Characterising Heavy Mineral Concentrate Grain Morphology and Mineralogy with Computer Vision

Melvin Hartley

Melinda Hodkiewicz
Mechanical Engineering
The University of Western Australia

Nigel Brand, Christabel Brand
CEED Client: Portable Spectral Services

Abstract

Traditional geological methods for quantifying grain morphology and mineralogy are often manual, time-consuming, and rely on subjective classification criteria. This case study delineates and characterises individual heavy mineral concentrate grains ranging from 300 to 100 microns in diameter within sample sizes ranging from hundreds to thousands of grains. We compare two segmentation algorithms designed to enhance the quantification of granule morphology: a) a Euclidean-based watershed (EBW) technique and a deep convolutional neural network, and b) the Segment Anything (SAM) model from Meta, for their grain segmentation accuracy. Empirical results indicate a marked improvement with the SAM model, showcasing a Dice Similarity Coefficient (DICE) of 0.974 compared to 0.665 for the watershed method, a Jaccard index of 0.949 versus 0.526, a Grain Count Accuracy (GCA) of 0.939 as opposed to 0.393, and an F-index of 0.974 compared to 0.665. We then integrate the SAM method with micro-X-ray Fluorescence (XRF) analysis to correlate mineralogical and morphological data on a grain-by-grain basis. This approach has been tested at scale in a robot-enabled commercial laboratory. It is a significant improvement in analytical precision and operational efficiency in joint characterisation of mineralogy and morphology of samples for geologists, geometallurgists, and process engineers in the environmental and mining industries.

1. Introduction

The characterisation of the morphology of individual grains provides insight into the deposition system and environment of the sample, weathering processes, alteration (such as grain pseudo morphing), and mineralization (Boggs et al., 2012). Additionally, the characterisation of individual grains facilitates the examination of their distinct chemical attributes, thereby yielding a more comprehensive insight into the mineral's composition. Geology is an observational science and in geological analysis traditional methodologies have relied on manual, on-site visual inspection and subjective classification of mineralogical samples. These approaches are time-consuming and are constrained by the human eye's limitations in discerning subtle colours and textures within the visible spectrum (W. J. Shim et al.,2017). Additionally, these methods do not make optimal use of specialised human expertise, which is an increasingly scarce resource.

Recent developments in scanning technologies and image analysis facilitate rapid, objective, and comprehensive data extraction, significantly enhancing the efficiency and reliability

of mineralogical data acquisition for use by geologists and geometallurgists. However, technological limitations persist, particularly in the accurate identification of grain borders and morphology when features overlap or are poorly defined (Balaram, 2021).

X-ray fluorescent analysis (XRF) is a common method of mineralogical quantification due to its non-destructive, non-contact nature. State-of-the-art analysis uses micro XRF for large-scale, robot-enabled, quantitative analysis of heavy mineral concentrates. The output is a detailed chemical analysis for each pixel in an image however there is, as yet, no published approach to directly link specific grains in the visual images with corresponding XRF mineral data. Addressing this issue is an industry imperative, as joint morphological and mineralogical characterisation of grains informs the geologist as to mineral deposit genesis and the geometallurgist considering mineral processing options. This case study addresses this need by presenting a segmentation model that enables grain-by-grain morphological quantification and a pipeline to associate micro XRF mineral data with specific grains in the image.

2. Process

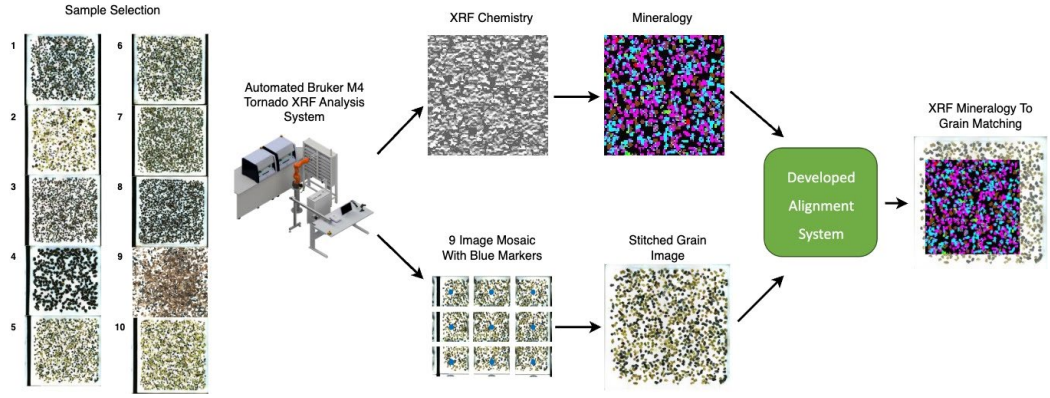

Figure 1 Spectral and image analysis with alignment.

