

# Technical Note: colab\_zirc\_dims: a Google-Colab-based ~~Toolset~~compatible toolset for ~~Automated and Semi-automated Measurement and semi-automated measurement~~ of ~~Mineral Grains~~mineral grains in LA-ICP-MS ~~Images Using Deep Learning Models~~images using deep learning models

Michael C. Sitar<sup>1</sup>, Ryan J. Leary<sup>2</sup>

<sup>1</sup>Department of Geosciences, Colorado State University, Fort Collins, CO, 80523-1482, USA

<sup>2</sup>Department of Earth and Environmental Science, New Mexico Institute of Mining and Technology, Socorro, NM, 87801, USA

10 *Correspondence to:* Michael C. Sitar (mcsitar@colostate.edu)

## Abstract

Collecting grain measurements for large detrital zircon age datasets is time-consuming, but a growing number of studies suggest such data are essential to understanding complex roles of grain size and morphology in grain transport and as indicators for grain provenance. We developed the colab\_zirc\_dims Python package to automate deep-learning-based segmentation and measurement of mineral grains from scaled images captured during laser ablation at facilities that use Chromium targeting software. The colab\_zirc\_dims package is implemented in a collection of ~~freely accessible, ready-to-highly interactive~~ Jupyter notebooks that can be run either on a local computer or installation-free via Google Colab. ~~These notebooks with~~also provide additional functionalities for dataset preparation and ~~for~~ semi-automated grain segmentation and measurement using a simple graphical user interface. Our automated grain measurement algorithm approaches human measurement accuracy when applied to a manually measured  $n = 5,004$  detrital zircon dataset, ~~but persistent errors.~~ Errors and uncertainty related to variable grain exposure necessitate semi-automated measurement for production of publication-quality ~~datasets.~~ We ~~measurements, but we~~ estimate that our semi-automated grain segmentation workflow will enable users to collect grain ~~measurements-measurement datasets~~ for large ( $n \geq 5,000$ ), applicable image datasets in under a day of work, ~~and we.~~ We hope that the colab\_zirc\_dims toolset allows more researchers to augment their detrital geochronology datasets with grain measurements.

## 25 ~~1~~ Introduction

~~Despite an increasing number of studies on the subject, the degree to which detrital geochronology datasets are affected by sample and mineral grain size remains unresolved. Several detrital zircon studies have documented substantial grain size-dependent mineral fractionation leading to biased detrital age spectra and erroneous provenance interpretations (e.g., Lawrence~~

et al., 2011; Ibañez-Mejía et al., 2018; Augustsson et al., 2018; Cantine et al., 2021), whereas several other studies have identified provenance-dependent grain size relationships in detrital samples but have found little evidence of age spectra biasing by selective transport processes (e.g., Muhlbauer et al., 2017; Leary et al., 2020, accepted). Because the number of studies characterizing grain size of detrital zircon datasets remains relatively small, especially compared to the number of studies employing detrital zircon geochronology, we likely lack the necessary volume and diversity of datasets to understand under which specific circumstances zircon transport processes will bias age spectra and interpreted provenance (Leary et al., accepted). A principal challenge in collecting such data has been that few automated approaches have been published, and manual collection of grain dimensions from large detrital datasets has been a barrier to such studies.

## 1 Introduction

Despite an increasing number of studies on the subject, the degree to which detrital geochronology datasets are affected by sample and mineral grain size remains unresolved. Several detrital zircon studies have documented substantial grain size-dependent mineral fractionation leading to biased detrital age spectra and erroneous provenance interpretations (e.g., Lawrence et al., 2011; Ibañez-Mejía et al., 2018; Augustsson et al., 2018; Cantine et al., 2021), whereas several other studies have identified provenance-dependent grain size relationships in detrital samples but have found little evidence of age spectra biasing by selective transport processes (e.g., Muhlbauer et al., 2017; Leary et al., 2020a, 2022). Because the number of studies characterizing grain size of detrital zircon datasets remains relatively small, especially compared to the number of studies employing detrital zircon geochronology, we likely lack the necessary volume and diversity of datasets to understand under which specific circumstances zircon transport processes will bias age spectra and interpreted provenance (Leary et al., 2022). A principal challenge in collecting such data has been that few automated approaches have been published (e.g. Scharf et al., 2022), and the time required to manually collect grain dimensions from large detrital datasets is a substantial barrier to widespread application of these methods (e.g. Leary et al., 2020a).

Zircon grains can be measured manually using analogue methods prior to LA-ICP-MS but doing so is prohibitively time consuming. Grains may also be imaged, characterized, and measured via scanning electron microscope before or after analysis, but this too incurs time and instrumentation costs that increase with sample size, and such analyses are not standard at most labs. Many LA-ICP-MS facilities using Teledyne-Photon machines laser ablation systems with proprietary Chromium (~~Teledyne Photon Machines, 2020~~)(Teledyne Photon Machines, 2020) targeting software save reflected light images of samples during analysis with scaling and shot location metadata files and provide these files to facility users. Images from these facilities may be full-sample mosaics captured prior to analyses or single, grain-centred per-shot images captured during ablation. The former are provided by the University of Arizona LaserChron Center (ALC) and the latter by the University of California, Santa Barbara (UCSB) Petrochronology Center. Many researchers who have not otherwise imaged their large-n detrital mineral datasets do have access to these files, and these can be used to locate and manually measure detrital mineral grains using the offline version of the Chromium targeting software (~~Leary et al., 2020~~)(Leary et al., 2020a).

Three limitations to manual grain measurement in Chromium (Leary et al., 2020a) are a) grains may be partially exposed or over-polished at the surfaces of epoxy mounts, so measurements are minimum, rather than true dimensions, b) this method is extremely time consuming, and c) this method can only produce one-dimensional (i.e., length) measurements. The first problem is inherent to reflected light images, but the latter two can be mitigated and solved, respectively, via automated two-dimensional grain-image segmentation and measurement of segmentation results. Deep learning methods, wherein training-optimizable models are used to algorithmically extract information from data (e.g., images) with minimal pre-processing (Alzubaidi et al., 2021; Alzubaidi et al., 2021), are at the cutting edge of accuracy in image segmentation and so allow grain image segmentation to be automated to a greater degree than other methods (e.g., thresholding).

~~We developed the colab\_zirc\_dims Python package, which contains code to automatically segment and measure mineral grains from Chromium-scaled LA-ICP-MS reflected light images using deep learning instance segmentation (i.e., where grains are treated as separate objects and distinguished from one another) models. Such models are computationally expensive to run and can be quite slow without a good graphics processing unit (GPU), so we implemented our code in Google Colab notebooks for maximally accessible public use. Google Colab is a free service that allows users to run Jupyter notebooks (i.e., Kluyver et al., 2016) on cloud-based virtual machines with high-end GPUs. Because its user interface is Colab notebook-based, colab\_zirc\_dims is not a per se application but a set of simplified, highly interactive scripts that rely on a backend of code in the colab\_zirc\_dims package. Deep learning-based techniques are increasingly geologic image segmentation tasks (e.g., fission track counting (Nachtergaele and De Grave, 2021), cobble measurement (Soloy et al., 2020), and photomicrograph grain segmentation (e.g., Bukharev et al., 2018; Jiang et al., 2020; Latif et al., 2022)), but the colab\_zirc\_dims package and processing notebooks represent, to the best of our knowledge, the first deep learning-based approach to per-grain detrital mineral separate measurement.~~

We developed the colab\_zirc\_dims Python package, which contains code to automatically segment and measure mineral grains from Chromium-scaled LA-ICP-MS reflected light images using deep learning instance segmentation (i.e., where grains are treated as separate objects and distinguished from one another) models. Such models are computationally expensive to run and can be quite slow without a good, code-compatible graphics processing unit (GPU). In order to maximize its accessibility, we implemented our code in Jupyter notebooks (i.e., Kluyver et al., 2016) that can be run either offline or online and installation-free using Google Colab (Sitar, 2022). Google Colab is a free service that allows users to run Jupyter notebooks on cloud-based virtual machines with variably high-end GPUs from the NVIDIA Tesla series (i.e., K80, T4, P100, and V100) that are allocated based on availability. Because its user interface is notebook-based, colab\_zirc\_dims is not a per se application but a set of simplified, highly interactive scripts that rely on a backend of code in the colab\_zirc\_dims package. Deep-learning-based techniques are increasingly applied to geologic image segmentation tasks (e.g., fission track counting (Nachtergaele and De Grave, 2021), cobble measurement (Soloy et al., 2020), and photomicrograph grain segmentation (e.g., Bukharev et al., 2018; Jiang et al., 2020; Latif et al., 2022)), but the colab\_zirc\_dims package and processing notebooks represent, to the best of our knowledge, the first deep-learning-based approach to per-grain detrital mineral separate measurement.

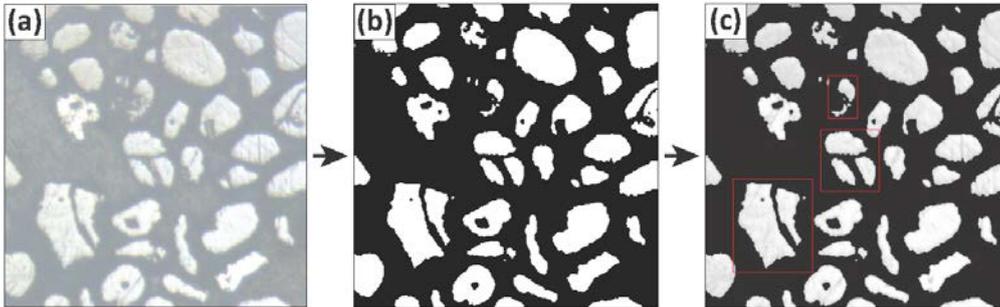
## 95 2 Established image segmentation techniques and related software

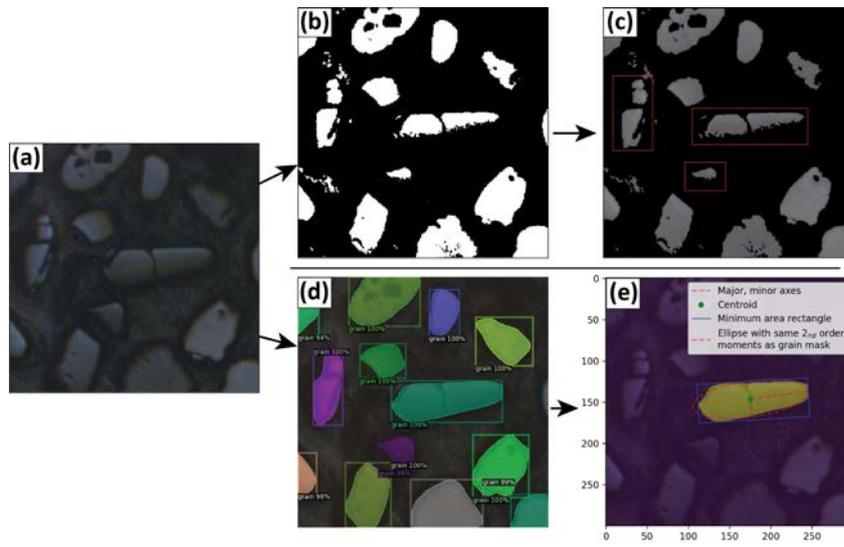
Automated segmentation of mineral grains in LA-ICP-MS images can be achieved with some success using relatively simple image segmentation techniques such as k-means clustering, edge detection, and intensity thresholding. ~~Otsu's thresholding method (i.e., Otsu, 1979), wherein image pixels are automatically segmented into background and foreground classes with Otsu's thresholding method (i.e., Otsu, 1979), wherein image pixels are automatically segmented into background and~~  
100 ~~foreground classes~~ via maximization of inter-class intensity variance, is particularly well-suited for reflected light images because mineral grains appear as a bright phase against an epoxy background (Fig. 1). Though grain segmentations produced through Otsu thresholding are often accurate, they tend to split single fractured grains into multiple sub-grains (~~Fig. 1c, A1~~) and can be wildly inaccurate where image artefacts affecting pixel intensity (e.g., anomalous bright spots; ~~Fig. A1~~) are present. These problems are common to automated segmentation techniques, and edge detection methods additionally contend  
105 with mis-segmentations along artefactitious edges where sub-image boundaries appear within larger, otherwise uniform mosaic images- (e.g., ~~Fig. A1~~). Because deep learning models can be optimized through training to ignore image artefacts and intra-grain fractures, they are likely the best available tool for achieving fully automated mineral grain segmentations with near-human accuracy.

Some existing software applications enable measurement of mineral grains in images with varying degrees of  
110 automation. The offline version of the Chromium LA-ICP-MS targeting application supports loading and viewing of scaled alignment images and shot locations; users can manually measure the axial dimensions of analysed grains using a ruler-like "measure" tool (Leary et al., 2020; Teledyne Photon Machines, 2020). The ZirconSpotFinder module of the MATLAB-based AgeCalcML application likewise supports loading and viewing of Chromium scaled LA-ICP-MS alignment images, but also  
115 implements semi-automated grain segmentation using user-selected thresholds, filtering of segmented grains by surface area, and export of area filtered shot lists (Sundell et al., 2020). AnalyZr, a new application designed specifically for measurement of zircon grains in images, combines Otsu thresholding with a novel boundary separation algorithm to automatically segment grains and allows users to edit the resulting segmentations before exporting automatically generated, grain-specific  
120 dimensional analyses (Scharf et al., 2022). Because AnalyZr supports loading of grain image .png files from any source with manual capture of image scale, it can be used to extract more detailed per-grain information (e.g., zonal area from cathodoluminescence images, unobscured grain dimensions from transmitted light images) than is obtainable using only reflected light images (Scharf et al., 2022). AnalyZr's manual scaling implementation (Leary et al., 2020a; Teledyne Photon  
125 Machines, 2020). The ZirconSpotFinder module of the MATLAB-based AgeCalcML application likewise supports loading and viewing of Chromium-scaled LA-ICP-MS alignment images, but also implements semi-automated grain segmentation using user-selected thresholds, filtering of segmented grains by surface area, and export of area-filtered shot lists (Sundell et al., 2020). AnalyZr, a new application designed specifically for measurement of zircon grains in images, combines Otsu thresholding with a novel boundary separation algorithm to automatically segment grains and allows users to edit the resulting segmentations before exporting automatically-generated, grain-specific dimensional analyses (Scharf et al., 2022). Analytical

spot identification and localization in AnalyZr is done manually through an interface that also allows input of spot-specific comments and qualitative internal grain zoning descriptors that persist into the program's exports (Scharf et al, 2022). Because AnalyZr supports loading of grain image .png files from any source with manual capture of image scale, it can be used to extract more detailed per-grain information (e.g., unobscured grain dimensions from transmitted light images) than is obtainable using only reflected light images (Scharf et al., 2022). AnalyZr's manual spot placement and scaling implementations and thresholding-based segmentation algorithm also, however, necessitate substantial human involvement in producing accurate grain segmentations and measurements. The colab\_zirc\_dims package and notebooks are likely better suited for rapid measurement of mineral grains in applicable (i.e., with Chromium-scaled images) large-n datasets due to their automated image loading, scaling, and generally accurate deep-learning-based automated segmentation capabilities.

**Figure 1. A visualization** of image thresholding segmentation using Otsu's method (Otsu, 1979) and its inherent problems in the context of reflected light detrital zircon grain images; (top row), and of the colab\_zirc\_dims segmentation and grain measurement process (bottom row). (a) An original, unaltered LA-ICP-MS reflected light image. (b) A binary image resulting from segmentation of the original image into foreground (white) and background (black) classes using Otsu's method. (c) The original image with "background" masked out using the binary image; fractured single grains that have been erroneously split into multiple grains are highlighted in red. Red highlights indicate single grains that have been erroneously eroded, segmented into multiple grains along fractures, or both. (d) The results (bounding boxes, probability scores, and masks) of instance segmentation of the original image using a Mask RCNN model (M-ST-C), as displayed by the Detectron2 "visualizer" module. (e) The resulting colab\_zirc\_dims verification image, scaled in  $\mu\text{m}$  and displaying the identified central grain mask (yellow), mask centroid (green), minimum-area circumscribing rectangle (blue), and ellipse with the same 2<sup>nd</sup> order moments as the grain mask along with its axes (red).





## 150 3 Methods

### 3.1 Dependencies

The `colab_zirc_dims` package was written in Python 3.8 and relies on some non-standard Python packages (Van Rossum, 2020). Pillow, OpenCV, and Matplotlib are respectively used for image loading, image display, and to create and save verification segmentation images; Matplotlib was additionally used to create figures for this manuscript (Umesh, 2012; Bradski, 2000; J. D. Hunter, 2007). NumPy is used for array operations and conversions, and pandas is used in some contexts for data organization and export (Harris et al., 2020; McKinney, 2010). The “measure” module of Scikit Image is used to produce unscaled dimensional analyses from segmented grain masks and to extract mask outlines for conversion into user-editable polygons (van der Walt et al., 2014). Interactivity in `colab_zirc_dims` processing notebooks is implemented using IPython (Pérez and Granger, 2007). Detectron2, which is a deep learning library that was developed by Facebook and is itself built on Google’s PyTorch, was used for model construction and training and is used to deploy models within `colab_zirc_dims` processing notebooks (Detectron2; Paszke et al., 2019).

Our notebook-based implementation The `colab_zirc_dims` package was written in Python 3.8 and relies on some non-standard Python packages (Van Rossum, 2020). Pillow and Matplotlib are respectively used for image loading and to create and save verification segmentation images; Matplotlib was additionally used to create figures for this manuscript (Umesh, 2012; J. D. Hunter, 2007). OpenCV (Bradski, 2000) is used to display images as well as to fit minimum-area circumscribing

rectangles to masks (e.g., Figs. 1e, A1c). NumPy is used for array operations and conversions, and pandas is used in some contexts for data organization and export (Harris et al., 2020; McKinney, 2010). The “.measure” module of Scikit-Image is used to produce unscaled dimensional analyses from segmented grain masks and to extract mask outlines for conversion into user-editable polygons (van der Walt et al., 2014). Interactivity in colab\_zirc\_dims processing notebooks is implemented using IPython (Pérez and Granger, 2007). Detectron2, which is a deep learning library that was developed by Facebook and is itself built on PyTorch, also developed by Facebook, was used for model construction and training and is used to deploy models within colab\_zirc\_dims processing notebooks (Detectron2; Paszke et al., 2019).

Local and online execution of the colab\_zirc\_dims package relies, obviously, on notebooks rely, respectively, on Jupyter and Google Colab to run. We recognize that Google Colab is Jupyter-style notebooks are an unconventional platform for final deployment of scientific computing algorithms and that #Google Colab in particular does have some significant disadvantages (e.g., runtimes will automatically disconnect if left idle for too long) versus deployment in a standalone, purpose-built local, or web-based application-based deployment. Nevertheless, we believe that Google Colab’s benefits in this use-case outweigh its disadvantages, especially with regards to accessibility. The colab\_zirc\_dims notebooks can be run using otherwise-expensive GPUs by anyone with a Google account, regardless of their local hardware or prior Python experience. We also mitigate potential connection-related issues by implementing automatic saving to Google Drive during online automated and semi-automated grain-image processing: if a user’s runtime disconnects, they can simply re-connect and resume work from the last sample processed before disconnection. The aforementioned timeout and connectivity problems will not affect the processing notebooks if they are run locally, and local notebook execution (i.e., Sitar, 2022, ‘Advanced Local Installation Instructions’) remains an option for users who are equipped with suitable hardware and either chafe against the constraints of Google Colab or are otherwise unable to access Google services.

## 3.2 Training-validation dataset

### 3.2 Deep learning models

#### 3.2.1 Mask RCNN

Mask RCNN is an instance segmentation (i.e., Figs. 2a, 2b) convolutional neural network (CNN) model architecture that was developed by Facebook and, when first released, achieved state-of-the-art results when trained and tested on the MS-COCO (“Common Objects in Context”) instance segmentation baseline dataset (He et al., 2018; Lin et al., 2015). Image inputs in Mask RCNN models initially pass through a backbone neural network, which extracts a feature map containing high-level information (e.g., possible objects) from the image (Ren et al., 2016; He et al., 2018). If the model includes a Feature Pyramid Network (FPN), this feature map will be enhanced via convolutional upsampling to improve resolution of small objects (Lin et al., 2016). The feature map is passed to a region proposal network, which generates variably-sealed anchor boxes encompassing the entire map before predicting whether each anchor box contains an actual (i.e., non-background) object and fitting bounding boxes to predicted regions of interest (ROIs) potentially containing objects (Ren et al., 2016; He et al., 2018). Lower-confidence ROIs are dropped at this stage if their degree of overlap with other possible ROIs for the same object is beyond a modifiable non-max suppression (NMS) threshold (Ren et al., 2016). The feature map is subsequently resampled for each remaining ROI by an “ROIAlign” layer to extract ROI-specific feature maps that are aligned pixel-to-pixel with the input image (He et al., 2018). Each

Formatted: Heading 2

Formatted: Heading 2

ROI-specific feature map is finally passed to both an object-detection head layer and a mask head network, which together produce the final model outputs: per-object classifications, fitted bounding boxes, and segmentation masks (He et al., 2018; Ren et al., 2016).

Components and branches of models available in Detectron2 are to a large extent modular and can be swapped with components or branches with different (e.g., newer) designs. We trained and tested Mask RCNN FPN models using two different backbone architectures: ResNet, and Swin (He et al., 2015; Liu et al., 2021). ResNet networks, which are standard for Detectron2 Mask RCNN implementations, use a variable number of residually connected blocks of batch-normalized convolutional layers to produce feature maps (He et al., 2015). ResNets are generally classified by their total number of convolutional layers (i.e., ResNet-101 has 101 layers); we tested Mask RCNN models with both ResNet-50 and ResNet-101 backbones (Table 1). Whereas ResNet is somewhat dated, the Swin-network architecture is near cutting-edge and achieved state-of-the-art results on several baseline datasets including MS-COCO when it was released last year (Liu et al., 2021). Swin models use transformer blocks to extract information (i.e., features) from iteratively upsampled and merged image patches within shifting self-attention windows (Liu et al., 2021). We trained and tested one otherwise standard Mask RCNN FPN model with an unofficial Detectron2 implementation (Ye et al., 2021) of Swin-T as a backbone (Table 1).

Figure 2. A visualization of the colab\_zirc\_dims deep learning-based mineral grain segmentation and measurement process. (a) The original image extracted from an ALC mosaic image, centred on a detrital zircon grain. (b) The results (bounding boxes, probability scores, and masks) of instance segmentation of the original image using a Mask RCNN model (M-R101-C), as displayed by the Detectron2 “visualizer” module. (c) The resulting colab\_zirc\_dims verification image, scaled in  $\mu\text{m}$  and displaying the identified central grain mask (yellow), mask centroid (green), and measured long and short grain axes (red).

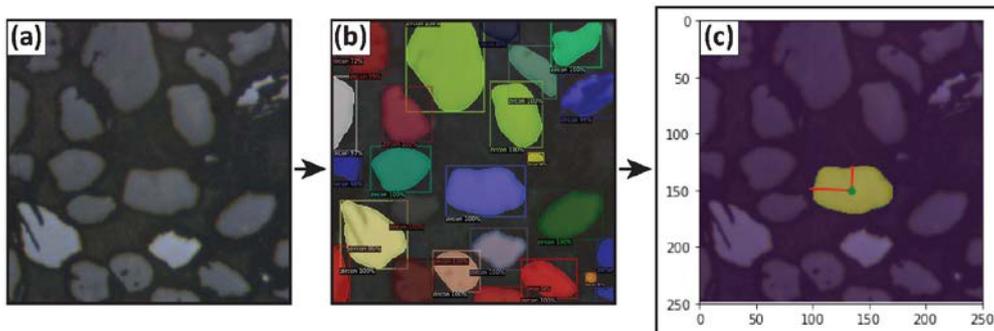


Table 1. Architectures, pre-training and training parameters, and number of iterations for trained model checkpoints picked for evaluation on the full Leary et al. (accepted) dataset.

We present “czd large”, a new training-validation dataset comprising 16,464 semi-automatically generated per-grain annotations in 1,558 LA-ICP-MS reflected light images of mineral grains (Table 1). Constituent images, which are sourced from both ALC and UCSB, were compiled via Chromium-metadata-informed (i.e., all images are non-overlapping in real-

world space) random selection. ALC source mosaic images (Table 1) were captured during analyses of detrital zircon from the Eagle and Paradox basins, USA; dates and Chromium-derived manual grain measurements resulting from these analyses were published by Leary et al. (2020a). UCSB images (Table 1) were captured during unpublished analyses of detrital zircon from units in east-central Nevada, USA. Automatic per-grain instance segmentations were generated using a Mask-RCNN Resnet-101 model trained on a smaller, manually annotated dataset compiled from the same sources (Table B1; Sitar, 2022, ‘Training Datasets’). These automatic segmentations were converted to the VGG image annotation format (Dutta and Zisserman, 2019) using a custom Python script, and annotations for every image were then manually reviewed and, where necessary, corrected or extended using the VIA Image Annotator (Dutta and Zisserman, 2019). Approximately 15% of the full dataset was split off into a validation subset via sample-stratified random selection (Table 1). We provide granular information (e.g., image sizes and scales, training versus validation set image and annotation distributions, etc.) about the dataset and a link to download it in the ‘Training Datasets’ subdirectory of our project GitHub page (Sitar, 2022).

**Table 1. A summary of the “czd\_large” dataset used to train the deep learning model presented in this manuscript for reflected light mineral grain segmentation. Please refer to Sitar (2022) for more in-depth information on the composition of the dataset.**

Source facility	Training set images	Validation set images	Training set annotations	Validation set annotations
ALC	1203	212	12923	2326
UCSB	121	22	1039	176
<b>Total</b>	<b>1324</b>	<b>234</b>	<b>13962</b>	<b>2502</b>
	<b>1558</b>		<b>16464</b>	

**Full-evaluated-model-name**

Abbreviated model name	Architecture	Base network	Pre-training image dataset	Training augmentation	Training iterations <sup>b</sup>	Model	MS COC	Yes	6000
				M-C	Mask RCN	ResNet-101	MS COC	Yes	6000
				M-R50-C	Mask RCN	ResNet-50	MS COC	Yes	6000
				M-R101-S	Mask RCNN	ResNet-101	MS COC	No	8000
				M-R50-S	Mask RCNN	ResNet-50	MS COC	Yes	4000

50_from_scratch_no_aug_s_4.0k	M-R50-S-NA	Mask RCNN	ResNet-50	None	No	4000
mask_rcnn_swint_7.0k	M-ST-C	Mask RCNN	Swin-T*	MS-COCO	Yes	7000
centermask2_4.0k	C-V-C	Center mask2	VoVNet V2	MS-COCO	Yes	4000

\*Unofficial Detectron2 implementation (Ye et al., 2021).

<sup>b</sup>Number of two-image training iterations completed for model selected for detailed evaluation on test dataset.

### 3.2.2 Centermask2

Centermask2 is a Detectron2-specific update to the Centermask instance segmentation-CNN-model architecture; Centermask is somewhat similar to Mask RCNN but differs in some important aspects (Lee and Park, 2020). In Centermask models, multi-scale feature maps extracted by backbone networks with FPN are passed directly to a Fully Convolutional One-Stage Object Detector (FCOS) model which identifies, classifies, and fits bounding boxes to objects without an intermediate (e.g., in Mask RCNN) anchor box prediction stage (Lee and Park, 2020; Tian et al., 2019). Object-specific masks are then predicted by a mask head which focuses on apparently information-rich pixels (i.e., spatial attention) from input images (Lee and Park, 2020). By reducing the number of steps involved in predicting and segmenting objects, Centermask models are able to out-perform Mask RCNN models on baseline datasets such as MS-COCO in accuracy, especially with regard to object bounding box predictions (Lee and Park, 2020). The standard backbone for Centermask models is VoVNetV2, which uses staged feature map downsampling (“one-shot aggregation”) with residual connections to predict features somewhat more accurately than ResNet (Lee and Park, 2020; Lee et al., 2019). We trained and tested one Centermask2 model equipped with an FPN VoVNetV2-99 (i.e., with 99 convolutional layers) backbone (Table 1).

### 3.2.3 Training

Deep learning models can be trained either from scratch with randomly initialized weights (i.e., training-optimizable parameters which determine model behaviour) or from a pre-trained baseline. Most of the weights in well-trained instance segmentation models will be broadly applicable to distinguishing objects (e.g., from 80 classes including “person”, “ear”, and “fork” in MS-COCO) from backgrounds, so pre-trained models can be adapted to segment new object classes by retraining only some (e.g., the last backbone convolutional layer) on new data (Lin et al., 2015; Alzubaidi et al., 2021). Such adaptation of pre-trained weights (i.e., “transfer learning”) generally allows retraining of models for new applications with smaller datasets, less training, or some combination thereof (Alzubaidi et al., 2021). We re-trained four models from MS-COCO pre-trained weights (Table 1) and trained three additional Mask-RCNN ResNet-FPN models from scratch (Table 1) to verify that pre-training was appropriate for achieving accurate grain segmentations (Detectron2; Lee and Park, 2020).

**Table 2. A summary of the dataset used to train the deep learning models presented in this manuscript for reflected light mineral grain segmentation.**

Set	Images	Zircon grains
Training	112 <sup>a</sup>	1555
Validation	31 <sup>b</sup>	480

<sup>a</sup> 18 from UCSB, 94 from ALC.

<sup>b</sup> 12 from UCSB, 19 from ALC.

We assembled a training dataset (with training and validation sub-sets; Table 2) of reflected light detrital zircon LA-ICP-MS images and manually segmented all mineral grains visible in the images using VIA Image Annotator (Dutta and Zisserman, 2019). ALC images (Table 2) were algorithmically extracted at varying, sample-dependent scales and resolutions (194 by 194 pixels to 398 by 398 pixels) from mosaic images captured during analyses of detrital zircon from the Eagle and Paradox basins, USA; dates and Chromium-derived manual grain measurements resulting from these analyses were published by Leary et al. (2020).

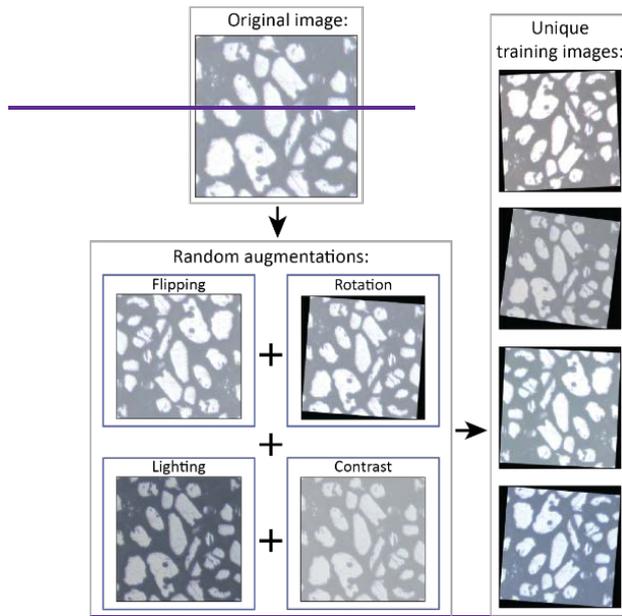
Some training and validation images contain likely detrital apatite grains in addition to zircon, and we segmented all visible mineral grains into a single “grain” class to avoid harming our models’ generalization abilities in the presence of varying image exposure and brightness levels. Models trained on “czd\_large” are consequently likely applicable to segmentation of all reflected-light bright-phase minerals but are unable to distinguish these minerals from one another. Both automatically and manually generated annotations are conservative with regards to interpreting grain extent; we only segmented areas where grains are exposed above the epoxy surface except in cases where larger subsurface extents are incontrovertibly apparent.

### 3.3 Deep learning models

Using the “czd\_large” dataset, we have trained several Detectron2-based instance segmentation models (i.e., configurations with trained weights) that can be applied in colab\_zirc\_dims processing notebooks. As of colab\_zirc\_dims v1.0.10, said models encompass several architectures and variations therein, including Mask-RCNN models with ResNet-FPN backbones, a Mask-RCNN model with a Swin-T backbone implemented using third-party code (Ye et al., 2021), and a Centermask2 model with a VovNetV2-99 backbone (Table B1). Given the rapid pace of progress in deep learning research and our own blundering yet continual progress in optimizing model hyperparameters for application in colab\_zirc\_dims, we expect that these models could be superseded by more performant models in the future. As such, we host our current models (i.e., configuration files and links to weights) and all explanatory information (i.e., training metrics, post-training evaluation metrics, and summary tables and diagrams) on a mutable ‘Model Library’ page within the project GitHub repository (Sitar, 2022). Users can refer to this page to learn more about the current selection of models, and to the linked Jupyter notebook files if they would like to train their own models using our training workflow. Models are loaded for application in local and Colab-based colab\_zirc\_dims processing notebooks through a dynamic selection and downloading interface; the current default model is a Mask-RCNN model with a Swin-T-FPN backbone (Table B1), which was selected due to its apparent low propensity for producing aberrantly over-interpretive segmentation masks (Sitar, 2022). This model is herein referred to a “M-ST-C” and was used to produce all measurements and segmentation images presented in the current study.

3.4UCSB images (Table 2) were captured during analyses of detrital zircon from units in east-central Nevada, USA and added to the dataset as-is, with respective resolutions and scales of 1280 by 1024 pixels and 0.43  $\mu\text{m}/\text{pixel}$ . Training and validation images were selected semi-randomly: most were picked at random using the Python 3.8 “Random” module (Van Rossum, 2020) and some additional images were hand-selected in order to ensure that models were exposed to common image artefacts and atypical grain morphologies during training. Although some training and validation images contain likely detrital apatite grains in addition to zircon, we segmented all visible mineral grains into a single “zircon” class to avoid harming our models’ generalization abilities in the presence of varying image exposure levels. Our models are consequently likely applicable to segmentation of all reflected light bright phase minerals but are unable to distinguish these minerals from one another. Because our dataset is small by deep learning standards, we used Detectron2-based random augmentations (Fig. 3) during training of all but one (control) model (Table 1) to ensure that models did not see the exact same image twice and so improve learning and mitigate potential overfitting.

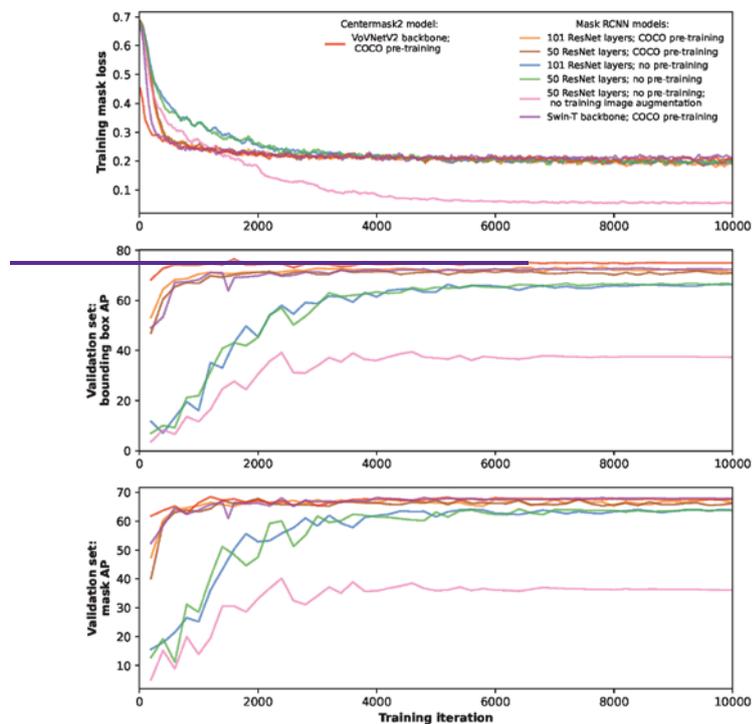
Figure 3. A graphical summary of the image augmentation method used during model training. Additional uniform (i.e., to 800 pixels for Mask RCNN ResNet models) and random (i.e., to 400 to 800 pixels for Mask RCNN Swin-T and Centermask2) short-edge resizing augmentations were also applied.



We trained each of our models in Google Colab using Detectron2 for at least 11,000 total two-image iterations with model-dependent learning rate schedules, all of which incorporated a 1,000 iteration warmup period and stepped 50% learning

rate reductions at variable (generally 1,000 iteration) intervals. MS-COCO style evaluations of the validation set (Fig. 4) were run every 200 iterations, and checkpoint model weights were saved at ~1000 iteration intervals for further evaluation against outside data. Losses during training were calculated using model default loss functions (He et al., 2018; Lee and Park, 2020; Detectron2); cross-entropy mask loss is a component of all model loss functions and so can be used for one to one inter-model training loss comparison (Fig. 4). We did not implement weight decay during training of any models except Mask RCNN with Swin-T backbone, where we used a starting weight decay parameter of 0.05 with 75% reduction at every stepped learning rate drawdown. Grain boundaries visible in reflected light images should not overlap, so we identified optimally reduced NMS thresholds (0.06 for Mask RCNN and 0.4 for Centermask2) for generating accurate single grain detections and masks and implemented these thresholds both during training and in colab\_zirc\_dims processing notebooks. We have followed the example of Nachtergaele and De Grave (2021) and include in the data repository for this manuscript our training dataset and Google Colab notebooks for training new Mask RCNN (ResNet and Swin-T backbone) and Centermask2 models to allow for future model improvements via training parameterization or implementation of new architectures.

**Figure 4. Plots showing values of training cross-entropy mask loss and validation set MS-COCO style AP metrics (mean average precision; calculated from multi-scale, multi-confidence-level intersection over union means) for bounding boxes and masks during model training.**



### 3.3 Dimensional analysis of mineral grains

The initial step in dimensional analysis of grains using colab\_zirc\_dims is standardized loading of grain images for segmentation such that differently formatted image-datasets can be processed using a single set of algorithms. Shot-centred single images (e.g., from UCSB) can be passed to models for segmentation as-is, but segmentation of grains from mosaic image datasets (e.g., from ALC) is performed on scaled, shot-centred sub-images extracted from mosaics using shot coordinate metadata. Grain-centred images are segmented by a deep learning model, and the resulting segmentations (e.g., [Fig. 2b](#), [Figs. 2d, A1c](#)) are passed to an algorithm that attempts to identify and return a “central” mask corresponding to the shot target grain LA-ICP-MS analysis (Fig. 2c). If no mask is found at the actual centre of the image, as may be the case in slightly misaligned images, the algorithm searches radially outwards until either a mask is identified or the central ~10% of the image has been checked. If a central grain is found, its dimensions are analysed using the “regionprops” function from the scikit image “measure” module (van der Walt et al., 2014), and theTo avoid erroneously returning significantly off-central (i.e., non-target)

grains, the algorithm is considered to have “failed” if it cannot find a grain mask after this search, and null values are returned for the spot instead of shape parameters. If a central grain is found, its dimensions are analysed using functions from OpenCV (Bradski, 2000) and the scikit-image “measure” module (van der Walt et al., 2014). The resulting measurements and properties are, where applicable, scaled from pixels to  $\mu\text{m}$  or  $\mu\text{m}^2$  using a Chromium-metadata-derived scale factor.

Successful grain-image processing by the colab\_zirc\_dims grain segmentation and measurement algorithm will return the following grain-mask properties: *area*, *convex area*, *eccentricity*, *equivalent diameter*, *perimeter*, *major axis length*, *minor axis length*, and *circularity*. *Circularity*, *long axis rectangular diameter*, *short axis rectangular diameter*, *best long axis length*, and *best short axis length*. Details on the derivation of all output grain-mask properties can be found on the ‘Processing Outputs’ section of the colab\_zirc\_dims GitHub page (Sitar, 2022), but some properties merit further discussion. *Circularity*, for instance, is calculated from scikit-image-derived area and perimeter measurements using Eq. (1); this is a notably simpler and likely less robust calculation than would be required for grain roundness (i.e., Resentini et al., 2018). Calculations for other returned properties can be found in the scikit image documentation (van der Walt et al., 2014). (i.e., Resentini et al., 2018).

**Equation 1:**

$$\text{Circularity} = \frac{4\pi * \text{Area}}{\text{Perimeter}^2}$$

Major and minor axis lengths are calculated from the moments of the grain mask image and reported axes thus correspond to “the length of the... axis of the ellipse that has the same normalized second central moments as the region” (van der Walt et al., 2014). These axial measurements will consequently fit exactly to perfectly elliptical and circular grain masks but may be more approximate in the cases of rectangular and irregularly shaped grains (e.g., Fig. 1e). Rectangular diameter measurements correspond to the long and short axes of the minimum area circumscribing rectangle (e.g., Fig. 1e) that can be fitted to a grain mask using the OpenCV minAreaRect function (Bradski, 2000). Minimum area rectangles will exactly fit to rectangular grain masks, but in the case of more equant grains may be grossly misaligned from the grain axes that a human researcher would interpret. The two types of calculated axial measurement parameters each have drawbacks. To split the difference, we implement “best” long and short axis measurement fields. These fields return either moment-based or rectangle-based axial measurements depending on whether each grain mask’s aspect ratio (i.e., moment-based long axis length divided by moment-based short axis length) is above or below an empirically chosen threshold of 1.8. Minimum-area bounding rectangles should trend towards coaxiality with moment-based axes with increasing aspect ratio, so rectangle-based measurements for grain masks with higher aspect ratios and moment-based measurements are returned for those with lower aspect ratios.

## 4 colab\_zirc\_dims Implementation

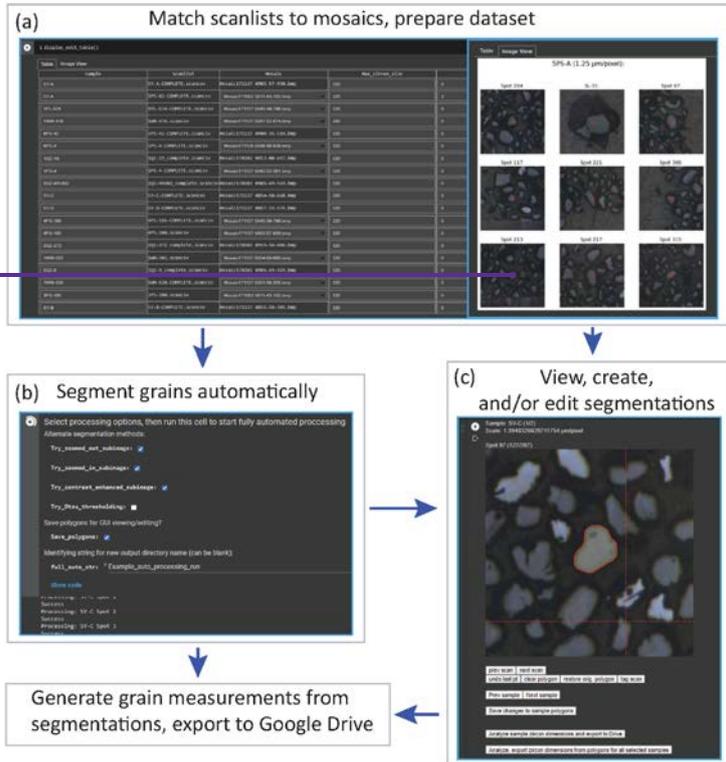
### 4.1 The colab\_zirc\_dims package

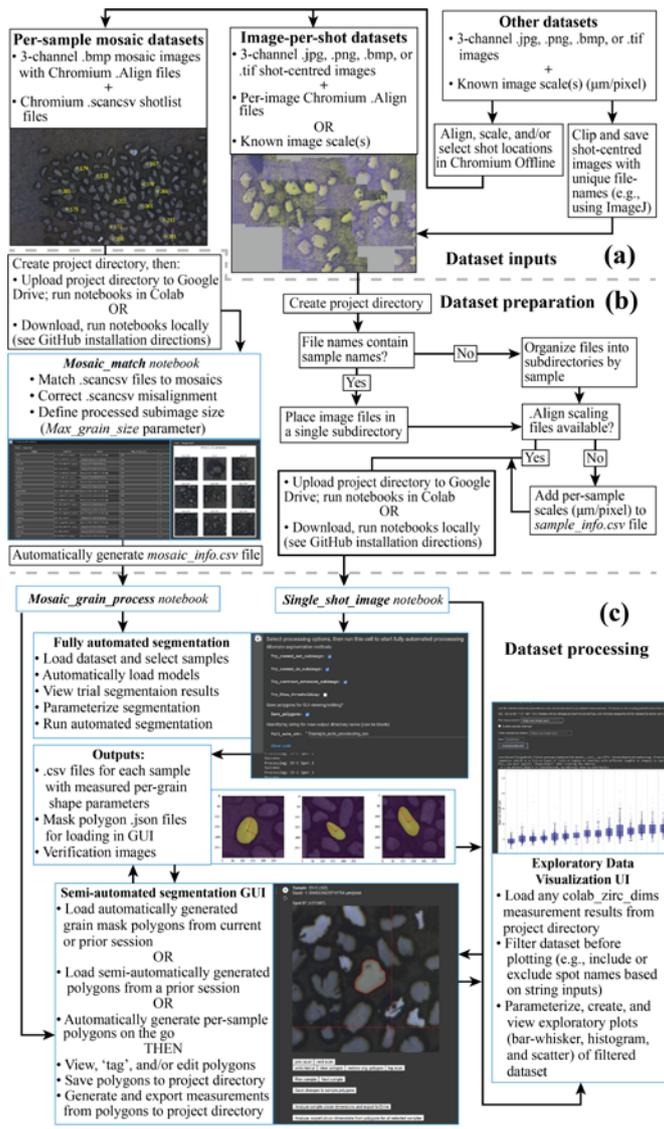
370 Code for loading and parsing Chromium alignment and shot list files, segmenting and measuring grains using deep learning  
models, and interacting with Colab notebooks using widgets is contained within the colab\_zirc\_dims package, which we have  
made available on the Python Package Index ([Python Package Index — PyPI, 2022](#)) for easy installation to local and virtual  
(e.g. [Python Package Index - PyPI, 2022](#)) for easy installation to local and virtual (i.e., Google Colab) machines. Some  
colab\_zirc\_dims modules (e.g., utilities for reading Chromium metadata files and basic segmentation functions) will work in  
375 local installations, but the package is only fully functional in Colab virtual runtimes without Detectron2 and other bulky  
dependencies, but these must be installed for full functionality.

### 4.2 Dataset organization

Before using colab\_zirc\_dims notebooks to automatically or semi-automatically measure grains, users must [upload](#) [set up a](#)  
[project folder containing](#) their dataset (i.e., image and metadata files) [as a](#). If users plan to use colab\_zirc\_dims in Google  
380 Colab, they must then upload their project folder [onto](#) Google Drive. ([Fig. 2A](#)). Required formats for colab\_zirc\_dims project  
folders are simple but necessarily differ slightly between dataset types (e.g., ALC mosaics or UCSB per-shot images), and  
they are thoroughly documented in the processing notebook for each type of dataset. Once a project folder has been created  
[and, optionally, uploaded](#) to a user's Google Drive, they can proceed either directly to Colab-notebook-based processing in  
the case of per-shot image datasets or to an additional, likewise notebook-based dataset preparation step in the case of mosaic  
385 image datasets- ([Fig. 2A](#)).

**Figure 52. A graphical summary of interfaces and workflow options available in colab\_zirc\_dims processing notebooks. (a) The  
interactive ALC mosaic dataset preparation (matching mosaic images to scansv shot lists, setting sub-image sizes, adjusting  
misalignment, etc.) interface. (b) The automated segmentation options interface for ALC datasets. (c) The GUI for semi-automated  
segmentation and editing of model- or human-generated segmentations**Tasks that are handled automatically or semi-automatically  
390 **by processing notebooks are shown in blue boxes. (a) A summary of possible dataset inputs to which can be processed or made  
processable with the provided notebooks. (b) Summary of the workflow for preparing datasets for fully or semi-automated  
segmentation. (c) Summary of possible workflows for automated or semi-automated grain measurement and for exploratory  
visualization of the resulting measurements.**





## 4.3 ~~Colab notebooks~~ Notebooks

### 4.3.1 Dataset preparation tools

As we note in Sect 3.34, segmentation and measurement of grains in mosaic image datasets requires extraction of shot-specific sub-images from larger mosaics using shot locations in corresponding .scanscv shot metadata files. Information on which mosaic file in a project folder matches which .scanscv file must consequently be provided by users for processing. Because  
400 deep learning models struggle to identify and segment grains when they cannot see all grain boundaries (e.g., if sub-images are smaller than grains), sub-image extraction also requires a user-provided, mosaic-specific sub-image size [parameter](#) (“[Max\\_grain\\_size](#)”) for accurate segmentations and measurements. Colab\_zirc\_dims processing notebooks read the aforementioned information from “mosaic\_info” .csv files stored in project folders. Though these “mosaic\_info” files can be created ~~and uploaded to project folders~~ manually, they can also be generated quickly and easily using the “*Mosaic\_Match*”  
405 colab\_zirc\_dims notebook (Fig. 5a2b) that we provide. The “*Mosaic\_Match*” notebook implements code that automatically finds matches between shot lists and mosaics in a project folder and allows ~~to~~ users to generate, modify, and export “mosaic\_info” tables (Fig 5a2b). Users can view sample shot locations and sub-images using a “Display” function (Fig. 5a2b), thus allowing interactive mis-alignment correction, adjustment of sub-image sizes, and, in cases where multiple possible  
410 mosaics could match a single .scanscv file, identification and selection of the correct mosaic from a dynamically populated dropdown menu. After exporting a “mosaic\_info” .csv file, users can proceed to fully or semi-automated segmentation and measurement of their dataset (Figs. 5b, 5e2b, 2c).

### 4.3.2 Fully automated segmentation and measurement

We provide ~~Colab~~ notebooks for automated and semi-automated processing of both mosaic image  
415 (“*Mosaic\_~~zirc~~grain\_process*”) and per-shot image (“*Single\_shot\_image\_~~zirc~~grain\_process*”) datasets; these notebooks are respectively currently set up to fully support processing of ALC and UCSB datasets but will likely work with datasets from other facilities sans-modification. The per-shot image notebook additionally supports loading and processing of any grain-centred reflected light grain images without Chromium scaling metadata, in which case users can provide custom per-sample scaling information in a .csv file or use a default scale of 1  $\mu\text{m}/\text{pixel}$ . [Researchers with datasets comprised of reflected light  
420 images that are not shot-centred and lack Chromium metadata can adapt \(i.e., Fig 2a\) their image datasets for use with colab\\_zirc\\_dims by either using Chromium Offline \(Teledyne Photon Machines, 2020\) to generate scaling and/or shot placement metadata or by manually cropping shot-centred images from mosaics \(e.g., using ImageJ’s “multicrop” function\). Such a workflow \(Fig. 2a\) will, however, bypass most of the automation in the colab\\_zirc\\_dims data loading process, and potential users are advised that collecting grain measurements using existing software \(i.e., AnalyZr; Scharf et al., 2022\) will  
425 likely be less arduous.](#)

Deep learning segmentation model weights are selected by users from a dropdown menu and downloaded to ~~Colab~~  
~~run~~ [times](#) virtual or local machines from an Amazon Web Services [S3](#) repository (provided by us) prior to model initialization

and processing. After weight file download and model initialization, users can select options for automated processing (Fig. 5b2c). These options include whether to attempt segmentation with various alternate methods (e.g., zooming out slightly or, 430 increasing image contrast before reapplying the model, or, as a last resort, using Otsu thresholding) if segmentation is initially unsuccessful, and whether to save polygons approximating model-produced masks for viewing or modification in the colab\_zirc\_dims graphical user interface (GUI) (Fig. 5b2c). During automated processing, per-grain dimensional analyses (Sect. 3.3) in per-sample .csv files along with are saved and exported to the user's project folder (Fig. 2c) alongside verification mask image .png files (e.g., Fig. 2e) are saved and exported to the user's Google Drive (Figs. 1e, A1c).

#### 435 4.3.3 Colab Notebook-based GUI for semi-automated segmentation and measurement

Because our deep learning models cannot fully match human accuracy when segmenting grains, we provide a simple, Colab-based GUI (Fig. 5c) extended from code in the Tensorflow Object Detection API (TensorFlow Developers, 2022). We provide 440 a simple, