A Deep Learning-enabled Approach for Plagioclase Textural Studies

● Norbert Toth *$^\alpha$ and ● John Maclennan$^\alpha$

$^\alpha$ University of Cambridge, Cambridge, UK.

ABSTRACT

Textural information, such as crystal size distributions (CSDs) or crystal aspect ratios are powerful tools in igneous petrography for interrogating the thermal history of rocks. They facilitate the investigation of crystal nucleation, growth and mixing as well as the cooling rate of the rock. However, they require large volumes of crystal segmentations and measurements that are often obtained with manual methods. Here a deep learning-based computer vision technique, termed instance segmentation, is proposed to automate the pixel-by-pixel detection of each plagioclase crystal in thin section images. Using predictions from a re-trained model the physical properties of the detected crystals, such as size and aspect ratio, can be rapidly generated to provide textural insights. The present segmentations are validated against published results from manual approaches to prove the method's accuracy. The power and efficiency of this automated approach is showcased by analysing an entire sample suite, segmenting over 48,000 crystals in only a matter of days. Widescale use of this method is expected to drive significant developments in the igneous petrography and related fields.

KEYWORDS: Deep Learning; Segmentation; Petrography; Timescales; MCMC.

1 INTRODUCTION

The textures of igneous rocks provide direct observational contraints upon the physics of magmatic processes [Cashman 1990; Higgins 2010; Askaripour et al. 2022]. Petrologists have long made use of quantitative descriptions of rock texture such as crystal size distributions (CSD's) [Cashman and Ferry 1988; Marsh 1988; Mangan 1990; Marsh 1998; Higgins 2000; Armienti 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 predictably with crystallization time [Lofgren 1974]. A simple parameterization 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 relationship provides a direct method of determining the timescales of cooling of plagioclase-bearing igneous rocks from simple two-dimensional observations alone.

---

*✉ nt398@cam.ac.uk

More recent work by Mangler et al. [2022] took this idea further and looked at plagioclase microlites 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] developed a method to determine most likely 3D crystal morphology from 2D slice observations. 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 image thresholding, where possible, to provide approximate crystal masks, or more often on manual crystal segmentation which is a labour intensive and time consuming task. To overcome this significant effort, a number of approaches have been proposed based on machine learning. The simplest, and perhaps least relevant in this case, were classification methods based on convolutional neural networks (CNN) to predict the rock type shown in images [Alférez 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 SEM 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 localization 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. The 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. Overall using the present method, a ten-fold speedup in segmentation time is reported only using CPU's for processing, 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. All images were acquired using both linear polarizing filters crossed at 90 degrees and circular polarizers to remove the crystal orientation dependence of birefringence colour intensity. 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 μm/pixel. The proprietary ZEISS file formats were exported as *.jpg* images with as little loss of quality from compression as possible. These files can now be used directly by the deep learning models.

2.2 Instance Segmentation Models

Instance segmentation is a subset of image segmentation methods that aim 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. 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. Based on the calculated features, the model proposes a large number of regions of interest. These regions are then evaluated based on a trained scoring strategy for the given application, and any region below a threshold score are discarded. For the 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*; note that these are strictly not probability values.

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. It is 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 most useful and in fact they are in the identical form as 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. 2019] and DetectoRS [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 DetectoRS 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. DetectoRS 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 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 and 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, bounding box detection and mask segmentation. 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 de-
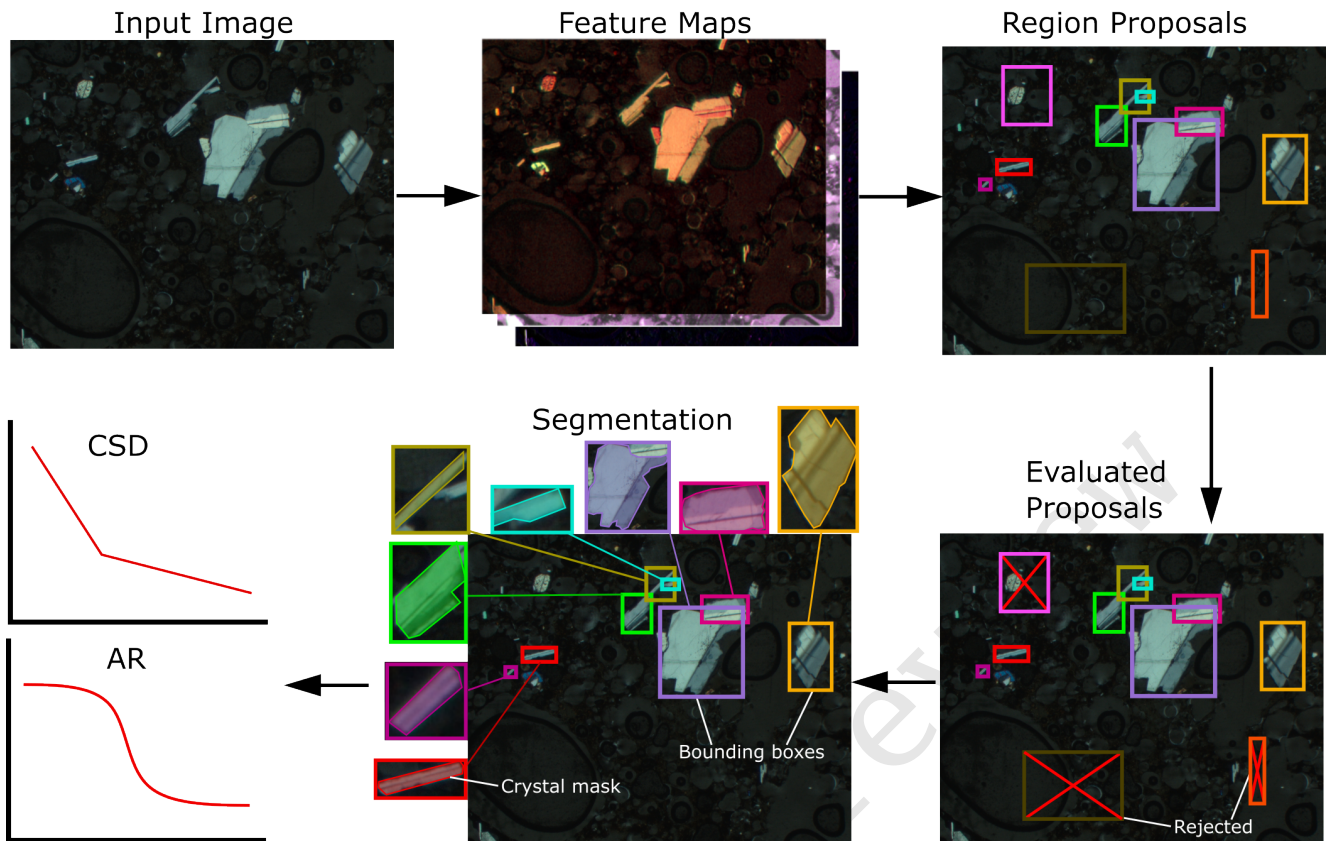
Figure 1: Overview of the proposed region-based deep learning method for plagioclase feldspar detection. Deep learning models necessarily transform the input into large stacks of feature maps that are used for further processing; next the model proposes regions of interests that are evaluated on whether 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 and aspect ratio distributions. Specific terms used to indicate the type of segmentation results such as *bounding box* or *mask* is indicated as appropriate for clarity.

fines the effective step size during the parameter 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 to 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 2(a). 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 no date]. 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 [XCRobert 2019] with a training-

test split of 80%. Examples of manually segmented plagioclase crystals are presented in Figure 2(b).

Model training is performed through the MMDetection open-source object detection toolbox [MMDetection Contributors 2018] that 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 Stochastic Gradient Descent (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 trained for 600 epochs. Testing scores for each model and learning rate during training are shown in Figure 2(c); there
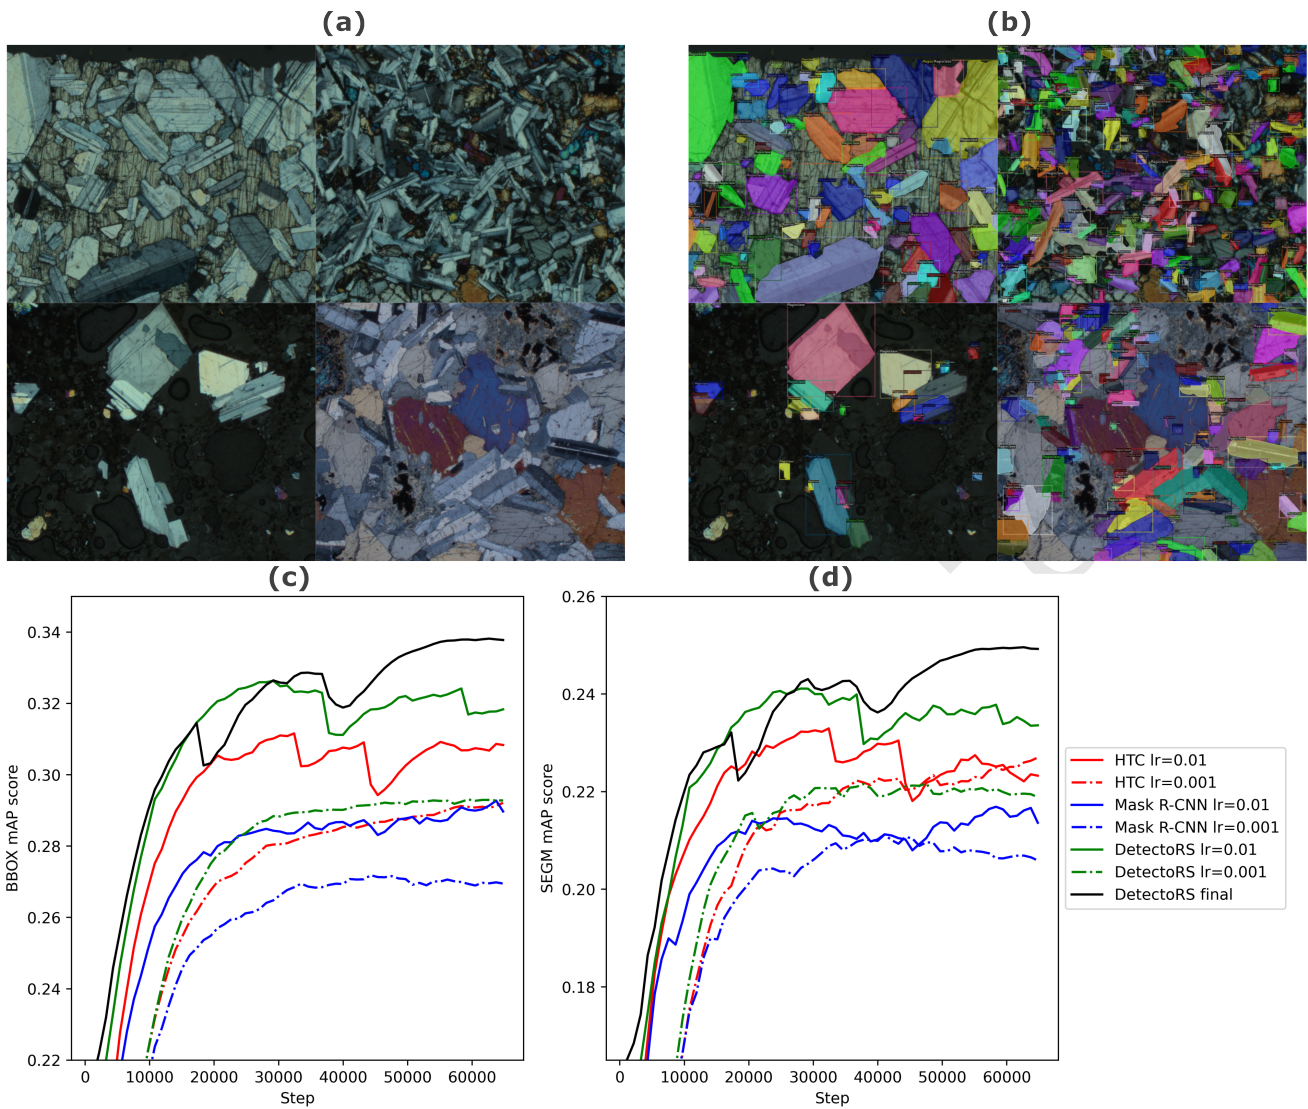
(a)

