

CNN-based 3D characterization and liberation analysis of lithium-bearing slag particles using correlative CT and SEM imaging

Tom Kirstein^{*1}, Cindytami Rachmawati², Kai Bachmann³, Erik Löwer², Ralf Ditscherlein²,
Orkun Furat^{1,4}, Urs A. Peucker², and Volker Schmidt¹

¹Institute of Stochastics, Ulm University, 89096 Ulm, Germany

²Institute of Mechanical Process Engineering and Mineral Processing, Technische Universität Bergakademie Freiberg, 09599 Freiberg, Germany

³Helmholtz-Zentrum Dresden-Rossendorf, Helmholtz Institute Freiberg for Resource Technology, 09599 Freiberg, Germany

⁴SDU Applied AI and Data Science Unit, University of Southern Denmark, 5230 Odense, Denmark

April 29, 2026

Abstract

Efficient recovery of critical raw materials such as lithium from metallurgical slags requires optimized liberation of target phases during comminution. To determine effective mechanical process parameters for target phase recovery, an in-depth understanding of the characteristics of slag particles is crucial. For this purpose, modern tomography techniques, such as computed tomography (CT), can provide high-resolution 3D images of micrometer-sized slag particles. However, analysis of such CT images poses challenges, such as insufficient grayscale contrast between mineral phases and partial-volume effects. This paper presents a scalable workflow for accurate phase- and particle-wise 3D characterization of particle systems by correlating 3D CT images with 2D mineral maps. For this purpose, high-resolution scanning electron microscopy (SEM) slices are registered in 3D CT images and used as ground truth to train 3D convolutional neural networks (CNNs) for the segmentation of individual particles and mineral phases. This approach addresses the principal challenges of obtaining CT-based mineralogical characterizations, allowing for the particle-wise 3D characterization of complex slag systems with minimum manual labeling effort. The trained CNNs are then applied to CT images of particle systems with different particle sizes (from 63 μm to 100 μm and from 100 μm to 250 μm) of a lithium-bearing slag with LiAlO_2 as the target phase. Although virtual cross-sections of the predicted 3D segmentations show excellent agreement with mineral liberation obtained from 2D validation SEM-EDS data, the derived 3D mineral liberation statistics differ significantly from 2D estimates. In particular, our results show that the 2D analyses significantly overestimate mineral liberation compared to the 3D characterization. By addressing this stereological bias, the correlative 3D characterization workflow provides essential insights required to tailor pyrometallurgical and mechanical processing parameters to improve the recovery of raw materials.

Keywords and Phrases: machine learning, correlative microscopy, convolutional neural network, lithium-bearing slag, engineered artificial mineral, 3D particle characterization

*Corresponding author: tom.kirstein@uni-ulm.de

1 Introduction

Efficient recovery of critical raw materials, such as lithium, nickel, and rare earth elements, is crucial to achieving sustainability and digitalization, as categorized by the European Union [1]. These raw materials are increasingly sought after for their applications in modern technologies, such as batteries, renewable energy systems, and functional composite structures in general [2]. However, many of these raw materials represent only a small fraction of the waste streams. Thus, the effectiveness of traditional recycling methods is limited, as they generally target the recovery of more common materials. As a result, the recovery of critical raw materials is often too expensive compared to extracting them through mining. For example, during so-called pyrometallurgical processes, critical raw materials can form various (artificial) mineral phases, which can accumulate in a separate molten layer, known as slag [3–5]. However, the phases that contain critical raw materials typically occur at low concentrations in the slag, making their extraction economically and technically challenging. For this reason, slag is often considered a byproduct and used, for example, in the cement industry for construction purposes, leaving valuable raw materials lost after the recycling process [6].

In the case of lithium-ion battery (LIB) recycling, pyrometallurgical processes are used to separate cobalt, nickel and copper as the valuable metals, whereas lithium as a light element is contained in the slag and forms several lithium phases [7]. When the slag solidifies, it forms either an amorphous fine structure or defined mineral phases, depending on the slag production conditions. To recover lithium from such slag, the target lithium phases must grow as slowly and unhindered as possible into compact crystal structures, also referred to as grains, as the slag solidifies. This growth can be controlled by adjusting the parameters of the pyrometallurgical processes, such as the chemical composition and the cooling temperature conditions, with the aim of growing crystals without inclusions of other phases to a size of a few hundred μm [8–11].

This crystallization behavior directly influences the subsequent comminution process, which is carried out to liberate the target phases by fragmenting the slag into discrete units, so-called particles. Here, improved crystal growth could enhance the efficiency of the comminution process, as larger grain sizes of the target phases can lead to a higher probability of liberation [12].

Typically, after the comminution process, a separation process is carried out to separate the particles into a valuable concentrate and waste tailings [13–15]. This separation is achieved by exploiting certain physical properties of the target phases, such as mass density, magnetizability and hydrophobicity. For the separation process to be effective, these exploited properties must be sufficiently different from the corresponding properties of the other mineral phases. Moreover, since the separation process acts on the particles, its efficiency depends heavily on the liberation of the target phases within these particles.

To achieve an economically feasible recovery of critical raw materials, the process parameters must be optimized not only with respect to the recovery of the target phases but also with respect to time and cost. For example, comminuting to a smaller particle size would lead to a higher probability of liberation, but the process also becomes significantly more time- and energy-intensive. In order to determine the effective mechanical process parameters for the recovery of the target phase, an in-depth understanding of the characteristics of slag particles is crucial [16]. For this purpose, so-called particle-wise descriptors are used, which quantify characteristics such as size, shape, and mineralogical composition of particles [17, 18]. Moreover, by combining multiple descriptors such as particle volume, sphericity, and elongation, as well as particle-wise volume fractions of the target phases, we can obtain a multivariate characterization of particles [19]. By analyzing such descriptor vectors, we can then optimize the slag production, comminution and separation processes. For instance, if liberation is too low, it indicates that crystallization or comminution was suboptimal. Moreover, by multivariate probabilistic modeling [15], the correlations and interdependencies of multiple descriptors can be quantified. This allows for the description of not only individual particles but also entire particle systems. To accurately compute particle descriptors, precise imaging techniques are essential, the choice of which depends on the material types and particle sizes. Modern tomography techniques, such as X-ray computed tomography (CT), are effective in obtaining high-resolution 3D images of micrometer-sized slag particles non-destructively [20]. In order to facilitate the recovery of critical raw materials, a primary objective is the mineralogical characterization of slag particles, which requires the identification of distinct mineral phases within the particles. However, mineralogical characterization using CT images poses

challenges [21], such as insufficient grayscale contrast between mineral phases and partial volume effects, where the measured intensity in a single volumetric pixel (voxel) is influenced by multiple mineral phases.

Although high-resolution 3D mineralogical information can be obtained using focused ion beam scanning electron microscopy with energy dispersive X-ray spectroscopy (FIB-SEM-EDS) [22], this approach is impractical due to its time-intensive nature. Alternatively, 2D SEM-EDS can provide mineralogical information in 2D images of planar sections through the material. Although previous studies, such as those reported in [19], have used stochastic models to correlate 2D mineralogical information with CT-derived features, the intergrowth of phases in slag particles renders 2D images insufficient as representative ground truth for entire 3D particles. Recent work in CT image analysis has demonstrated that deep learning approaches, especially convolutional neural networks (CNNs), can achieve highly accurate segmentations of CT images into multiple mineral phases. In [23] a 2D CNN was trained on glutenite CT images using ground truth derived from corresponding SEM-EDS images. Furthermore, in [24], a computational pipeline has been described that registers 2D SEM-EDS images in 3D CT images to create ground truth for the training of both 2D and 3D CNN models. The resulting segmentations were stable and significantly outperformed conventional naive Bayes and feature-based classifiers. These studies illustrate that CNN-based methods can effectively utilize 2D mineralogical information to segment phases in 3D CT images, even when the grayscale contrast between different phases is low. However, to obtain particle-wise descriptors, we need not only the identification of the mineral phases, referred to as phase-wise segmentation, but also the identification of the individual particles, referred to as particle-wise segmentation. Due to the high spatial resolution of the SEM-EDS images, particles are mostly well-separated in these images with sharp boundaries. Thus, SEM-EDS images can also be used as ground truth for particle-wise segmentation, allowing us to almost entirely avoid manual labeling, which is time intensive and error prone.

To effectively employ 2D SEM-EDS images for neural network training, we first register them in CT images using rigid image registration [25]. The combined data then serves as a training dataset for two neural networks: one for phase-wise segmentation and the other for particle-wise segmentation of CT images. Specifically, we utilize so-called U-nets, which were originally designed to analyze biomedical CT images; see [26]. Although a phase-wise segmentation could, in principle, solve both tasks, as is the case for SEM-EDS images, the lower resolution of the CT images makes it difficult to separate interfacing particles. Therefore, a dedicated particle-wise segmentation network is used to focus almost exclusively on detecting these interface regions, greatly simplifying training and leading to much more reliable particle separation. In Figure 1, an overview of the workflow is depicted for 3D particle-wise characterization by training neural networks on registered SEM-EDS images. Using ground truth derived from SEM-EDS images for both segmentation networks, the characterization workflow is highly scalable, with minimum manual annotations required.

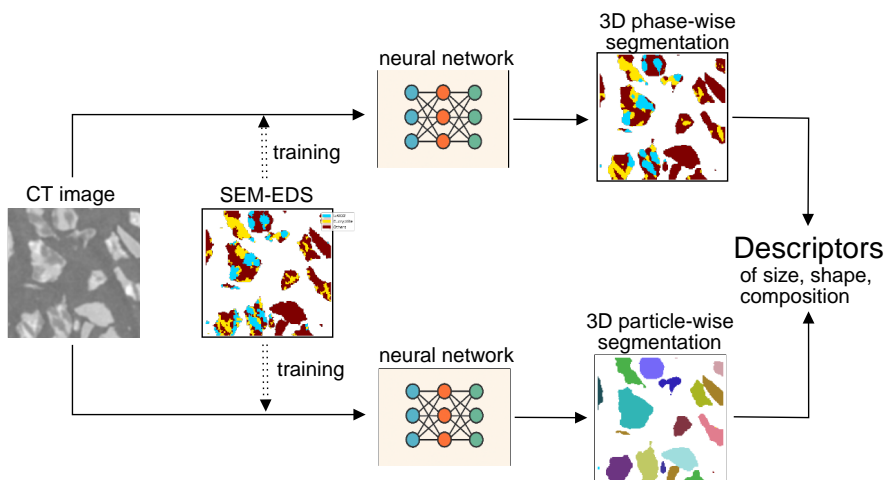


Figure 1: Workflow for phase-wise and particle-wise 3D characterization.

The predicted particle-wise segmentation allows us to obtain a comprehensive quantification of

particle characteristics. For example, by combining both segmentations, we can compute the volume fraction of the target phase for every individual particle. This allows us to derive so-called mineral liberation curves, quantifying the fraction of the target phase accessible to recovery processes. Whereas conventional mineral liberation curves derived from 2D slices inherently overestimate liberation due to stereological biases [27, 28], our workflow generates these based on volumetric data and should not suffer from stereological bias.

The remainder of this paper is organized as follows. In Section 2, we give a comprehensive overview of the materials and methods considered in this paper. Then, in Section 3, we evaluate the precision of the segmentation networks and the resulting particle-wise characterization. In particular, we quantify the discrepancy between the mineral liberation curves obtained by using the proposed 3D characterization workflow and 2D approaches. In Section 4, we discuss the advantages and limitations of the proposed characterization workflow and give an outlook on future improvements. Finally, Section 5 concludes.

2 Materials and Methods

2.1 Materials and sample preparation

The materials used in this study come from a synthetic slag produced by the Institute of Metallurgy at RWTH Aachen University, where a mixture of Li_2O , MnO , Al_2O_3 , SiO_2 , and CaO is heated to 1450°C and then cooled at a temperature rate of 50 K h^{-1} . The chemical compounds mentioned above are chosen as the slag system because they form the main component of the slag that occurs when recycling LIBs pyrometallurgically. Upon cooling, the slag solidified and crystallized as minerals, including lithium aluminate (LiAlO_2), gehlenite ($\text{Ca}_2\text{Al}(\text{AlSiO}_7)$), eucryptite (LiAlSiO_4), and glaucochroite (CaMnSiO_4). Apart from forming minerals, the slag also formed a minor amount of amorphous phases. The target phase is lithium aluminate with gehlenite as the main matrix. The slag is comminuted with a jaw crusher and further milled with a centrifugal ball mill to a particle size below 1 mm, where the size of almost 50 % of the particles is below $100\ \mu\text{m}$. Two size fractions from $63\ \mu\text{m}$ to $100\ \mu\text{m}$ and from $100\ \mu\text{m}$ to $250\ \mu\text{m}$ are prepared by sieving. These two size fractions are chosen as optimal particle sizes for this study, because particle sizes below $63\ \mu\text{m}$ will be more difficult to characterize using CT imaging, and particle sizes above $250\ \mu\text{m}$ will be too coarse to obtain a representative sample for image analysis. A rotary splitter is utilized to obtain approximately 2 g of representative subsamples from these two size fractions.

In order to prepare the size fractions stated above for CT imaging, each of them is mixed with carbon black nanoparticles (carbon black powder, Lamp black 101, Orion) using a volume ratio of 80:20, where the graphite nanoparticles will act as a spacer between the particles to be analyzed. Moreover, the low density of graphite will result in a distinct contrast between graphite nanoparticles and the sample material within acquired image data. Epoxy resin (EpoThin 2, Buehler) is then mixed with the sample-graphite mixture to create a slurry. A syringe, connected to a polymeric tube with an inner diameter of 2 mm, is used to draw the slurry mixture into the tube and left to cure overnight at room temperature for the hardening process. This particle sample preparation method follows the step-by-step procedure as implemented by [29]. Afterwards, the hardened particle sample is removed from the tube and glued onto a sample holder for the CT measurement. The resulting samples will be referred to as epoxy cylinder sample, see Figure 2a.




