice denoised with the proposed TV denoising technique. Segmentations are presented as well, where the
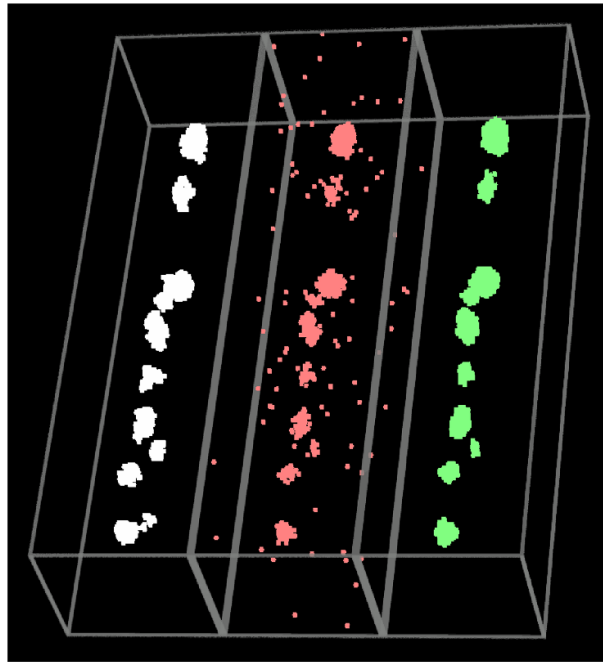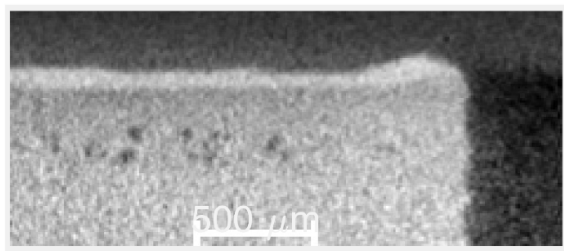
Figure 17. A 3D visualization of the pore segmentations, with the ground truth solution in white, the best scoring segmentation on the simulated data set in red, and the segmentation of the denoised solution with $\lambda = 2000$ and threshold $36\,000$ in green.
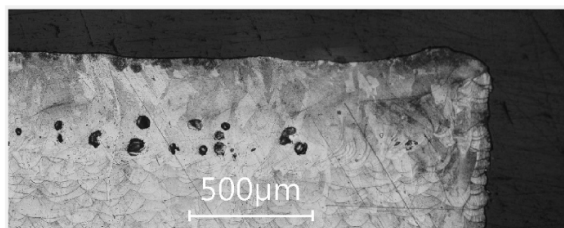
(a) Cut image

(b) Segmented cut image

(c) Original CT

(d) Segmented original CT

(e) Denoised CT

(f) Segmented denoised CT

Figure 18. The high-resolution image of the cut and polished object (top row), the original X-CT slice that corresponds maximally with this cut (middle row) and the same slice after denoising (bottom row). The right columns shows the segmentation results.

thresholds were manually selected based upon the image histograms and the resulting segmentation.

Quantitative results are hard to obtain in this type of experiment, as the employed metric is very sensitive to geometrical issues such as alignment and scaling, and threshold selection. Qualitatively, the denoised X-CT image looks visually smoother with better defines borders around the complex pores. Also the segmentation shows less false positives, while retaining at least the larger pores visible in the cut image.

6. Conclusions

We have presented a TV denoising technique for denoising X-CT images of additive manufactured parts with the goal of improving pore segmentation. The technique is based upon a popular optimization scheme developed by Chambolle and Pock [15], which already finds other applications in image denoising and medical X-CT imaging. The parameter space of the algorithm is explored within the context of the pore segmentation problem, an optimal parameter set is derived, and robustness against deviations from this set is demonstrated.

Several experiments are presented, where low- and high-quality X-CT scans are compared, a simulated X-CT data set is used to assess the technique, and destructive testing is employed to obtain a ground-truth 2D cut through the object. These results clearly indicate that pore segmentation improves in quality after application of the TV denoising technique, assessed by the Dice coefficient against a ground-truth segmentation. As the proposed algorithm can be easily implemented and can be sped up by parallellization, practical implementation in additive manufacturing quality processing chains should be relatively straightforward.

Possible future work is the inclusion of the proposed TV denoising technique into the X-CT reconstruction algorithm itself. This approach can be accommodated by the Chambolle–Pock algorithm, and has already been shown to be feasible in medical X-CT image reconstruction. The advantage then would be to have a single-step algorithm where the optimization happens against the actually measured X-CT sinogram, as compared to the current two-step approach of reconstruction followed by denoising. We expect such an integrated algorithm to perform superior to the currently proposed version.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgment

ORCID iDs

Rob Heylen ⬛ https://orcid.org/0000-0002-8341-9397
Aditi Thanki ⬛ https://orcid.org/0000-0002-6834-6900
Dries Verhees ⬛ https://orcid.org/0000-0002-5584-4601
Domenico Iuso ⬛ https://orcid.org/0000-0001-7919-7878
Jan De Beenhouwer ⬛ https://orcid.org/0000-0001-5253-1274
Jan Sijbers ⬛ https://orcid.org/0000-0003-4225-2487
Ann Witvrouw ⬛ https://orcid.org/0000-0003-3693-9290
Han Haitjema ⬛ https://orcid.org/0000-0001-7955-875X
Abdellatif Bey-Temsamani ⬛ https://orcid.org/0000-0002-8198-0965