(b)

(c)

(d)

Figure 2: Training dataset and overview of trained models' performance. (a) Example micrographs from the training deteset, (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 maximized for best performance. Smoothing was applied using an exponential moving average method.

are no results for models with learning rates of 0.1 as their SGD never converged and were rendered unusable. 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 outperfom 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 where it is once again stepped down to 0.0001 for the last 100 epochs. All other SGD parameters are kept fixed. This is done to potentially 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 2(c), the final DetectoRS model outperforms all previously trained ones and is therefore the best one to implement for all subsequent work. One must note this may not be the most optimised - 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 - overall equating to three days of computational time.

2.4 Inference Software Implementation

The bulk of the software was written 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 utilized to analyse the results and build directly onto the inference routines in the MMDetection API. 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 100GB or more. Therefore, a custom tiling strategy is employed during inference; a simple schematic is shown in Figure 3(a). 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 3(b). 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 3(b), 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.

Using the trained DetectoRS model and the software outlined above, an example segmentation image is shown in Figure 4(a) 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. This shows 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.

3 EXAMPLE USAGE

With the model trained on ground truth labels for plagioclase feldspar crystals in basaltic composition igneous rocks, an example use of the method is showcased to analyse natural volcanic samples. Segmentation results will be used to produce CSD's and aspect ratio distributions to interrogate the samples' textural history. To assess the accuracy of the present approach, the plots generated will be compared to published manually collected results. The following section is merely a guide and by no means an exhaustive description of the method's capabilities.

3.1 The Skuggafjöll Eruption

Skuggafjöll is a NE-SW-striking hyaloclastite ridge between the Mýrdalsjökull 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

(a)

(b)

$$\text{IoU} = \frac{\text{Intersection}}{\text{Union}} = \frac{\textcolor{green}{A_G}}{\textcolor{green}{A_G} + \textcolor{blue}{A_B} + \textcolor{red}{A_R}}$$
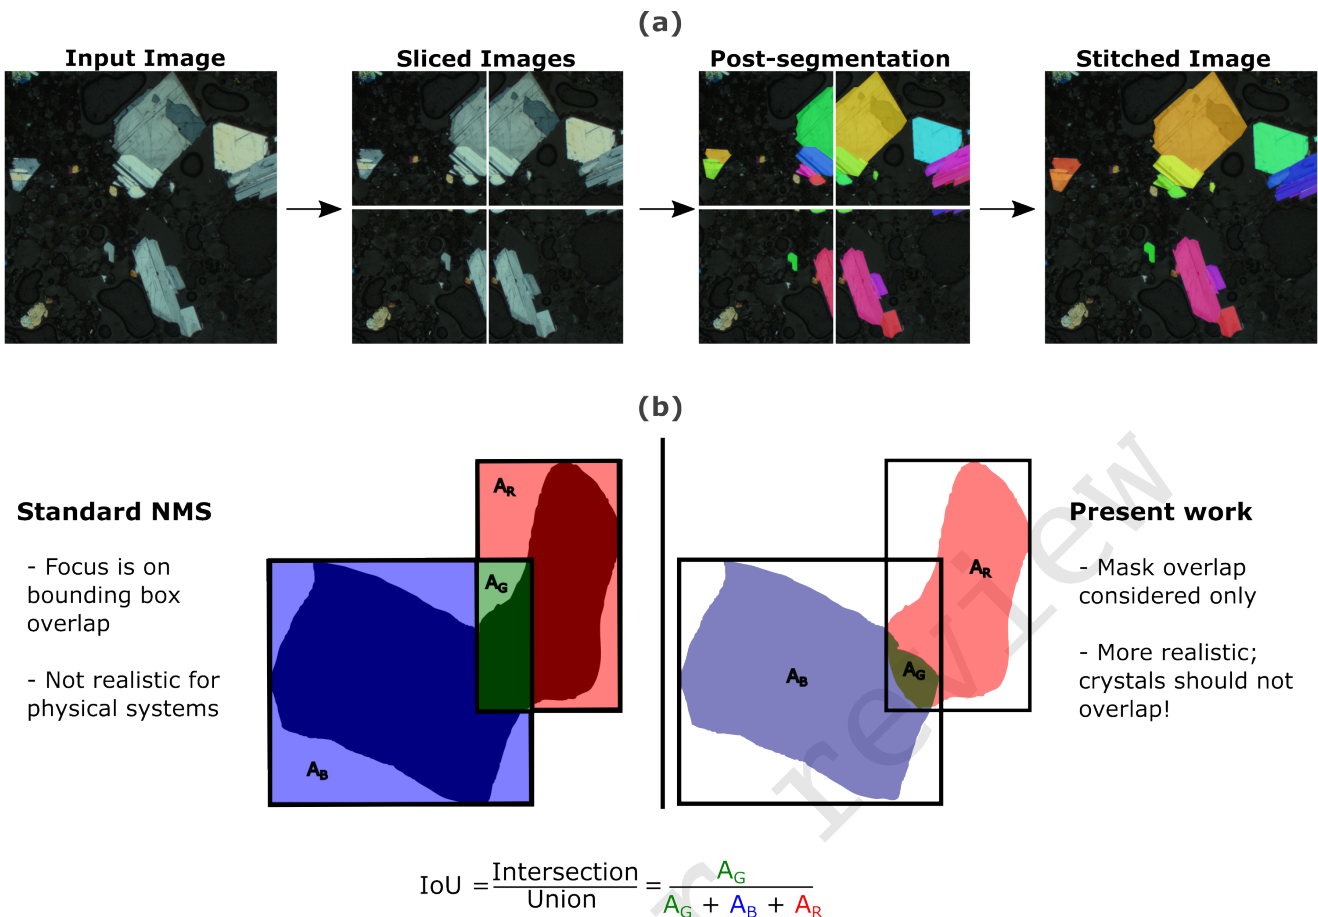
Figure 3: Overview of the approach developed for detection on large area section scans and the post-processing step for removing overlapping segmentations. (a) The input image is initially sliced into the desired tiles with overlaps between them; the individual tiles are used for the detection to reduce the computational power required, which are stitched together to produce the final segmented image.(b) Schematic overview of the NMS approach used in the present work compared to the standard algorithms most often utilized for object detection. Inference-over-union (IoU) scores are introduced as the primary way to calculate overlaps in images.