Our approach is described in three parts, 1) sample preparation, 2) segment model selection process and 3) aligning the high-resolution grain image with the XRF mineral composition analysis.

2.1 Sample Selection and Image Preparation

First, a set of 10 test samples were selected using a design of experiments approach to cover variations in sample characteristics such as colour, morphology, and size. The goal being to challenge the robustness and flexibility of the segmentation algorithms. The samples are processed by an automated robotic system. Each sample of heavy mineral concentrate grains ranging from 300 to 100 microns in diameter is accurately weighed and placed on a white Teflon puck. The sample is transferred to a Bruker M4 Tornado XRF analyser for detailed scanning. This scanning system produces 1) XRF chemistry result, and 2) a high resolution nine-image mosaic. The XRF data is then used to generate mineralogy maps. The location of the camera and therefore the coordinates of the centre for each image, as well as the bounding coordinates of the XRF scanning path, are recorded in the XRF coordinate system and documented in a log file.

As seen in Figure 1 the scanned XRF mineral data captures only a subset of the larger sample of grains, necessitating an alignment system to accurately identify the portions of each sample corresponding to the regions scanned by the XRF analysis. A referencing system was developed that utilises a distinct marker identifiable in both the contour map image and the scanned XRF mineralogy data. This marker was used to establish a scale between the two corresponding domains. A distinct blue marker is placed centrally in each of the nine mosaic images to facilitate precise image alignment. These images are then stitched together using a Binary Robust Invariant Scalable Keypoints (BRISK) algorithm. Following the stitching process, an object detection algorithm that utilises colour and circularity thresholds is deployed to identify the locations of the nine markers within the stitched grain image. This provides a set of new coordinates indicating the location of the markers within the stitched image coordinate space.

2.2 Segmentation Model Assessment

The baseline method used a conventional water segmentation technique using gradient and thresholding mechanisms along with a Euclidean distance-based marker system to define distinct regions within the image. The Euclidean distance transform is applied to a pre-processed image to help generate markers at the centres of objects, which serve as the starting points for the flooding process (Chen et al., 2004).

The challenger method uses DCNNs, specifically the SAM model from Meta AI Research. The SAM was chosen for its zero-shot performance and its training on a dataset of over 11 million images (Kirillov et al., 2023). Post-processing steps applied to the segmentation masks generated by the SAM model include identifying contours of the masks, isolating the largest contour, and assessing the area enclosed by the contour against predefined minimum and maximum thresholds, which the user sets to specify the desired size range of the grains of interest.

For both approaches the detected contours of each grain are then used to identify the bounds within the coloured grain image for grain-by-grain colour analysis. This analysis averages the RGB values of all the pixels within the coloured image that are bounded by a grain's contour. The contours for each identified grain are used to generate the morphological characterisation of the grains. The morphological statistics calculated for each grain are area, perimeter, ellipticity, compactness, roundness and extent. Assessment of each approach is made using specificity, precision, accuracy, negative predictive value (NPV), and the F-Index scores plus two metrics commonly used in segmentation performance quantification, the Dice Similarity Coefficient (DICE) and the Jaccard (JAC) index. In addition, for each segmentation model, the particle count was taken and compared to the verified manual count of the grains. Adopting a similar approach to the Cell Count Accuracy (CCA) method a Grain Count Accuracy (GCA) metric was established. Details of the results are described in the results section. The SAM model was selected for the next phase of the work. The contours for each identified grain are then used to generate 1) the morphological characterisation of the grains (area, perimeter, ellipticity, compactness, roundness, and extent), and 2) averages for the RGB values of all the pixels within the coloured image that are bounded by a grain's contour.