References

[1] Debroy T *et al* 2018 Additive manufacturing of metallic components–process, structure and properties *Prog. Mater. Sci.* **92** 112–224

[2] Yap C Y, Chu a C K, Dong Z L, Liu Z H, Zhang D Q, Loh L E and Sing S L 2015 Review of selective laser melting: materials and applications *Appl. Phys. Rev.* **2** 041101

[3] Boley C D, Khairallah S A and Rubenchik A M 2015 Calculation of laser absorption by metal powders in additive manufacturing *Appl. Opt.* **54** 2477–82

[4] Wu Y-C, San C-H, Chang C-H, Lin H-J, Marwan R, Baba S and Hwang W-S 2018 Numerical modeling of melt-pool behavior in selective laser melting with random powder distribution and experimental validation *J. Mater. Process. Technol.* **254** 72–78

[5] Zhang X, Yocom C J, Mao B and Liao Y 2019 Microstructure evolution during selective laser melting of metallic materials: a review *J. Laser Appl.* **31** 031201

[6] Vrancken B, Thijs L, Kruth J-P and Van Humbeeck J 2012 Heat treatment of Ti6Al4V produced by selective laser melting: microstructure and mechanical properties *J. Alloys Compd.* **541** 177–85

[7] Zhang B, Li Y and Bai Q 2017 Defect formation mechanisms in selective laser melting: a review *Chin. J. Mech. Eng.* **30** 515–27

[8] Spierings A, Schneider M and Eggenberger R 2011 Comparison of density measurement techniques for additive manufactured metallic parts *Rapid Prototyp. J.* **17** 380–6

[9] Carmignato S, Dewulf W and Leach R 2018 *Industrial X-Ray Computed Tomography* (Berlin: Springer)

[10] Thompson A, Maskery I and Leach R K 2016 X-ray computed tomography for additive manufacturing: a review *Meas. Sci. Technol.* **27** 072001

[11] Bellens S, Vandewalle P and Dewulf W 2021 Deep learning based porosity segmentation in X-ray CT measurements of polymer additive manufacturing parts *Proc. CIRP* **96** 336–41

[12] Rodrí-guez-Sánchez A, Thompson A, Körner L, Brierley N and Leach R 2020 Review of the influence of noise in x-ray computed tomography measurement uncertainty *Precis. Eng.* **66** 382–91

[13] Hansen D C and Sorensen T 2018 Fast 4D cone-beam CT from 60 s acquisitions *Phys. Imaging Radiat. Oncol.* **5** 69–75

[14] Ritschl L, Sawall S, Knaup M, Hess A and Kachelriess M 2012 Iterative 4D cardiac micro-CT image reconstruction using an adaptive spatio-temporal sparsity prior *Phys. Med. Biol.* **57** 1517–25

[15] Chambolle A and Pock T 2011 A first-order primal-dual algorithm for convex problems with applications to imaging *J. Math. Imaging Vis.* **40** 120–45

[16] Mory C and Jacques L 2014 A modified 4D ROOSTER method using the Chambolle–Pock algorithm *Proc. 3rd Int. Conf. on Image Formation in X-Ray CT* pp 191–3

[17] Taubmann O, Haase V, Lauritsch G, Zheng Y, Krings G, Hornegger J and Maier A 2017 Assessing cardiac function from total-variation-regularized 4D C-arm CT in the presence of angular undersampling *Phys. Med. Biol.* **62** 2762–77

[18] Nikitin V, Carlsson M, Andersson F and Mokso R 2019 Four-dimensional tomographic reconstruction by time domain decomposition *IEEE Trans. Computat. Imaging* **5** 409–19

[19] Condat L 2013 A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms *J. Optim. Theory Appl.* **158** 460–79

[20] De Samber B, Renders J, Elberfeld T, Maris Y, Sanctorum J, Six N, Liang Z, De Beenhouwer J and Sijbers J 2021 FleXCT: a flexible X-ray CT scanner with 10 degrees of freedom *Opt. Express* **29** 3438–57

[21] Danckaers F, Huysmans T, Lacko D, Ledda A, Verwulgent S, Van Dongen S and Sijbers J 2014 Correspondence preserving elastic surface registration with shape model prior *2014 22nd Int. Conf. on Pattern Recognition* (IEEE) pp 2143–8

[22] Otsu N 1979 A threshold selection method from gray-level histograms *IEEE Trans. Syst. Man Cybern.* **9** 62–6

[23] Fuchs P, Kröger T, Dierig T and Garbe C 2019 Generating meaningful synthetic ground truth for pore detection in cast aluminum parts *Proc. Conf. on Industrial Computed Tomography (ICT2019)* (*Padua, Italy*)

[24] Thanki A, Goossens L, Mertens R, Probst G, Dewulf W, Witvrouw A, and Yang S 2019 Study of keyhole-porosities in selective laser melting using X-ray computed tomography (available at: www.ndt.net/iCT2019) pp 1–7

[25] Bircher B A, Wyss S, Gage D, Küng A, Körner C and Meli F 2021 High-resolution x-ray computed tomography for additive manufacturing: towards traceable porosity defect measurements using digital twins *Joint Special Interest Group Meeting Between Euspen and ASPE Advancing Precision in Additive Manufacturing Inspire AG* (*St. Gallen, Switzerland*)