well as a combination of X-ray flourescence, ICP-MS 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 through 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 present Deep Learning instance segmentation 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 restitched using the custom tiling subroutine. Any instances with a detection score below 0.8 or with mask overlap above 0.5 were removed. 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 5 shows the comparison between the CSD's 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. There is impressive similarity in the results of the new automated and old manual methods, despite the possibility of bias inherent to each technique. Prior work has established the CSD's for these exact samples show characteristic crystal lengths (L_D) of 70-100 µm for the population of small macrocrysts and 700-1030 µm for the large population [Neave et al. 2014]. The present automated segmentation approach yields very similar values of 80-100 µm and 590-1090 µm for the small and large macrocryst populations respectively. Likewise the pronounced in-
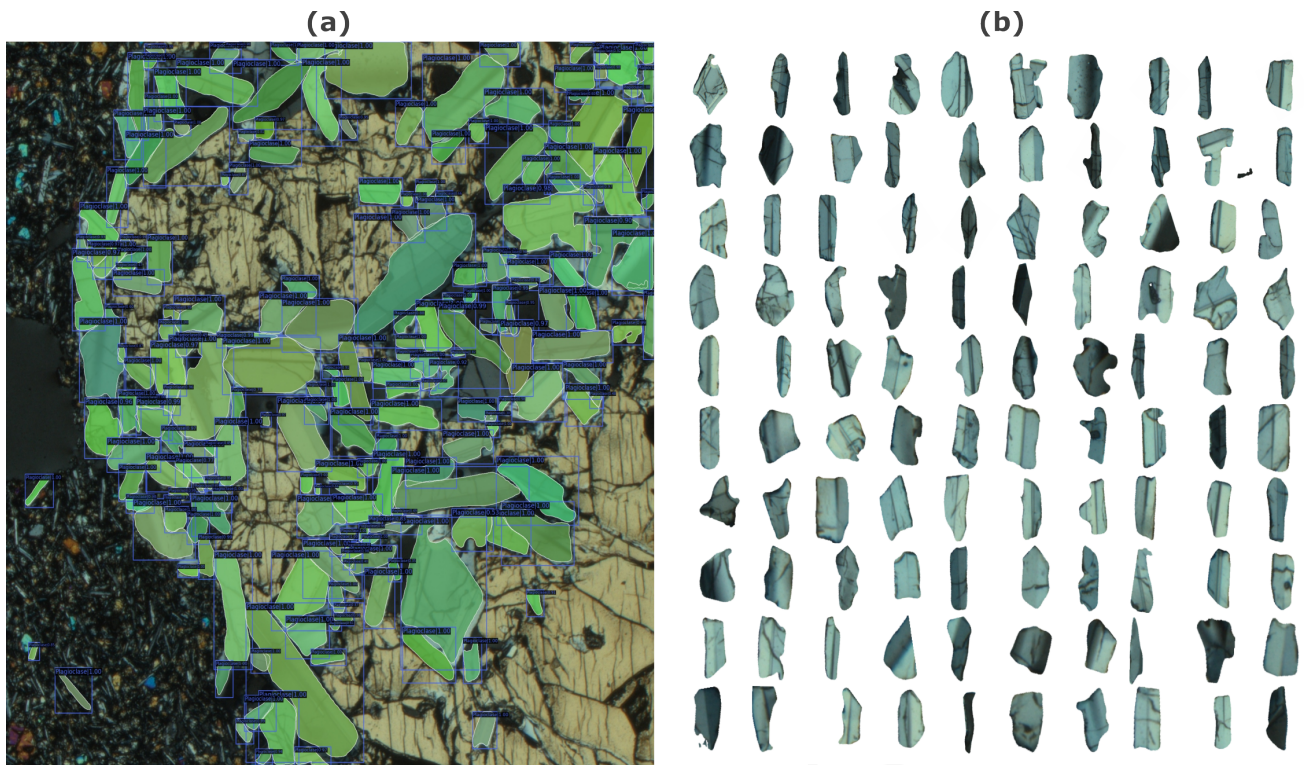
Figure 4: Example segmentation on a test image of touching ophitic plagioclase crystals. (a) - Segmented image showing the segmented crystal masks, and (b) - the 100 largest crystals lined up in descending order for further inspection.
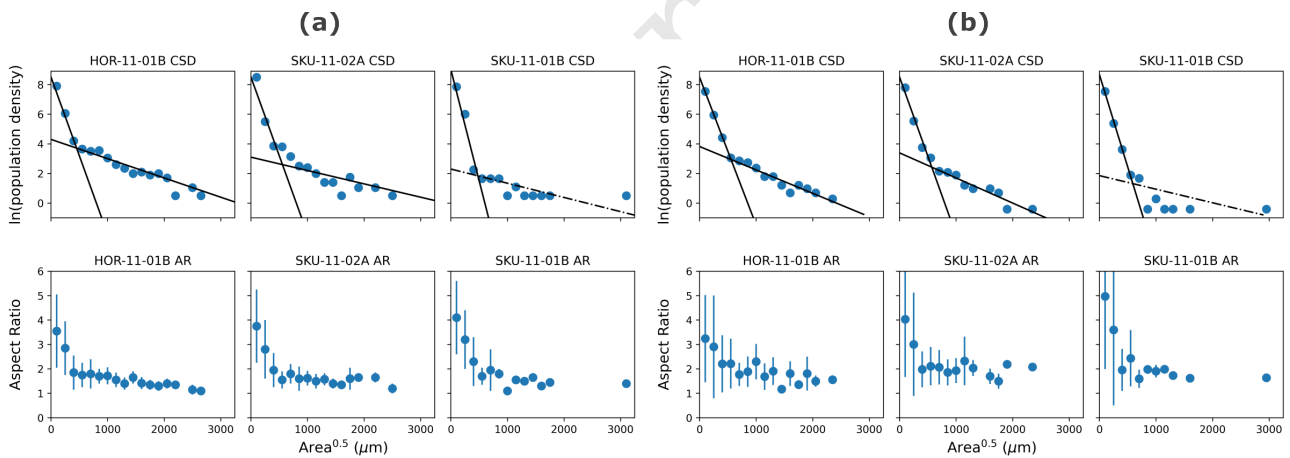

Figure 5: Comparison of Skuggafjöll crystal size and aspect ratio distributions acquired using the present method for three selected samples with manual segmentation data available from Neave et al. [2014]. (a) Figures from Neave et al. [2014]. (b) Data acquired from DL approach. The data sets show exceptional matching across the two methods used, serving as useful validation for the method presented.

flections in the CSD's 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 aspect ratio plots at the approximately identical crystal sizes, with the large macrocrysts composed of low aspect ratio (∼2) with the smaller crystals showing a continuously increasing trend with decreasing crystal size. It is worth noting that 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. 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 GPU's will reduce this to around 30 minutes per section or less. 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 segmentation 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 $100\,\mu m$ square root crystal area. CSDs for all samples are shown in Figure 6(a) with piece-wise linear regression calculated using a single breakpoint for kinked CSDs and simple linear regression for the two samples showing straight line CSDs. This is done using the Python package *piecewise-regression* that uses an iterative approach to find breakpoints whilst using a sampling method to find global optima. All fitted parameters are reported with suitable statistical information including the standard errors from the sampling procedure. Figure 6(b) shows the plots between each fitted parameter with one standard error deviation shown where appropriate. α_1 and α_2 are the gradients of the CSDs as shown; 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 -1.7×10^{-2}, suggesting average L_D of $60\,\mu m$ square root area. α_2's have a mean of -2.2×10^{-3} meaning average L_D is approximately $450\,\mu m$ for the large macrocrysts. Breakpoints cluster around $400\,\mu m$ square root area, but show significant variability with a standard error of $64\,\mu m$. These values tend to be on the smaller end of the values discussed previously and found by Neave et al. [2014] for the three selected samples. This can be attributed to the automatic approach perhaps not excluding some of the likely anomalous points in the CSDs that a human would during linear regression. However, the errors arising from this most likely form a significant part of the uncertainty in the fitted parameters as determined by the model's sampling approach. Overall such a regression model may explore a larger range of possible fits than one where the breakpoints and outliers are identifed on a subjective manual basis. Statistically it is expected the compilation of results from all 27 samples is a more accurate representation of the plagioclase population as a whole than one can expect from a subset of only three samples.

Nevertheless, 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. This is shown best by the distribution of breakpoints as samples with a higher proportion of large phenocrysts show a change in population at a smaller crystal size. Figure 6(b) does show some scatter in the parameters, however the majority of the 27 samples do

plot in similar positions within error. The scattered points also clearly show significantly larger errors, showing the model is significantly less certain in the proposed fit.

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 $100\,\mu m$ 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 from Holness [2014] were imaged in panorama mode utilizing circular polarised optics as described for the previous set of samples discussed. Input images were sliced to 1000×1000 px with 250 px overlap. The NMS threshold was 0.5 IoU overlap as before. The resulting segmentations were used to calculate crystal sizes and aspect ratios.

