Key Engineering Materials
ISSN: 1662-9795, Vol. 1047, pp 191-204

Submitted: 2026-01-19
Accepted: 2026-02-11
Online: 2026-04-13

Volumetric Anomaly Detection in LPBF by Segmentation and Classification of Exposure Optical Tomography (EOT) Image Stacks

Karsten Scheibe[1,a*], Johann Albers[1,b], Felix Jensch[1,c], Chukwuemeka Okolo[1,d], and Sebastian Härtel[1,e]

Department of Hybrid Manufacturing[1], Brandenburg University of Technology Cottbus–Senftenberg, Konrad-Wachsmann-Allee 17, 03046 Cottbus, Germany

[a]Karsten.Scheibe@b-tu.de, [b]Johann.Albers@b-tu.de, [c]Felix.Jensch@b-tu.de, [d]okolo@b-tu.de, [e]Sebastian.Haertel@b-tu.de

Keywords: LPBF, Exposure OT, Image Segmentation, Anomaly Detection, Machine Learning

Abstract. In Laser Powder Bed Fusion (LPBF), ensuring reproducible part quality remains a major challenge despite the availability of high-resolution in-situ monitoring systems, such as Exposure Optical Tomography (EOT). While EOT provides detailed layer-wise optical information, most existing approaches focus on single-layer analysis or real-time process control and do not exploit the full volumetric information contained in the acquired data.

This work presents a modular framework for the volumetric reconstruction and post-process analysis of EOT image data. Sequential EOT images are processed using volumetric component segmentation (VCS) and fused into a three-dimensional OT Image Cube, forming a central volumetric data structure called LPBF Cube. Each voxel encodes spatial and radiometric information and can optionally be augmented with additional process metadata.

Based on this representation, high-resolution two-dimensional slices are rendered along arbitrary orientations using profile-based slicing strategies for planar, cylindrical, and complex geometries. These slices enable intuitive, part-level inspection of laser exposure history and spatial process variations. The framework is validated using geometric and radiometric analyzes, demonstrating good agreement with nominal CAD geometry and a clear correlation between EOT-derived emission values, laser energy input, and local cross-sectional area.

The proposed approach extends the use of EOT data beyond layer-wise monitoring toward comprehensive, volumetric part inspection and provides a practical basis for geometry-aware quality assessment in LPBF, particularly for prototyping and post-build evaluation.

Introduction

Laser Powder Bed Fusion (LPBF) is one of the most mature and widely adopted technologies in metal Additive Manufacturing. Despite its industrial relevance, ensuring reproducible part quality remains a significant challenge. Process-induced variations—such as fluctuations in melt pool dynamics, laser–powder interactions, or insufficient fusion may lead to defects including porosity, lack of fusion (LoF), and local overheating zones [1]. To mitigate such deviations, the development of robust in-situ monitoring technologies has become increasingly important.

Exposure Optical Tomography (EOT) represents a promising sensing modality that captures high-resolution transmission or emission maps for each processed layer during LPBF [2]. These images encode information about the optical behavior of the powder bed and the solidification process; local intensity variations often correlate with process anomalies and may reveal indicators of melt pool instability, powder distribution irregularities, or localized defects [3]. Consequently, numerous studies investigate signature-based EOT or OT to analyze individual layers or to support machine-learning-based defect detection [4–7]. Moreover, several studies have reported a strong correlation between EOT-derived features and defect indications obtained from X-ray micro-computed tomography (μCT), demonstrating the suitability of EOT data for defect assessment, for example, in [8].

However, the vast majority of existing research focuses on layer-wise monitoring aimed primarily at real-time process control, anomaly detection, or parameter optimization. A comprehensive perspective on the entire manufactured part—across all layers and the full build volume—has received much less attention. However, for prototyping, post-build evaluation, and the identification of design- or manufacturing induced weak points in critical regions that may result in LoF or keyhole formation, which are commonly identified as pores, a holistic part-level interpretation is crucial. In many cases, only minor adjustments to the part design or the build planning, particularly the scan or laser path strategy, are sufficient to achieve improved process stability and part quality while maintaining identical functional performance.

In contrast to conventional layer-centric approaches, this work introduces a volumetric methodology in which EOT images are not analyzed in isolation but fused into a unified three-dimensional data representation. Starting from the complete stack of EOT images, a 3D OT Image Cube is reconstructed, from which 3D Data (i.e., point clouds or meshed vertices) in standard industrial formats (e.g., STL or PLY) are derived. This geometric representation forms the basis for subsequent segmentation and classification tasks, enabling structural inspection of the printed component.

Furthermore, the volumetric representation allows the generation of arbitrary high-resolution 2D cross-sectional slices extracted from the 3D OT Cube. These renderings facilitate detailed analysis of the laser exposure history and the spatial distribution of potential process anomalies across multiple adjacent layers. As a result, the proposed approach provides a powerful tool for part-level quality assessment—particularly valuable in prototyping and design validation contexts.

Thesis Statement

The volumetric reconstruction of a 3D OT Image Cube from sequential EOT layer images enables the generation of arbitrary high-resolution 2D cross-sections, significantly facilitating the identification of design-induced weak points as well as process-related anomalies throughout the entire part. This capability goes well beyond the limitations of purely layer-based in-situ monitoring approaches. By integrating all EOT layers into a unified three-dimensional voxel matrix, a central volumetric data structure—the LPBF Cube—is formed. Each voxel stores not only its spatial coordinates and local emission value but may optionally contain additional process metadata such as timestamps, scan strategy, laser power, or acoustic process signatures. This enriched volume model establishes the foundation for a holistic, multi-perspective analysis of the manufactured component.

This article underscores the following key contributions:

1. Reconstruction of a 3D Model and development of various tools for comprehensive analysis of OT Image Data, including derivation of STL/PLY data and volumetric data structures for downstream evaluation.
2. High-Resolution 2D Image Rendering from the 3D OT Image Cube, enabling flexible visualization and targeted analysis of laser exposure history and spatial anomalies.

Methodology for Processing EOT Data

The proposed framework comprises three main components:

1. data acquisition and preprocessing, including initial volumetric component segmentation;
2. volumetric reconstruction and 3D model generation; and
3. segmentation, classification, and analytical evaluation.

The overall architecture is designed to support both geometric reconstruction of the part and process-oriented quality assessment.