Figure 2: Correlative preparation procedure of CT and SEM-EDS samples: (a) hardened mixture of epoxy, sample and graphite, called an epoxy cylinder sample used for CT; (b) each epoxy cylinder sample is sliced into five disk pieces; (c) a total of ten pieces are obtained from two epoxy cylinder samples with two different size fractions, $63\ \mu\text{m}$ to $100\ \mu\text{m}$ and $100\ \mu\text{m}$ to $250\ \mu\text{m}$; (d) embedding the cut pieces into epoxy resin; (e) the resulting embedded sample which will later be re-embedded in epoxy, ground and polished to have the epoxy grain mount for SEM-EDS. The size of underlying square (b-d) is 10 mm.

After CT measurement, each epoxy cylinder sample is prepared for SEM-EDS analysis. The cylinder is sectioned along its axis into five shorter cylindrical segments, which are then embedded into epoxy resin inside a 20 mm diameter Teflon mold cup for both size fractions, see Figures 2b to 2e. This type of sample preparation provides five slices of the epoxy cylinder for each size fraction to be analyzed with SEM-EDS, to obtain a sufficient amount of particle data for the 2D-3D data correlation. The sample is then re-embedded in epoxy resin, ground, and polished to meet standard SEM-EDS preparation requirements. The resulting sample is called an epoxy grain mount, with a diameter of 25 mm.

2.2 CT and SEM-EDS based automated mineralogy

The epoxy cylinder sample was scanned using CT (Zeiss Xradia 510 Versa, Carl Zeiss). A short introduction to CT and its application to particle characterization can be found, e.g., in [30] and [31]. The scan parameters are listed in Table 1. The measurement is carried out using Stitch-mode to allow multiple scans at different heights of the sample to capture the whole sample. The 3D reconstruction is done using a filtered back projection algorithm in the Zeiss Reconstructor version 16.1.13038.

Table 1: Scan parameters for the CT measurement.

parameter	value	parameter	value
source position (mm)	-15	number of projections	1601
detector position (mm)	29	scan angle range ($^{\circ}$)	360
magnification lens	4X	exposure time (s)	1.5
acceleration voltage (kV)	80	camera binning	2
electrical power (W)	7	number of stitches	5
filter (Zeiss Standard)	LE1	voxel size (μm)	2.3

After CT measurement, the epoxy cylinder sample is used to prepare an epoxy grain mount, as explained in Section 2.1. Following these steps, automated mineralogy measurement by Mineral Liberation Analysis (MLA) software is carried out using an SEM device (FEI Quanta 650F, Thermo Fisher Scientific) equipped with two EDS detectors (Bruker). Before this measurement, the polished surface was coated with a thin carbon layer to minimize surface charging and improve the accuracy of the EDS measurement. The measurement is carried out at an acceleration voltage of 25 kV, using a backscattered electron (BSE) image resolution of $1\ \mu\text{m}$ per pixel and an X-ray acquisition step size of $6\ \mu\text{m}$ for EDS. Based on the collected EDS spectra, a mineral list is created and then used to classify the mineral phases in the sample by MLA. The resulting SEM-EDS image is a 2D false color image with a resolution of $1\ \mu\text{m}$ per pixel, where each mineral is assigned to a different color.

2.3 Phase-wise and particle-wise segmentation

Ground truth derived from 2D SEM-EDS images accurately labels mineral phases in 2D slices of the corresponding CT image. Moreover, these SEM-EDS images also provide high-resolution details on the

shapes and sizes of cross-sections of particles. However, as this type of data is in 2D, it cannot fully capture the morphological 3D aspects of particles such as their volume and mineralogical composition.

Due to the complex internal structure of the slag particles, see Figure 3, the mineral liberation observed in a single 2D slice of a particle may not be representative for the entire particle. This discrepancy stems from fundamental stereological biases that arise when using 2D slices for the computation of particle-wise descriptors of size, shape and mineral liberation. For example, 2D analysis can introduce systematic bias, leading to an underestimation of particle size [32] and an overestimation of mineral liberation [33]. Moreover, other particle descriptors also suffer from stereological biases when estimated using 2D data [27]. Although the specific magnitude of the bias varies, empirical studies such as those reported in [34, 35] have observed significant discrepancies. These stereological biases can be avoided when particle characterization is based on 3D data given by CT images. However, obtaining a particle-wise characterization from CT images is challenging, especially regarding the mineralogical composition of particles. This is due to factors such as partial volume effects and inhomogeneous grayscale levels caused by beam hardening artifacts. Moreover, the polychromatic X-ray spectrum combined with the cone-beam geometry of the X-ray projections produces artifacts in the outer regions inside the epoxy cylinder sample (see Figure 3d), leading to gray value gradients. This can also be seen in Figure 3e where the average gray value (or CT intensity) is higher in the outer part of the sample.




Figure 3: Example of a 2D slice (size fraction from $63\ \mu\text{m}$ to $100\ \mu\text{m}$) from the 3D CT image (a) and classified 2D SEM-EDS image (b), where different minerals are assigned with different colors; together with magnifications from inner (c) and outer (d) areas of the cylinder. The average CT intensity over the distance from the boundary of the epoxy cylinder sample is shown in (e), where this boundary was manually segmented for five slices of the CT image and the visualized results were then aggregated over these slices.

We now show how 2D SEM-EDS images can be used to segment mineral phases and individual particles in 3D CT images, i.e., to assign a specific mineral phase and particle label to each voxel. The segmentations obtained in this way will then be used to compute particle-wise descriptor vectors to quantify various characteristics of the particles, as described in Section 2.4 below. To determine these segmentations, we utilize neural networks, which have to be trained using ground truth. Note that for the methodology and training process described in the following, we focus exclusively on the image data from the finer size fraction (from $63\ \mu\text{m}$ to $100\ \mu\text{m}$). The coarser fraction (from $100\ \mu\text{m}$ to $250\ \mu\text{m}$) is held

out entirely for training and serves as an additional dataset in Section 3 to evaluate the robustness of the trained networks to different particle sizes. While providing comprehensive 3D ground truth using technologies like FIB-SEM-EDS is possible, it is costly and often redundant, particularly in neighboring slices with little variation. Instead, as described in Section 2.1, we obtained a few SEM-EDS images that are sparsely distributed in the 3D CT images to capture diverse, non-redundant data across the sample, ensuring each image provides unique information. In the following, we explain how such a 2D ground truth can be utilized for the training of both (segmentation) networks.

2.3.1 Obtaining ground truth from SEM-EDS images

To utilize these types of images effectively for training neural networks, we first formalize the data and segmentation frameworks.

A 3D image is considered as an array $I = (I_x, x \in \Omega) \in \mathbb{R}^{H \times W \times D}$, where $H, W, D \in \mathbb{N} = \{1, 2, \dots\}$ are the spatial dimensions of the image. The component I_x then represents the grayscale value at voxel $x \in \Omega = [H] \times [W] \times [D]$, where $[H] = \{1, \dots, H\}$. Similarly, a particle-wise (or phase-wise) segmentation can be defined as a labeled image $S = (S_x, x \in \Omega) \in L^{H \times W \times D}$, where $L = \{-1, 0, 1, \dots, c\}$ and $c \in \mathbb{N}$ is the number of particles or (non-background) mineral phases. Here, $S_x = 0$ indicates that voxel x belongs to the background, i.e., it does not belong to any particle (or mineral phase) and $S_x = i$ for $i \in \{1, \dots, c\}$ means that voxel x belongs to the i -th particle (or mineral phase). Furthermore, $S_x = -1$ indicates that the label of the particle (or mineral phase) at voxel x is unknown. In the following, we refer to S as a complete segmentation if there are no voxels with unknown labels, i.e., if $S_x \geq 0$ for all voxels $x \in \Omega$, otherwise we refer to it as an incomplete segmentation.

As mentioned above, the SEM-EDS images provide ground truth for 2D slices of the cylindrical epoxy sample. However, to obtain these incomplete ground truth segmentations, we first need to register the SEM-EDS images in the CT volume. This process involves determining the translations and rotations of the slices to ensure that they align correctly with the CT image. For that purpose, the SEM-EDS images are first downsampled by a factor of 2.3 to align their pixel size with that of the CT images. The SEM-EDS images are then sectioned via overlapping tiles into smaller cutouts of size 180×180 pixels (compared to the downsampled SEM-EDS image sizes of about 1000×1000), and the registration algorithm detailed in [25] is applied to obtain translation and rotation for each cutout. In the present paper, the SEM-EDS images were sectioned because larger cutouts could not be accurately aligned with the CT image. This could be due to difficulties finding optimal rotations and translations or due to slight curvatures of the imaged surfaces that have been exposed by the sample preparation for SEM imaging. By applying the optimal translations and rotations to each cutout, we directly obtain an incomplete ground truth phase-wise segmentation $S^M \in \{-1, 0, \dots, c^M\}^{H \times W \times D}$, where $c^M \in \mathbb{N}$ denotes the number of mineral phases identified in the SEM-EDS images. We grouped the identified mineral phases into four distinct classes ($c^M = 3$ non-background classes): background, LiAlO_2 , eucryptite, and a combined class comprising gehlenite and the other phases. This grouping simplifies the segmentation task while preserving the distinction between the target phase (LiAlO_2), other lithium-bearing phases (eucryptite) and the other phases (gehlenite and other phases).

Obtaining an incomplete ground truth particle-wise segmentation requires additional processing steps to ensure that individual particles can be distinguished. Due to the high resolution of the original SEM-EDS images, most particles that appear to touch each other in the lower-resolved CT image are well-separated by background in the original SEM-EDS images, but this spatial separation may be lost in the downsampling step. Therefore, before downsampling to obtain the ground truth for an incomplete particle-wise segmentation, individual particles are first identified in the high-resolution SEM-EDS images by means of a connected component algorithm [36]. The high-resolution particle labels obtained in this way are then downsampled to the CT resolution using nearest-neighbor interpolation. Subsequently, the translation and rotation parameters, as determined during the registration of the 2D cutouts, are applied to align this particle-wise segmentation with the CT volume. During this process, voxels from different particles may become neighbors. Therefore, if a voxel belonging to a particle has a voxel belonging to a different particle as neighbor, we set the labels of those voxels to the background class 0. This produces an incomplete particle-wise segmentation $S^P \in \{-1, 0, \dots, c^P\}^{H \times W \times D}$, where $c^P \in \mathbb{N}$ denotes the number of particles identified in the original SEM-EDS images and all particles are

spatially separated by at least one pixel of background.

The incomplete ground truth segmentations S^M and S^P will be used for the training of two neural networks, one for phase-wise segmentation and one for particle-wise segmentation. However, directly using S^P as the ground truth for the training of the particle-wise segmentation network would be inefficient, as the specific label values assigned to the particles are arbitrary and the total number of labels is variable, i.e., the total number of particles can vary. Therefore, we aim to reduce the task to a classification problem with a fixed number of classes. More specifically, to train the neural network, we use the following incomplete foreground-background segmentation $S^F \in \{-1, 0, 1\}^{H \times W \times D}$, given by

$$S_x^F = \begin{cases} 1, & \text{if } S_x^P \geq 1, \\ 0, & \text{if } S_x^P = 0, \\ -1, & \text{if } S_x^P = -1, \end{cases} \quad (1)$$

for each $x \in \Omega$, which allows us to distinguish between the background and the particles. Moreover, since we ensured that neighboring particles are separated by the background, this simplification does not result in a loss of information. The distinct particle instances can be fully recovered from S^F , by simply applying a connected component algorithm.

In the next section, we introduce a network architecture that can utilize such incomplete ground truth segmentation for training.

2.3.2 Neural network architecture

In order to obtain complete segmentations of the CT image, as defined in Section 2.3.1, we utilize two neural networks: one for phase-wise segmentation and one for particle-wise segmentation. Both networks utilize an identical underlying architecture based on the U-net, originally proposed for 2D biomedical images [37] and later extended to 3D [26]. These U-nets have an encoder-decoder architecture, where features are processed at different spatial resolutions through subsequent down- and upsampling layers, as shown in Figure 4.

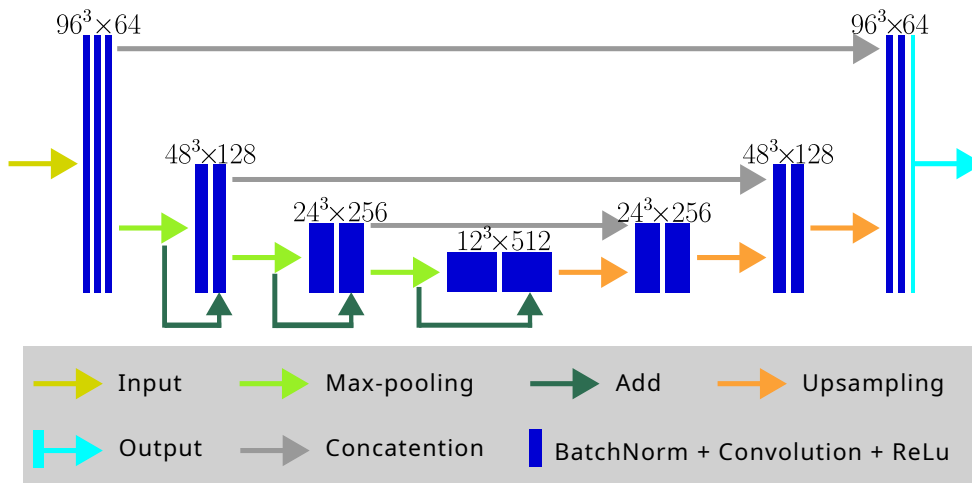


Figure 4: U-net architecture used for phase-wise and particle-wise segmentations. The width of the convolutional layer represents the number of feature channels, the height represents the spatial dimensions. Each convolutional layer (dark blue) comprises subsequent batch normalization, $3 \times 3 \times 3$ convolutions and ReLu activation layers. The first convolutional layer after the input layer has 32 filters ($96^3 \times 32$). The output layer applies a softmax activation function for phase-wise segmentation and the sigmoid activation function for particle-wise segmentation.