Figure 7 shows a plot of published aspect ratio data [Holness 2014; Holness et al. 2017] compared to the values obtained in the present study. 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 Neave et al. [2014]. The observed differences arise from the differences in aspect ratio measurements used by the two approaches. The manual approach utilized by 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 required for such analytical work [Higgins 2000]. The fact that variation in the two sets of values is clearly systematic ($r = 0.815$), and not due to random segmentation
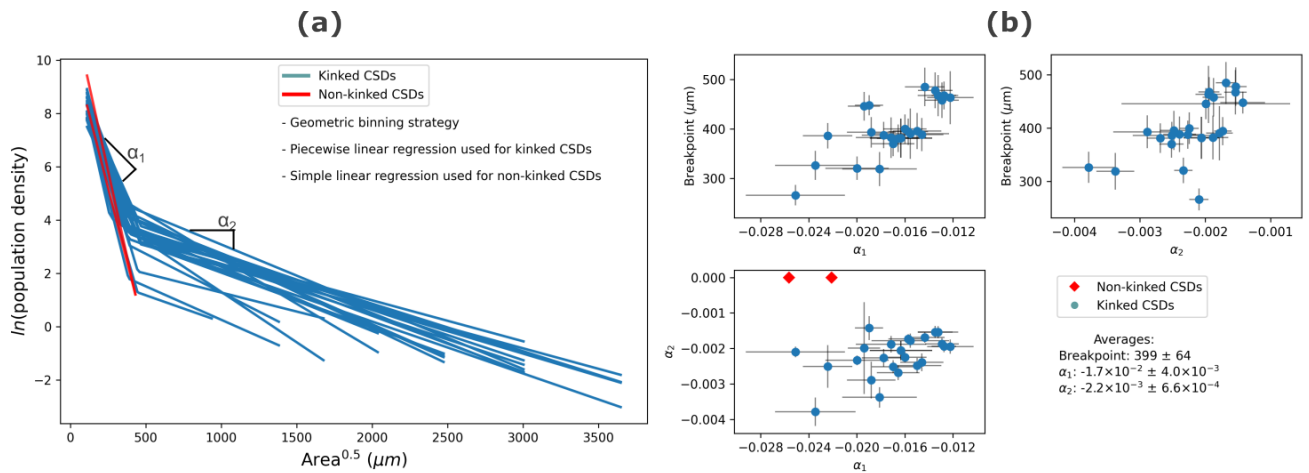
(a)

(b)

Figure 6: Crystal size distribution of plagioclase for the entire Skuggafjöll suite of thin sections, excluding hyaloclastites, as segmented using the DL method presented (a) Piecewise linear regression modelling of all but two CSDs with a single breakpoint [Pilgrim 2021], the remaining two exhibit straight CSDs and ordinary linear regression was performed; (b) Plots showing the relationships between the three parameters obtained from the regression: the breakpoint and the two different gradient values as shown. The two samples without any macrocrysts present that show straight line CSDs have gradients on the steeper end of distribution, but they are by no means anomalous. Quantification of the mean and 1 standard error spread of the fitted parameters was calculated using a sampling approach to take into account the spread in the sample means and the uncertainty in each fitted value.
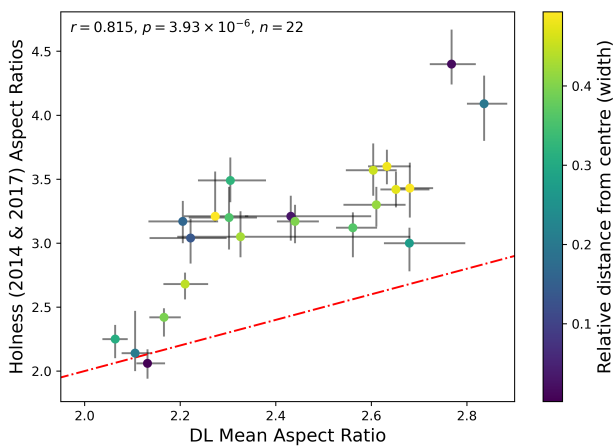

Figure 7: Comparison of aspect ratio values obtained from Holness [2014] and using the present approach. The red line shows the 1:1 line. The 95% confidence intervals were calculated using the identical bootstrapping method from the previous work. Pearson's correlation coefficient indicated shows a strong correlation does exist between the two datasets, although they clearly do not lie on the 1:1 line.

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 samples close to the margins of sills exhibit anomalous grain shapes of significantly lower aspect ratios than expected from the calculated crystallisation times. 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 utilized through Markov-chain Monte Carlo (MCMC) [Andrieu et al. 2003] using the Python packages *Bilby* and *emcee* [Foreman-Mackey et al. 2013; Ashton et al. 2019]. 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 $[-10, 0]$ and $[0, 25]$ for m and c respectively, drawing 20,000 samples to construct the posterior joint distribution after the initial burn-in period. The calibration is shown on 8(a) plotted with the data with the posterior joint distribution for m and c plotted on (b). This method is robust to noise, takes account of the uncertainty in the observations and provides confidence intervals for the calibration.