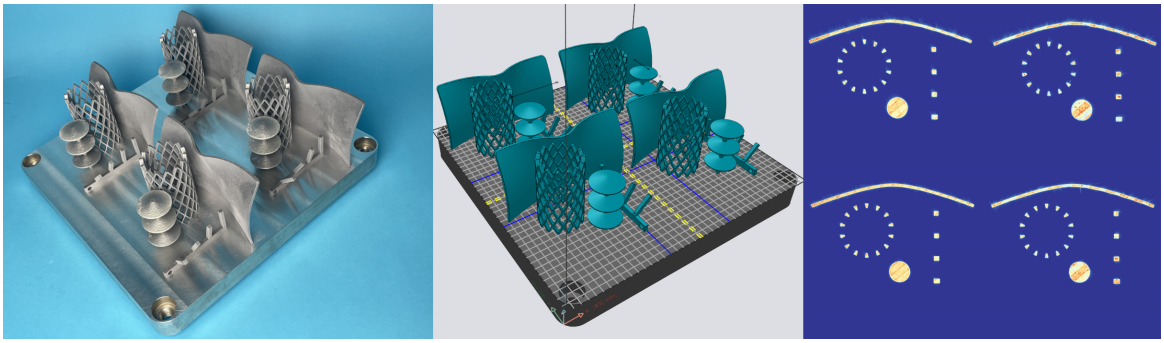

Fig. 1: Complementary views of the input dataset: a representative photograph of the finished part (left), the corresponding CAD model defining the intended geometry (center), and a typical EOT slice acquired during the LPBF process (right).

Acquisition and Pre-Processing of EOT Data

The components were manufactured on an EOS M400-4 L-PBF system, which has a build volume of 400 × 400 × 400 mm³. It is equipped with four 400 W Yb fiber lasers, each operating at a wavelength of 1070 nm. Each laser covers a working area of 250 × 250 mm² on the substrate plate, resulting in overlapping regions between the individual quadrants. All components were produced using identical parameters: a laser power of 370 W, an exposure speed of 1260 mm/s, a hatching distance of 0.18 mm, and a layer thickness of 80 µm. Nitrogen was used as a protective gas during fabrication, and the powder material was applied using a polymer lip. AlSi10Mg powder served as the base material for all manufactured components.

The component geometries were chosen to intentionally induce local overheating during production, in order to demonstrate the full potential of the method presented in this paper. The four components were fabricated once in each of the four laser working areas, with only one laser used per quadrant and slightly varied scan strategies and laser power between the quadrants. Additional parameter variations were investigated in separate builds for further studies but these results are outside the scope of the present work.

The L-PBF system is equipped with an exposure optical tomography (EOT) system, in which an sCMOS camera captures images in the near-infrared wavelength range across the entire build area. These images record radiation emissions originating from the laser–material interaction.

The EOT system provides two acquisition modes: a maximum gray-value mode and an integral gray-value mode. In this work, the integral gray-value mode is used, which accumulates the emitted radiation over the exposure time of each layer. While the recorded signal does not represent temperature directly, the temporal integration inherently captures thermal effects such as heat accumulation and heat flow within the material. Accordingly, the EOT signal reflects the effective energy input rather than instantaneous peak emissions.

Nevertheless, layer-integrated EOT signals have been shown to correlate strongly with the applied laser energy, local heat accumulation, and melt pool stability, making them highly suitable for component-level and process-consistent radiometric analyses. Accordingly, the EOT signal is interpreted in this work as a relative, process-dependent indicator of the effective thermal energy state rather than as an absolute physical quantity such as temperature or total energy.

Throughout the paper, and in particular in the radiometric validation section, a clear distinction is made between laser power applied per unit area and the effectively integrated energy input per layer, enabling a consistent interpretation of radiometric and process-related quantities.

The foundation of this work is a multi-modal dataset that combines conventional post-build inspection data, the corresponding CAD representation of the build job, and layer-wise EOT images acquired during fabrication. Figure 1 illustrates these three complementary perspectives: the left im-

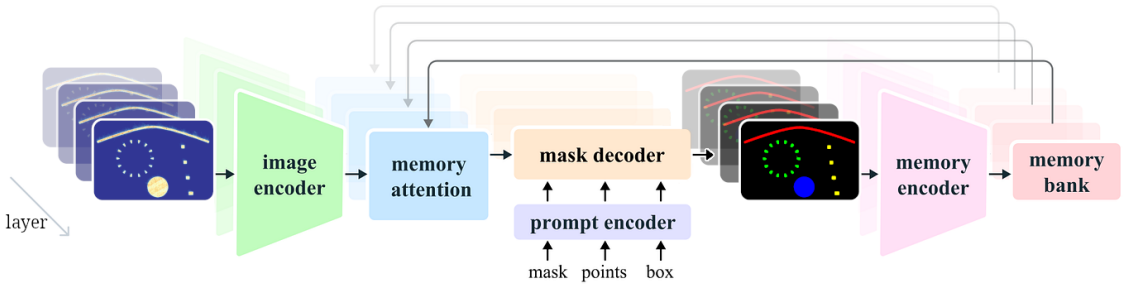

Fig. 2: Integration of the SAM2 architecture into the proposed framework for prompt-based volumetric component segmentation of memory-driven sequential EOT data in LPBF.

age shows a representative photograph of the completed part obtained after removal of the construction site and surface inspection; the center image depicts the underlying CAD model that defines the intended geometry of the component; and the right image provides a typical EOT slice captured during the LPBF process, representing the spatial distribution of optically emitted radiation for a single processed layer. For each manufactured layer, a corresponding EOT image is recorded, typically resulting in several thousand images per build (commonly between 1,000 and 5,000 layers per part). Before volumetric reconstruction, raw images undergo a standardized preprocessing pipeline that includes gray-value normalization, noise reduction, and geometric alignment to ensure spatial consistency across the stack. Additionally, the layer index is synchronized with the build-height information, enabling an unambiguous mapping of each image to its physical position within the final volumetric model.

Image Segmentation and Classification

Segment Anything Model 2 (SAM2) is the second generation of Meta's Segment Anything framework, designed to provide robust, promptable segmentation across a wide variety of visual domains. SAM2 introduces significant architectural and performance improvements over the original SAM, including a more efficient image encoder, a persistent memory mechanism for handling video or sequential imagery, and stronger generalization capabilities, particularly for non-natural or highly technical image data. As a prompt-driven model, SAM2 can generate high-quality object masks based on points, boxes, or coarse masks, making it well-suited for segmenting complex LPBF and EOT imagery where clear object boundaries are difficult to obtain with classical image-processing pipelines. Its ability to handle challenging illumination conditions, texture-poor regions, and irregular geometries makes SAM2 an effective tool for segmentation of volumetric components in additive manufacturing applications.