2.3 Aligning the High Resolution Grain Image With the XRF Mineral Composition Analysis

The coordinates in the image space corresponding to the detected markers are combined with the XRF process marker coordinates to calibrate an affine transformation across the two coordinate systems. This transformation is subsequently applied to the bounding coordinates provided by the XRF coordinate system, yielding an equivalent set of bounding coordinates within the stitched image space. These new bounding coordinates delineate the scanned area within the stitched image as defined by the affine transformation. The corresponding region of the grain image is then cropped and processed through our segmentation algorithm to obtain the segmentation masks of the grains. When these segmentations are superimposed onto the XRF mineralogy map, a translational discrepancy emerges between the detected grain positions and the mineral data points. This discrepancy is attributable to the incident angle of the XRF beam during analysis, which causes the XRF chemistry to project an offset representation of where the grains are located. A cross-correlation technique using fast Fourier transforms is utilised to account for the translational discrepancies between the XRF mineral data points and the locations of the detected grains. The process begins by generating a binarised image representation of the XRF mineral dataset, simplifying the depiction of minerals to indicate the presence or absence at specific locations.

The identified region of interest from the stitched grain image, determined through the previous affine transformation, is processed through our segmentation model to determine the locations and masks of the grains. These segmentation masks are then binarised to indicate the presence or absence of grains at specific locations. Now, the XRF mineral data is captured at a significantly lower resolution, represented by a 90 by 90 grid of scanned data points, which is substantially coarser than the approximately 1250-pixel resolution of the binarised grain segmentation. This binarised grain segmentation is subsequently dilated using OpenCV's erode function with a 15 by 15-pixel kernel. This dilation effectively enlarges the black areas of the image (represented by zeros), reducing the visibility of finer features within the higher-resolution grain segmentation. This adjustment helps the grain segmentation more closely resemble the lower-resolution, binarised XRF mineral dataset. By aligning the resolution characteristics of both datasets through dilation, the comparability is enhanced, which is ideal for effective cross-correlation analysis.

Following the refinement of grain image alignment with the XRF mineral data, the mineral composition analysis is conducted using the segmentation masks of the grains to assign pixels to corresponding grains For each pixel within a grain, an assignment is made to the corresponding XRF scanned point. Each scanned point captures individual elemental counts and the resultant mineral. The frequency of these pixel assignments to various XRF scanned points is systematically accumulated for each grain, culminating in a detailed mineral and elemental composition profile for each grain within the sample.

3. Results and Discussion

Table 1 illustrates the averaged segmentation performance results for ten different samples for the following 1) particle count, DICE, JAC and GCA measures, and FI scores. The SAM model consistently outperformed the Euclidean-based Watershed (EBW) in grain segmentation accuracy across all samples and performance metrics. The EBW model achieved its best results in scenarios with minimal particle aggregation and where the grain colours distinctly contrasted

with the background. Regarding particle count there is a significant improvement in SAM over the EBW model but it is not flawless so there is more work to be done in this area.

	Euclidean-based Watershed model					Segment Anything model					True Counts
	Part. Count	DICE	JAC	GCA	F-Index	Part. Count	DICE	JAC	GCA	F-Index	
Sample1	410	0.703	0.542	0.375	0.703	545	0.977	0.956	0.953	0.977	568
Sample2	276	0.511	0.343	0.273	0.511	413	0.975	0.951	0.943	0.975	431
Sample3	466	0.703	0.542	0.378	0.703	645	0.969	0.939	0.923	0.969	673
Sample4	201	0.924	0.859	0.759	0.924	204	0.978	0.958	0.958	0.978	213
Sample5	381	0.765	0.620	0.468	0.765	536	0.972	0.945	0.925	0.972	560
Sample6	476	0.650	0.482	0.330	0.650	648	0.958	0.919	0.902	0.958	691
Sample7	551	0.763	0.616	0.451	0.763	683	0.973	0.947	0.937	0.973	717
Sample8	585	0.810	0.681	0.519	0.810	706	0.969	0.939	0.922	0.969	743
Sample9	889	0.174	0.095	0.051	0.174	1020	0.993	0.985	0.979	0.993	1028
Sample10	486	0.648	0.479	0.329	0.648	688	0.977	0.955	0.950	0.977	712
Average		**0.665**	**0.526**	**0.393**	**0.665**		**0.974**	**0.949**	**0.939**	**0.974**	

Table 1 Comparison of results for the Euclidean-based Watershed model and the Segment Anything model

In scenarios with dense particle aggregation and where grain colours closely resembled the background, the watershed model exhibited significant limitations in performance. Conversely, the SAM model demonstrated superior segmentation capabilities under these challenging conditions. For example, in the analysis of test sample 9, the SAM model achieved a DICE coefficient of 0.993, markedly higher than the 0.174 recorded for the watershed model, as detailed in Table 1. Similarly, the JAC and GCA indexes for the SAM model were 0.985 and 0.979 respectively, compared to only 0.051 and 0.174 for the watershed model. This substantial difference in performance metrics demonstrates the robustness of the SAM model in handling this complex segmentation task compared to the traditional Waterfall model for this specific task.

