

MinDet1: A deep learning-enabled approach for plagioclase textural studies

✉ Norbert Toth*^α and  John Maclennan^α

^αUniversity of Cambridge, Cambridge, UK.

ABSTRACT

Quantitative textural attributes, such as crystal size distributions (CSDs) or crystal aspect ratios, provide important constraints on the thermal history of rocks. They facilitate the investigation of crystal nucleation, growth, and mixing as well as the cooling rate. However, they require large volumes of crystal segmentations and measurements often obtained with manual methods. Here, a deep learning-based technique—instance segmentation—is proposed to automate the pixel-by-pixel detection of plagioclase crystals in thin-section images. Using predictions from a re-trained model, the physical properties of the detected crystals (size and aspect ratio) are rapidly generated to provide textural insights. These are validated against published results from manual approaches to demonstrate the accuracy of the method. The power and efficiency of this approach is showcased by analysing an entire sample suite, segmenting over 48,000 crystals in a matter of days. The approach is available as MinDet1 software for users with moderate expertise in Python. Widespread use of MinDet may facilitate significant developments in igneous petrography and related fields.

KEYWORDS: Deep Learning; Segmentation; Petrography; Timescales.

1 INTRODUCTION

The textures of igneous rocks provide direct observational constraints on the physics of magmatic processes [Cashman 1990; Higgins 2010b; Askaripour et al. 2022]. Petrologists have long made use of quantitative descriptions of rock texture such as crystal size distributions (CSDs) [Cashman and Ferry 1988; Marsh 1988; Mangan 1990; Marsh 1998; Higgins 2000; Armi-enti 2008] and crystal shapes [Jerram 2003; Hersum and Marsh 2006; Holness 2014] as powerful methods for interrogating the thermal history of rocks and the timescales of processes affecting them. Due to the relative ease of direct observation, detailed petrographic descriptions of igneous rocks, often made using thin-sections, are among the first steps in most petrological studies. In some cases the entire observational record from a study is petrographic. The understanding gained from this primary analysis enables petrologists to target the most useful parts of the rock for further, often more resource-intensive, analyses. Therefore, the development of rapid and accurate textural tools may allow for further streamlining of future studies.

Plagioclase feldspar is the most abundant mineral in the Earth's crust [Smith and Brown 1988]. As a major primary igneous phase over a wide compositional range, it has found extensive use as a reliable tracer for determining igneous thermal history [Cashman 1993; Higgins 1996; Neave et al. 2017; Holness et al. 2020] with crystal shapes shown to vary with crystallisation time [Lofgren 1974]. A simple parameterisation of crystallisation time, t and mean crystal aspect ratio, A , was developed by Holness [2014] where $\log_{10}(t)$, is linearly related to A . This calibration was based on observations from a set of basaltic sills with varying thicknesses, quantifying prior observations that apparently tabular feldspars record slower cooling than crystals with higher aspect ratios. This relation-

ship provides a direct method of determining the timescales of cooling of plagioclase-bearing igneous rocks from simple two-dimensional observations alone, albeit that this calibration is restricted to sills and other magmatic environments with impinging plagioclase networks. More recent work by Mangler et al. [2022] took this idea further and looked at plagioclase microclites in intermediate melts to study how plagioclase crystals change their crystal shape in three dimensions as a function of size during crystallisation. As part of this work, Mangler et al. [2022] built on and corrected previous methods used to determine most likely 3D crystal morphology from 2D slice observations [Morgan and Jerram 2006]. This method makes use of forward models of ideal cuboidal crystal slicing.

All large-scale textural studies are faced with a common problem: What is the best approach to generate statistically acceptable quantities of data with the least observational effort (e.g. time spent manually tracing digital photomicrographs). So far almost all work has been based either on thresholding RGB or grayscale values for images acquired using optical or electron microscopes, to provide approximate crystal masks, or more often on manual crystal segmentation which is a labour-intensive and time-consuming task. Electron microscopes are a rich source of semantically strong datasets using EDS (energy-dispersive X-ray spectroscopy) or WDS (wavelength-dispersive X-ray spectroscopy) detectors to capture X-rays that can be directly related to mineral chemistry. Image segmentation methods have long been successfully utilised for these datasets using a wide variety of techniques [Ortolano et al. 2014; Lormand et al. 2018; Sheldrake and Higgins 2021; Tung et al. 2023; Visalli et al. 2023]. However optical microscopy is significantly less expensive and less labour-intensive in terms of sample preparation. It is therefore vital that similar methods are developed for optical image segmentation as well.

*✉ nt398@cam.ac.uk

Superpixel segmentation methods for human-computer collaborative tasks are readily employed for optical image segmentation [Yu et al. 2023], but they still require extensive human effort for large-scale studies as they heavily rely on stacked images with different light polarisation. To overcome this significant effort, a number of approaches have been proposed based on machine learning. One of the first examples of such work using optical microscope images are a set of classification methods based on convolutional neural networks (CNN) to predict the rock type shown in images [Alf rez et al. 2021; Patro et al. 2022]. Such a task is merely a demonstration of past capabilities, but it is not of obvious quantitative use in petrological or petrographic analysis. One promising optical method was developed for sandstone grain segmentation using the popular U-Net semantic segmentation model [Das et al. 2022], but these models can only predict phase labels and individual grains are not segmented. The most effective methods for phase segmentation in rock sections have been based on electron microscopy rather than optical microscopy, where EDS-derived compositional maps were used to classify pixels as mineral phases either using machine learning classification [Leichter et al. 2022] or through a database matching approach [Neave et al. 2017; Bell et al. 2020; Schulz et al. 2020]. Even this sort of scanning electron microscopy approach cannot automatically separate touching grains because phase classification based on pixel chemistry can also only provide semantic labels.

Here, we present a new automated method that capitalises upon modern advances in Deep Learning (DL) computer vision technologies. Specifically, region-based instance segmentation models [Hafiz and Bhat 2020] are employed as they are able to generate both semantic (object class) and localisation labels thereby not just segmenting plagioclase from everything else, but also providing individual crystal segmentation. These results may then be used for appropriate quantitative textural analyses. Instance segmentation can therefore overcome the limitations of semantic-only segmentation techniques and are ripe for use in geological applications. We refer to the approach as MinDet and provide links to the software in the [Data Availability](#) section. This first version of the software, MinDet1, is exclusively developed for plagioclase detection. The MinDet1 models presented here are trained on datasets of manually segmented plagioclase crystals and the ability of these trained models to rapidly generate accurate large-scale data is demonstrated in the following sections. MinDet provides a ten-fold speedup in segmentation time using CPUs for processing (as opposed to GPUs which would be even faster), whilst maintaining comparable accuracy to manual segmentation methods. Rapid insights into the crystallisation timescales of different plagioclase populations are for the first time possible. This is demonstrated using over 38,000 segmented crystals from an entire sample suite of volcanic igneous rocks from Iceland.

2 METHODS

2.1 Image acquisition

All images used for training and inference for the example presented here were acquired using a ZEISS Axio Imager.A2

optical microscope. The images were acquired using both linear polarising filters crossed at 90 degrees and circular polarisers to remove the crystal orientation dependence of birefringence colour intensity [Higgins 2010a]. Each image was acquired with identical light sources to eliminate any effects on colour. Large thin-section scans were performed using the panorama imaging mode where each tile was acquired using identical imaging conditions. Images were captured using one of two different objective lenses as deemed suitable for the crystal size within the sample; spatial resolution for these images were either 1.38 or 2.76 $\mu\text{m}/\text{pixel}$. The proprietary ZEISS file formats were exported as .jpg images with as little loss of quality from compression as possible. These .jpg files can be used directly by the deep learning models.

2.2 Instance segmentation models

Instance segmentation is a subset of image segmentation methods that aims to detect each instance of an object as well as its boundaries within an image. Popular implementations involve multi-stage deep convolutional neural network (CNN) models trained on ground-truth labels. Convolutional neural networks are ideally suited for image processing as they can effectively combine both image channel (e.g. RGB) and spatial information. Furthermore, deep neural networks, as opposed to more conventional machine learning models, go a step further with the ability to learn their own features through their optimisation process [LeCun et al. 2015]. This diminishes the need for precise feature engineering, where the user attempts to create new features manually, as the model will automatically learn the best set of features for the given task as long as they are trained well. Therefore all images in the present work are input to the model in RGB colour space as acquired directly from the microscope camera with no other features identified manually. Any other colour space can be reproduced by the model's feature extraction stage if it is deemed suitable during model training. A key disadvantage of this approach is that neural networks inherently produce black box models and it is not possible to justify all of the model's actions. Explainable AI is a major area of ongoing research within the deep learning community.

A schematic view of the operation of such models for the segmentation of plagioclase feldspar is shown in [Figure 1](#). Input images are first processed by deep convolutional layers to generate a stack of feature maps. These feature maps, as explained above, are a result of complex mathematical functions applied to all image channels with additional spatial information from the convolutional operation over a range of lengthscales. In the models used in the present manuscript this is done by pooling spatial information together from up to 32^2 pixel^2 area which allows the model to learn from more regional contexts, but it often leads to loss of information at the smallest lengthscales. The effect of this pooling operation will be the subject of discussion in the performance evaluation section. Based on these features, composed of RGB and spatial information, the model proposes a large number of regions of interest. The calculated regions are then evaluated based on a trained scoring strategy for the given application, and any region below a threshold score is discarded. For the

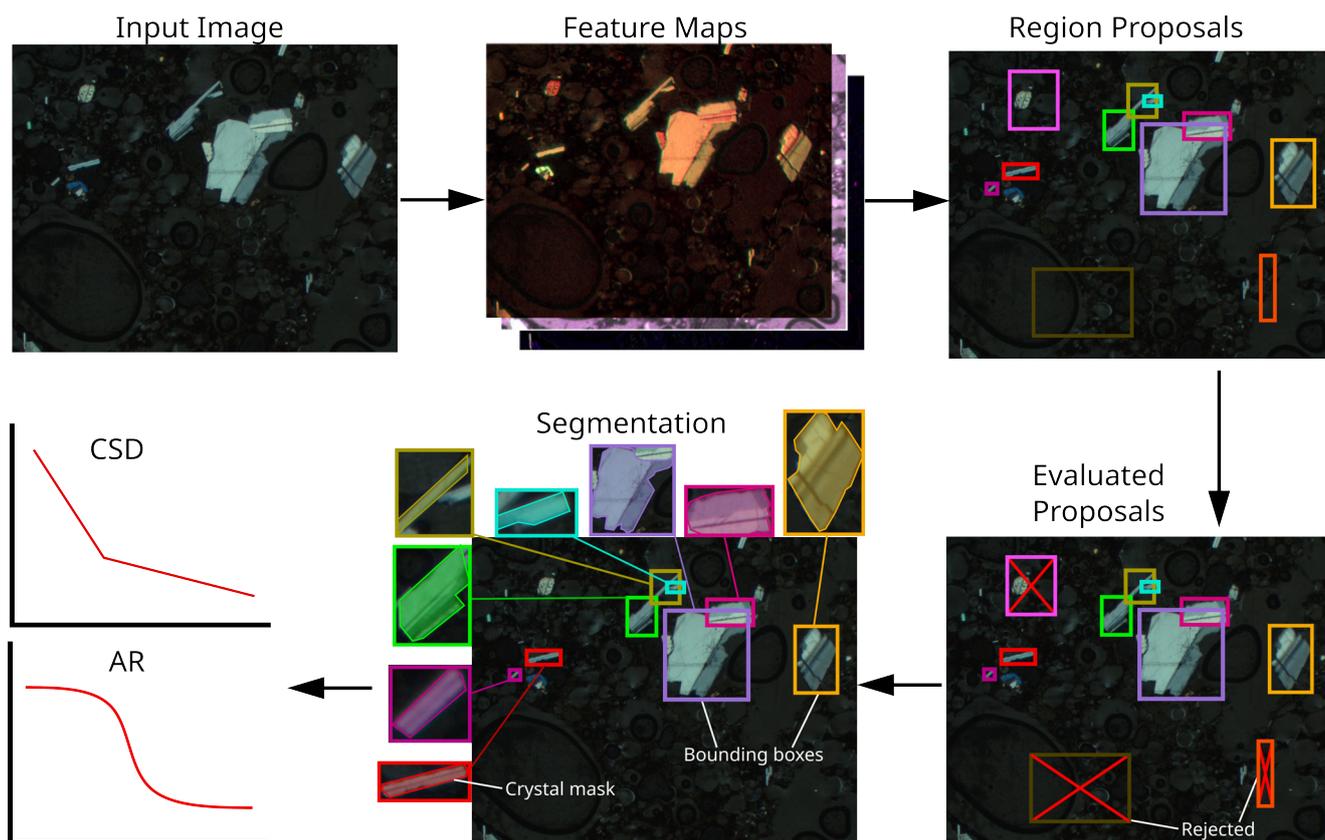


Figure 1: Overview of the proposed region-based deep learning method for plagioclase feldspar detection. Deep learning models necessarily transform the input into large stacks of feature maps that are used for further processing; next the model proposes regions of interests that are evaluated on whether there are any plagioclase crystals present. Regions that survive the evaluation process are taken individually and each crystal present is segmented out as binary phase masks which can be turned into textural information such as crystal size (CSD) and aspect ratio (AR) distributions. Specific terms used to indicate the type of segmentation results such as *bounding box* or *mask* is indicated as appropriate for clarity.