Figure 2 illustrates the SAM2 architecture and highlights its integration into the proposed framework, showing how prompt-based interaction and sequential EOT data are combined to perform volumetric component segmentation in the context of LPBF process monitoring [9].

Volumetric Component Segmentation (VCS). A central challenge in processing EOT data is the reliable isolation of the additively manufactured component from background structures such as the powder bed, platform boundaries, optical artifacts, or process noise. Traditional image-processing techniques—such as edge detection, threshold-based binarization, morphological operations (e.g., erosion or closing), or region-growing methods—have been applied in earlier studies to extract component contours. However, these approaches often suffer from sensitivity to local contrast variations, powder accumulation, spatter-induced artifacts, and layer-to-layer fluctuations, which can lead to fragmented or incomplete segmentation masks. To robustly address these limitations, the present work employs Volumetric Component Segmentation (VCS) based on modern deep-learning-driven foundation models such as SAM2. This model enables highly stable, geometry-aware separation of the component from its surroundings, even in the presence of significant noise or non-uniform illumination condi-

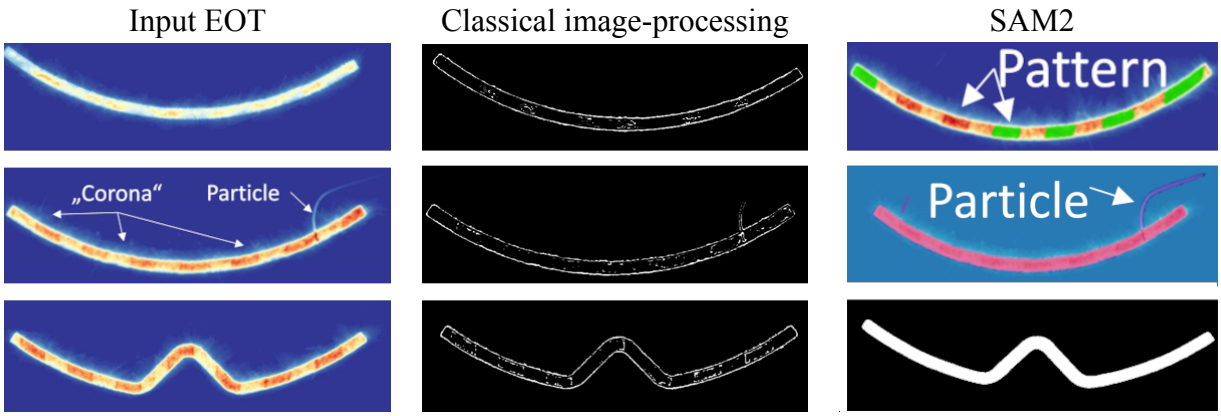

Fig. 3: Representative EOT slices of the curved wall example from Fig. 1, illustrating typical segmentation challenges.

tions typical for EOT imagery. SAM2 provides consistent per-slice masks that accurately delineate the part boundaries, thereby enabling the generation of a clean, binary segmentation volume.

Figure 3 illustrates these challenges by comparing classical image-processing results with representative examples of EOT slices. From top to bottom, the figure presents three characteristic cases. In each row, the left column shows the raw EOT input, the center column displays the corresponding mask obtained using classical image-processing techniques, and the right column shows the segmentation generated by SAM2, which yields robust and high-quality results across all scenarios.

The top row represents a comparatively clean and simple example, where both classical processing and SAM2 perform well. However, the classical pipeline exhibits slight instability due to pronounced patterning effects, whereas SAM2 can be explicitly prompted to recognize and classify these patterns as distinct regions if required (green patches detected by SAM2; two are disabled for improved visual clarity). It is important to note that such patterns should not be interpreted as defects as long as they originate from optical emission effects rather than from actual process irregularities. These patterning effects are primarily caused by process-related optical phenomena. For example, process emissions such as smoke or plume dynamics can disturb the optical signal, and incident laser power may be partially absorbed and re-emitted at different wavelengths (e.g., Raman scattering). To mitigate these influences (which cannot be entirely eliminated), a narrow near-infrared bandpass filter (center wavelength 900 nm, bandwidth 25 nm) is used in the EOT system, effectively suppressing plasma emissions and back-reflected laser light at 1060 nm. The observed patterns typically consist of two stripe-like patterns corresponding to antiparallel laser scan directions (180° apart).

The middle row highlights typical challenges already described above—most notably acquisition artifacts (e.g., spatter-induced particle) and the “corona” around the edges of the part (e.g., very small particles and condensate), where sharp boundary separation is often difficult for classical methods. Due to the high energy density, the melt pool reaches the evaporation point of certain alloying elements and as a consequence, the vaporized metal is ruptured out of the melt pool. Vaporized metal is cooled down very quickly and condensates, forming particles with very small diameters.

The bottom row shows the discussed odd–even pattern, for which simple threshold-based approaches may fail, while SAM2 still yields an accurate segmentation.

The resulting component mask is subsequently used for voxel-space reconstruction, where either the entire segmented region or only its convex hull can be considered, depending on the desired level of geometric fidelity. This mask defines the volumetric envelope of the reconstructed LPBF Cube and ensures that downstream analyzes operate exclusively within the relevant part volume.

Optional Segmentation and Process-State Classification. Beyond component-level segmentation, additional segmentation and classification of internal regions may be performed as an extended step in the analysis pipeline. Modern methods such as SAM2-based region segmentation, U-Net architec-

tures, self-organizing maps (SOM)[6], or clustering-based techniques can be applied to partition the EOT data into semantically meaningful subregions. Each segmented region can then be assigned to a predefined process class, for example: Normal Consolidation, LoF, Overheating, or other condition-specific categories reported in prior work. When applied voxel-wise across the entire 3D OT Image Cube, this classification produces a fully annotated volumetric representation of the component, capturing both geometric and process-state information.

This optional phase enables a richer interpretation of spatial defect patterns and their relationship to part geometry, but is conceptually distinct from the primary VCS step, which focuses solely on isolating the component volume. In addition, this extended segmentation stage is particularly beneficial for arbitrary 2D slice rendering, where class-aware slice extraction provides clearer, context-sensitive visualization of internal structures and defect regions.

Construction of a 3D OT Image Cube

In previous work, the construction and inspection of volumetric optical-tomography data has often relied on generic visualization tools such as the ImageJ *3D Viewer* plugin. While these tools provide a convenient means for rendering stacked EOT slices, they quickly reach their limits even for small to medium-sized parts due to out-of-memory constraints. Moreover, such approaches typically operate on the entire EOT acquisition without explicitly separating the part from the surrounding measurement domain. As a result, the generated volume represents a global EOT cube rather than a part-specific reconstruction, causing relevant structural details to be obscured by non-informative data and remaining identifiable—if at all—only through manual slicing. Although sufficient for basic visualization, such volumes lack semantic or structural linkage between rendered pixels and therefore do not represent a coherent, component-aware data structure.

By contrast, once individual components—or even sub-components—are separated using the described machine-learning-based segmentation methods, additional volumetric processing steps become straightforward to implement. In particular, volumetric region growth algorithms can be applied directly within the segmented voxel space, enabling the extraction of connected regions and the quantitative evaluation of geometric properties such as shape, size, and volume. These metrics can subsequently be compared against nominal CAD data to assess deviations in geometry and structural integrity and to derive indicators of manufacturing effects such as distortion or warping.

Figure 4 is organized into two rows and three columns. The top-left panel schematically illustrates the generation of component masks from a stack of approximately 2,000 sequential EOT images using volumetric component segmentation. In the center column, example segmentation results are shown for two representative layers: Layer 0 (top) and Layer 1500 (bottom). Here, individual components are visualized as color-coded classes, with the curved wall shown in red, the cylindrical lattice structure in green, the benchmark tree element in blue, and a sloped structure in yellow. The dotted profiles shown in Layer 1500 are not required for the three-dimensional reconstruction but are included here for illustration purposes and will be explained in detail in the context of high-resolution 2D slice rendering.

The top-right panel presents an aggregated view of the four individual components, displayed according to their segmentation-class masks. The bottom-left panel highlights the lattice cylinder as an example of a single sub-component. In the bottom-right panel, the same four components are shown, but the color corresponds to the original EOT emission values rather than the class-based segmentation.

In the bottom-left and bottom-right panels, the components are rendered using different color schemes to emphasize distinct aspects of the EOT data. For the lattice cylinder shown in the bottom-left panel, the rendering represents the integrated emission values of the original EOT images (without applying a jet color-map), providing a direct visualization of the raw layer-wise signal distribution.

Both panels exhibit slightly bluish shading along the boundary of the components. As shown in the zoomed view of the top-left panel (red arrows), these regions correspond to low-emission con-

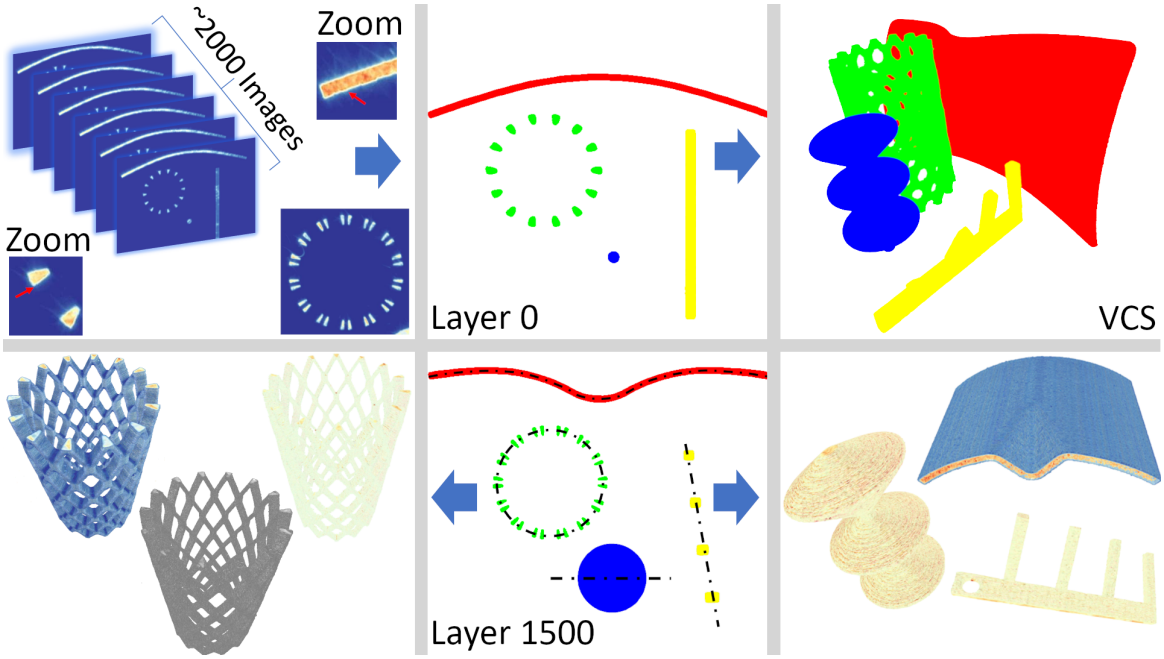

Fig. 4: Illustration of volumetric component segmentation and reconstruction from EOT image stacks, showing class-based masks at different layers and component-specific volumetric representations.

tour scans, which are intentionally exposed with reduced energy and therefore appear radiometrically "cold". These effects must not be confused with process-induced LoF defects. In EOT data, however, both phenomena may exhibit very similar low-intensity signatures. In addition, previously described corona-related artifacts may occur in the vicinity of component edges; these arise from optical effects such as scattering, condensate, or stripe-border effects and represent a further independent source of low-signal regions.

To avoid false-positive LoF detection at the component boundaries, the VCS-derived segmentation mask is first eroded by a small number of pixels. This purely geometric preprocessing step removes the contour region from the subsequent analysis. Only after this erosion step is a radiometric threshold applied to identify low-emission regions within the remaining interior of the component. This sequential combination of geometric mask erosion and radiometric thresholding ensures that intentionally low-energy contour exposures and corona-related edge artifacts are excluded, while potential LoF-related low-emission regions inside the bulk material are preserved. As a result, LoF detection becomes more robust and less sensitive to boundary-related artifacts, without sacrificing sensitivity to process-induced defects in the component interior.

Building on the volumetric reconstruction and VCS-based component segmentation, the following example demonstrates how radiometric information can be exploited for intuitive part-level analysis. Figure 5 illustrates the potential to identify overheated regions within the volume of reconstructed components. For narrow-band EOT images, such regions can often be highlighted without the need for elaborate algorithms. The three views on the left show a benchmark lattice cylinder rendered as a VCS-derived point cloud, where the base geometry is displayed in medium gray and regions exceeding specified emission thresholds are emphasized. In the right lattice cylinder view, the remaining gray points therefore represent regions well below the mean energy level. Regions far below the mean could alternatively be highlighted using a contrasting color (e.g., blue), which is not shown here.

The rightmost panel of Fig. 5 demonstrates the same concept applied to a different component geometry. In contrast to the lattice cylinder shown in the left panels, this example represents a fully solid component, whereas the lattice cylinder consists primarily of a thin, hollow lattice ring. In this case, the VCS-derived segmentation masks are exploited not only for radiometric thresholding but also for efficient geometric processing.

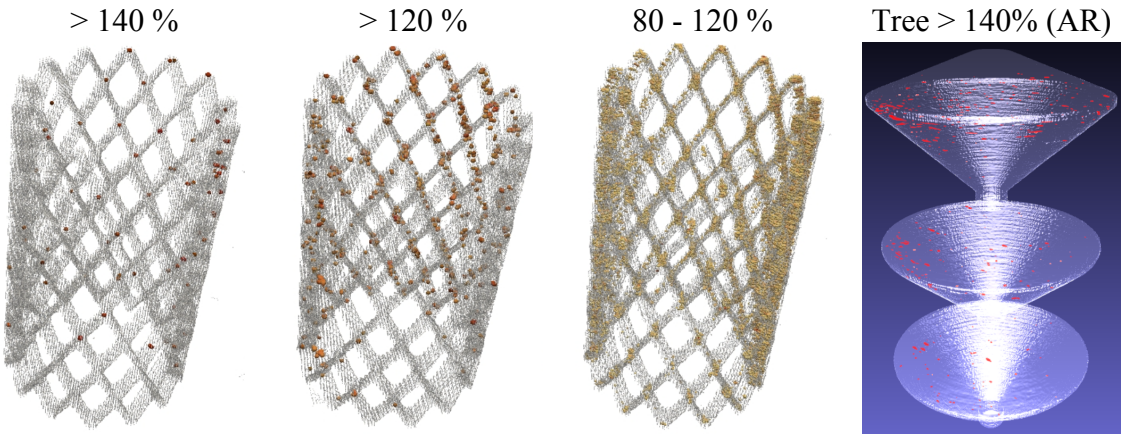

Fig. 5: VCS-based renderings with additional thresholding of EOT emission relative to its mean. Left: three point-cloud views of a lattice cylinder. Right: an alternative transparent surface rendering (AR), derived from VCS masks using erosion and convex hull extraction of the benchmark tree.

The component boundary is determined through mask erosion, after which a convex hull is computed directly on the segmented volume. Compared to operating in the full point cloud, this approach significantly simplifies hull generation and subsequent meshing. The resulting surface mesh is rendered using a transparent, glass-like material, while regions exceeding a defined emission threshold are overlaid as red points within the enclosed volume. This combined visualization allows simultaneous inspection of the internal high-energy regions and the geometry of the enclosed component.

Together, these examples demonstrate that even simple threshold-based analyzes applied to segmented volumetric data can reveal meaningful spatial patterns of elevated energy input. When combined with mask-based geometric processing, the approach enables intuitive and efficient visualization of both radiometric and structural information, complementing the segmentation- and cube-based volumetric analyzes described above.

Finally, this volumetric representation enables a wide range of downstream analyzes that go beyond purely visual inspection. In addition to geometric validation and qualitative evaluation, the reconstructed volume provides a consistent basis for the quantitative evaluation of structural features, defect-prone regions, and process-dependent variations within the component. Because volumetric data are derived directly from the layer-resolved EOT acquisition and are constrained by the VCS segmentation, all subsequent analyzes remain inherently component-aware and spatially consistent across the entire build volume. The reconstructed model can be exported in standard industrial formats such as STL or PLY, enabling direct interoperability with established analysis and inspection tools, including MeshLab, CloudCompare, and numerical simulation environments. This allows, for example, comparison with nominal CAD geometries, mesh-based deviation analysis, or integration of the reconstructed component into downstream simulation workflows. At the voxel level, the LPBF Cube acts as a unified container that combines geometric, radiometric, and process-related information. Each voxel in the reconstructed volume stores at least the following information:

1. spatial coordinates (x, y, z),
2. EOT emission values derived from the optical tomography images, and
3. optionally, additional process metadata such as timestamp, scan strategy, laser power, or acoustic signals.

By aggregating this information into a single volumetric data structure, the resulting 3D representation forms the central volumetric *LPBF Cube*. It provides a flexible and extensible foundation for component-level inspection, multi-modal data fusion, and future process-aware analysis approaches in LPBF.

High-Resolution 2D Slice Rendering

While the construction of the 3D OT Image Cube primarily enables volumetric visualization and component-level inspection, many process-related phenomena and defect patterns remain difficult to interpret in purely three-dimensional renderings. For detailed analysis, engineers and process experts, therefore, rely on two-dimensional cross-sectional views that provide direct insight into internal structures and local variations of energy input. Building upon the volumetric representation introduced in the previous section, high-resolution 2D slices are generated by performing virtual cuts through the volume along arbitrary orientations, including standard Cartesian planes (xy, yz, xz) as well as freely defined slicing directions. These slices enable detailed analysis of laser energy distributions, facilitate the identification of vertically extended defect structures, and provide intuitive cross-sectional representations suitable for engineering assessment. All slicing operations described in the following are performed on the VCS-derived masks (cf. dotted lines in Fig. 4), ensuring that all extracted views are strictly constrained to the segmented component geometry. However, the rendering itself is performed directly on the original EOT images rather than on the reconstructed point clouds, which may suffer from resolution loss due to point reduction or meshing operations.

Profile-Based Slicing for simple component geometries, linear slicing profiles are sufficient to generate planar cross-sections that are well aligned with the part geometry. While comparable slicing functionality is available in most common CAD environments, these basic slicing operations are implemented as lightweight “gadgets” within the proposed framework to enable direct interaction with the EOT-based volumetric data. Such profile-based slicing is particularly suitable for components whose cross-sectional shape remains invariant along the build direction with respect to the slicing profile. In these cases, the profile can be defined in various ways, including freehand drawing, spline-supported curves, or polylines. A representative example is the blue benchmark tree shown in Figure 4: although the structure exhibits gradual tapering along the build direction, the underlying slicing profile remains unchanged across layers, allowing consistent extraction of 2D cross-sections throughout the component height directly in the VCS mask.