More specifically, both the encoder and the decoder are composed of subsequent convolutional blocks, which comprise two convolutional layers with kernel size $3 \times 3 \times 3$. Each convolutional layer is then followed by a batch normalization layer and a ReLU activation function. In the encoder, these convolutional blocks double the number of feature channels and are followed by a $2 \times 2 \times 2$ max-pooling

layer to halve each of the spatial dimensions. In the decoder, the number of feature channels is halved in each convolutional block and the spatial dimensions are doubled using a bilinear upscaling layer.

Additionally, so-called skip connections concatenate the output of each max-pooling layer with the input of the corresponding convolutional block (in the decoder) with the same spatial dimensions. Note that while these two sets of features share the same spatial dimensions, they represent the input images at different levels of abstraction. The features provided by the skip connections are extracted earlier in the network with the high spatial resolution necessary to resolve fine details, such as particle boundaries. By contrast, the features upsampled from the deeper layers are extracted from a larger field of view, providing larger context, but with lower spatial resolution due to previous downsampling operations.

As shown in Figure 4, we modify the standard 3D U-net architecture with the addition of so-called residual connections [38] in the downsampling layers. Specifically, the input of each downsampling layer is added to its output, resulting in a residual 3D U-net, similar to the U-nets used in [19, 39]. As stated in [38], these connections can make the network easier to optimize and allow for a larger number of layers.

The outputs of these residual 3D U-nets are multi-channel images, where the three spatial dimensions are of the same size as the input images. More precisely, a residual 3D U-net with trainable parameters $\theta \in \Theta$ describes a function $f_\theta: \mathbb{R}^{H \times W \times D} \rightarrow [0, 1]^{H \times W \times D \times k}$, with the number of output channels $k \in \mathbb{N}$ and spatial dimensions $H, W, D \in \mathbb{N}$. The residual 3D U-net comprises three subsequent max-pooling layers, each halving the spatial dimensions, and thus H, W, D have to be divisible by 8. Except for this constraint, the spatial dimensions of the input image can be freely chosen, independently of the parameter constellation θ . Moreover, the number of output channels k depends on the specific segmentation task defined in Section 2.3.1. For the phase-wise network, we set $k = c^M + 1$ representing the identified mineral phases and the background. In contrast, for the particle-wise network, we use a single output channel $k = 1$ to efficiently represent the foreground and background.

In the final layer, we apply a specific activation function to constrain the network output to the interval $[0, 1]$. For the phase-wise network, we use a voxel-wise softmax function, whereas for the particle-wise network, we utilize a sigmoid function. The output of the network f_θ then represents a probability map, where each value indicates the likelihood that a voxel belongs to a specific class. From these probabilities, we can obtain complete segmentations, as described in Section 2.3.1, by applying a threshold (for the particle-wise network) or selecting the channel with the highest probability (for the phase-wise network). However, for these segmentations to be accurate, the network’s trainable parameters θ must first be optimized during so-called training. One challenge in this optimization is the incomplete nature of our ground truth segmentations, which contain undefined regions where no valid ground truth exists. Therefore, we introduce voxel weights in the next section, which allow us to mask these unknown regions and focus the training process exclusively on the regions with known ground truth data.

2.3.3 Weight maps

To train segmentation networks on incomplete ground truth segmentations, we employ so-called voxel weights, i.e., arrays of the form $\Lambda = (\Lambda_x, x \in \Omega) \in [0, \infty)^{H \times W \times D}$, which are utilized during training to determine the influence each voxel has on the training process. In our approach, the voxel weights assign a weight of zero to voxels with a label of '-1', directing the network to learn exclusively from voxels with known ground truth during training. The voxel weights $\Lambda^M \in \{0, 1\}^{H \times W \times D}$, for training the phase-wise segmentation network, are given by

$$\Lambda_x^M = \begin{cases} 1, & \text{if } S_x^M \geq 0, \\ 0, & \text{if } S_x^M = -1, \end{cases}$$

for each voxel $x \in \Omega = [H] \times [W] \times [D]$. The particle-wise segmentation network faces a more nuanced set of challenges to achieve accurate segmentations. Most importantly, incorrectly segmented voxels in the regions between interfacing particles can cause these particles to merge incorrectly when identifying connected components. To address this critical issue, the voxel weights $\Lambda^F \in [0, \infty)^{H \times W \times D}$ for training the particle-wise segmentation network are designed to exhibit increased weights for voxels in regions

between interfacing particles. More specifically, the voxel weights Λ^F , which rely on the foreground-background segmentation S^F defined in Eq. (1), are given by

$$\Lambda_x^F = \begin{cases} D(S^F)_x + \lambda^{\text{bg}}, & \text{if } S_x^F = 0, \\ \lambda^f, & \text{if } S_x^F > 0, \\ 0, & \text{if } S_x^F = -1, \end{cases} \quad (2)$$

for each $x \in \Omega$, where the weighting function $D : \{-1, 0, 1\}^{H \times W \times D} \rightarrow [0, \infty)^{H \times W \times D}$ results in increased weights for voxels between particles that may appear as interfacing in the CT image. More precisely, the image $D(S^F)$ takes its maximum value at those voxels where the two nearest particles are at a distance of one voxel unit, and it is set to 0 if the second nearest particle is more than five voxel units away; see [37] for more details. The constants $\lambda^{\text{bg}} = 0.013$ and $\lambda^f = 0.04$ were chosen to ensure that the total weight assigned to the background and foreground voxels is approximately equal. Furthermore, these values ensure that foreground voxels are weighted significantly higher than background voxels for which $D(S^F)$ takes a value of 0. The voxel weights Λ^M and Λ^F allow for a more meaningful comparison of the network outputs and the incomplete ground truth segmentations during the training of the neural networks.

2.3.4 Training procedure

Using the ground truth segmentations and corresponding voxel weights, as introduced in Sections 2.3.1 and 2.3.3, we can train the phase-wise and particle-wise segmentation networks. For this purpose, the network's predictions are iteratively compared to the corresponding ground truth by means of a so-called loss function [40]. In each iteration, a so-called training step, this loss, i.e., the value of the loss function, is used to update the network's trainable parameters θ using stochastic gradient descent.

More specifically, in each training step, the network is evaluated on a batch of randomly chosen input images $I' = (I'_x, x \in \Omega') \in \mathbb{R}^{H' \times W' \times D'}$ with spatial dimensions $H', W', D' \in \mathbb{N}$, where $H' < H$, $W' < W$ and $D' < D$. Each of these input images is determined by a randomly chosen spatial offset $y \in \Omega$ such that $I'_x = I_{x+y}$ for each $x \in \Omega' = [H'] \times [W'] \times [D']$, where $I = (I_x, x \in \Omega)$ is the CT image introduced in Section 2.3.1. For each input I' , the network yields a voxel-wise class probability map $P' = (P'_{x,i}) \in [0, 1]^{H' \times W' \times D' \times (c+1)}$ such that

$$\sum_{i=1}^{c+1} P'_{x,i} = 1 \quad \text{for each } x \in \Omega' = [H'] \times [W'] \times [D'],$$

where $c \in \mathbb{N}$ is the number of particles or (non-background) mineral phases and $P'_{x,i}$ denotes the predicted probability that voxel x belongs to class $i - 1$.

Note that, while the phase-wise network f_θ^M directly outputs a probability map over the $c^M + 1 = 4$ channels, the particle-wise network f_θ^P outputs a single channel representing the foreground probabilities. To obtain a multi-channel probability map $P' \in [0, 1]^{H' \times W' \times D' \times 2}$ for the particle-wise network, the network output $f_\theta^P(I')$ is concatenated with its complement, i.e., the probability map P' is given by

$$P'_{x,i} = \begin{cases} 1 - f_\theta^P(I')_{x,1}, & \text{if } i = 1, \\ f_\theta^P(I')_{x,1}, & \text{if } i = 2, \end{cases}$$

for each $x \in \Omega'$.

The predicted probability map P' is then compared to the corresponding ground truth by means of the so-called weighted focal loss FL. This loss is given by

$$\text{FL}(\tilde{P}, \Lambda') = - \sum_{x \in \Omega'} \sum_{i=1}^{c+1} \Lambda'_x \cdot (1 - \tilde{P}_{x,i})^\gamma \cdot \log(\tilde{P}_{x,i}),$$

where the array $\tilde{P} \in [0, 1]^{H' \times W' \times D' \times (c+1)}$, defined as

$$\tilde{P}_{x,i} = \begin{cases} P'_{x,i}, & \text{if } S'_x = i - 1, \\ 1 - P'_{x,i}, & \text{otherwise,} \end{cases}$$

for each $x \in \Omega'$, measures how closely P' aligns with the ground truth. Here, S' and Λ' denote the corresponding ground truth and voxel weights with the same spatial offset as the input image, defined by $S'_x = S_{x+y}$ and $\Lambda'_z = \Lambda_{x+y}$, respectively. The exponent $\gamma \geq 0$ is the focal parameter, where a higher value of γ reduces the relative loss for well-classified voxels. This increases the contribution of misclassified voxels to the loss function, allowing the model to learn more effectively from novel features. In this work, we set $\gamma = 2$. In contrast, for two classes, setting $\gamma = 0$ would reduce focal loss to the weighted binary cross-entropy loss.

Finally, to minimize the loss, we compute the gradient $\nabla_{\theta} \text{FL}(\tilde{P}, \Lambda')$ of the loss function with respect to the network’s trainable parameters θ . This gradient is then utilized by the ADAM optimizer [41] with learning rate 0.0001 to iteratively update the trainable parameters of the network.

Before training, the input images were converted to an 8-bit integer format to reduce memory consumption and accelerate the training process without any perceptible loss in segmentation quality. The phase-wise segmentation network was trained for a total of 150 epochs, where we defined an epoch as 100 individual training steps with a batch size of 4. In each training step, the input was transformed by a series of augmentations to improve generalization [42]. Specifically, the input images I' , as well as their corresponding voxel weights Λ' and ground truth S' , were randomly rotated around each axis by multiples of 90 degrees. Additionally, the grayscale values of each input image were uniformly shifted by an integer randomly chosen between -10 and 10. This intensity shift corresponds to the observed difference in the average CT intensity between the inner and outer regions of the epoxy cylinder sample illustrated in Figure 3.

For training of the particle-wise segmentation network, we performed an additional data augmentation step using elastic transformations, as described in [43]. Moreover, the particle-wise segmentation network was trained for two training runs, the first consisting of 100 epochs and the second one consisting of 150 epochs. The reason for having two training runs is as follows: After the first training run, the trained network was applied to a larger subvolume of the CT image, and some obvious segmentation errors were identified. By manually labeling voxels associated with incorrectly segmented regions, the training data was refined for the second training run of the network. More precisely, let $\mathcal{E} \subset \Omega$ denote the set of voxels manually labeled as segmentation errors, and let $\hat{S} \in \{0, 1\}^{H' \times W' \times D'}$ be the binarized prediction of the network after the first run. The refined ground truth segmentation $S^{\text{F, new}} \in \{0, 1\}^{H' \times W' \times D'}$ is then given by assigning the opposite value of the initial prediction to the manually labeled regions:

$$S_x^{\text{F, new}} = \begin{cases} 1 - \hat{S}_x, & \text{if } x \in \mathcal{E}, \\ S_x^{\text{F}}, & \text{if } x \in \Omega \setminus \mathcal{E}. \end{cases}$$

These errors were generally of two types: regions between particles (where the new label becomes 0) and darker regions inside particles (where the new label becomes 1). To ensure that the network effectively corrects these specific issues, the voxel weights Λ^{F} , as defined in Eq. (2), is updated accordingly. The refined voxel weights $\Lambda^{\text{F, new}} \in [0, \infty)^{H' \times W' \times D'}$ are given by

$$\Lambda_x^{\text{F, new}} = \begin{cases} \max(\Lambda^{\text{F}}), & \text{if } x \in \mathcal{E} \text{ and } S_x^{\text{F, new}} = 0, \\ \delta^{\text{bal}}, & \text{if } x \in \mathcal{E} \text{ and } S_x^{\text{F, new}} = 1, \\ \Lambda_x^{\text{F}}, & \text{if } x \in \Omega \setminus \mathcal{E}, \end{cases} \quad (3)$$

where $\delta^{\text{bal}} = \max(\Lambda^{\text{F}}) \frac{|\{x \in \mathcal{E}: S_x^{\text{F, new}} = 0\}|}{|\{x \in \mathcal{E}: S_x^{\text{F, new}} = 1\}|}$ is chosen such that the total weight assigned to the manually labeled particle regions equals the total weight assigned to the manually labeled background regions.

Utilizing these manually annotated regions during the second training run successfully eliminated the most visible segmentation errors, i.e., the artificial merging of neighboring particles and the generation of holes within single particles. Note that while identifying these segmentation errors was trivial for the particle-wise segmentation, the same could not be done for the phase-wise segmentation, since manually distinguishing between LiAlO_2 and eucryptite is not feasible.

2.3.5 Network inference

Once the trainable parameters θ of the phase-wise and particle-wise networks are optimized, the trained networks are applied to a full CT image. Due to GPU memory constraints, each CT image is processed

using an overlap-tile strategy [37], which produces the predicted probability maps $\hat{P}^M \in [0, 1]^{H \times W \times D \times 4}$ and $\hat{P}^P \in [0, 1]^{H \times W \times D \times 2}$ for mineral phases and particles, respectively. For the CT volumes considered in this study, inference was computationally feasible on a NVIDIA RTX 4080. For example, applying the phase-wise or particle-wise network to a $500 \times 500 \times 500$ voxel subvolume required approximately 250s of runtime. Of this runtime, approximately 45s were spent on the GPU-based evaluation of the neural networks, whereas the remaining time was dominated by the overlap-tile strategy, including tile extraction, stitching and associated data handling. Therefore, the runtime to obtain network predictions is reasonable compared to sample preparation and CT measurements.