particular method described in this paper, regions with low scores would be those not containing plagioclase. Surviving regions of interest are passed to the segmenter network which generates the final segmentation results. These results are also scored from 0 to 1 to quantify the model's confidence in the resulting segmentation. From here on these will be referred to as *detection scores*. There is a temptation to use these as probability values but this should be avoided because this model, in contrast to Bayesian Neural Networks, is not conditioned over a distribution of parameters. The score simply reflects the confidence of the model in its results as a score between 0 and 1.

The two different segmentation results produced by instance segmentation models are also highlighted in Figure 1—bounding boxes and crystal masks. Bounding boxes show the best-fit rectangular region for an instance of an object. Bounding box approaches are used universally in object detection tasks where exact pixel-by-pixel location and shape of objects are not important, but their number and approximate location within an image is of interest. For the present application, the image masks generated are more useful and are identical in form to segmentation generated through most manual methods. Masks show the segmented boundaries of

objects, thereby pointing to the exact pixel-by-pixel locations. Crystal shapes and sizes may then be rapidly calculated using well-established image processing techniques [Schneider et al. 2012].

There are a large number of different instance segmentation models and architectures developed and available for custom training. In the present work region-based multi-stage segmentation models are considered only, more specifically the three models trained are Mask R-CNN [He et al. 2017], HTC [Chen et al. 2019a] and DetectorRS [Qiao et al. 2020]. These were chosen as Mask R-CNN is one of the most popular and well-tested models that performed exceptionally in all image segmentation competitions even though it is a relatively simple model. HTC and DetectorRS are later modifications to Mask R-CNN, hence they are expected to perform better. HTC specifically makes use of a significantly more complex segmentation stage to improve on the final results. DetectorRS goes one step further by also altering the model backbone—the stage that generates the stack of feature maps. Their philosophy relies on the idea that improved feature production should improve the results of any subsequent stage of the model. Whilst the modified models were shown to perform better on standard datasets such as ImageNet [Deng et al. 2009] or COCO [Lin et

al. 2014], it is not guaranteed they will outperform Mask R-CNN on the rock samples used for training here.

2.3 Inference Software Implementation

We designed the bulk of the software written for this paper to provide a user-friendly interface between the deep learning model and its geological application. The main aim is to make the method accessible for researchers with moderate Python experience as well as enabling its incorporation into pre-existing subroutines or methodologies. This approach will allow other researchers to custom-build on top of the basic textural analysis performed in this paper. Popular and reliable Python tools for image processing from the `scikit-image` [Van der Walt et al. 2014] library are heavily used to analyse the results and build directly onto the inference routines in the `MMDetection` application programming interface. High-throughput use of the present method is enabled by custom tiling and stitching routines of entire thin-section scans as well as a custom post-detection processing algorithm that was developed for more realistic detection results.

Panorama images of entire thin-sections imaged at close to micrometer resolution have image sizes on the order of 20,000×20,000 pixels. Segmentation of these large images with the deep learning models, known as inference, would require exceedingly large amounts of memory, on the order of 100 GB or more. Therefore, a custom tiling strategy is employed during inference; a simple schematic is shown in Figure 2A. Each panorama image is sliced into separate overlapping tiles according to the user's preferred size, for example 1,000×1,000 pixels with 100 pixel overlap on each side. The model performs inference on each tile image in series and the results are stitched together resulting in the same panorama image as before, but with each plagioclase crystal segmented.

We employ a very simple stitching strategy, with the core aim being to avoid over-segmenting crystals that lie across multiple tiles. Each tile is taken in turn to be stitched together row-wise until we have a complete set of all rows stitched together. Entire rows are then stitched together successively to form the final image. During the stitching process itself we simply take the overlapping regions and re-label the overlapping crystal masks such that they are consistent across the tile boundary. This eliminates any over-segmentation, which we see as a more significant issue than under-segmentation.

Object detection pipelines make use of a standard algorithm in their post-detection steps to remove overlapping segmentations. This step is often done with a method known as non-maximum suppression (NMS) [Neubeck and Van Gool 2006]; an example is shown in Figure 2B. The standard NMS algorithm looks at overlap of bounding boxes of detected objects by means of calculating Intersection over Union (IoU) scores and removes bounding boxes, except the one with the highest detection score, above a critical IoU threshold. It is by nature an iterative process, with overlap between every detection box necessarily calculated. A development of NMS was proposed [Bodla et al. 2017] which does not remove objects straight away, but instead decays the detection scores as a function of bounding box overlap. This approach is called soft-NMS [Bodla et al. 2017]. It includes only a very subtle

change to standard NMS and was found to increase overall segmentation performance for bounding box detection scores significantly.

For crystal segmentation, however, the main interests lie in the crystal masks predicted by the model and not the bounding boxes. Therefore overlapping bounding boxes are not necessarily undesirable in this case and instead one performs NMS over instance mask IoU. This is equivalent to requiring that two crystals cannot occupy the same pixel. The NMS algorithm developed for the `SOLov2` model [Wang et al. 2020] goes some way towards achieving this goal by considering predicted masks only and decaying each mask's scores as according to their overlap with other masks. In the present approach, instead of decaying detection scores the algorithm computes mask overlap, as shown in Figure 2B, and any detection with overlap above a critical threshold is suppressed regardless of prior detection score. This method is a more realistic approach as simple score decay can still give rise to significant overlaps without mask suppression due to high detection scores. Any remaining overlap that falls under the threshold for suppression is assigned to the mask with the highest detection score as it shows a higher level of confidence by the model.

Using a trained model, explained in the subsequent sections, and the software outlined above, an example segmentation image is shown in Figure 3A with the crystals outlined using the segmented crystal masks. Each crystal from the image can be extracted using these crystal masks, as well as for measuring sizes, shapes, and other properties, to allow for inspection of the segmented crystals. This is shown for the 100 largest crystals from the segmented image lined up in decreasing size order with their long axes aligned vertically to allow for further interrogation of them if necessary; each crystal was scaled to fit into each grid. This provides a visual example of model performance that is not necessarily apparent from numerical metrics alone. Visually, the results show good crystal localisation with limited errors. Furthermore, the figure shows that such data-driven approach can go well beyond just providing crystal shape and size information due to the wealth of other available data, such as the original image itself. These images of individual crystals could potentially be lines of future development, aided by the rapid segmentation provided by the deep learning method, although these will not be discussed in the present publication.

2.4 Evaluation Metric

As previously outlined, instance segmentation aims to solve a complex set of tasks that involve object localisation, classification, and segmentation at the same time; a combination of classical object detection and semantic segmentation. Owing to the complexity of the task at hand, the models and their evaluation are necessary also complex. Classical metrics used with machine learning models for classification tasks are based on the number of true positives (TP), false positives (FP)

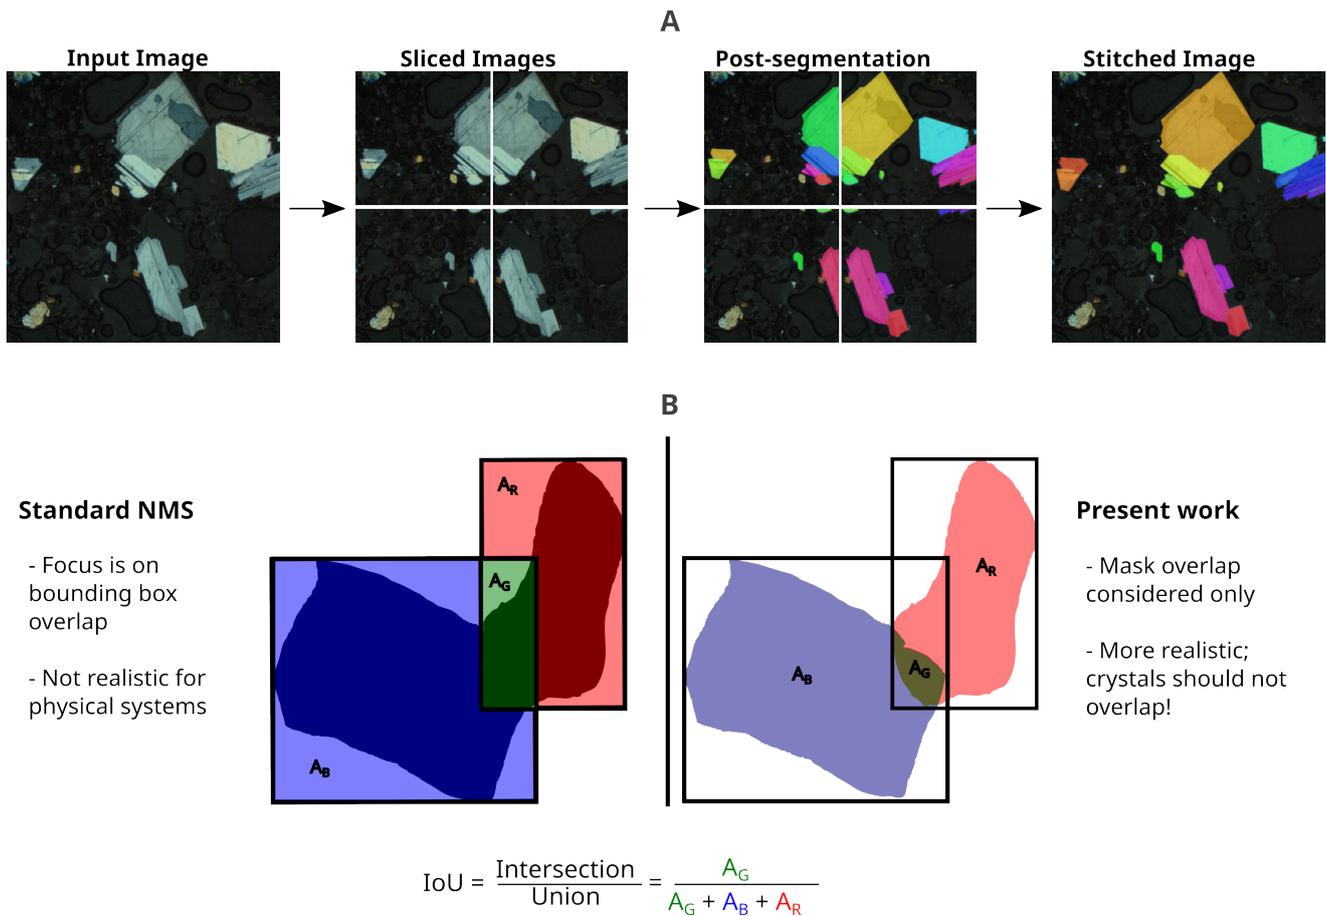


Figure 2: Overview of the approach developed for detection on large area section scans and the post-processing step for removing overlapping segmentations. [A] The input image is initially sliced into the desired tiles with overlaps between them; the individual tiles are used for the detection to reduce the computational power required, which are stitched together to produce the final segmented image. [B] Schematic overview of the NMS (Non-Maximum Suppression) approach used in the present work compared to the standard version most often used for object detection. Intersection-over-union (IoU) scores are introduced as the primary way to calculate overlaps in images. Standard NMS algorithm uses bounding boxes to check for object overlap and suppresses highly overlapped segmentations according to that. In the present case, bounding boxes are not a suitable metric, therefore mask-based NMS is employed where IoU of masks are calculated to support suppression of objects. Overlapping segmentations are in general assigned to the object with the highest detection score; this one is removed from any other object. The subscripts *G*, *B* and *R* refer to the green, blue and red colours of the areas shown.

and false negatives (FN) as shown:

$$\begin{aligned} P &= \frac{TP}{TP + FP} \\ R &= \frac{TP}{TP + FN} \\ F &= \frac{2 \times P \times R}{P + R} \end{aligned} \quad (1)$$

where *P* is precision, *R* is recall and *F* is the F1 score—a combination of precision and recall into a single metric. Often model performance is judged by F1 score alone as an easy to use classification measure indicating the best trade off between model precision and recall. The closer F1 score is to 1 the better the performance is judged to be, with anything below 0.7 often cited as a bad performance.