3.3 Application of Timescales

One can make use of the above calibration to get a sense of the growth conditions of plagioclase crystals, assuming their growth can be related to cooling in a body akin to the sills described by Holness [2014]. To demonstrate this, the segmentation results of the whole suite of Skuggafjöll will be used;
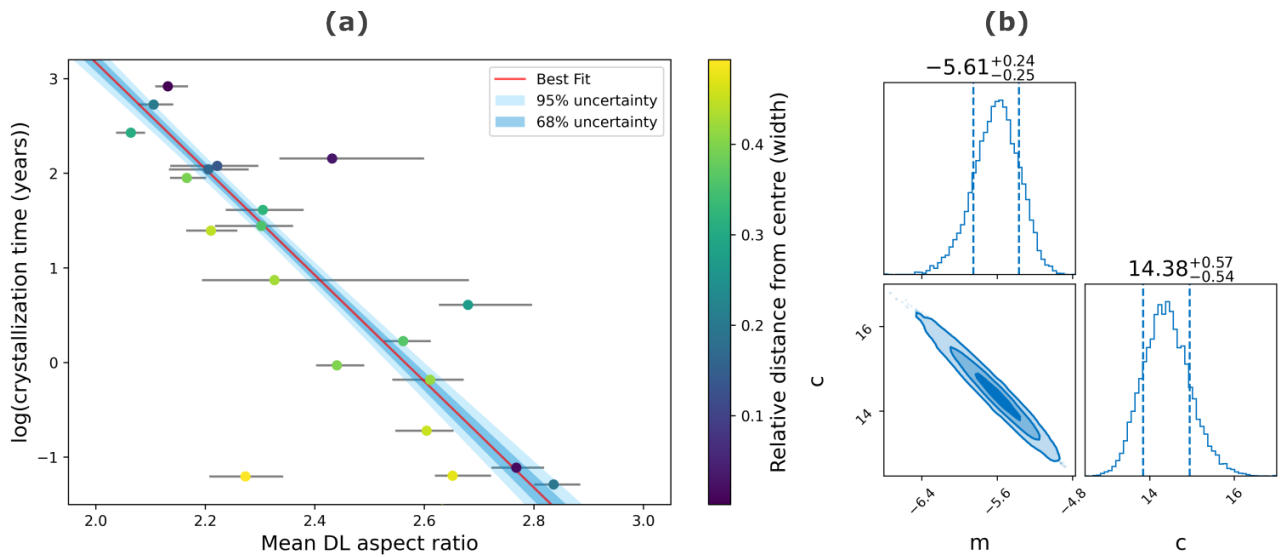
(a)

(b)

Figure 8: 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.
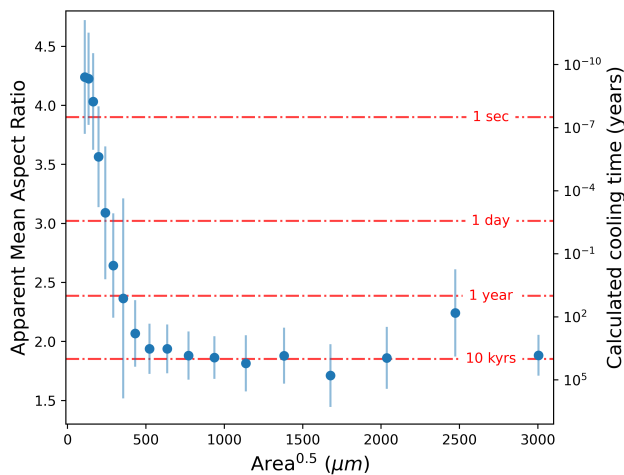

Figure 9: Plagioclase shape information from the whole suite of Skuggafjöll samples, showing the average aspect ratio for a given size bin with ±1SD error bars in blue. Data points in red show calculated characteristic cooling times using the calibration from the previous section, also showing ±1σ error bars.

collated aspect ratio values are shown in Figure 9 showing both the average and the spread of aspect ratio measurements across all the samples. As discussed previously, large macrocrysts, larger than 500-600 μm in square root area, consistently show aspect ratios ~2, with the smaller population of crystals showing an increasing trend between from ~2 to ~4.5 with decreasing square root area. This trend may indicate changing growth conditions in the final stages of cooling prior to eruption. As the crystals entrained in the erupted lavas are unlikely to come from fully crystallised sill-like bodies, such as those used for the timescale calibration above, the times

calculated will be referred to as *effective cooling times* and will not be interpreted as exact timescales.

Time estimates are also shown in Figure 9 using the previous calibration, taking account of the spread in aspect ratio values across all the samples showing ±1σ uncertainty. Care was taken to propagate all uncertainties to final estimates. The large macrocryst population shows effective crystallisation times on the order of 1,000-10,000 years, whereas the highest aspect ratio (smallest in size) crystals show extremely short timescales on the order of a second or so. However, anything above aspect ratio of 3 is an extrapolation of the calibration and the calculated time values should not be taken for granted as the linear model may no longer apply due to finite crystal growth rates. The fastest reported growth rates of plagioclase, producing skeletal crystals, from cooling experiments is around 10^{-3} mm s^{-1} [Arzilli et al. 2019] which is not fast enough to account for the calculated times. The growth of non-skeletal plagioclase is orders of magnitude slower [Moschini et al. 2023]. However, below aspect ratio of 3 the linear model is indeed applicable and the uncertainty in exact fit at those higher values is taken care of by modelling uncertainties. The further the model strays from the calibration data the larger the inherent modelling uncertainties as evident even from Figure 8.

Furthermore, one can take the CSD results from Figure 6 to gain independent timescale information. Although the calculated CSD's are not stereologically corrected, it has been shown previously that order of magnitude arguments can be made without the need for such corrections [Neave et al. 2017; Cashman 2020]. Slopes of CSD for the smaller crystal populations were found in the range -0.012 and -0.020 μm^{-1} and taking the range of growth rates as $10^{-6} - -10^{-7}$ mm s^{-1} [Neave et al. 2017] the calculated residence time ranges from 0.6 days to 0.6 years. Given the large uncertainties, these time-

cales are broadly consistent with the those estimated from crystal shapes.