Aligning the XRF analysis and the image is a two-stage process because the affine transformation on its own revealed a pronounced bias towards the Y-axis. Specifically, the mean percentage shift on the Y-axis is significantly higher at 4.76%, compared to a much smaller mean shift on the X-axis of only 0.30%. This indicates a dominant translational shift predominantly in the vertical direction. An explanation for the observed bias in the Y-axis shifts could be attributed to parallax error, which arises from the angled positioning of the XRF camera relative to the vertical axis. Specifically, when scanning larger grains, the XRF signals may seem to originate from a specific point on the detection plate; however, they emanate from a point closer to the camera, located at the top of the grain. This can be seen in Figure 3.

The impact on our analysis is mitigated by a cross-correlation method as previously described. This technique is applied to each sample, effectively correcting any translational discrepancies, whether they arise from parallax or other sources. The effectiveness of this correction is demonstrated in the subsequent Figures 3, 4, which highlight the alignment improvements between the XRF data and the grain images. The assessment is a qualitative one, future work will look at how to quantify this alignment.

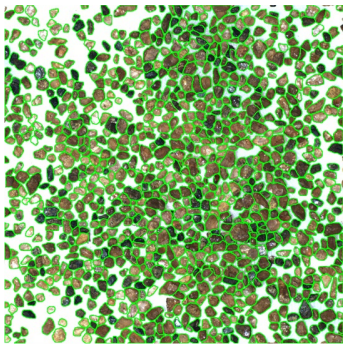

Figure 2: Segment Anything model segmentation results.

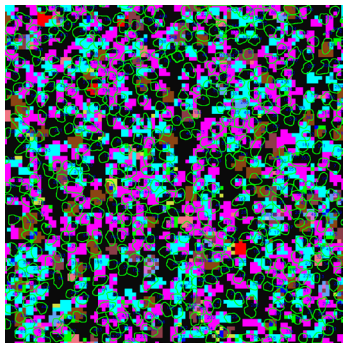

Figure 3: Segmentation data overlaid on the micro XRF mineral data.

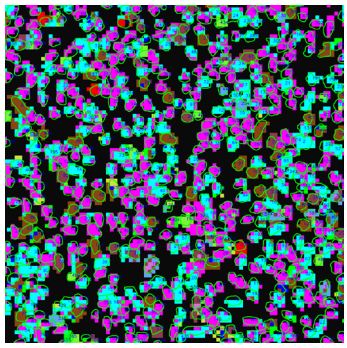

Figure 4: After cross-correlation alignment.

4. Conclusions and Future Work

This study has developed and compared two segmentation models, demonstrating that the SAM model significantly outperforms the Euclidean-based watershed method in delineating grain morphology within dense, homogeneous mineral aggregates and has an acceptable performance for this use case. We note that other groups involved in the mineral and grain processing sectors are also proposing SAM-based image processing 1) for particle size distribution over conventional mechanical sieve analysis and 2) with fine tuning for digital rock image physics.

5. Acknowledgements

The author extends sincere gratitude to several individuals whose support was invaluable throughout this project. Special thanks to the entire team at Portable Spectral Services (PSS), particularly Dr. Nigel Brand and Christabel Brand, for their invaluable insights and assistance. Appreciation is also due to Melinda Hodkiewicz, my supervisor, whose guidance and expertise have been fundamental to the success of this research.

6. References

Balaram, V. (2021). Current and emerging analytical techniques for geochemical and geochronological studies. Geological Journal, 56(5), 2300–2359.

Boggs, S., et al. (2012). Principles of sedimentology and stratigraphy. Biblioteca Hern´an Malo Gonz´alez.

Chen, Q., Yang, X., & Petriu, E. M. (2004). Watershed segmentation for binary images with different distance transforms. Proceedings of the 3rd IEEE international workshop on Haptic, Audio and Visual environments and their applications, 2, 111–116.

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C.Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W. -Y., et al. (2023). Segment anything. Proceedings of the IEEE/CVF International Conference on Computer Vision, 4015–4026.

Shim, W. J., Hong, S. H., & Eo, S. (2017). Identification methods in microplastic analysis: A review. Analytical Methods, 9(9), 1384–1391.