For object detection tasks the gold standard evaluation metric is the COCO mean average precision score (mAP) introduced alongside the COCO object detection data set [Lin et al.

2014]. Although the naming is misleading, it is not an average measure of precision alone but it is a metric used to describe the entire precision-recall curve as a single value between 0 and 1: the closer it is to 1, the better the model performance. Furthermore, complications arise when one starts to consider what should count as TP in instance segmentation, such as how much overlap between predicted and real object is required. This can be measured by IoU of either bounding boxes or segmentation masks with ground-truth labels. For COCO this is dealt with by calculating precision-recall curves at IoU threshold scores ranging from 0.5 to 0.95 and averaging them to calculate the AP score. When instances of multiple classes are present and segmented, the AP scores of each class are averaged to get the final mAP score; in the present case only a single class is dealt with so AP and mAP scores are identical. Specifically for instance segmentation, mAP scores can be calculated for both bounding box and mask segmentation

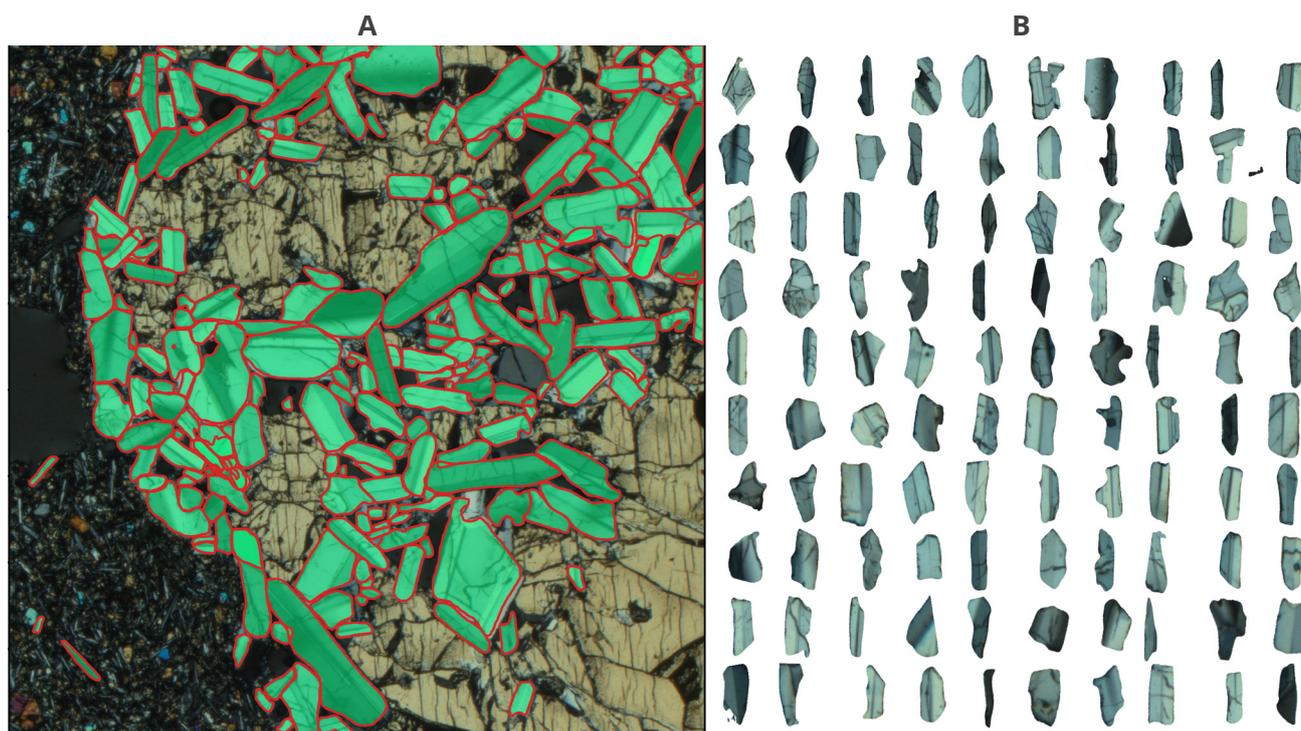


Figure 3: Example segmentation on a test image of touching ophitic plagioclase crystals. [A] Segmented image showing the segmented crystal masks, and [B] The 100 largest crystals lined up in descending order for further inspection. All crystals were scaled to fit into each grid position and do not show true relative sizes.

performance in the identical way to evaluate the two different aspects of object segmentation. A mAP score of 1 shows perfect object localisation with perfect precision and recall, but further quantitative interpretation of the mAP metric is not as straightforward as for the aforementioned classical machine learning metrics. This is largely down to the large number of parameters that need to be averaged over and hence; they have largely been used as relative performance measures only.

A schematic example of calculating mAP score is shown in Figure 4 for use with segmentation masks for a single class of objects. The table in Figure 4B shows the seven predicted masks ranked in order of decreasing model confidence—or prediction score—alongside their respective IoU value with ground-truth labels. Precision and recall values are calculated sequentially in order of decreasing prediction score for given IoU thresholds, taking all preceding detection masks into account at all times. Masks with IoU greater than the threshold are deemed “true positive” and all others “false positive”. Recall is defined as the proportion of all positive segmentations ranked above the given prediction score and precision as the proportion of all segmentations above that rank that are true positives [Everingham et al. 2009]. These are plotted to form the precision-recall curves in Figure 4C; the numbers denote the precision and recall values at a given rank. AP values for a given IoU threshold are calculated by averaging over all precision and recall values; it’s a measure of area under the curve. When the IoU threshold is set very low, such as 0.05, the model shows excellent performance as almost all detections are classified “true positive”, i.e. $AP = 1$. As the IoU threshold is decreased, the AP value also decreases as low IoU masks

are now classified as false positives, leading to lower precision at high recall values. The calculated AP values form the AP-IoU curve shown in Figure 4D. COCO mAP is determined by averaging over AP scores above IoU threshold of 0.5. This method necessarily misses out the left-hand side of the curve where segmentation is made easy by low thresholds. Therefore COCO mAP highly rewards models that are very precise at object localisation and is therefore a powerful metric.

Finally, we would like to stress that as important as quantified metrics are, such metrics may not provide the clarity of model performance as actual visualisation of results, such as Figure 3.

2.5 Training

The term *training* is used to define the optimisation procedure of models in machine learning. The training of neural networks is primarily done using iterative methods called Stochastic Gradient Descent (SGD) and backpropagation on a defined objective function, termed *loss function*, to find the best fit parameters for a given task. In the case of instance segmentation a multi-task loss function is defined as $L = L_{cls} + L_{box} + L_{mask}$ to direct model performance for each task of classification (L_{cls}), bounding box detection (L_{box}), and mask segmentation (L_{mask}). Note in the present case that because there is only one class segmented, plagioclase, L_{cls} is not an important term. To configure the SGD process one must choose a configuration of the learning mechanism external to the machine learning models, universally referred to as *hyperparameters*. The most important of them is the *learning rate* which defines the effective step size during the parameter

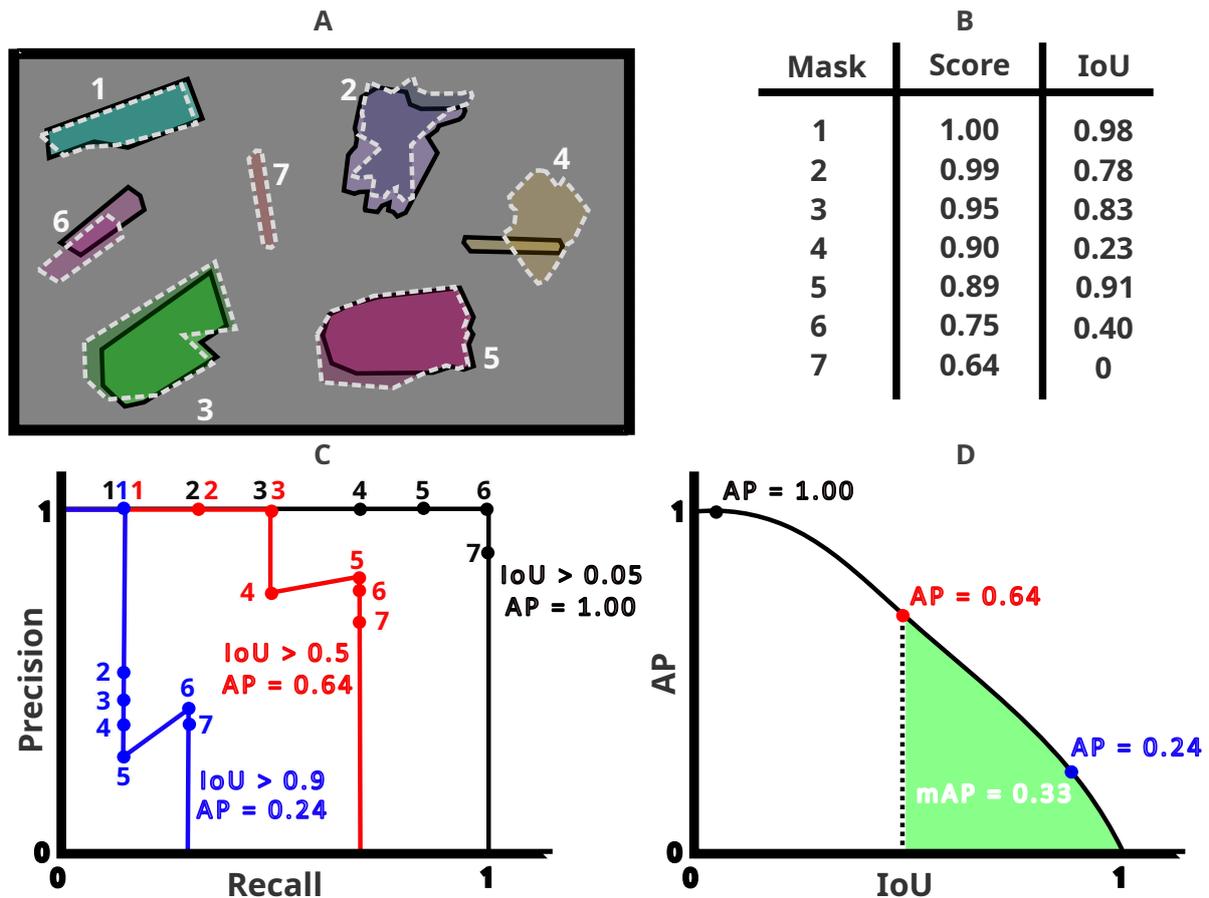


Figure 4: Schematic object detection example using segmentation masks to localise objects and calculate AP scores based on IoU threshold. [A] Schematic scene with objects (black bold outlines) and detections (white dotted outlines); [B] table of detections ordered according to detection score with IoU calculated between detected and “real” masks. [C] Precision-recall curves for selected IoU thresholds showing the plotted points for each mask; [D] Interpolated AP-IoU threshold curve with mAP score calculated as shown. Note how the AP score decreases significantly as the IoU threshold is made tougher.

update step of SGD. Larger learning rates tend to converge to minima in fewer steps, but if too large then SGD can overshoot the minimum or lead to divergent behaviour and fail to converge; if learning rates are too small, SGD will take significantly longer to converge. Therefore it is vital to find the best values for learning rate during training by setting up identical training runs and varying the learning rate or some other hyperparameter to find the best performing model and training regime configuration.