From these predictions, we reconstruct the final predictions for the complete segmentations. However, since the two networks operate independently of each other, they may produce different foreground regions. That is, some voxels may be assigned to a particle but not to a mineral phase, and vice versa. Moreover, due to the heavy weighting of the interface regions (via the voxel weights Λ^F as defined in Eq. (2)), the particle-wise segmentation network is encouraged to erode the boundary of the particles to ensure their spatial separation. Therefore, relying on this output for particle-wise characterization would result in biases, e.g., an underestimation of the particle size. To resolve these issues, we utilized the following reconstruction strategy to unify the foreground based on the predictions of the phase-wise network while using the predictions of the particle-wise network to separate individual particles.

First, we obtain the predicted phase-wise segmentation by assigning each voxel to the class with the highest probability. More specifically, the predicted phase-wise segmentation $\hat{S}^M \in \{0, \dots, 3\}^{H \times W \times D}$ is given by

$$\hat{S}_x^M = \operatorname{argmax}_{i \in \{1, \dots, 4\}} \hat{P}_{x,i}^M - 1,$$

for each $x \in \Omega$. We consider the foreground $\Omega_{\text{fg}} = \{x \in \Omega : \hat{S}_x^M > 0\}$ defined by this segmentation to be the reference mask for the union of particles.

Second, we utilize the particle-wise network output to generate markers for individual particles. Because the particle-wise network was trained to predict wide gaps between touching particles, a direct thresholding of \hat{P}^P produces well-separated, albeit eroded, particles. More specifically, the predicted foreground-background segmentation $\hat{S}^F \in \{0, 1\}^{H \times W \times D}$ given by

$$\hat{S}_x^F = \operatorname{argmax}_{i \in \{1, 2\}} \hat{P}_{x,i}^P - 1,$$

for each $x \in \Omega$, can be used to obtain a set of (particle) markers \mathcal{M} by means of a connected components algorithm, as described in [36]. These markers are considered as the initial

Finally, to recover the particles from these markers while ensuring consistency with the phase-wise segmentation, we apply a marker-based watershed transform [44]. This can be described as a region growth starting from each individual marker $m \in \mathcal{M}$. These markers then grow simultaneously into the set of unassigned foreground voxels $\Omega_{\text{fg}} \setminus \bigcup_{m \in \mathcal{M}} m$, such that every voxel in the foreground Ω_{fg} is assigned to exactly one marker, and voxels that are not in the foreground Ω_{fg} are assigned to no marker. This yields the predicted particle-wise segmentation \hat{S}^P , where the union of all particles corresponds exactly to the foreground region determined by the phase-wise segmentation, i.e., $\Omega_{\text{fg}} = \{x \in \Omega : \hat{S}_x^M > 0\} = \{x \in \Omega : \hat{S}_x^P > 0\}$, ensuring that every voxel assigned to a mineral phase is also assigned to a specific particle, and vice versa. Finally, to exclude particles that cannot be reliably characterized [45], any predicted particle comprising fewer than 100 voxels is removed from \hat{S}^P .

2.4 Multivariate characterization of particles

To quantify characteristics such as size, shape, and mineralogical composition of particles, we compute descriptor vectors based on the phase-wise and particle-wise segmentations. In the following, let $S^M \in \{-1, 0, \dots, c^M\}^{H \times W \times D}$ and $S^P \in \{-1, 0, \dots\}^{H \times W \times D}$ denote phase-wise and particle-wise segmentations, respectively. These may represent either the incomplete ground truth derived in Section 2.3.1 or the predictions for the complete segmentations obtained in Section 2.3.5.

We then identify the set of individual particles $\mathcal{P} = \{P_1, \dots, P_N\}$, where $N > 0$ denotes the total number of particles. Here, the i -th particle is given by $P_i = \{x \in \Omega : S_x^P = i\}$, i.e., the set of voxels associated with the label $i \in \{1, \dots, N\}$ in the particle-wise segmentation S^P . For each particle $P \in \mathcal{P}$,

characteristics such as size, shape and mineralogical composition can be quantified. More specifically, so-called descriptor vectors are computed, allowing us to transform the complex image data of each particle into interpretable, low-dimensional descriptor vectors. Since ground truth descriptors can only be computed on 2D SEM-EDS images, we also introduce the corresponding 2D descriptors when they differ from their 3D counterparts. In the 2D case, we consider the planar section $P' = P \cap E$, where the particle $P \in \mathcal{P}$ results from an (incomplete) segmentation and $E \subset \Omega$ is a discretized plane.

2.4.1 Particle size and shape

The size of a particle $P \in \mathcal{P}$ is quantified by the volume-equivalent and area-equivalent diameters $D_D(P)$ and $D'_D(P')$, respectively, given by

$$D_D(P) = \left(\frac{6|P|}{\pi} \right)^{1/3} \quad \text{and} \quad D'_D(P') = \left(\frac{4|P'|}{\pi} \right)^{1/2},$$

where $|\cdot|$ denotes cardinality.

To quantify how closely the particle shape resembles that of a sphere or circle, we compute the sphericity $D_S(P)$ and circularity $D'_S(P')$, given by

$$D_S(P) = \frac{\pi^{1/3}(6|P|)^{2/3}}{|\partial P|} \quad \text{and} \quad D'_S(P') = \frac{4\pi|P'|}{(|\partial P'|)^2},$$

where $\partial P \subset P$ denotes the surface voxels of P and $\partial P' \subset P'$ is the set of those surface voxels of P that intersect with the discretized plane $E \subset \Omega$. These descriptors take values between 0 and 1, with 1 indicating a perfect sphere (or circle). Note that while $|\partial P|$ is an intuitive approximation of the surface area of P , it can be improved by introducing corrective weights [46]. The notation based on the cardinality of sets of voxels to quantify surface areas and perimeters is to maintain notational simplicity only, while we computed surface area and perimeter using implementations of [46] and [47], respectively.

Although most of the particles in the considered dataset exhibit an irregular shape, some particles display an elongated or even plate-like shape. To differentiate between these shapes, we apply an oriented minimum-volume bounding box algorithm for particles, as implemented in [48], to determine the lengths of the third longest, second longest and longest axes, denoted by a_3 , a_2 and a_1 , respectively. We then compute the flatness $D_F(P)$ and elongation $D_E(P)$ of a particle $P \in \mathcal{P}$, given by

$$D_F(P) = \frac{a_3(P)}{a_2(P)} \quad \text{and} \quad D_E(P) = \frac{a_2(P)}{a_1(P)}, \quad (4)$$

which take values closer to 0 for plate-like or elongated particles.

2.4.2 Mineral liberation

In addition to computing the descriptors mentioned in Section 2.4.1 that quantify the size and shape of the particles, we can also utilize a phase-wise segmentation S^M to characterize the mineralogical composition of the particles.

The volume liberation $D_{VL}(P)$ and surface liberation $D_{SL}(P)$ of a particle $P \in \mathcal{P}$ are given by

$$D_{VL}(P) = \frac{|P^{c_T}|}{|P|} \quad \text{and} \quad D_{SL}(P) = \frac{|\partial P \cap P^{c_T}|}{|\partial P|},$$

where $P^{c_T} = \{x \in P: S_x^M = c_T\}$ is the set of voxels of particle P associated with the target phase $c_T \in \{1, \dots, c_M\}$ and $\partial P \subset P$ denotes the set of surface voxels of P . These descriptors quantify the fraction of the target phase within the particle and its surface. In general, while values of $D_{VL}(P)$ and $D_{SL}(P)$ close to 0 or 1 are often desired for separation processes, a large number of slag particles contains the target phase along with other mineral phases. For such particles, intergrowth between these phases

critically affects liberation efficiency during comminution. To quantify this, we define the intergrowth-boundary area $D_{\text{IB}}(P)$ as the area of the shared interface between valuable and non-valuable mineral phases, given by

$$D_{\text{IB}}(P) = \frac{1}{2} \left(|\partial(P \cap P^{\text{cT}})| + |\partial(P \setminus P^{\text{cT}})| - |\partial P| \right), \quad (5)$$

for each particle $P \in \mathcal{P}$. This value can then be normalized to compare particles of different sizes. For example, the Amstutz-Giger index, originally introduced in [49], normalizes the intergrowth boundary with the sum of the intergrowth boundary and the outer boundary of the particle. Since the Amstutz-Giger index was originally introduced for line segments, rather than particles, we define a corresponding particle-wise descriptor, denoted by $D_{\text{AG}}(P)$ and given by

$$D_{\text{AG}}(P) = \frac{D_{\text{IB}}(P)}{D_{\text{IB}}(P) + |\partial P|}, \quad (6)$$

for each $P \in \mathcal{P}$. By definition, the intergrowth boundary $D_{\text{IB}}(P)$ is smaller than the surface area of the minority phase, i.e., the phase with the lower volume. However, for particles with a volume liberation close to 0 or 1, the volume of the minority phase is low, and thus its surface area also tends to be lower. Therefore, a small but highly intergrown inclusion could still lead to a low Amstutz-Giger index. That is, while a high Amstutz-Giger index reliably indicates that a particle is highly intergrown, it only does so for particles with a volume liberation $D_{\text{VL}}(P)$ close to 0.5. In order to identify highly intergrown particles with a lower $D_{\text{VL}}(P)$ as well, the intergrowth boundary $D_{\text{IB}}(P)$ can be compared with the minimum possible interphase boundary [50]. However, computing such a minimum boundary for 3D particles is challenging. Instead, we compare the intergrowth boundary with the surface area of a sphere with the same volume, which leads to an intergrowth descriptor that is similar to the sphericity of particles. The normalized intergrowth boundary $D_{\text{NIB}}(P)$ and its 2D counterpart $D'_{\text{NIB}}(P')$, where $P' = P \cap E$ is the planar section of a particle $P \in \mathcal{P}$ and the discretized plane $E \subset \Omega$, are given by

$$D_{\text{NIB}}(P) = \frac{\pi^{1/3}(6|P^{\text{cT}}|)^{2/3}}{D_{\text{IB}}(P)} \quad \text{and} \quad D'_{\text{NIB}}(P') = \frac{2\sqrt{\pi}|P',\text{cT}|}{D_{\text{IB}}(P')}, \quad (7)$$

where $D_{\text{IB}}(P')$ denotes the intergrowth boundary length of the planar section $P' \subset P$, which is computed using Eq. (5). Note that smaller values of $D_{\text{NIB}}(P)$ and $D'_{\text{NIB}}(P')$ indicate that a particle has more highly intergrown mineral phases, and both descriptors can also take values greater than 1 if a large fraction of the target phase is located on the surface of the particle.

Although the liberation and intergrowth descriptors stated above can be evaluated in a particle-wise manner, they are often aggregated to characterize the entire particle system. A commonly used descriptor of this type is the degree of liberation, which typically quantifies the fraction of the target mineral phase contained in particles that exceed a specific volume (or surface) liberation threshold, e.g., 0.9. However, relying on a single liberation threshold provides an incomplete picture, as different separation processes may require different liberation thresholds for efficient recovery.

We instead compute so-called liberation curves, which quantify the degree of liberation of all particles in the set of particles \mathcal{P} over the entire range of possible liberation thresholds. More specifically, for each liberation threshold $\tau \in [0, 1]$, the sets $\mathcal{V}_\tau, \mathcal{S}_\tau \subset \Omega$, given by

$$\mathcal{V}_\tau = \bigcup_{\substack{P \in \mathcal{P} \\ D_{\text{VL}}(P) \geq \tau}} P^{\text{cT}} \quad \text{and} \quad \mathcal{S}_\tau = \bigcup_{\substack{P \in \mathcal{P} \\ D_{\text{SL}}(P) \geq \tau}} P^{\text{cT}},$$

describe the set of those voxels that are associated with the target phase c_{T} in particles with a volume liberation and a surface liberation of at least τ , respectively. Thus, $|\mathcal{V}_0|$ and $|\mathcal{S}_0|$ represent the total volume of the target phase in the population of particles, i.e., $\mathcal{V}_0 = \mathcal{S}_0 = \bigcup_{P \in \mathcal{P}} P^{\text{cT}}$. The cumulative volume and surface liberation curves, denoted by $L_{\text{V}}, L_{\text{S}} : [0, 1] \rightarrow [0, 1]$, respectively, are then given by

$$L_{\text{V}}(\tau) = \frac{|\mathcal{V}_\tau|}{|\mathcal{V}_0|} \quad \text{and} \quad L_{\text{S}}(\tau) = \frac{|\mathcal{S}_\tau|}{|\mathcal{S}_0|}, \quad (8)$$

for each $\tau \in [0, 1]$.