These results support the arguments of Neave et al. [2014] that the large plagioclase macrocrysts represent disaggregated mush material likely forming a large reservoir as indicated by the crystal shapes and therefore their calculated effective crystallisation times. Disaggregation of this crystal mush has most likely occurred less than a year prior to eruption. Meanwhile the smaller plagioclase crystals most likely grew post- or syn-disaggregation.

4   Conclusions

A toolbox based on a deep learning algorithm is presented for the textural analyses of petrological thin sections, focussing on plagioclase feldspar. 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 sample 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 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 charcterization 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 measurements 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

Acknowledgements

Data Availability

The training dataset, model configuration and 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/PlagDetect (https://doi.org/10.5281/zenodo.7851939).

Copyright notice

References

Alférez, 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). *Machine Learning* 50(1/2), pages 5–43. DOI: 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.

Arzilli, F., G. L. Spina, M. R. Burton, M. Polacci, N. L. Gall, M. E. Hartley, D. D. Genova, B. Cai, N. T. Vo, E. C. Bamber, S. Nonni, R. Atwood, E. W. Llewellin, R. A. Brooker, H. M. Mader, and P. D. Lee (2019). "Magma fragmentation in highly explosive basaltic eruptions induced by rapid crystallization". *Nature Geoscience* 12(12), pages 1023–1028. DOI: 10.1038/s41561-019-0468-6.

Ashton, G., M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov, F. Hernandez Vivanco, J. Harms, M. E. Lower, G. D. Meadors, D. Melchor, E. Payne, M. D. Pitkin, J. Powell, N. Sarin, R. J. E. Smith, and E. Thrane (2019). "BILBY: A User-friendly Bayesian Inference Library for Gravitational-wave Astronomy". 241(2), 27, page 27. DOI: 10.3847/1538-4365/ab06fc.

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.

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), egaa047. 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". In: *2017 IEEE International Conference on Computer Vision (ICCV)*. IEEE. DOI: 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.

Cashman, K. V. (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.

– (2020). "Crystal Size Distribution (CSD) Analysis of Volcanic Samples: Advances and Challenges". *Frontiers in Earth Science* 8. DOI: 10.3389/feart.2020.00291.

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.

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 (2019). "Hybrid task cascade for instance segmentation". In: *IEEE Conference on Computer Vision and Pattern Recognition*.

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.

Deng, J., W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei (2009). "Imagenet: A large-scale hierarchical image database". In: *2009 IEEE conference on computer vision and pattern recognition*. Ieee, pages 248–255.

Foreman-Mackey, D., D. W. Hogg, D. Lang, and J. Goodman (2013). "emcee: The MCMC Hammer". 125(925), page 306. DOI: 10.1086/670067.

Hafiz, A. M. and G. M. Bhat (2020). "A Survey on Instance Segmentation: State of the art". *CoRR* abs/2007.00047.

He, K., G. Gkioxari, P. Dollár, and R. Girshick (2017). "Mask R-CNN". In: *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988. DOI: 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.

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.

– (2000). "Measurement of crystal size distributions". *American Mineralogist* 85(9), pages 1105–1116. DOI: 10.2138/am-2000-8-901.

– (2010). "Textural coarsening in igneous rocks". *International Geology Review* 53(3-4), pages 354–376. DOI: 10.1080/00206814.2010.496177.

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.

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). egac107. DOI: 10.1093/petrology/egac107.

– (2014). "The effect of crystallization time on plagioclase grain shape in dolerites". *Contributions to Mineralogy and Petrology* 168(5). DOI: 10.1007/s00410-014-1076-5.

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.

Jakobsson, S. P. and M. T. Gudmundsson (2008). "Subglacial and intraglacial volcanic formations in Iceland". *Jökull* 58, pages 179–196.

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.

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.

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.

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.

Lofgren, G. (1974). "An experimental study of plagioclase crystal morphology; isothermal crystallization". *American journal of Science* 274(3), pages 243–273.

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.

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.

Marsh, B. D. (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.

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.

MMDetection Contributors (2018). *OpenMMLab Detection Toolbox and Benchmark*.

Moschini, P., S. Mollo, A. Pontesilli, M. Nazzari, C. Petrone, S. Fanara, A. Vona, M. Gaeta, C. Romano, and P. Scarlato (2023). "A review of plagioclase growth rate and compositional evolution in mafic alkaline magmas: Guidelines for thermometry, hygrometry, and timescales of magma dynamics at Stromboli and Mt. Etna". *Earth-Science Reviews* 240, page 104399. DOI: https://doi.org/10.1016/j.earscirev.2023.104399.

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.

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.

Neubeck, A. and L. Van Gool (2006). "Efficient non-maximum suppression". In: *18th international conference on pattern recognition (ICPR'06)*. Volume 3. IEEE, pages 850–855.

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.

Pilgrim, C. (2021). "piecewise-regression (aka segmented regression) in Python". *Journal of Open Source Software* 6(68), page 3859. DOI: 10.21105/joss.03859.

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.

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.

Smith, J. V. and W. L. Brown (1988). *Feldspar Minerals*. Springer Berlin Heidelberg. DOI: 10.1007/978-3-642-72594-4.

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, e453.

Wada, K. (no date). *Labelme: Image Polygonal Annotation with Python*. DOI: 10.5281/zenodo.5711226.

Wang, X., R. Zhang, T. Kong, L. Li, and C. Shen (2020). "SOLOv2: Dynamic and Fast Instance Segmentation". In: *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.

XCRobert (2019). *Labelme2coco: A lightweight package for converting your labelme annotations into COCO object detection format*.