The training dataset was compiled from over 20 igneous rock thin-sections of basaltic compositions of varying textures from porphyritic volcanic rocks to gabbros with varying mineral assemblages and plagioclase abundance; see examples in Figure 5A. Some tolerance was given to varying acquisition parameters to ensure no crystals were over-saturated in any section. Plagioclase crystals were segmented individually by tracing the edges of crystals using the open-source Python software Labelme [Wada 2016]. The final set of training images consisted of over 8000 segmented crystals. The resulting annotations, stored in separate JSON files for each image, were converted to the popular COCO [Lin et al. 2014] format using the software labelme2coco [Xu 2019] with a training-test split of 80 %. Examples of manually segmented plagioclase

crystals are presented in Figure 5B. No validation data is used in the present manuscript as the intention is not to introduce a benchmarking dataset or mAP score as these are not necessarily absolute performance metrics and no model development was performed.

Model training is performed through the MMDetection open-source object detection toolbox [Chen et al. 2019b], which provides flexible and easy-to-use implementations of the most popular object detection, semantic segmentation, instance segmentation, and many more types of models. All three instance segmentation models considered in the present work are implemented and readily available for training on a custom dataset. Training was performed using SGD [Kiefer and Wolfowitz 1952] with variable learning rates and fixed default values for momentum and gradient decay. Data augmentation is readily performed by MMDetection and throughout the training procedure. Random resize, random flipping, and random cropping were performed during all training procedures to effectively enlarge the size of the training dataset and help avoid overfitting models.

To optimise the training of the final model, a methodical experiment was set up across all three models that were trained using three different learning rates: 0.1, 0.01, and 0.001, each

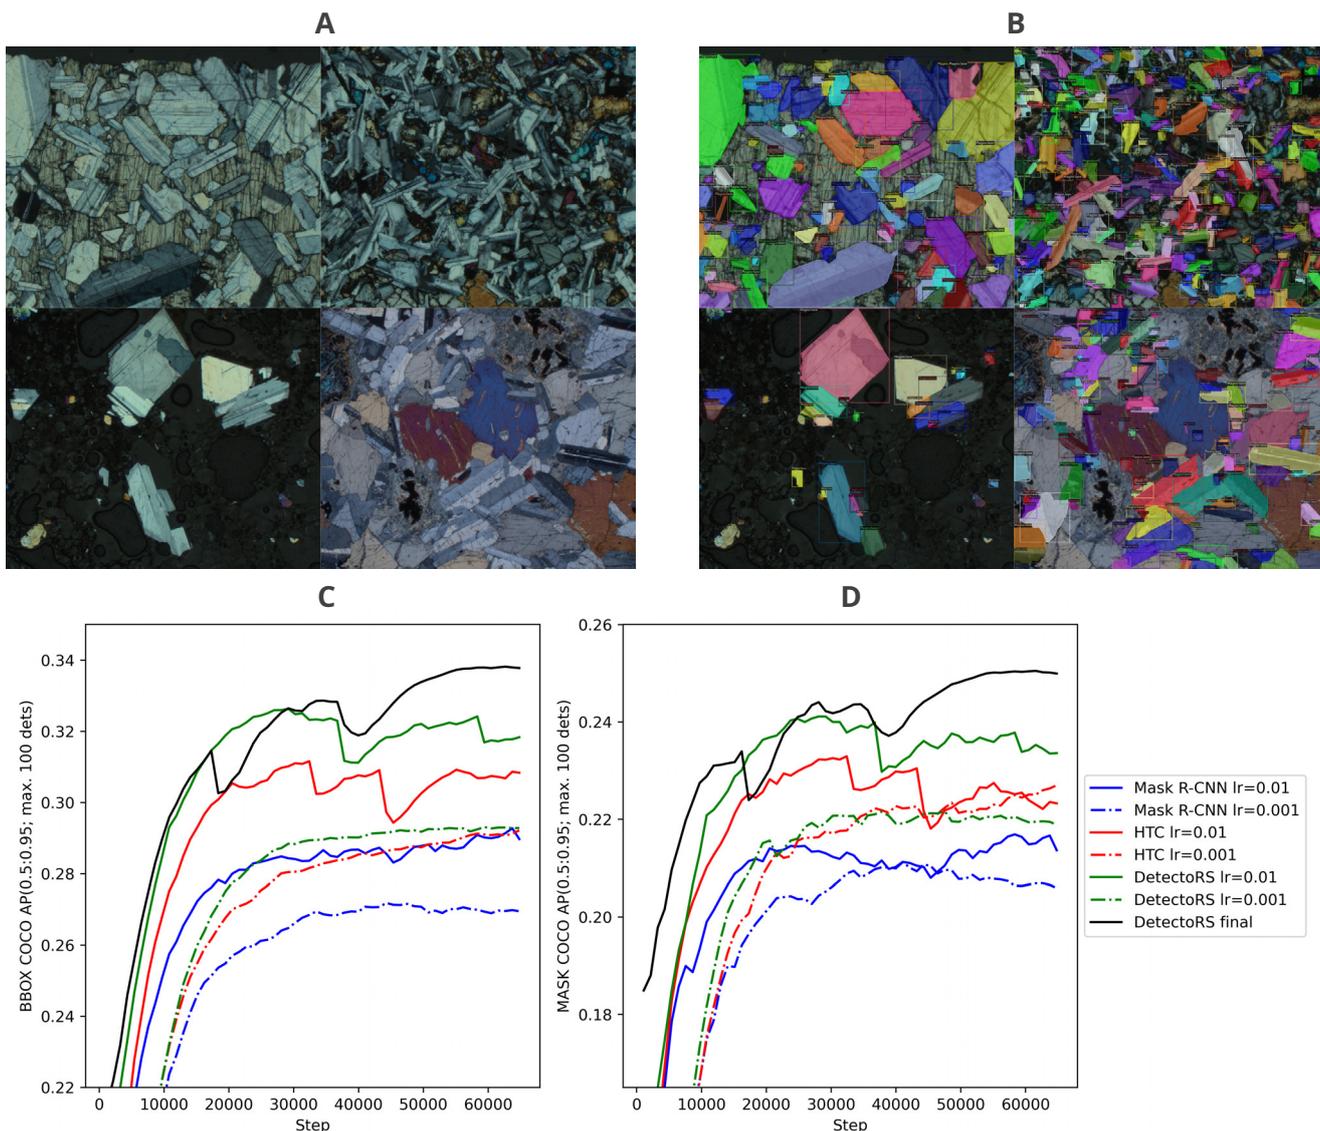


Figure 5: Training dataset and overview of trained models' performance. [A] Example micrographs from the training dataset, [B] manually segmented examples of the ground-truth labels used in the training data, and [C]–[D] testing scores of the three different region-based instance segmentation models trained with various learning rates evaluated for bounding box and crystal mask accuracy respectively. The implemented model was trained using stepped learning rates based on this result and chosen to be DetectoRS model due to its high performance in tests—this model, once trained, clearly outperforms all others previously trained. Note mAP scores (mean average precision) are a popular measure of performance in object detection. Here these are maximised for best performance. Smoothing was applied using an exponential moving average method.

trained for 600 epochs. Testing mAP scores for each model and learning rate during training are shown in Figure 5C; there are no results for models with learning rates of 0.1 as their SGD never converged and were rendered unusable. Note that MMDetect only calculates COCO mAP using the first 100 detections, so for images with over 100 objects present this caps the maximum recall for those specific samples. Whilst it does decrease the AP scores for those samples, this is irrespective of the model used and the mAP scores can still be used qualitatively to compare performance. No NMS and prediction score threshold was used during testing.

There are two clear trends from the results as both model type and learning rate are seen to make a difference in model

performance. First of all, with the same learning rate, Mask R-CNN is outperformed by both other models and HTC is always outperformed by DetectoRS; this is no surprise as one is a development of the other. Secondly, learning rate, *lr*, is seen to play a crucial role where models trained with *lr* = 0.01 always outperform those with *lr* = 0.001 once testing losses have stabilised.

DetectoRS was therefore chosen as the model used in all subsequent segmentation work with a variable learning rate over its 600 epochs of training time, starting at 0.01 for the first 400 epochs which is then reduced by a factor of 10 to 0.001 for the next 100 epochs whereupon it is once again stepped down to 0.0001 for the last 100 epochs. All other SGD parameters

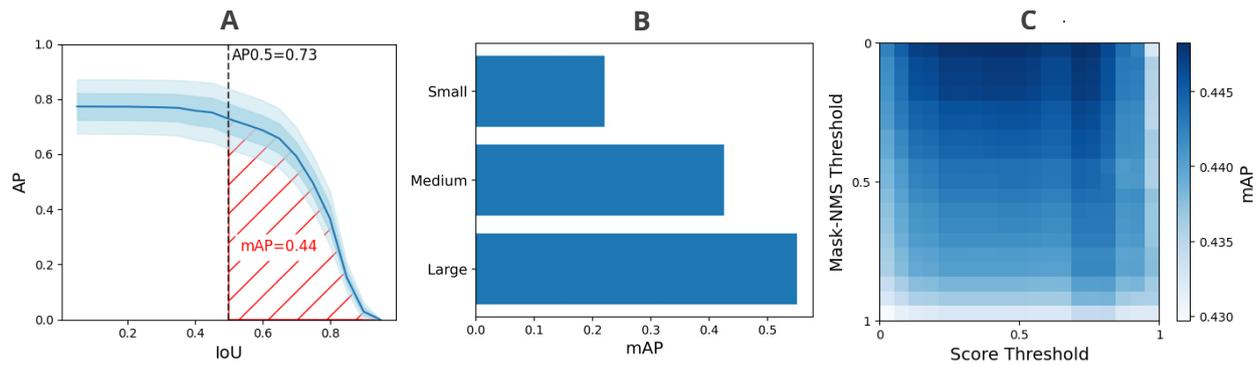


Figure 6: Variation of AP values on the testing dataset with respect to IoU threshold, crystal size, NMS and detection score threshold. AP confidence intervals are showing the variability between different images. IoU thresholds are averaged over between 0.5 and 0.95 when it is not a free variable [B]–[C] and likewise for NMS and detection score thresholds between 0.05 and 0.95 [A]–[B].

are kept fixed. This is done to enable SGD to find more optimum minima closer to (or at) the global minimum as lower learning rates reduce the chance of skipping over such minima. This training strategy had clearly worked as shown on Figure 5C, the final DetectoRS model outperforms all previously trained ones and is therefore the best one to implement for all subsequent work.

2.6 Model Performance

Before applying the final trained model to acquire data from real sections, it is imperative that we understand the limitations and behaviour of the model for a range of parameters a user can change during inference. Therefore a set of experiments were set up on the testing dataset to evaluate how AP scores are affected by IoU threshold, grain size, detection score threshold, and mask-NMS threshold. In these experiments, AP was calculated using the COCO conventions, but using all detections and not only the top 100 to better reflect its use on real samples and remove the effect of capping the maximum recall value. All AP and IoU values are calculated for segmentation masks only as these are the most important part for petrographic purposes (such as calculating crystal shapes and sizes).