Since volume (and surface) liberation is expected to correlate with particle size, we also computed size-based liberation diagrams to investigate this relationship across different particle size ranges. For this purpose, the set of particles \mathcal{P} is divided into bins based on the volume-equivalent diameter D_D and, for each bin, a LiAlO₂ volume-weighted average volume (and surface) liberation is computed. More specifically, let $b_0 = 0 < b_1 < \dots < b_n$ be the size thresholds for some $n \in \mathbb{N}$, where $b_n \geq \max_{P \in \mathcal{P}} D_D(P)$. Each particle is then assigned to one of the bins $B_1, B_2, \dots, B_n \subset \mathcal{P}$, where $B_k = \{P \in \mathcal{P} : D_D(P) \in (b_{k-1}, b_k]\}$ for each $k \in \{1, \dots, n\}$. The volume-weighted average volume and surface liberation of the k -th size bin, denoted by $\bar{D}_{\text{VL}}(B_k)$ and $\bar{D}_{\text{SL}}(B_k)$, respectively, are then given by

$$\bar{D}_{\text{VL}}(B_k) = \frac{\sum_{P \in B_k} |P^{c_T}| \cdot D_{\text{VL}}(P)}{\sum_{P \in B_k} |P^{c_T}|} \quad \text{and} \quad \bar{D}_{\text{SL}}(B_k) = \frac{\sum_{P \in B_k} |P^{c_T}| \cdot D_{\text{SL}}(P)}{\sum_{P \in B_k} |P^{c_T}|}, \quad (9)$$

for each $k \in \{1, \dots, n\}$. Using these size-based liberation diagrams, we can identify optimal particle size ranges that maximize liberation while considering the trade-offs between processing costs and recovery efficiency.

3 Results

This section presents the results of the workflow for the 3D morphological and mineralogical characterization of slag systems, focusing on descriptors that quantify the size, shape and mineralogical composition of particles. For that purpose, trained networks were applied to CT images of particle systems with different size fractions (from 63 μm to 100 μm , and from 100 μm to 250 μm , the latter being held out entirely during training), and descriptor vectors were computed based on the resulting segmentations, as described in Section 2.4. However, the accuracy of this characterization is highly dependent on the accuracy of the underlying phase-wise and particle-wise segmentations. Thus, in a first step, we evaluate the suitability of the combined segmentation approach considered in this paper.

3.1 Evaluation of segmentation accuracy

We first evaluated under- and over-segmentation resulting from the particle-wise segmentation in a CT subvolume of size 150^3 voxels. For this purpose, markers were manually placed inside each sufficiently large particle visible in the subvolume. In total, 106 particles were identified in this way. A predicted particle label was counted as under-segmented if it contained markers from more than one manually identified particle. Since a predicted particle label without a marker may either indicate over-segmentation or a particle that was missed during manual annotation, these labels were additionally inspected visually to confirm over-segmentation. Using these criteria, one under-segmented particle and one over-segmented particle were identified in the predicted particle-wise segmentation.

Furthermore, the predicted phase-wise segmentation $\hat{S}^M = (\hat{S}_x^M, x \in \Omega)$ and particle-wise segmentation $\hat{S}^P = (\hat{S}_x^P, x \in \Omega)$ were compared to the ground truth segmentation $S = (S_x, x \in \Omega)$ derived from 2D SEM-EDS images, as shown in Figure 5.

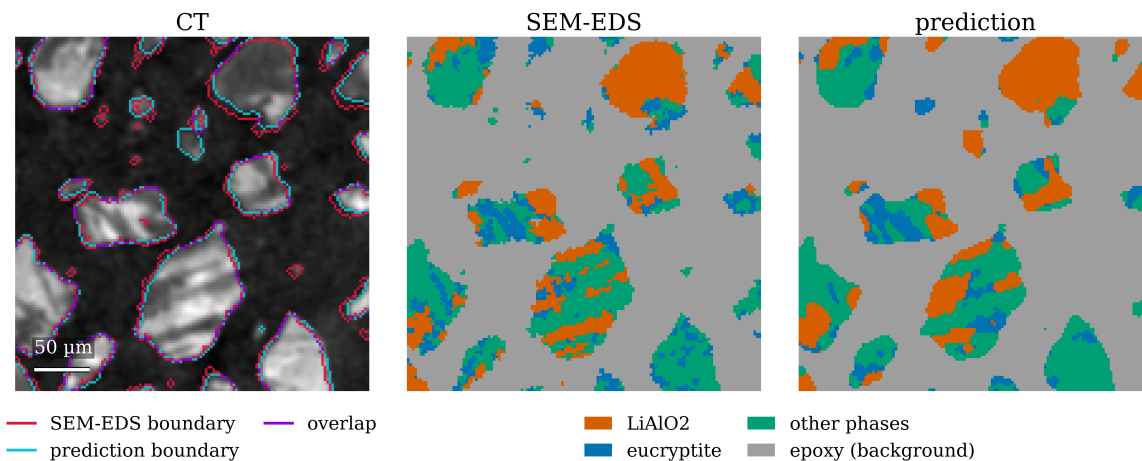


Figure 5: Validation of predicted 3D segmentations against ground truth for a cutout not used for network training.

This evaluation is restricted to the set of voxels $\Omega_{\text{val}} = \{x \in \Omega: S_x \neq -1\}$, i.e., those voxels with known ground truth. We consider the Dice coefficient DICE and the accuracy ACC, given by

$$\text{DICE} = \frac{2 \sum_{x \in \Omega_{\text{val}}} \mathbb{I}(\hat{S}_x^{\text{M}} > 0) \cdot \mathbb{I}(S_x > 0)}{\sum_{x \in \Omega_{\text{val}}} (\mathbb{I}(\hat{S}_x^{\text{M}} > 0) + \mathbb{I}(S_x > 0))} \quad \text{and} \quad \text{ACC} = \frac{\sum_{x \in \Omega_{\text{val}}} \mathbb{I}(\hat{S}_x^{\text{M}} = S_x)}{|\Omega_{\text{val}}|},$$

where $\mathbb{I}(\cdot)$ denotes the indicator.

We computed a Dice coefficient of 0.90, indicating strong agreement between the predicted particle boundaries and the ground truth. The phase-wise segmentation achieves an overall classification accuracy of 0.89 across the four groups of mineral phases: lithium aluminate (target phase), eucryptite, epoxy (background) and other phases. However, the ground truth is significantly more fine-grained than the predicted phase-wise and particle-wise segmentations.

To quantify the effect of this smoothing of fine details on particle-wise characterization, we compared particle-wise descriptors derived from the registered cutouts of 2D SEM-EDS images, with a size of 180×180 voxels, to those derived from the corresponding planar sections of the predicted segmentations, as described in Section 2.4. For this purpose, we computed the mean absolute error MAE, the bias BIAS, and the coefficient of determination R^2 between the vectors of the ground truth descriptors $y = (y_1, \dots, y_N)^\top$ and the predicted descriptor vectors $\hat{y} = (\hat{y}_1, \dots, \hat{y}_N)^\top$, given by

$$\text{MAE}(y, \hat{y}) = \frac{1}{N} \sum_{i=1}^N |\hat{y}_i - y_i|, \quad \text{BIAS}(y, \hat{y}) = \frac{1}{N} \sum_{i=1}^N (\hat{y}_i - y_i), \quad \text{and} \quad R^2(y, \hat{y}) = 1 - \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{\sum_{i=1}^N (y_i - \bar{y})^2},$$

respectively, where $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$ is the mean of the ground truth descriptor values. Furthermore, we computed the Spearman rank correlation coefficient ρ , given by

$$\rho(y, \hat{y}) = \frac{\sum_{i=1}^N (r_i - \bar{r})(\hat{r}_i - \bar{\hat{r}})}{\sqrt{\sum_{i=1}^N (r_i - \frac{N+1}{2})^2 \sum_{i=1}^N (\hat{r}_i - \frac{N+1}{2})^2}},$$

where $r_i \in \{1, \dots, N\}$ denotes the rank of y_i , defined as the position of y_i in the ordered set $\{y_1, \dots, y_N\}$ when sorted in ascending order for each $i \in \{1, \dots, N\}$. Similarly, \hat{r}_i denotes the rank of \hat{y}_i within the ordered set of predicted values.

To obtain the individual components y_i and \hat{y}_i for a given 2D descriptor, denoted by \mathcal{D} , we first identify a set of $N \in \mathbb{N}$ matched particle pairs $\{(P_1, \hat{P}_1), \dots, (P_N, \hat{P}_N)\}$. Specifically, for each $i \in \{1, \dots, N\}$, the set $P_i \subset \Omega$ denotes a particle in the ground truth segmentation, and $\hat{P}_i \subset \Omega$ denotes the corresponding particle in the predicted 3D segmentation having the largest spatial overlap with P_i . Then, the components $y_i = \mathcal{D}(P_i)$ and $\hat{y}_i = \mathcal{D}(\hat{P}_i \cap E_i)$ denote the descriptor values for the particle in the ground truth and the corresponding planar section of the predicted segmentations, respectively.

Here, $E_i \subset \Omega$ is the discretized plane defined by the registration of the specific SEM-EDS cutout that contains the particle P_i .

Moreover, to reliably estimate the segmentation performance, the registered cutouts from one SEM-EDS image were used as a test dataset, i.e., they were only used for evaluation and not for training. In Table 2, we quantified the accuracy of the descriptors introduced in Section 2.4 on the training and test datasets separately. Predictions and ground truth of these particle-wise descriptors are also shown in Figure 6. Note that while the volume-equivalent diameter D_D has a lower bound due to the size filtering step described in Section 2.3.5, this is not necessarily the case for the area-equivalent diameter D'_D . Moreover, the normalized intergrowth boundary D'_{NIB} defined in Eq. (7) can only be evaluated for particles where the intergrowth boundary length D_{IB} is larger than 0, both for the predicted segmentations and the ground truth SEM-EDS image. To prevent unbounded outlier values, we truncate D'_{NIB} at 3, i.e., any value exceeding 3 is set to 3.

Table 2: Mean absolute error, bias, coefficient of determination R^2 and Spearman rank correlation ρ for particle descriptors introduced in Section 2.4. These scores are computed for particles intersecting with the SEM-EDS images used for training (left column) and a held-out SEM-EDS image not used for training (right column).

descriptor	training dataset				test dataset			
	MAE	BIAS	R^2	ρ	MAE	BIAS	R^2	ρ
area-equivalent diameter (pixels)	0.935	-0.546	0.980	0.986	0.760	-0.457	0.990	0.994
volume liberation	0.076	0.017	0.740	0.868	0.078	0.014	0.748	0.883
surface liberation	0.102	0.014	0.610	0.828	0.083	0.029	0.765	0.877
intergrowth boundary (pixels)	6.564	-6.091	0.840	0.895	8.865	-6.145	0.788	0.906
normalized intergrowth boundary	0.257	0.230	0.707	0.983	0.418	0.361	0.324	0.950
Amstutz-Giger index	0.132	-0.097	0.582	0.819	0.128	-0.041	0.641	0.801
elongation	0.065	-0.003	0.708	0.851	0.051	0.002	0.739	0.867
perimeter (pixels)	7.982	-7.093	0.945	0.983	9.946	-9.841	0.929	0.993
circularity	0.176	0.149	-0.570	0.855	0.154	0.147	0.513	0.894

In general, the mean absolute errors for the descriptors considered in Table 2 are similar for the training and test datasets, indicating limited overfitting. Although the size D'_D of the particles is reliably estimated in the predicted segmentation, the shapes of the particles are simplified, resulting in a significant overestimation of the circularity D'_S and the normalized intergrowth boundary D'_{NIB} . For circularity, in particular, this leads to a negative R^2 score on the training dataset. However, the Spearman rank correlation coefficient ρ remains high, e.g., $\rho > 0.95$ for the normalized intergrowth boundary D'_{NIB} on the test dataset. This indicates that the relative ordering of the particles in terms of their shape and intergrowth is preserved, despite the significant bias introduced by the smoothing of fine details. That means that, when comparing multiple particle systems, we can determine which particle system contains more highly intergrown particles. Moreover, this systematic bias could be reduced by introducing a corrective term, however, deriving a robust 3D corrective term using only 2D ground truth data remains challenging.

Volume liberation D_{VL} is estimated with a mean absolute error of 0.078, indicating that the proportions of the mineral phases within the particles are accurately preserved. Surface liberation D_{SL} is also well-estimated in the predicted segmentation with a mean absolute error of 0.083. Although the segmentation of surface voxels is significantly more challenging due to partial volume effects, the MAE is similar to that for volume liberation.

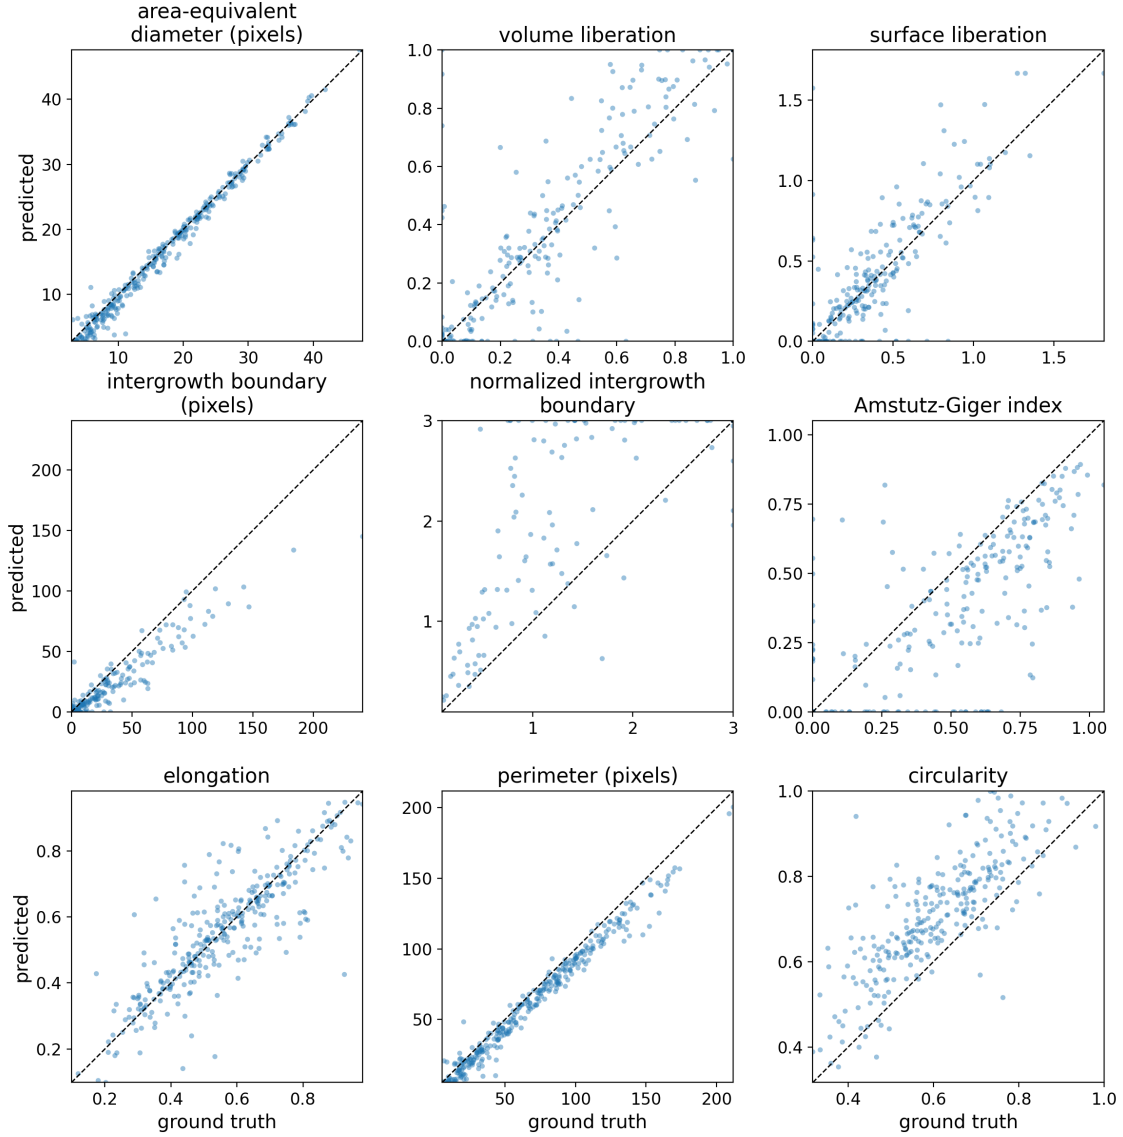


Figure 6: Scatterplots of predicted and ground truth particle-wise descriptors for particles intersecting with SEM-EDS images.

3.2 Particle-wise characterization

Following the comparison of descriptors that have been computed from 2D slices of the segmentations with registered SEM-EDS ground truth, we now characterize the slag particle system using descriptor vectors of the 3D counterparts computed from the predicted 3D segmentations. These vectors provide a low-dimensional representation of the particles' 3D morphology and mineralogy, as well as interdependencies between the components of these vectors.

Figure 7 shows number-based and volume-weighted histograms of particle size. Due to the large number of small particles, the number-based histograms are concentrated at the lower end of the size range. Although a large fraction, 44.77% of all particles, is below a size of 50 μm , they only constitute 2.83% of the total particle volume. Therefore, to obtain a representative characterization of the material, we will only consider volume-weighted histograms. Although the particles were sieved to the size range from 63 μm to 100 μm , particles with a volume-equivalent diameter outside of this range constitute 18.7% of the total particle volume, including a significant number of particles with size greater than 115 μm .




Figure 7: Number-based (left) and volume-weighted (right) histograms of volume-equivalent diameter resulting from the predicted 3D particle-wise segmentation.

As shown in Figure 8, it turns out that the largest particles ($>100\ \mu\text{m}$) are more elongated, with an average elongation D_E below 0.5. In general, the volume-equivalent diameter D_D and the elongation D_E are negatively correlated with a Pearson correlation coefficient of -0.225, while the volume-equivalent diameter D_D and the flatness D_F are positively correlated with a Pearson correlation coefficient of 0.156.

It is important to note that the observed shift towards more elongated shapes for larger particles could also be a result of under-segmentation, as two combined particles would generally exhibit more elongated shapes. Although under-segmentation was extremely uncommon throughout the segmented volume, it could still lead to a small number of large and elongated particles.




Figure 8: Volume-weighted bivariate histograms of particle shape and size descriptors resulting from the 3D particle-wise segmentation: Elongation and volume-equivalent diameter (left), flatness and volume-equivalent diameter (right). The colors represent the weighted relative frequency.

In order to investigate the potential under-segmentation of larger particles and to visualize the shape differences, we have depicted projections of particles along their thinnest axis in Figure 9. Meaning that if the fitted minimum volume oriented bounding box was aligned with the axes, the figure would show the largest 2D projection of the particle and its elongation. The top row shows the 10 largest particles, without visible under-segmentation, and in the bottom row there are 10 randomly chosen particles in the size range from $63\ \mu\text{m}$ to $100\ \mu\text{m}$. Similarly to the shift in the histograms, we can observe a clear visual difference between the shapes of the largest particles and those of the 10 randomly chosen particles, especially with respect to their elongation.




Figure 9: The 10 largest particles (top row) and particles randomly sampled within the size range from 63 μm to 100 μm (bottom row), together with their volume-equivalent diameter and elongation. The depicted elongation values are those obtained from minimum volume oriented bounding boxes.

In Figure 10 (left), a scatterplot is shown that visualizes the results obtained for volume and surface liberation, where the descriptors D_{VL} and D_{SL} exhibit a strong positive correlation of 0.988. This means that particles with a larger volume of the target phase generally also expose more of it at the surface. Moreover, very few particles are fully liberated for surface or volume liberation. When looking at the scatterplot of the two intergrowth descriptors in Figure 10 (right), the normalized intergrowth boundary D_{NIB} (given in Eq. (7)) and the Amstutz-Giger index D_{AG} (given in Eq. (6)) provide complementary information regarding the intergrowth of mineral phases. Although the two descriptors do not exhibit a strong correlation for particles with a low Amstutz-Giger index, this highlights their specific utility for different particle compositions. Particles with very low or very high values of D_{VL} necessarily exhibit a low Amstutz-Giger index. However, these particles may still be characterized as highly intergrown by the normalized intergrowth boundary due to normalization by the volume of the target phase. In contrast, particles with a high Amstutz-Giger index are also described as highly intergrown by the normalized intergrowth boundary.

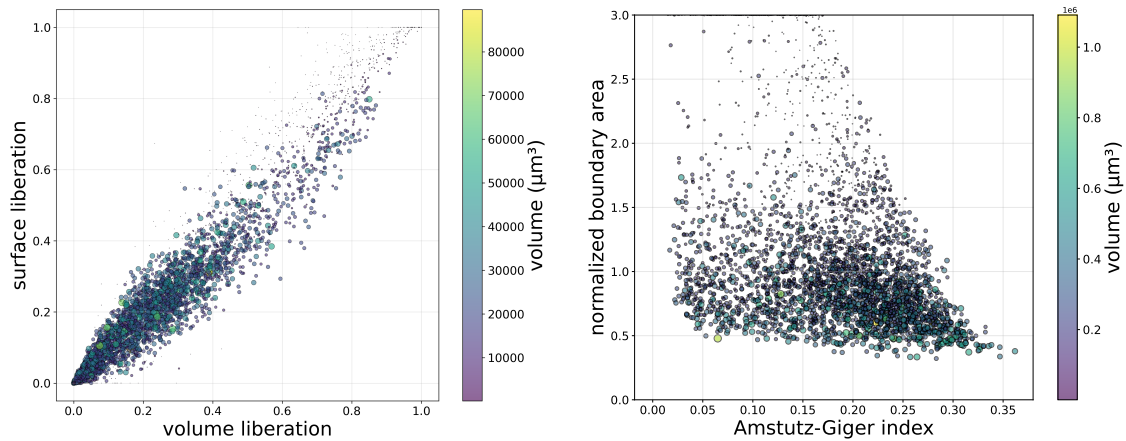


Figure 10: Scatterplots for mineral liberation and intergrowth descriptors resulting from the combined 3D segmentations: Valuable mineral fraction and surface liberation (left) and normalized intergrowth boundary and Amstutz-Giger index (right).

In Section 3.3, we compute the liberation curves introduced in Section 2.4.2 to aggregate the mineralogical descriptors considered above across the entire particle system to characterize the material's degree of liberation.

3.3 Liberation analysis

As shown in Figure 10 (left), only a small number of larger particles exhibit high volume or surface liberation. Although this is in line with visual observations of 3D CT images, the 2D SEM-EDS

images suggest otherwise. Therefore, before investigating the 3D liberation curves, we first compare the liberation curves $L_V^{\text{SEM-EDS}}$ and $L_V^{2\text{D-Pred}}$ resulting from SEM-EDS images and slices of the predicted segmentations, respectively. This allows us to identify biases and inaccuracies in the predicted phase-wise segmentation.

More specifically, the liberation curve $L_V^{\text{SEM-EDS}}$ is computed by applying Eq. (8) to the set of all connected components in the original high-resolution SEM-EDS images. Analogously, $L_V^{2\text{D-Pred}}$ is computed by considering the particles from 50 planar sections of the predicted 3D segmentations. The resulting liberation curves describe the liberation of 2D sections of a large number of particles from the same population and should match closely if the predicted segmentations are accurate.

As shown in Figure 11 (left), the liberation curves $L_V^{\text{SEM-EDS}}$ and $L_V^{2\text{D-Pred}}$ are in close agreement across the entire range of liberation thresholds $\tau \in [0, 1]$. This shows that the neural networks accurately predict the fractions of the target phase across particles of varying mineralogical compositions without introducing systematic biases at specific liberation thresholds. Therefore, the combination of phase-wise and particle-wise segmentation provides a reliable basis for the computation of these liberation curves.

However, a particle’s mineralogical composition observed in a single slice may not be representative of the entire particle due to complex intergrowth structures. This assumption, which largely motivated the use of 3D segmentations in the first place, is especially relevant when investigating the liberation of particles. Therefore, in order to quantify this effect, we compared $L_V^{\text{SEM-EDS}}$ with the liberation curve $L_V^{3\text{D-Pred}}$, which is computed by applying Eq. (8) to the full 3D particles identified in the predicted segmentations. As shown in Figure 11 (right), the liberation curve $L_V^{3\text{D-Pred}}$ resulting from the predicted 3D segmentations consistently shows lower liberation for all thresholds than those observed in 2D. Hence, whereas a significant number of particles is highly liberated in both the SEM-EDS images and individual slices of the 3D segmentations, this is not the case when considering the full 3D volume.

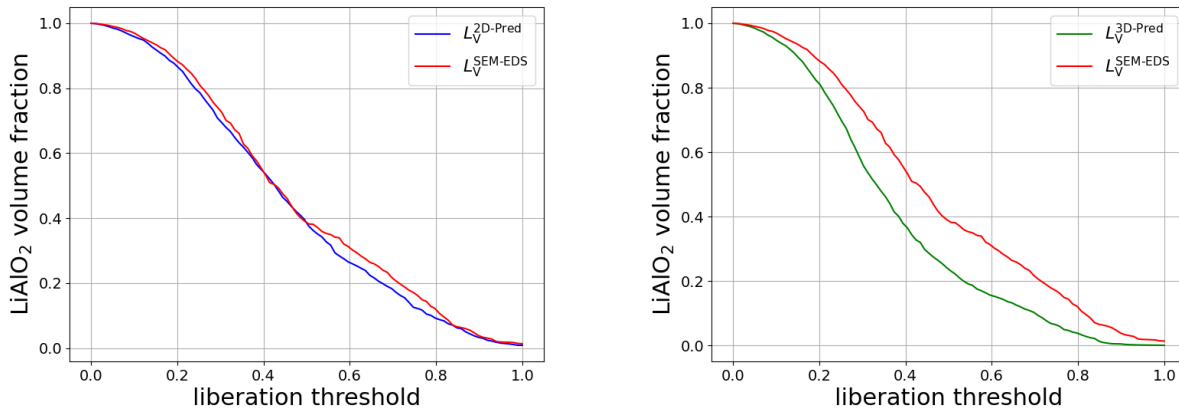


Figure 11: Comparison of predicted and ground truth liberation curves computed for LiAlO_2 using Eq. (8): for 2D slices of the predicted 3D segmentation and 2D ground truth (left), and for the complete predicted 3D segmentation and 2D ground truth (right).

Analogously, we also computed the surface liberation curves for the three datasets, denoted by $L_S^{\text{SEM-EDS}}$, $L_S^{2\text{D-Pred}}$, and $L_S^{3\text{D-Pred}}$. Figure 12 (left) shows that the surface liberation $L_S^{2\text{D-Pred}}$ computed from slices of the segmentations is significantly shifted compared to $L_S^{\text{SEM-EDS}}$. This is likely a result of the fundamental difficulty in correctly segmenting voxels directly at the surface of particles due to partial volume effects. To address these difficulties, we applied a morphological erosion of one voxel [44] to the particles and recomputed their surface liberation. Figure 12 (left) shows that the curve derived from the eroded 2D segmentation aligns closely with the curve derived from the eroded SEM-EDS images. Furthermore, the SEM-EDS curve itself remains largely unchanged after erosion, indicating that the erosion does not fundamentally alter the surface liberation in the ground truth but rather filters out segmentation artifacts at the surface. Therefore, we utilize this erosion for the 3D analysis. Figure 12 (right) displays the surface liberation curve computed from the eroded 3D particles compared to the

eroded SEM-EDS baseline. Similarly to volume liberation results, the 3D surface liberation indicates a lower degree of liberation across all thresholds compared to the 2D estimates.

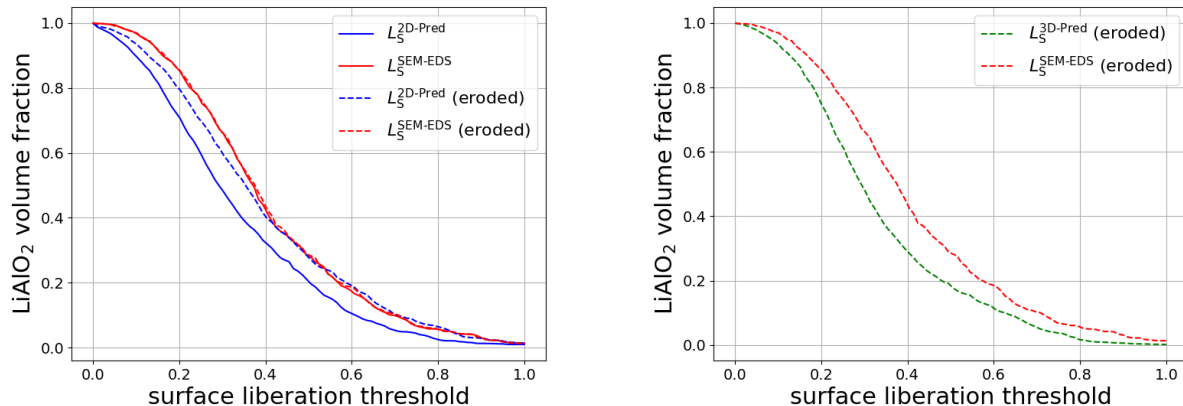


Figure 12: Comparison of predicted and ground truth surface liberation curves computed for LiAlO_2 using Eq. (8): for 2D slices of the predicted 3D segmentation and 2D ground truth (left), and for the complete predicted 3D segmentation and 2D ground truth (right). Dotted lines indicate that the surface liberation curve was computed using eroded particles.

In Section 4 we will discuss alternative approaches that could be applied to datasets where slight erosion is not acceptable.

3.4 Characterization of additional particle systems

As described in Section 2.1, individual CT measurements were conducted for two particle systems derived from the same slag source: one with a particle size range from $63\ \mu\text{m}$ to $100\ \mu\text{m}$ and the other with a size range from $100\ \mu\text{m}$ to $250\ \mu\text{m}$.

In the previous sections, neural networks were only trained and evaluated on the particle system with smaller particles. Although obtaining 3D phase-wise and particle-wise segmentations of a single CT volume is helpful, segmenting additional particle systems can provide further insight. For example, while we have observed a shift towards more elongated shapes for particles larger than $100\ \mu\text{m}$ in Figure 8, it is not clear if this is a result of the sieving process or breakage. Moreover, we also gain insight into the robustness of the segmentation method itself, allowing for better application to entirely different particle systems in the future.

When the networks trained on the size fraction from $63\ \mu\text{m}$ to $100\ \mu\text{m}$ are applied to the size fraction from $100\ \mu\text{m}$ to $250\ \mu\text{m}$, the fraction of LiAlO_2 in the predicted phase-wise segmentation is 0.21, which is similar to that of the SEM-EDS images at 0.22. Moreover, as shown in Figure 13 (left), the observed elongation in the range from $100\ \mu\text{m}$ to $250\ \mu\text{m}$ for this particle system appears similar to that of the other particle system within its respective sieving range from $63\ \mu\text{m}$ to $100\ \mu\text{m}$, despite the significant difference in particle sizes. For example, the volume-weighted mean elongation for both particle systems is 0.72. When looking at Figure 13 (right), we can observe that for larger particles, the variability of volume liberation reduces slightly. To further investigate size and liberation, we consider the size-based liberation diagrams introduced in Section 2.4.2. Since we have particle-wise data available for both particle systems, we can compute the individual volume-weighted average volume liberation for arbitrarily chosen size ranges, as given by Eq. (9).




Figure 13: Volume-weighted bivariate histograms for the particle system with particles in the size range from 100 μm to 250 μm : elongation and equivalent diameter (left) and valuable mineral fraction and equivalent diameter (right).

In Figure 14 we compare the size-based liberation diagrams resulting from the segmentations of both particle systems. It turns out that both size-based liberation diagrams are largely similar in the 40 μm to 125 μm size range, which is consistent with the shared origin and processing of the slag. In that range, the volume-weighted liberation of LiAlO_2 decreases from 55% to 30%, after which it plateaus around 21%. Since this is the overall volume fraction of the target phase, we do not expect to see further changes as the particle size increases. The only significant disagreement between the two diagrams occurs for particles with sizes from 0 μm to 30 μm .




Figure 14: Top: Size-based liberation histograms for the particle systems with smaller (blue) and larger (red) particles. Bottom: Histograms of the number of particles in the respective size bins used for the computation of the size-based liberation diagrams.

This investigation of the relationship between particle size and liberation could be used to identify an optimal target particle size range. Here, the term “optimal” may refer to a trade-off between average liberation and other factors such as costs, emissions and time constraints, which would also need to be considered. In the following section, we discuss the results and give an outlook on further improvements and applications of the presented methods.

4 Discussion

In Section 3, we evaluated the accuracy of the segmentation workflow presented and the resulting particle-wise characterization for both particle systems. The proposed approach allows for a scalable 3D characterization using correlative CT and SEM-EDS images. The resulting phase-wise segmentation achieved an accuracy of 89.23%. The network is able to distinguish mineral phases based on features observed in CT images, despite similar CT intensities for LiAlO_2 and eucryptite.

Although the voxel-wise accuracy is a good indicator of the performance of the phase-wise segmentation network, the evaluation of particle-wise descriptors provides more meaningful insights into the accuracy of the characterization workflow. Descriptors that do not rely on fine details, such as elongation and equivalent diameter, were estimated with high precision. However, descriptors such as perimeter and intergrowth-boundary length were consistently underestimated, indicating that the network tends to smooth out fine details. Consequently, shape descriptors that are influenced by these fine details, such as circularity, suffer with respect to their accuracy. This behavior is likely driven by two factors. For one, partial volume effects at the boundaries of mineral phases (and particles) fundamentally limit the segmentation of fine details. Moreover, minor misalignments of the SEM-EDS images create inconsistent ground truth for fine details, discouraging the network from learning these features. However, the high rank-correlation coefficients for all considered descriptors suggest that the segmentations still capture relative morphological differences between particles, allowing for the reliable identification of highly intergrown particles.

When analyzing the particle system with a particle size range from $63\ \mu\text{m}$ to $100\ \mu\text{m}$, we identified a significant number of particles that were larger than $100\ \mu\text{m}$. As shown in Figure 9, these particles are not an indication of under-segmentation. Instead, as shown in Figure 8, these larger particles tend to be elongated, with an almost cylindrical shape. Because their thickness and width are smaller than the aperture size of the sieve, these particles can end up in the finer fraction [51, 52].

A significant advantage of this workflow is the ability to quantify stereological biases. For this purpose, we computed liberation curves, as described in Section 2.4.2, using particles identified in the 2D SEM-EDS images and the predicted 3D segmentations. The liberation curves resulting from 2D slices of the predicted 3D segmentation showed excellent agreement with those obtained from the SEM-EDS images. However, when using the full 3D volume for this analysis, a clear discrepancy emerged. The liberation curves resulting from 3D images consistently indicated significantly lower liberation than those resulting from 2D slices of the same segmentation. Regarding surface liberation, the segmentation results initially exhibited significant bias compared to the 2D SEM-EDS ground truth. However, as demonstrated in Section 3.3, applying a morphological erosion of one voxel was sufficient to correct this bias. This suggests that while the surface voxels are prone to misclassification due to partial volume effects, voxels close to (but not at) the surface are significantly less affected by these errors. Although this erosion step did not affect the surface liberation computed with the particles in the SEM-EDS images, as shown in Figure 12, this observation would not hold for particle systems where the particles are only coated by a thin layer of the target phase. In such cases, additional loss terms may be necessary for the training of the neural network. For example, a discriminator could be used in an adversarial setting to encourage segmentation networks to learn to predict fine surface details that can be discouraged by the pixel-wise loss function [53, 54].

To evaluate whether the trained particle-wise network can be reliably applied to particle systems of different size fractions, we investigated its sensitivity to changes in particle size relative to the image voxel size. Since rescaling a CT image changes the number of voxels representing a particle, upsampling generates larger particles while downsampling generates smaller ones. Note that this sensitivity analysis was not performed for the phase-wise segmentation, as scaling the image also scales the grains, which generates unrealistic internal structures. Using bilinear interpolation, we applied scaling factors

between 0.5 and 3.0 to the same 150^3 -voxel subvolume that was used to quantify segmentation errors in Section 3.1. The rescaled images were processed by the network without retraining and the results were compared against the predictions at the original resolution. As shown in Figure 15, for scaling factors between 0.75 and 1.75 the volume-weighted errors (missed, under-segmented, and over-segmented particles) remained below 0.5%. However, performance degraded at higher scaling factors, where the network increasingly merged neighboring particles. At the highest considered scale factor of 3.0, over-segmentation remained negligible, but roughly 13% of the reference volume was under-segmented and 16% was missed entirely.

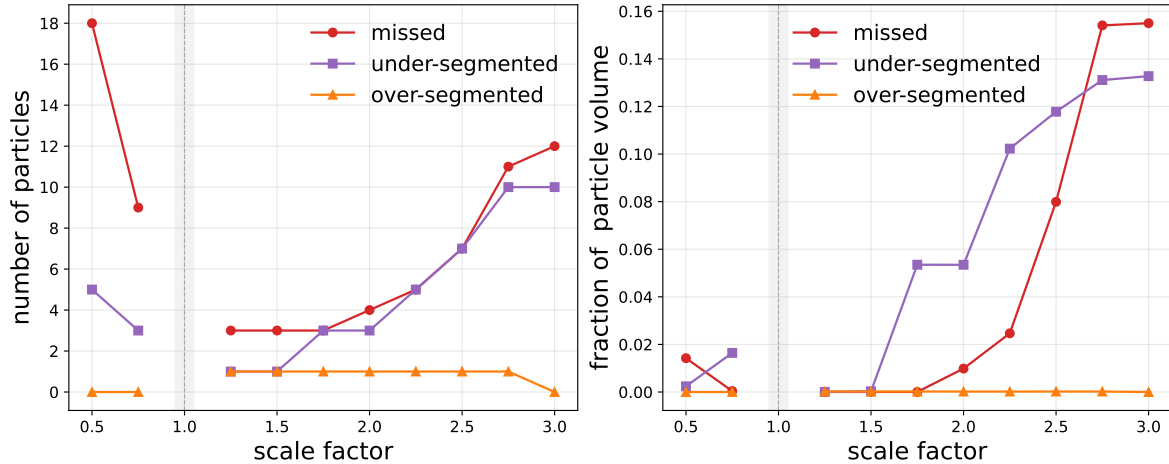


Figure 15: Sensitivity of the particle-wise network to changes in particle size relative to voxel size, evaluated by rescaling a 150^3 -voxel sub-volume by factors between 0.5 and 3.0: number of reference particles that were missed, under- or over-segmented (left) and corresponding fractions of the total reference particle volume (right).

This robustness is important for the analysis of the second slag particle system with a size range from $100\ \mu\text{m}$ to $250\ \mu\text{m}$, as its particles are roughly twice as large as those in the training data. Although trained exclusively on the finer fraction (with particle sizes from $63\ \mu\text{m}$ to $100\ \mu\text{m}$), the networks performed similarly on the second slag particle system with respect to visual inspection, and the overall LiAlO_2 fraction aligned well with that of the ground truth. Limitations arise from the imaged CT volume itself, as it inherently features fewer, larger particles, yielding less robust statistics. Moreover, the grayscale gradient in CT intensity (similar to Figure 3e) required cropping, resulting in a smaller field of view, further reducing the number of particles. In addition, the particles themselves are significantly larger, and thus a more significant fraction of particles intersected the boundary of this smaller field of view and was excluded from subsequent analysis. Since larger and more elongated particles are more likely to intersect with the cropped image boundary, a shift in the respective histograms is expected. The size-based liberation diagrams, shown in Figure 14, for both particle systems showed strong agreement. The only notable disagreement was observed for particles smaller than $30\ \mu\text{m}$. Although this may point to segmentation limitations at the resolution limit for small particles [45], these particles fall significantly outside the sieving ranges of both datasets and do not represent a significant fraction of the overall particle volume. However, it should be noted that applying these trained networks to particle systems originating from a different slag could require retraining due to differences in grain morphology.

Despite the accuracy of the predicted segmentations and the robustness of the trained networks, the workflow is currently constrained by the quality and usability of the image registration process. This registration process relies on an initial binarization of the CT image to align the SEM-EDS images. In the present paper, standard Otsu thresholding was used, which proved insufficient for some slices. As a result, two of the five SEM-EDS images were not successfully registered, leading to less training data. Moreover, it remains difficult to fully quantify registration errors in the successfully registered slices, as discrepancies could stem from either registration errors or inaccuracies in the initial Otsu binarization step itself. Although the SEM-EDS sample preparation is arranged based on the height

of the CT epoxy cylinder sample, providing additional markers on the CT sample could be beneficial for the registration. For example, small cuts or holes could be made at different heights of the epoxy cylinder sample, which would be easily distinguishable in the CT images. Alternatively, particles with higher X-ray attenuation, such as metallic particles, or distinct shapes, such as glass spheres, could be added at distinct heights to provide easily identifiable landmarks compared to slag particles.

Future research will focus on applying the proposed method to slags produced under varying process conditions, such as different cooling rates. Although we were able to distinguish between different particles and mineral phases accurately, the segmentation of individual grains within these particles remains a challenge. A reliable grain-wise segmentation would allow for a more indepth grain-wise characterization, providing further insight into grain growth.

5 Conclusion

In this paper, we presented a scalable workflow for the phase-wise and particle-wise 3D characterization of engineered artificial mineral particle systems using correlative CT and SEM-EDS imaging. By registering 2D SEM-EDS images with 3D CT images, we generated high-quality ground truth to train convolutional neural networks for both phase-wise and particle-wise segmentation. This approach was applied to two particle systems derived from the same lithium-bearing slag: a finer fraction (with particle sizes from 63 μm to 100 μm) used for training and evaluation, and a coarser fraction (with particle sizes from 100 μm to 250 μm) used to assess the robustness of the trained networks.