Figure 6 shows, on the left, the rendered benchmark tree at an approximately isotropic resolution. This isotropy is achieved by exploiting the inherently higher physical resolution in the build direction (80 µm), while the in-plane (xy) resolution is enhanced using high-quality interpolation based on the SciPy library (mode=’reflect’, anti_aliasing=True, preserve_range=True). The left rendering preserves the original radiometric values of the EOT data, whereas the right rendering applies median smoothing to mitigate layer-wise pattern variations caused by alternating laser scan directions along the build axis. As a result, the latter representation emphasizes the cumulative laser energy distribution across the entire component within the selected slice. Optionally, individual slices can either be analyzed independently or aggregated across multiple layers and processed using the algorithms introduced earlier, which is particularly beneficial for identifying vertically extended phenomena such as keyhole-related features.

Cylindrical and Revolved Profiles for rotationally symmetric structures, such as lattice cylinders, cylindrical slicing profiles are employed. These profiles can be unwrapped into two-dimensional representations, enabling continuous inspection of circumferential features within a single image. To facilitate this process, a lightweight three-point selection gadget is implemented within the proposed framework, allowing the slicing profile to be interactively defined directly on the VCS-derived masks and subsequently rendered as a high-resolution 2D slice.

The primary advantage of this representation lies in the ability to visualize the entire unwrapped surface of the component. In contrast to conventional planar cuts in CAD environments—where only a narrow set of pixels intersecting the cutting plane, or at best all pixels behind the plane, are displayed—the unwrapped view provides complete circumferential coverage. Figure 6 illustrates this approach in the center panel, showing an unwrapped lattice cylinder in which all lattice intersection points are fully represented. The coordinate system of the unwrapped image is defined consistently

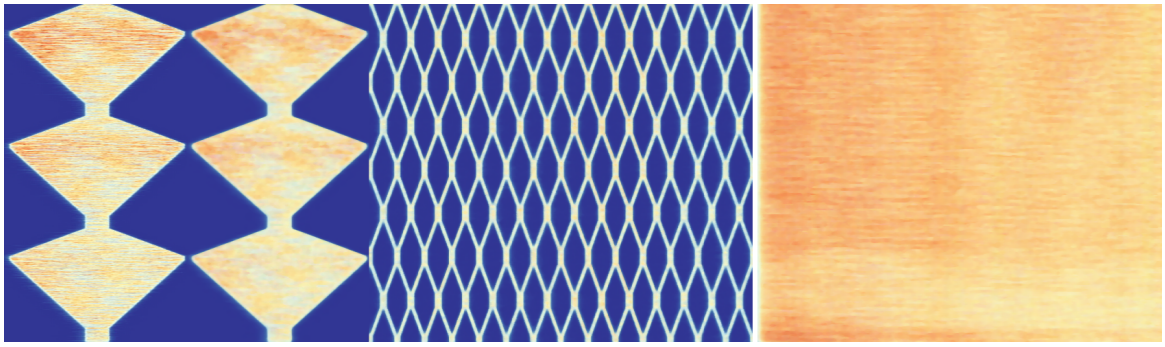

Fig. 6: High-resolution 2D slice renderings obtained from VCS-derived profiles, illustrating planar slicing of a benchmark tree (left), cylindrical unwrapping of a lattice structure (center), and skeleton-based slicing of a curved wall (right).

with the physical build geometry. Along the vertical axis (y-axis), the coordinate corresponds directly to the build direction and is metrically scaled according to the thickness of the layer (80 µm). Along the horizontal axis (x-axis), the coordinates are expressed in an angular space, preserving the angular fidelity in polar coordinates. Consequently, a constant arc length is maintained along the slicing profile. The angular resolution in x is adapted to achieve an approximate isotropy with the higher resolution in the build direction and corresponds to approximately 0.25° per pixel, which is equivalent to an arc length of 80 µm in the original rendered images.

Skeleton-Based Slicing for Complex Geometries such as curved walls, skeleton-based slicing strategies are applied [10, 11]. This approach relies on the extraction of the medial axis of the segmented component, which provides a compact representation of its topology and central geometry. Given a VCS binary segmentation mask $M(x, y)$, the skeleton $S \subset M$:

$$S = \left\{ p \in M \;\middle|\; \exists\, q_1, q_2 \in \partial M,\ q_1 \neq q_2,\ d(p, q_1) = d(p, q_2) = \min_{q \in \partial M} d(p, q) \right\}$$

is defined as the set of all points p inside the segmentation mask M for which there exist at least two distinct boundary points q_1 and q_2 on ∂M such that the distance from p to each of these boundary points is equal and corresponds to the minimum distance from p to the boundary. This formulation yields a one-pixel-wide centerline that captures the essential shape of the component and serves as a curved slicing path. Here, ∂M denotes the boundary of the segmentation mask M, i.e., the set of all pixels in M that are adjacent to the background and thus form the interface between the segmented component and its surrounding domain.

Based on the medial axis extracted from the VCS-derived masks, a curved slicing profile is generated and dynamically adapted across the build layers to account for geometric evolution along the build direction. In such cases, the geometry of the components may change significantly with increasing build height, for example, due to gradual deformation or flaring (outward bulging) toward the top layers (see Fig. 4, center column, comparing the red wall in Layer 0 and Layer 1500). Figure 6 illustrates the resulting cut of the curved wall in the rightmost panel, shown as an orthogonal projection along the derived curved profile. As with the unwrapped cylindrical representations, this approach offers the key advantage that all pixels associated with the slicing profile are projected into a single 2D image, rather than only those intersecting a planar cutting surface. This enables continuous inspection of complex, non-planar structures either by following a single centerline-based slicing profile or by aggregating all profile samples between the minimum and maximum extents across successive build layers into a single two-dimensional representation, while preserving the full spatial context of the component.

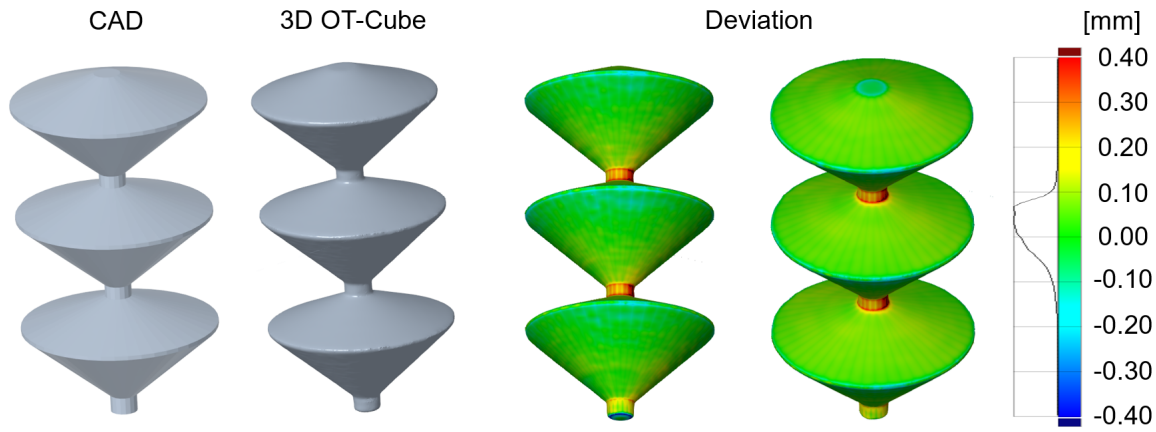

Fig. 7: Geometrical validation of the reconstructed 3D OT-Cube against the nominal CAD geometry.