The results are shown in Figure 6 with [A] showing the variation of AP with respect to the IoU threshold for positive detection; in this case all AP values are averaged over all scores and NMS thresholds between 0.05 and 0.95. Below 0.5 IoU, the model shows very good performance with AP close to 0.8: not far off perfect precision and recall. There is a significant drop off in performance at 0.75 IoU. This shows that segmented crystal masks are not perfect, but a limit of 0.75 IoU is generally considered a good performance as it is used as a strict performance metric. One should bear this in mind when applying the model to tasks where exact mask shapes are crucial, such as determining crystal habits; further training on more labelled data will help improve this aspect. Figure 6B shows the variation of mAP, averaged over IoU greater than 0.5, mask-NMS, and detection score thresholds, over different object sizes. A standard definition of small, medium and large object sizes from the COCO metric are used: denoting

small objects as less than 32^2 pixel area, medium as between 32^2 and 96^2 pixel area, and large as anything over 96^2 pixel area. Clearly, segmentation results are best for the largest set of crystals with medium not far behind, but model performance drops off significantly for small crystals. This is due to the deepest, semantically very strong, layer of the network using a 32-fold downsampling of the input image, therefore anything smaller than 32^2 pixel area may go missing. It is advised for users to make sure crystals of interest appear larger in their images or they use this threshold in subsequent analyses as an effective detection limit.

Finally, Figure 6C shows the variation of AP, averaged over IoU greater than 0.5 as in COCO, with respect to NMS and detection score thresholds. There is very little variation observed between AP of 0.43 and 0.45 depending on the exact thresholds used. One must note that best performance for this set of samples was observed for low NMS thresholds (<0.5). Score threshold on the whole made little difference unless set too low (<0.1) decreasing model precision or too high (>0.9) decreasing recall significantly. These values depend exactly on the dataset used and users are encouraged to tune these parameters through visual inspection for best performance.

It is important to note that the version of DetectoRS used in this paper may not be the best optimised for the segmentation task performed here. Various parameters, especially training dataset size, may influence the final state and therefore the accuracy of the model which would require further work. Overall, the collection and segmentation of the training set required three weeks of manual work, whilst each training run lasted 6–8 hours, cumulatively representing three days of computational time.

3 EXAMPLE USAGE

With the MinDet model trained on ground-truth labels for plagioclase feldspar crystals in basaltic composition igneous rocks, an example use of the MinDet method is showcased to analyse natural volcanic samples. Segmentation results were used to produce CSDs and aspect ratio distributions to interrogate the samples' textural history. To assess the accuracy of the present approach, the plots generated can be compared to

published manually collected results. The following section is merely a guide and by no means an exhaustive description of the MinDet method's capabilities.

3.1 The Skuggafjöll Eruption

Skuggafjöll is a NE–SW-striking hyaloclastite ridge between the Mjrdalsjökull and Vatnajökull glaciers in southern Iceland [Jakobsson and Gudmundsson 2008]. Its lower slopes are composed of pillow lavas occasionally intercalated with hyaloclastite. The first complete textural and chemical description of hand-specimens from the edifice was carried out by Neave et al. [2014] using standard manual point counting and segmentation methods for the textural studies as well as a combination of X-ray fluorescence, inductively coupled plasma mass spectrometry, and electron microanalytical tools for their chemical analyses. It was shown that geochemical variability in whole-rock samples was best accounted for by accumulation of plagioclase and olivine through a mush disaggregation process prior to eruption. This model was supported by textural work on three samples showing two distinct crystal populations present. The first population is composed of large disaggregated macrocrysts of plagioclase and olivine which were part of a crystal mush. The second population is composed of smaller clinopyroxene, plagioclase and olivine crystals that is thought to have grown rapidly in an event possibly associated with destabilisation of the mush and magma storage zones at depth prior to eruption.

The manually acquired segmentation data from the Neave et al. [2014] study provide an ideal validation dataset for the MinDet method. The thin-sections used in the previous study were mapped in panorama mode using circular polarised optical illumination and plagioclase crystals were segmented using the trained DetectoRS model on a CPU-only cluster computer. 1000×1000 pixel tiles with 250 px overlap were sliced out from the panorama images, they were individually segmented and re-stitched using the custom tiling subroutine. Any instances with a detection score below 0.8 or with mask overlap above 0.5 were removed using NMS. Dimensions and shapes of individual shape masks were calculated using a best-fit ellipse approach and extracted from the final segmentations to be used to determine textural properties directly.

Figure 7 shows the comparison between both approximate and corrected CSDs, and aspect ratio distributions from the previous work using the manual segmentation approach [Neave et al. 2014] and the new DL-based method presented here. Best fit lines are the approximate CSDs are plotted using two separate linear regressions for each population identified in each plot; this is an identical procedure to that applied in Neave et al. [2014]. This was purposefully done to allow for the most transparent comparison possible with the published dataset. The present work is not meant to imply which CSD approach is most appropriate. There is impressive similarity in the results of the new automated and old manual methods, despite the possibility of bias inherent to each technique. This is shown by the close similarities in lines of best fit for approximate CSDs, points of inflections and aspect ratio values.

The characteristic crystal length (L_D) of a population of crystals can be found from regressions through CSDs, where

L_D is defined as $-1/\alpha$ with α being the slope. [Higgins and Chandrasekharam 2007]. Prior work has established that the approximate CSDs for these Skuggafjöll samples show L_D , modelled as square root of crystal area, of 70–100 μm for the population of smaller macrocrysts and 700–1030 μm for the larger population [Neave et al. 2014]. The present automated segmentation approach with MinDet yields very similar values of 80–100 μm and 590–1090 μm for the small and large macrocryst populations respectively, as calculated from the slopes in Figure 7. Likewise the pronounced inflections in the approximate CSDs are calculated to occur at square root crystal areas of 555–590 μm , in agreement with the reported approximate value of 600 μm previously [Neave et al. 2014]. A change in crystal populations is also evident in the mean aspect ratio plots at the approximately identical crystal sizes, with the large macrocrysts exhibiting low aspect ratio (~ 2) with the smaller crystals showing a continuously increasing trend with decreasing crystal size. Note there are some deviations between the two sets of aspect ratio values, but they are within the error bars as indicated. Both the manual [Neave et al. 2014] and the present work used a best-fit ellipse method to calculate these crystal parameters that help to eliminate any effects other than segmentation method. If a different measurement method was applied the presence of systematic differences may be expected.

Stereologically corrected CSDs were calculated using the data from Neave et al. [2014] and the methods from Higgins [2000] as shown in the bottom row of Figure 7. Best fit crystal habits were determined using the methods indicated [Morgan and Jerram 2006; Mangler et al. 2022]. All CSDs show very close matches, for both gradients and points of inflection, regardless of the crystal habits determined. As all samples are composed of mixed populations all determined crystal habits showed poor fits, therefore subtle variations in the distribution of slice shapes can cause the discrepancies in determined crystal habits. As explained in a previous section, and in Figure 6, mAP scores are not 1 and therefore subtle differences in shape compared to manual segmentations are expected. These may be enough to create the discrepancies within the crystal habits determined. However, this is not a focus of the present study and requires further work. This discrepancy may be eliminated using a larger, more extensive training set for such studies; note that this does involve an increased volume of manual work.

The major advantage of the present automated approach lies in the time-saving in processing large datasets that are unreasonably time-consuming when performed manually. The above manual example from Neave et al. [2014] is reported to have taken the authors approximately 12 hours of segmentation time for a single phase per thin-section scan. Using an automated approach eliminates the human time for segmentation, with the DL method taking 1–2 hours of computational time per section for segmentation using CPU only; the availability of GPUs will reduce this to around 30 minutes or less per thin-section. This example serves as a useful validation for the accuracy and applicability of such automated methods, given the models are trained correctly.

Furthermore, as the most time-consuming part of the seg-

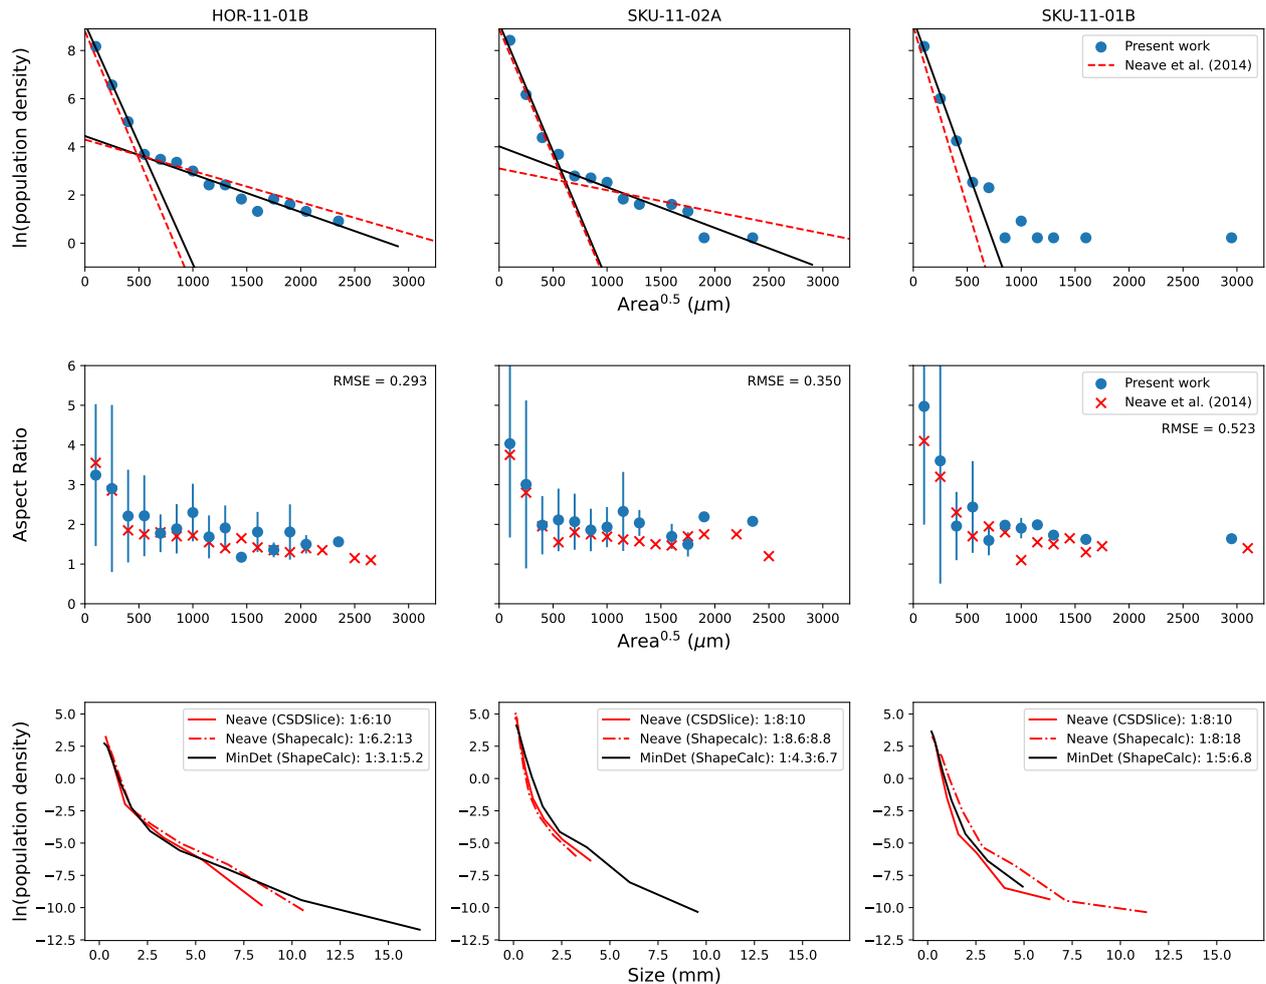


Figure 7: Comparison of `MinDet` automated results from this study with manual segmentation from Neave et al. [2014] for Skuggafjöll crystal size and aspect ratio distributions. [top] Comparison of `MinDet` output (blue symbols, black lines) with Neave et al. [2014] manual segmentation (red lines). The method for converting crystal shapes into approximate CSD was that described in Neave et al. [2014] in this case. A best-fit line is not provided for crystals >600 μm in SKU11-01B as most bins are underfilled, containing either one or no crystals at all. These panels are provided for data comparison purposes only. [middle] Comparison of calculated mean aspect ratios using the method of Neave et al. [2014]; RMSE values are shown for each sample between the present and Neave et al. [2014] datasets. Bins with no error bars contain only a single crystal measurement. [bottom] Comparison of `MinDet` output (black lines) with manual segmentation from Neave et al. [2014] (red lines) using stereologically corrected CSDs. The method for stereological corrections by Higgins [2000] was used in this case. Note that S/I values were determined to ±0.04–0.06 error with I/L showing ±0.21–0.25 error; as all samples are composed of mixed populations all determined crystal habits showed poor fits.