A major advantage of the proposed workflow is that it does not require manual labeling of mineral phases in the image data. Due to the insufficient grayscale contrast between the target phase LiAlO_2 and other phases, such as eucryptite, manual labeling of these phases is extremely difficult. Using registered SEM-EDS images as ground truth for the training of convolutional neural networks, we obtained accurate phase-wise and particle-wise segmentations of the 3D image data.

The resulting 3D characterization provided insight into the features of the slag particles without stereological biases. In particular, the computed liberation curves demonstrated that the 2D analysis significantly overestimates the liberation of the target phase LiAlO_2 compared to the 3D analysis. This discrepancy highlights the need for 3D characterization to determine effective mechanical process parameters. Furthermore, the proposed workflow showed high robustness when applied to a second particle system with a different size fraction (from 100 μm to 250 μm , compared to training data with a size range from 63 μm to 100 μm) without retraining the networks.

Future research will focus on applying this workflow to slag particle systems produced under varying process conditions, such as different cooling rates, to optimize crystal growth for better liberation. While the current study grouped gehlenite and other non lithium-bearing phases, future work could aim to distinguish between some of these phases to provide a more comprehensive mineralogical characterization. Additionally, we aim to improve the registration of SEM-EDS images to further reduce manual effort and allow grain-wise segmentation of the particle systems.

Acknowledgments

This research is partially funded by the German Research Foundation (DFG) through the research projects 470552553 and 470322626 within the priority programs SPP 2315 “Engineered Artificial Minerals (EnAMs): A Geo-Metallurgical Tool to Recycle Critical Elements from Waste Streams: Synthesis, Characterization, Metallurgical and Mechanical Processing.”

References

- [1] European Parliament and Council. Regulation (EU) 2024/1252 of 11 April 2024 establishing a framework for ensuring a secure and sustainable supply of critical raw materials (European Critical Raw Materials Act). *Official Journal of the European Union*, L 1252/1, 3 May 2024, 2024. URL <http://data.europa.eu/eli/reg/2024/1252/oj>. Accessed 4 July 2025.

- [2] E. Comission. *Critical Raw Materials for Strategic Technologies and Sectors in the EU: A Foresight Study*. Publications Office of the European Union, 2020.
- [3] H. Dang, Z. Chang, H. Zhou, S. Ma, M. Li and J. Xiang. Extraction of lithium from the simulated pyrometallurgical slag of spent lithium-ion batteries by binary eutectic molten carbonates. *International Journal of Minerals, Metallurgy and Materials*, 29:1715–1721, 2022.
- [4] T. Elwert, K. Strauss, T. Schirmer and D. Goldmann. Phase composition of high lithium slags from the recycling of lithium ion batteries. *World of Metallurgy - ERZMETALL*, 65:163–171, 2012.
- [5] C. Liu, J. Lin, H. Cao, Y. Zhang and Z. Sun. Recycling of spent lithium-ion batteries in view of lithium recovery: A critical review. *Journal of Cleaner Production*, 228:801–813, 2019.
- [6] G. Harper, R. Sommerville, E. Kendrick, L. Driscoll, P. Slater, R. Stolkin, A. Walton, P. Christensen, O. Heidrich, S. Lambert, A. Abbott, K. Ryder, L. Gaines and P. Anderson. Recycling lithium-ion batteries from electric vehicles. *Nature*, 575:75–86, 2019.
- [7] M. Sommerfeld, C. Vonderstein, C. Dertmann, J. Klimko, D. Oráč, A. Miškufová, T. Havlík and B. Friedrich. A combined pyro-and hydrometallurgical approach to recycle pyrolyzed lithium-ion battery black mass. Part 1: Production of lithium concentrates in an electric arc furnace. *Metals*, 10:1069, 2020.
- [8] C. W. Bale, P. Chartrand, S. Degterov, G. Eriksson, K. Hack, R. B. Mahfoud, J. Melançon, A. Pelton and S. Petersen. Factsage thermochemical software and databases. *Calphad*, 26:189–228, 2002.
- [9] S. Chakrabarty, H. Li, T. Schirmer, S. Hampel, U. Fittschen and M. Fischlschweiger. Non-equilibrium thermodynamic modelling of cooling path dependent phase evolution of Li_2iSiO_3 from $\text{Li}_2\text{O-SiO}_2$ melt by considering mixed kinetic phenomena and time-dependent concentration fields. *Scripta Materialia*, 242:115922, 2024.
- [10] C. Rachmawati, J. Weiss, H. I. Lucas, E. Löwer, T. Leißner, D. Ebert, R. Möckel, B. Friedrich and U. A. Peuker. Characterisation of the grain morphology of artificial minerals (EnAMs) in lithium slags by correlating multi-dimensional 2D and 3D methods. *Minerals*, 14:130, 2024.
- [11] A. Schnickmann, S. Hampel, T. Schirmer and U. E. A. Fittschen. Formation of lithium-manganates in a complex slag system consisting of $\text{Li}_2\text{O-MgO-Al}_2\text{O}_3\text{-SiO}_2\text{CaO-MnO}$ — A first survey. *Metals*, 13:2006, 2023.
- [12] L. G. Leon, K. J. Hogmalm and M. Bengtsson. Understanding mineral liberation during crushing using grade-by-size analysis—A case study of the Penuota Sn-Ta mineralization, Spain. *Minerals*, 10:164, 2020.
- [13] E. Baawuah, C. Kelsey, J. Addai-Mensah and W. Skinner. Comparison of the performance of different comminution technologies in terms of energy efficiency and mineral liberation. *Minerals Engineering*, 156:106454, 2020.
- [14] T. Leißner, K. Bachmann, J. Gutzmer and U. Peuker. MLA-based partition curves for magnetic separation. *Minerals Engineering*, 94:94–103, 2016.
- [15] S. Weber, O. Furat, T. Kirstein, T. Leißner, U. A. Peuker and V. Schmidt. Computational workflow for the characterization of size, shape, and composition of particles and their separation behavior during processing. *Powders*, 4:1, 2025.
- [16] J. Ohser and F. Mücklich. *Statistical Analysis of Microstructures in Materials Science*. J. Wiley & Sons, 2000.
- [17] T. Allen. *Powder Sampling and Particle Size Determination*. Elsevier, 2003.

- [18] S. J. Blott and K. Pye. Particle shape: A review and new methods of characterization and classification. *Sedimentology*, 55:31–63, 2008.
- [19] O. Furat, T. Kirstein, T. Leißner, K. Bachmann, J. Gutzmer, U. Peuker and V. Schmidt. Multidimensional characterization of particle morphology and mineralogical composition using CT data and R-vine copulas. *Minerals Engineering*, 206:108520, 2024.
- [20] E. Maire and P. J. Withers. Quantitative X-ray tomography. *International Materials Reviews*, 59:1–43, 2014.
- [21] M. Buchmann, N. Borowski, T. Leißner, T. Heinig, M. A. Reuter, B. Friedrich and U. A. Peuker. Evaluation of recyclability of a WEEE slag by means of integrative X-ray computer tomography and SEM-based image analysis. *Minerals*, 10(4):309, 2020.
- [22] Z. H. Weng, L. Van Zwieten, E. Tavakkoli, M. T. Rose, B. P. Singh, S. Joseph, L. M. Macdonald, S. Kimber, S. Morris, T. J. Rose, B. S. Archanjo, C. Tang, A. E. Franks, H. Diao, S. Schweizer, M. J. Tobin, A. R. Klein, J. Vongsvivut, S. L. Y. Chang, P. M. Kopittke and A. Cowie. Microspectroscopic visualization of how biochar lifts the soil organic carbon ceiling. *Nature Communications*, 13(1):5177, 2022.
- [23] P. Sun, L. Ge, Y. Liu, L. Bingke and X. Nie. Modeling of multi-mineral-component digital core based on Res-UNet. *Journal of Geophysics and Engineering*, 20:483–493, 2023.
- [24] I. Varfolomeev, V. Svinin and I. Yakimchuk. 3D multiclass digital core models via microCT, SEM-EDS and deep learning. *E3S Web of Conferences*, 366:01003, 2023.
- [25] O. Furat, T. Leißner, R. Ditscherlein, O. Sedivy, M. Weber, K. Bachmann, J. Gutzmer, U. Peuker and V. Schmidt. Description of ore particles from X-ray microtomography (XMT) images, supported by scanning electron microscope (SEM)-based image analysis. *Microscopy and Microanalysis*, 24:461–470, 2018.
- [26] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox and O. Ronneberger. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In S. Ourselin, L. Joskowicz, M. R. Sabuncu, G. Unal and W. Wells, editors, *Medical Image Computing and Computer-Assisted Intervention*, pages 424–432. Springer, 2016.
- [27] S. N. Chiu, D. Stoyan, W. S. Kendall and J. Mecke. *Stochastic Geometry and Its Applications*. J. Wiley & Sons, 3rd edition, 2013.
- [28] D. Uliana and C. Ulsen. Mineral liberation by 3D X-ray microtomography and SEM-based image analysis in low-grade iron ores with different mineralogy and texture. *Minerals Engineering*, 222:109150, 2025.
- [29] R. Ditscherlein, T. Leißner and U. A. Peuker. Preparation strategy for statistically significant micrometer-sized particle systems suitable for correlative 3D imaging workflows on the example of X-ray microtomography. *Powder Technology*, 395:235–242, 2022.
- [30] R. D. Hanna and R. A. Ketcham. X-ray computed tomography of planetary materials: A primer and review of recent studies. *Geochemistry*, 77:547–572, 2017.
- [31] P. Withers, C. Bouman, S. Carmignato, V. Cnudde, D. Grimaldi, C. Hagen, E. Maire, M. Manley, A. Du Plessis and S. Stock. X-ray computed tomography. *Nature Reviews Methods Primers*, 1:18, 2021.
- [32] E. E. Underwood. *Quantitative Stereology*. Addison-Wesley Publishing Company, 1970.
- [33] G. Barbery. *Mineral Liberation: Measurement, Simulation and Practical Use*. Éditions GB, 1991.
- [34] A. Siddique, J. R. Godinho, J. Sittner and L. Pereira. Overcoming stereological bias: A workflow for 3D mineral characterization of particles using X-ray micro-computed tomography. *Minerals Engineering*, 201:108200, 2023.

- [35] S. Spencer and D. Sutherland. Stereological correction of mineral liberation grade distributions estimated by single sectioning of particles. *Image Analysis & Stereology*, 19:175–182, 2000.
- [36] J. Hoshen and R. Kopelman. Percolation and cluster distribution. I. Cluster multiple labeling technique and critical concentration algorithm. *Physical Review B*, 14:3438–3445, 1976.
- [37] O. Ronneberger, P. Fischer and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W. M. Wells and A. F. Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241. Springer International Publishing, 2015.
- [38] K. He, X. Zhang, S. Ren and J. Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 770–778. IEEE, 2016.
- [39] Z. Zhang and Q. Liu. Road extraction by deep residual U-Net. *IEEE Geoscience and Remote Sensing Letters*, 15:749–753, 2018.
- [40] Z. Botev, D. P. Kroese and T. Taimre. *Data Science and Machine Learning: Mathematical and Statistical Methods*. Chapman and Hall/CRC, 2nd edition, 2025.
- [41] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In Y. Bengio and Y. LeCun, editors, *Track Proceedings of the 3rd International Conference on Learning Representations*, pages 1–15, 2015.
- [42] C. Shorten and T. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6:1–48, 2019.
- [43] P. Simard, D. Steinkraus and J. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Seventh International Conference on Document Analysis and Recognition*, pages 958–963, 2003.
- [44] P. Soille. *Morphological Image Analysis: Principles and Applications*. Springer, 2003.
- [45] J. G. Behnson, K. Black, J. E. Houghton and R. H. Worden. A review of particle size analysis with X-ray CT. *Materials*, 16(3):1259, 2023.
- [46] K. Schladitz, J. Ohser and W. Nagel. Measuring intrinsic volumes in digital 3D images. In A. Kuba, L. G. Nyúl and K. Palágyi, editors, *Discrete Geometry for Computer Imagery*, pages 247–258. Springer, 2006.
- [47] K. Benkrid, D. Crookes and A. Benkrid. Design and FPGA implementation of a perimeter estimator. *Proceedings of the Irish Machine Vision and Image Processing Conference*, pages 51–57, 2000.
- [48] G. Barequet and S. Har-Peled. Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. *Journal of Algorithms*, 38:91–109, 2001.
- [49] G. C. Amstutz and H. Giger. Stereological methods applied to mineralogy, petrology, mineral deposits and ceramics. *Journal of Microscopy*, 95:145–164, 1972.
- [50] L. Pérez-Barnuevo, E. Pirard and R. Bolibar. Textural descriptors for multiphasic ore particles. *Image Analysis & Stereology*, 31:175–184, 2012.
- [51] P. C. Bartley, B. E. Jackson and W. C. Fonteno. Effect of particle length to width ratio on sieving accuracy and precision. *Powder Technology*, 355:349–354, 2019.
- [52] S. Durand, B. E. Jackson, W. C. Fonteno and J. C. Michel. Particle size distribution of growing media constituents using dynamic image analysis: Parametrization and comparison to sieving. *Oil Science Society of America Journal*, 87:767–780, 2023.

- [53] P. Luc, C. Couprie, S. Chintala and J. Verbeek. Semantic segmentation using adversarial networks. In *NIPS Workshop on Adversarial Training*, pages 1–12, 2016.
- [54] Y. Zheng, M. Yang, M. Wang, X. Qian, R. Yang, X. Zhang and W. Dong. Semi-supervised adversarial semantic segmentation network using transformer and multiscale convolution for high-resolution remote sensing imagery. *Remote Sensing*, 14(8), 2022.