Validation

To assess the reliability and applicability of the proposed framework, a validation is performed with respect to both geometric accuracy and radiometric consistency. The validation focuses on evaluating how well the reconstructed volumetric representations derived from EOT data correspond to the nominal part geometry and whether the preserved optical information enables meaningful interpretation of process-related effects. Accordingly, the validation is divided into a geometric and a radiometric analysis.

Geometrically Validation aims to quantify the deviation between the reconstructed 3D OT-Cube and the nominal CAD geometry of the manufactured components. For this purpose, the volumetric reconstruction obtained from the EOT data is converted into a surface representation and spatially aligned with the corresponding CAD model. A deviation analysis is then performed by computing point-wise distances between the reconstructed geometry and the reference CAD surface.

Figure 7 illustrates the geometric validation workflow. The left side shows the input data, consisting of the original CAD model and the reconstructed 3D OT-Cube. On the right, a color-coded deviation map visualizes the local geometric differences between both models, with the deviation magnitude expressed in millimeters. This comparison allows systematic evaluation of dimensional fidelity and highlights regions where process-induced distortions or reconstruction inaccuracies may occur.

Radiometrical Validation analyzes the relationship between the optical emission information captured by the EOT system and the applied accumulated laser energy during the LPBF process. To this end, the benchmark tree geometry is evaluated using aggregated profile-based slicing, in which all slicing profiles spanning the local component cross-section are accumulated rather than relying on a single, arbitrarily chosen centerline. This approach enables a truly volumetric analysis by capturing the full radiometric response of the component across successive build layers while simultaneously reducing local scan-pattern-induced artifacts. By integrating all profile samples, the resulting two-dimensional representations reflect the cumulative energy input within the component and allow systematic investigation of both global trends along the build direction and geometry-dependent variations in laser exposure. As a result, correlations between radiometric emission, applied laser energy, and segmented cross-sectional area can be analyzed in a robust and physically meaningful manner.

Figure 8 illustrates the results of this analysis. On the left, high-resolution 2D slice renderings of the benchmark tree are shown, obtained by accumulating all slicing profiles across the local component cross-section, as described in Section *High-Resolution 2D Slice Rendering*. In contrast to centerline-based approaches, this representation integrates the radiometric contribution of the entire profile region and thus captures the full volumetric emission response associated with each individual build layer.

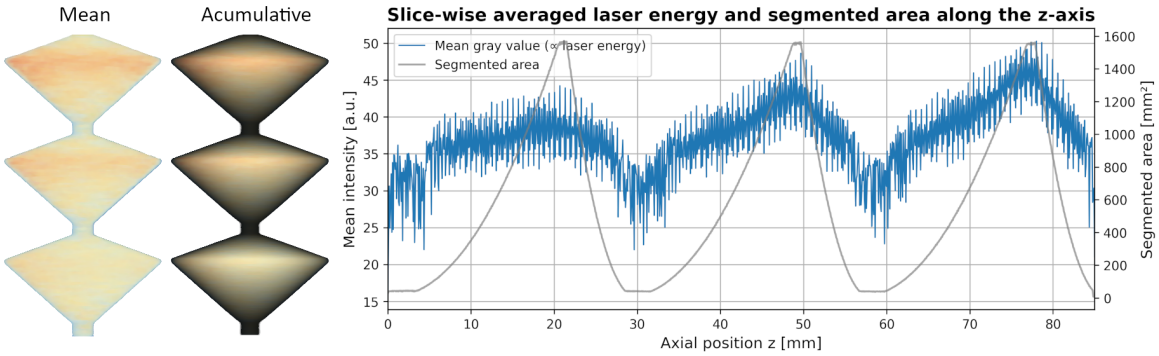

Fig. 8: Radiometric validation based on aggregated profile slicing of the benchmark tree, showing high-resolution 2D renderings and slice-wise emission statistics along the build direction.

Two complementary radiometric representations are shown. The first representation normalizes the accumulated emission by the effective profile extent, yielding a geometry-aware measure of the mean radiometric emission per unit area within each layer. Although evaluated on a per-layer basis, this quantity represents an area-normalized radiometric response rather than the total energy deposited within the layer. The second representation depicts the accumulated radiometric energy per layer without normalization by the number of contributing samples, thereby emphasizing the total effective energy input into the component cross-section. This accumulated view amplifies geometry-induced effects and highlights regions of increased thermal loading that are not apparent in purely area-normalized representations. On the right, the slice-wise radiometric measures are plotted as a function of the axial build position z. Both radiometric representations reveal a gradual increase in emission intensity along the build height, indicating a systematic process-related trend over the entire component. This build-height-dependent increase is observed consistently in this analysis and is also evident in Fig. 5 (Tree), where an increasing number of regions with elevated emission levels are detected toward higher build layers. With increasing build height, such regions become more frequent and spatially extended, which is consistent with an increased likelihood of local overheating phenomena that may manifest as keyhole-related defects.

In addition, the segmented cross-sectional area derived from the VCS masks is shown for each slice, revealing a clear correlation between local geometry and radiometric response. Layers with increased segmented area consistently coincide with elevated radiometric values, reflecting the higher effective energy input and enhanced thermal accumulation required to process larger contiguous cross-sections. Although one of the radiometric measures is evaluated in a normalized manner, the resulting mean emission value is not invariant with respect to the segmented cross-sectional area. This behavior does not arise from image discretizations effects but instead reflects the underlying physics of energy accumulation in contiguous material regions. For larger connected cross-sections, heat dissipation through component boundaries becomes less effective, while repeated laser exposure within the same layer leads to increased thermal superposition. Consequently, the integrated EOT emission represents the effective, thermally accumulated energy input per unit area rather than the nominal laser power density alone.

The radiometric trends observed in this analysis show very good agreement with results obtained from the EOS Melt Pool Monitoring (MPM) system [12]. In particular, regions of elevated and reduced energy input identified in the EOT-based evaluation correspond closely to the melt pool signatures reported by the MPM data. This consistency confirms that the proposed volumetric EOT analysis captures physically meaningful information related to the applied laser input and overall process stability. Based on the observed correlation between VCS-derived geometric information and radiometric emission values, the results further indicate that an inverse relationship between component geometry and required laser power could, in principle, be exploited to derive corrective laser power profiles.

A detailed discussion of such inverse correction strategies is beyond the scope of this work and is therefore addressed as an outlook in the summary.

Summary

This work introduced a volumetric framework for processing and analyzing Exposure Optical Tomography data acquired during the LPBF process. By reconstructing sequential EOT images into a component-aware 3D OT Image Cube, the proposed approach enables holistic, part-level inspection that extends beyond conventional layer-wise monitoring. The integration of volumetric component segmentation ensures that all subsequent analyzes are constrained to the actual component geometry, while high-resolution 2D slice rendering allows flexible visualization of internal structures along arbitrary slicing directions.

The presented slicing strategies—ranging from planar and cylindrical profiles to skeleton-based slicing for complex geometries—provide intuitive and information-dense representations that are well suited for engineering analysis. Geometric and radiometric validation demonstrate that the reconstructed volume preserves both dimensional fidelity and meaningful radiometric trends correlated with laser energy input and part geometry. As a result, the framework offers a practical foundation for post-process quality assurance, defect localization, and design-related weakness analysis in LPBF, with particular relevance for prototyping workflows.

Beyond the presented validation, the demonstrated correlation between volumetric component segmentation and radiometric emission data opens promising perspectives for future work. In particular, the inverse relationship between component geometry and laser energy input suggests that geometry-aware laser power correction profiles could be derived directly from sliced CAD data. By mapping local cross-sectional properties of connected components to expected radiometric responses, it becomes conceivable to predict EOT-like emission distributions prior to manufacturing and to compute mean correction factors for individual components or connected regions, but this has to be evaluated first.

Such an approach could enable geometry-adaptive process parameterization and complement existing monitoring systems by providing predictive, model-based laser power adjustments. While this work focuses on post-process analysis and validation, the presented framework establishes the necessary foundation for future investigations toward closed-loop or pre-compensated LPBF process control.

Acknowledgements

The author gratefully acknowledges EOS for the excellent calibration of the sCMOS camera used in the Exposure Optical Tomography (EOT) system, which significantly facilitated the quantitative evaluation of the acquired data. The resulting consistency between recorded optical signals and actual process conditions—summarized by the principle “What you CAM is what you get”—formed a reliable foundation for the analyzes presented in this work.

Special thanks are extended to the Additive Minds Academy of EOS for providing fundamental training in Laser Powder Bed Fusion and for the comprehensive documentation and knowledge transfer related to the EOSTATE MeltPool Monitoring (MPM) system, which supported the interpretation and validation of radiometric process data.

In addition, this work substantially benefited from the availability of open-source and public software. In particular, the Python ecosystem and its extensive libraries—including OpenCV, NumPy, SciPy, Matplotlib, and Open3D—provided essential tools for image processing, numerical analysis, and volumetric data handling. The author acknowledges the contributions of the open-source community and the developers who make these tools freely available for research and development.

The use of open AI-based tools is also gratefully acknowledged, especially the Segment Anything Model 2 (SAM2) developed by Meta and distributed through the broader research community, including platforms such as Hugging Face. In addition, ChatGPT was used as a supportive tool for translating assistance in refining the text during manuscript preparation.

Finally, the author acknowledges the availability of free software tools such as MeshLab, ImageJ, and GIMP, which were used for visualization, inspection, and figure preparation throughout this work.

References

[1] E. Duong: *Laser Powder Bed Fusion Process Monitoring Using Multisensory*: PhD thesis, RWTH Aachen (2023).

[2] Mohr, G.; Altenburg, S.J.; Ulbricht, A.; Heinrich, P.; Baum, D.; Maierhofer, C.; Hilgenberg, K.: *In-Situ Defect Detection in Laser Powder Bed Fusion by Using Thermography and Optical Tomography—Comparison to Computed Tomography.* Metals (2020), 10, 103. https://doi.org/10.3390/met10010103

[3] Breese, P., Becker, T., Oster, S., Altenburg, S., Metz, C., Maierhofer, C.: *Aktive Laserthermografie im L-PBF-Prozess zur in-situ Detektion von Defekten.* DGZfP Jahrestagung (2022), Kassel, May, Germany. https://www.ndt.net/?id=27067

[4] Aydogan, B.; Chou, K.: *Review of In Situ Detection and Ex Situ Characterization of Porosity in Laser Powder Bed Fusion Metal Additive Manufacturing.* Metals (2024), 14, 669. https://doi.org/10.3390/met14060669

[5] Herzog, T., Brandt, M., Trinchi, A. et al.: *Process monitoring and machine learning for defect detection in laser-based metal additive manufacturing.* J Intell Manuf 35, 1407–1437 (2024). https://doi.org/10.1007/s10845-023-02119-y

[6] Osazee E., Katayoon T., Ehsan T.: *Optical tomography and machine learning for in-situ defects detection in laser powder bed fusion: A self-organizing map and U-Net based approach* Additive Manufacturing, V78 (2023), https://doi.org/10.1016/j.addma.2023.103894.

[7] Atwya, M., Panoutsos, G.: *In-situ porosity prediction in metal powder bed fusion additive manufacturing using spectral emissions: a prior-guided machine learning approach.* J Intell Manuf 35, 2719–2742 (2024). https://doi.org/10.1007/s10845-023-02170-9

[8] Lim, Y., Wong and Tan: *Correlation of In-Process Monitoring Data and Defects in X-Ray CT for AM Swirler* Whitepaper EOS. `https://www.eos.info/`

[9] A. Kirillov, N. Ravi, D. Xu, S. Azadi, H. Xu, T. Darrell, K. He, R. Girshick, and P. Dollár: *Segment Anything Model 2 (SAM2)*. Meta AI Research, 2024. Available at: `https://github.com/facebookresearch/segment-anything-2`.

[10] T. Y. Zhang and C. Y. Suen: *A fast parallel algorithm for thinning digital patterns.* Communications of the ACM, Volume 27, Number 3 (March 1984).

[11] T.-C. Lee, R.L. Kashyap and C.-N. Chu: textitBuilding skeleton models via 3-D medial surface/axis thinning algorithms. Computer Vision, Graphics, and Image Processing, 56(6):462-478, (1994).

[12] L. Fuchs, C. Eischer: *In-process monitoring systems.* Whitepaper EOS. `https://www.eos.info/`