mentation process can be eliminated, it is now possible to perform rapid plagioclase crystal segmentation for the entire suite of 27 thin-sections of samples of pillow basalt from Skuggafjöll. These were imaged in panorama mode and processed almost identically to the previous examples, with a size threshold of 100 μm square root crystal area. Stereological corrections were applied, as according to Higgins [2000], using the best fit crystal shapes as determined using ShapeCalc [Mangler et al. 2022]. CSDs for all samples are shown in Figure 8(a) with Bayesian piece-wise linear regression calculated using a single breakpoint for kinked CSDs and Bayesian linear regression for the two samples showing straight line CSDs. This is

done using the Python package `emcee` [Foreman-Mackey et al. 2013] using Markov-chain Monte Carlo (MCMC) [Andrieu et al. 2003]. Uniform priors are set for all parameters; the bounds for these were set as [−20, 0] for α_1 and α_2 , [−20, 20] for $\ln(n^2)$ and [0.5, 5] for break points. Their joint posterior distributions are fitted from the 90,000 samples drawn by the MCMC procedure after discarding the initial burn-in period. This approach allows one to take into account the uncertainties in the CSD density values and report the fitted parameters with their uncertainties. Figure 8B shows the plots between each fitted parameter with one standard deviation shown where appropriate. α_1 and α_2 are the gradients of the CSDs as shown;

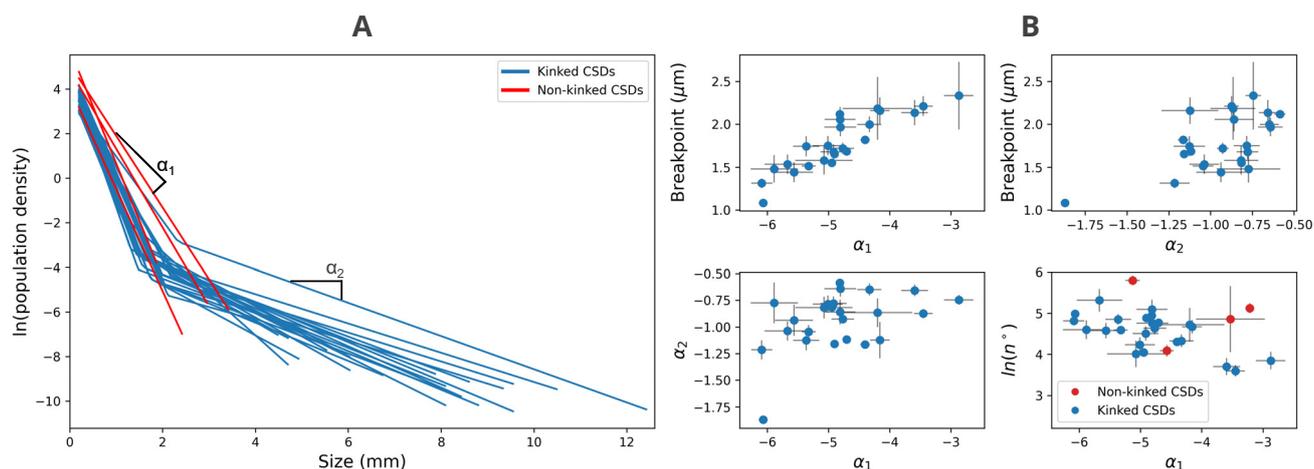


Figure 8: Stereologically corrected crystal size distribution of plagioclase for the entire Skuggafjöll suite of thin-sections, excluding hyaloclastites, as segmented using the DL method presented [A] Bayesian piece-wise regression modelling of all but four CSDs with a single breakpoint, the remaining two exhibit straight CSDs and bayesian linear regression was performed instead. These were performed using an MCMC approach; [B] Plots showing the relationships between the four parameters obtained from the regression: the breakpoint and the two different gradient values and $\ln(n^\circ)$, the natural logarithm of the nucleation density as shown. Quantification of the mean and 1 standard deviation spread of the fitted parameters was calculated fitted posterior joint distributions therefore taking into account the uncertainties in the stereological correction. The two samples showing highly uncertain α_2 and break point values were borderline between straight line and kinked CSDs, therefore they were processed with the piece-wise model.

the straight line CSDs do not have a breakpoint and their gradients are treated as α_1 , with α_2 naturally set as 0. α_1 values have a mean of -5.3 mm^{-1} , suggesting average L_D of $190 \text{ }\mu\text{m}$ length. α_2 's have a mean of -1.37 mm^{-1} meaning average L_D is approximately $730 \text{ }\mu\text{m}$ for the large macrocrysts. Breakpoints cluster around 1.5 mm length, but show significant variability with a standard deviation of $150 \text{ }\mu\text{m}$. A strong correlation may be observed between the α_1 and breakpoint results that are absent from all the other plots. Values of α_2 and n are constant within their respective uncertainties. This suggests that breakpoint values are controlled by α_1 , the slope of the small crystal population's distribution. Note the power of this method is displayed well in Figure 8 showing two samples where manually segmentation produced slightly kinked CSDs but MinDet results show poor fits for α_2 and break point values. This poor fit is shown by very high uncertainties in Figure 8B, though these samples do overlap with the other values. This shows they are described sufficiently well by both straight line and kinked CSD models without having to decide which model may fit better and therefore biasing the procedure.

The above results reveal that the vast majority of these samples share a history of textural evolution driven by two distinct stages of cooling and crystallisation as also observed in prior bulk rock chemical analyses, with the major changes relating to the proportion of macrocrysts present [Neave et al. 2014]. The macrocryst proportions are controlled by the amount of disaggregating mush entrained within each sample and their distribution is a constant across all samples. The major variable across the sample suite is the slope of the fine grained phenocryst population which is controlled by the timescale of growth of these crystals. This example from Skuggafjöll

demonstrates the power of large-scale crystal segmentation. In less than 48 hours of computational time, it was possible to segment large-area scans for 27 thin-sections, resulting in over 38,000 crystal segmentations over the $100 \text{ }\mu\text{m}$ size threshold. Based on the manual segmentation times mentioned previously, such a task would require around 300 hours of active manual segmentation time—unreasonably long for the vast majority of studies. Such holistic analyses may provide a much clearer picture of the magmatic history in this case as well as increasing the reliability of the conclusions drawn from them by reducing the uncertainty from only being able to analyse a subset of the available samples. It is hoped that automated segmentation will allow more petrologists to make use of such data-driven analyses and help advance related fields significantly more swiftly.

3.2 Crystal Shape and Crystallisation Time

As shown by Holness [2014], the mean aspect ratio, A , of plagioclase grains measured in dykes and sills can be calibrated with their crystallisation time, t , with a simple linear parameterisation: $\log_{10}(t) = mA + c$, where m and c are parameters to be fitted. A selection of samples that Holness [2014] used to build the calibration were imaged for this study in panorama mode using circular polarised optics as described for the previous set of samples discussed. Input images were sliced to $1000 \times 1000 \text{ px}$ with 250 px overlap. The NMS threshold was set at 0.5 IoU overlap and the score threshold at 0.8 as before. The resulting segmentations were used to calculate crystal sizes and aspect ratios.

Figure 9 shows a plot of published aspect ratio data [Holness 2014; Holness et al. 2017] compared to the values obtained in the present study. Bootstrap sampling was used to obtain the

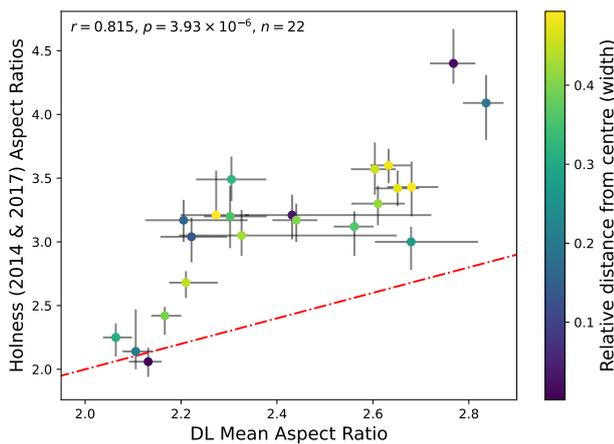


Figure 9: Comparison of aspect ratio values obtained from Holness [2014] and using the present approach. The red line shows the 1:1 line. The 95 % confidence intervals ($\pm 2\sigma$) were calculated using the identical bootstrapping method from the previous work. Pearson's correlation coefficient (r) shows a strong correlation does exist between the two datasets, although they clearly do not lie on the 1:1 line.

confidence intervals shown as performed by Holness [2014]. This is suitable as the distribution of aspect ratios in each section are only samples from the true population distribution. By re-sampling, as done using the bootstrap method, the variance of possible means can be quantified and therefore the confidence level in having obtained the true population mean is also quantified. The 1:1 line is plotted for illustrative purposes. A systematic difference in the values from automatic and manual characterisation is clear from the plot. This is in contrast to the previous section where automatic and manual segmentation values were in good agreement. The observed differences arise from the differences in aspect ratio measurements used by the two approaches. The manual approach of Holness [2014] involves picking and drawing on perpendicular sets of long and short axes for all visible grains in a section. This approach is significantly faster than manually segmenting crystal masks [e.g. Neave et al. 2014]. Nevertheless, the measured lengths and aspect ratios rely upon different sets of assumptions and biases on where to place such axes which may be ambiguous for irregularly shaped crystals. Such assumptions may be cause of the mismatch between the results of Holness [2014] and the DL automatic segmentation method. This is in contrast to the results of the previous section where the crystal-outline segmentation results of Neave et al. [2014] were in good agreement with our DL approach. Rather than invalidating the observations of Holness [2014], this mismatch highlights the importance of consistency, in terms measurement techniques employed, required for such analytical work [Higgins 2000]. The fact that covariation in the two sets of values is clearly systematic (Pearson's correlation coefficient: $r = 0.815$), and not due to random segmentation errors, one may be confident the DL-based method is also capturing real signals and real information from the thin-section scans. Similarly to the previous work, it is found that the one sample

taken from close to the margin of a sill exhibits anomalous grain shape of significantly lower average aspect ratio than expected from the calculated crystallisation times; this refers to the yellow data point with 2.3 aspect ratio and lower than 10^{-1} years crystallisation time. This is attributed to extensive through-flow of magma within the body of intrusion providing additional heat to the margins, melting back some of the original chilled margin [Holness 2022]. This change in the heat balance leads to longer cooling times in the areas of such melt-back leading to the observed anomalous grain shapes. These anomalous samples were removed from the regression analysis described below in order to mitigate the effects of physical processes not included in the original cooling calculations.

The results presented above were used to develop a calibration of crystallisation time as a function of aspect ratio where the aspect ratio is calculated on the basis of best-fit ellipses to whole crystal area segmentations. For robust fitting of crystal shape and crystallisation time, a Bayesian linear regression approach was used to find the mean crystallisation time as a function of aspect ratio, using MCMC once again, using the Python package emcee [Foreman-Mackey et al. 2013]. The model fitted is analogous to the calibration used by Holness [2014]: $\log_{10}(t) = mA + c$, where t is crystallisation time, A is aspect ratio and m and c are parameters to be fitted. Uniform priors were set as $[-20, 20]$ and $[0, 40]$ for m and c respectively, drawing 250,000 samples to construct the posterior joint distribution after the initial burn-in period is removed. The calibration is shown on Figure 10A plotted with the data with the posterior joint distribution for m and c plotted on Figure 10B. This method is robust to noise, takes account of the uncertainty in the observations and provides confidence intervals for the calibration. However, one must keep in mind that this is still only a simple physical model that only takes sliced crystal aspect ratios into account and does not consider the true three dimensional shape of crystals. One further caveat is that this model is only fitted to rocks from intrusive bodies that had undergone complete crystallisation and therefore it may not be applicable for crystals from other settings, such as those grown free floating in liquid.

4 CONCLUSIONS

The MinDet toolbox, using deep learning-based image segmentation, is presented for the textural analyses of petrological thin-sections. In this paper the method is focussed on plagioclase feldspar and called MinDet1. Instance segmentation models were trained on manually segmented images of plagioclase crystals in samples of basaltic composition. These images were acquired using circular polarising optics on a petrographic microscope. The best-performing segmentation model was used for the analysis of new samples to showcase the method's potential. This automated approach was shown to be able to replace days of manual segmentation work per sample with only 1–2 hours of computational time; such time saving can enable petrologists to analyse more samples and larger areas. In addition, relevant software was developed in conjunction to interface the acquired images with the segmentation models and to post-process the final crystal masks into a petrographically useful form. This includes the development

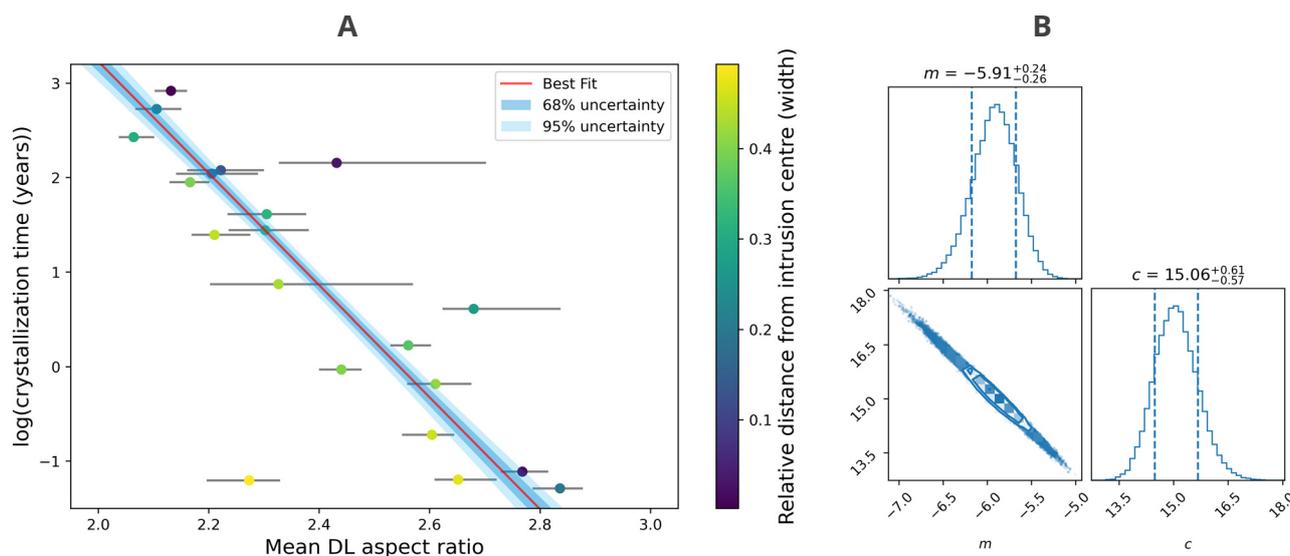


Figure 10: Calibration for expected time of crystallisation of plagioclase based on population shape distribution. [A] Plot of $\log_{10}(t)$ against crystal aspect ratio showing the best fit relationship and confidence intervals for the fitting based on the MCMC samples generated. [B] Posterior joint distribution of the two parameters of the model $\log_{10}(t) = mA + c$, where t is crystallisation time, A is aspect ratio and m and c are fitted parameters.

of a new mask-NMS algorithm to deal with overlapping segmentation proposals as well as the incorporation of popular Python libraries for image analysis and data processing.

The trained model was applied to two well-characterised sample sets from published material. First, thin-sections of sub-glacial pillow basalts, whose textures had been quantified through a slow manual approach [Neave et al. 2014], were used to validate the segmentation results of plagioclase crystals. The resulting CSDs and aspect ratio distributions for the three samples were an almost identical match between the manual and new automated method. Textural characterisation involving all other samples from this Skuggafjöll suite was presented in addition to a novel approach to detect variability using a piecewise segmentation method. The main power of such automated methods was presented making use of the large-scale study of Holness [2014] to create a new probabilistic calibration of sill and dyke crystallisation time and mean apparent crystal aspect ratio.

It is now possible to rapidly obtain plagioclase textural information for large sample suites with minimal user intervention required. This quantitative textural data can be used to constrain the timescale of igneous processes. The present method may now readily be used alongside well-established high-throughput timescale measurement approaches such as diffusion chronometry to better understand igneous processes and improve the quality of timescale estimates overall. More data-driven approaches for textural quantification must be developed to realise the full potential of this approach. The above work must also be extended to further mineral groups to allow for full-scale automated petrographic descriptions. Such methods will undoubtedly require vast amounts of training data, though it is technologically feasible as shown in the present work.

AUTHOR CONTRIBUTIONS

NT conceived of the presented idea, developed the software, produced the training dataset and performed the computations. JM provided the physical samples used for training and the Skuggafjöll sample set. All authors discussed the results and contributed to the final manuscript. JM supervised the project.

ACKNOWLEDGEMENTS

This research was funded by a NERC studentship awarded to N. Toth. The authors would like to thank Marian Holness for helpful advice on textural analyses and providing the samples used in Section 3.2. This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk). The authors would like to thank Martin Mangler, Zoja Vukmanovic, Michael Higgins and Thomas Sheldrake for their detailed and constructive comments, and Jamie Farquharson for his guidance and editorial handling.

DATA AVAILABILITY

The training dataset, model parameters and supporting data can be found at <https://doi.org/10.5281/zenodo.7930016>. The software code is freely available and installable at <https://github.com/norberttoth398/MinDet> (<https://doi.org/10.5281/zenodo.10061725>). Software documentation is available at <https://mindet.readthedocs.io/en/latest/?badge=latest>.

COPYRIGHT NOTICE

© The Author(s) 2024. This article is distributed under the terms of the [Creative Commons Attribution 4.0 International License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

REFERENCES

- Alfárez, G. H., E. L. Vázquez, A. M. Martínez Ardila, and B. L. Clausen (2021). “Automatic classification of plutonic rocks with deep learning”. *Applied Computing and Geosciences* 10, page 100061. DOI: <https://doi.org/10.1016/j.acags.2021.100061>.
- Andrieu, C., N. de Freitas, A. Doucet, and M. I. Jordan (2003). “An Introduction to MCMC for Machine Learning”. *Machine Learning* 50(1/2), pages 5–43. DOI: [10.1023/a:1020281327116](https://doi.org/10.1023/a:1020281327116).
- Armienti, P. (2008). “Decryption of Igneous Rock Textures: Crystal Size Distribution Tools”. *Reviews in Mineralogy and Geochemistry* 69(1), pages 623–649. DOI: [10.2138/rmg.2008.69.16](https://doi.org/10.2138/rmg.2008.69.16).
- Askaripour, M., A. Saeidi, P. Mercier-Langevin, and A. Rouleau (2022). “A Review of Relationship between Texture Characteristic and Mechanical Properties of Rock”. *Geotechnics* 2(1), pages 262–296. DOI: [10.3390/geotechnics2010012](https://doi.org/10.3390/geotechnics2010012).
- Bell, S., K. Joy, J. Pernet-Fisher, and M. Hartley (2020). “QEM-SCAN as a method of semi-automated crystal size distribution analysis: insights from Apollo 15 mare basalts”. *Journal of Petrology* 61(4), page 47. DOI: doi.org/10.1093/petrology/egaa047.
- Bodla, N., B. Singh, R. Chellappa, and L. S. Davis (2017). “Soft-NMS — Improving Object Detection with One Line of Code”. *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE, pages 5562–5570. DOI: [10.1109/iccv.2017.593](https://doi.org/10.1109/iccv.2017.593).
- Cashman, K. V. (1990). “Textural constraints on the kinetics of crystallization of igneous rocks”. *Reviews in Mineralogy and Geochemistry* 24(1), pages 259–314.
- (1993). “Relationship between plagioclase crystallization and cooling rate in basaltic melts”. *Contributions to Mineralogy and Petrology* 113(1), pages 126–142. DOI: [10.1007/bf00320836](https://doi.org/10.1007/bf00320836).
- Cashman, K. V. and J. M. Ferry (1988). “Crystal size distribution (CSD) in rocks and the kinetics and dynamics of crystallization”. *Contributions to Mineralogy and Petrology* 99(4), pages 401–415. DOI: [10.1007/bf00371933](https://doi.org/10.1007/bf00371933).
- Chen, K., J. Pang, J. Wang, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Shi, W. Ouyang, C. C. Loy, and D. Lin (2019a). “Hybrid Task Cascade for Instance Segmentation”. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4969–4978. DOI: [10.1109/CVPR.2019.00511](https://doi.org/10.1109/CVPR.2019.00511).
- Chen, K., J. Wang, J. Pang, Y. Cao, Y. Xiong, X. Li, S. Sun, W. Feng, Z. Liu, J. Xu, Z. Zhang, D. Cheng, C. Zhu, T. Cheng, Q. Zhao, B. Li, X. Lu, R. Zhu, Y. Wu, J. Dai, J. Wang, J. Shi, W. Ouyang, C. C. Loy, and D. Lin (2019b). “MMDetection: Open MMLab Detection Toolbox and Benchmark”. *arXiv preprint arXiv:1906.07155*.
- Das, R., A. Mondal, T. Chakraborty, and K. Ghosh (2022). “Deep neural networks for automatic grain-matrix segmentation in plane and cross-polarized sandstone photomicrographs”. *Applied Intelligence* 52(3), pages 2332–2345. DOI: [10.1007/s10489-021-02530-z](https://doi.org/10.1007/s10489-021-02530-z).
- Deng, J., W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei (2009). “ImageNet: A large-scale hierarchical image database”. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. DOI: [10.1109/CVPR.2009.5206848](https://doi.org/10.1109/CVPR.2009.5206848).
- Everingham, M., L. V. Gool, C. K. I. Williams, J. Winn, and A. Zisserman (2009). “The Pascal Visual Object Classes (VOC) Challenge”. *International Journal of Computer Vision* 88(2), pages 303–338. DOI: [10.1007/s11263-009-0275-4](https://doi.org/10.1007/s11263-009-0275-4).
- Foreman-Mackey, D., D. W. Hogg, D. Lang, and J. Goodman (2013). “emcee: The MCMC Hammer”. *125(925)*, page 306. DOI: [10.1086/670067](https://doi.org/10.1086/670067).
- Hafiz, A. M. and G. M. Bhat (2020). “A survey on instance segmentation: state of the art”. *International Journal of Multimedia Information Retrieval* 9(3), pages 171–189. DOI: [10.1007/s13735-020-00195-x](https://doi.org/10.1007/s13735-020-00195-x).
- He, K., G. Gkioxari, P. Dollár, and R. Girshick (2017). “Mask R-CNN”. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988. DOI: [10.1109/ICCV.2017.322](https://doi.org/10.1109/ICCV.2017.322).
- Hersum, T. G. and B. D. Marsh (2006). “Igneous microstructures from kinetic models of crystallization”. *Journal of Volcanology and Geothermal Research* 154(1-2), pages 34–47. DOI: [10.1016/j.jvolgeores.2005.09.018](https://doi.org/10.1016/j.jvolgeores.2005.09.018).
- Higgins, M. D. (1996). “Magma dynamics beneath Kameni volcano, Thera, Greece, as revealed by crystal size and shape measurements”. *Journal of Volcanology and Geothermal Research* 70(1-2), pages 37–48. DOI: [10.1016/0377-0273\(95\)00045-3](https://doi.org/10.1016/0377-0273(95)00045-3).
- (2000). “Measurement of crystal size distributions”. *American Mineralogist* 85(9), pages 1105–1116. DOI: [10.2138/am-2000-8-901](https://doi.org/10.2138/am-2000-8-901).
- (2010a). “Imaging birefringent minerals without extinction using circularly polarized light”. *The Canadian Mineralogist* 48(1), pages 231–235. DOI: [10.3749/canmin.48.1.231](https://doi.org/10.3749/canmin.48.1.231).
- (2010b). “Textural coarsening in igneous rocks”. *International Geology Review* 53(3-4), pages 354–376. DOI: [10.1080/00206814.2010.496177](https://doi.org/10.1080/00206814.2010.496177).
- Higgins, M. D. and D. Chandrasekharam (2007). “Nature of Sub-volcanic Magma Chambers, Deccan Province, India: Evidence from Quantitative Textural Analysis of Plagioclase Megacrysts in the Giant Plagioclase Basalts”. *Journal of Petrology* 48(5), pages 885–900. DOI: [10.1093/petrology/egm005](https://doi.org/10.1093/petrology/egm005).
- Holness, M. B. (2014). “The effect of crystallization time on plagioclase grain shape in dolerites”. *Contributions to Mineralogy and Petrology* 168(5), page 1076. DOI: [10.1007/s00410-014-1076-5](https://doi.org/10.1007/s00410-014-1076-5).

- Holness, M. B. (2022). “Microstructural Evidence for the Fluid Dynamical Behaviour in Vertically and Laterally Propagated Dykes of the British and Irish Paleogene Igneous Province”. *Journal of Petrology* 63(11). DOI: [10.1093/ petrology/egac107](https://doi.org/10.1093/petrology/egac107).
- Holness, M. B., C. Morris, Z. Vukmanovic, and D. J. Morgan (2020). “Insights Into Magma Chamber Processes From the Relationship Between Fabric and Grain Shape in Troctolitic Cumulates”. *Frontiers in Earth Science* 8. DOI: [10.3389/ feart.2020.00352](https://doi.org/10.3389/feart.2020.00352).
- Holness, M. B., J. A. Neufeld, A. J. Gilbert, and R. Macdonald (2017). “Orientation of Tabular Mafic Intrusions Controls Convective Vigour and Crystallization Style”. *Journal of Petrology* 58(10), pages 2035–2053. DOI: [10.1093/ petrology/egx081](https://doi.org/10.1093/ petrology/egx081).
- Jakobsson, S. P. and M. T. Gudmundsson (2008). “Subglacial and intraglacial volcanic formations in Iceland”. *Jökull* 58, pages 179–196. DOI: [10.33799/ jokull2008.58.179](https://doi.org/10.33799/jokull2008.58.179).
- Jerram, D. A. (2003). “Quantifying the Building Blocks of Igneous Rocks: Are Clustered Crystal Frameworks the Foundation?” *Journal of Petrology* 44(11), pages 2033–2051. DOI: [10.1093/ petrology/egg069](https://doi.org/10.1093/ petrology/egg069).
- Kiefer, J. and J. Wolfowitz (1952). “Stochastic Estimation of the Maximum of a Regression Function”. *The Annals of Mathematical Statistics* 23(3), pages 462–466. DOI: [10.1214/ aoms/1177729392](https://doi.org/10.1214/aoms/1177729392).
- LeCun, Y., Y. Bengio, and G. Hinton (2015). “Deep learning”. *Nature* 521(7553), pages 436–444. DOI: [10.1038/ nature14539](https://doi.org/10.1038/nature14539).
- Leichter, A., R. R. Almeev, D. Wittich, P. Beckmann, F. Rottensteiner, F. Holtz, and M. Sester (2022). “Automated Segmentation of Olivine Phenocrysts in a Volcanic Rock Thin Section Using a Fully Convolutional Neural Network”. *Frontiers in Earth Science* 10. DOI: [10.3389/ feart.2022.740638](https://doi.org/10.3389/feart.2022.740638).
- Lin, T.-Y., M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár (2014). *Microsoft COCO: Common Objects in Context*. DOI: [10.48550/ARXIV.1405.0312](https://doi.org/10.48550/ARXIV.1405.0312).
- Lofgren, G. (1974). “An experimental study of plagioclase crystal morphology; isothermal crystallization”. *American journal of Science* 274(3), pages 243–273. DOI: [10.2475/ajs.274.3.243](https://doi.org/10.2475/ajs.274.3.243).
- Lormand, C., G. F. Zellmer, K. Németh, G. Kilgour, S. Mead, A. S. Palmer, N. Sakamoto, H. Yurimoto, and A. Moebis (2018). “Weka Trainable Segmentation Plugin in ImageJ: A Semi-Automatic Tool Applied to Crystal Size Distributions of Microlites in Volcanic Rocks”. *Microscopy and Microanalysis* 24(6), pages 667–675. DOI: [10.1017/ s1431927618015428](https://doi.org/10.1017/s1431927618015428).
- Mangan, M. T. (1990). “Crystal size distribution systematics and the determination of magma storage times: The 1959 eruption of Kilauea volcano, Hawaii”. *Journal of Volcanology and Geothermal Research* 44(3-4), pages 295–302. DOI: [10.1016/0377-0273\(90\)90023-9](https://doi.org/10.1016/0377-0273(90)90023-9).
- Mangler, M. F., M. C. S. Humphreys, F. B. Wadsworth, A. A. Iveson, and M. D. Higgins (2022). “Variation of plagioclase shape with size in intermediate magmas: a window into incipient plagioclase crystallisation”. *Contributions to Mineralogy and Petrology* 177(6), page 64. DOI: [10.1007/s00410-022-01922-9](https://doi.org/10.1007/s00410-022-01922-9).
- Marsh, B. D. (1988). “Crystal size distribution (CSD) in rocks and the kinetics and dynamics of crystallization”. *Contributions to Mineralogy and Petrology* 99(3), pages 277–291. DOI: [10.1007/bf00375362](https://doi.org/10.1007/bf00375362).
- (1998). “On the Interpretation of Crystal Size Distributions in Magmatic Systems”. *Journal of Petrology* 39(4), pages 553–599. DOI: [10.1093/ petroj/39.4.553](https://doi.org/10.1093/ petroj/39.4.553).
- Morgan, D. J. and D. A. Jerram (2006). “On estimating crystal shape for crystal size distribution analysis”. *Journal of Volcanology and Geothermal Research* 154(1-2), pages 1–7. DOI: [10.1016/j. jvolgeores.2005.09.016](https://doi.org/10.1016/j.jvolgeores.2005.09.016).
- Neave, D. A., I. Buisman, and J. Maclennan (2017). “Continuous mush disaggregation during the long-lasting Laki fissure eruption, Iceland”. *American Mineralogist* 102(10), pages 2007–2021. DOI: [10.2138/am-2017-6015CCBY](https://doi.org/10.2138/am-2017-6015CCBY).
- Neave, D. A., J. Maclennan, M. E. Hartley, M. Edmonds, and T. Thordarson (2014). “Crystal Storage and Transfer in Basaltic Systems: the Skuggafjöll Eruption, Iceland”. *Journal of Petrology* 55(12), pages 2311–2346. DOI: [10.1093/ petrology/egu058](https://doi.org/10.1093/ petrology/egu058).
- Neubeck, A. and L. Van Gool (2006). “Efficient Non-Maximum Suppression”. *18th International Conference on Pattern Recognition (ICPR'06)*. Volume 3, pages 850–855. DOI: [10.1109/ICPR.2006.479](https://doi.org/10.1109/ICPR.2006.479).
- Ortolano, G., L. Zappalà, and P. Mazzoleni (2014). “X-Ray Map Analyser: A new ArcGIS® based tool for the quantitative statistical data handling of X-ray maps (Geo- and material-science applications)”. *Computers & Geosciences* 72, pages 49–64. DOI: [10.1016/j. cageo.2014.07.006](https://doi.org/10.1016/j.cageo.2014.07.006).
- Patro, S., D. Jhariya, M. Sahu, P. Dewangan, and P. Dhekne (2022). “Igneous rock classification using Convolutional neural networks (CNN)”. *IOP Conference Series: Earth and Environmental Science* 1032(1), page 012045. DOI: [10.1088/1755-1315/1032/1/012045](https://doi.org/10.1088/1755-1315/1032/1/012045).
- Qiao, S., L.-C. Chen, and A. Yuille (2020). “DetectoRS: Detecting Objects with Recursive Feature Pyramid and Switchable Atrous Convolution”. *arXiv preprint arXiv:2006.02334*.
- Schneider, C. A., W. S. Rasband, and K. W. Eliceiri (2012). “NIH Image to ImageJ: 25 years of image analysis”. *Nature Methods* 9(7), pages 671–675. DOI: [10.1038/nmeth.2089](https://doi.org/10.1038/nmeth.2089).
- Schulz, B., D. Sandmann, and S. Gilbricht (2020). “SEM-Based Automated Mineralogy and Its Application in Geo- and Material Sciences”. *Minerals* 10(11), page 1004. DOI: [10.3390/ min10111004](https://doi.org/10.3390/min10111004).
- Sheldrake, T. and O. Higgins (2021). “Classification, segmentation and correlation of zoned minerals”. *Computers & Geosciences* 156, page 104876. DOI: [10.1016/j. cageo.2021.104876](https://doi.org/10.1016/j.cageo.2021.104876).
- Smith, J. V. and W. L. Brown (1988). *Feldspar Minerals*. Springer Berlin Heidelberg. ISBN: 9783642725944. DOI: [10.1007/978-3-642-72594-4](https://doi.org/10.1007/978-3-642-72594-4).
- Tung, P.-Y., H. A. Sheikh, M. Ball, F. Nabiei, and R. J. Harrison (2023). “SIGMA: Spectral Interpretation Using Gaussian Mixtures and Autoencoder”. *Geochemistry, Geophysics, Geosystems* 24(1). DOI: [10.1029/2022gc010530](https://doi.org/10.1029/2022gc010530).



- Van der Walt, S., J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu (2014). “scikit-image: image processing in Python”. *PeerJ* 2, page 453.
- Visalli, R., M. Giuffrida, and M. Viccaro (2023). “Unraveling Textural and Chemical Features in Volcanic Rocks Through Advanced Image Processing: A Case Study From the 2019 Paroxysmal Eruptions of Stromboli”. *Geochemistry, Geophysics, Geosystems* 24(3). DOI: [10.1029/2022gc010774](https://doi.org/10.1029/2022gc010774).
- Wada, K. (2016). *Labelme: Image Polygonal Annotation with Python*. DOI: [10.5281/zenodo.5711226](https://doi.org/10.5281/zenodo.5711226).
- Wang, X., R. Zhang, T. Kong, L. Li, and C. Shen (2020). “SOLOv2: Dynamic and Fast Instance Segmentation”. *Advances in Neural Information Processing Systems*. Edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Volume 33. Curran Associates, Inc., pages 17721–17732.
- Xu, R. (2019). “Labelme2coco: A lightweight package for converting your labelme annotations into COCO object detection format.” *Github*. [Software] Available at <https://github.com/XCRobert/Labelme2Coco>.
- Yu, J., F. Wellmann, S. Virgo, M. von Domarus, M. Jiang, J. Schmatz, and B. Leibe (2023). “Superpixel segmentations for thin sections: Evaluation of methods to enable the generation of machine learning training data sets”. *Computers & Geosciences* 170, page 105232. DOI: [10.1016/j.cageo.2022.105232](https://doi.org/10.1016/j.cageo.2022.105232).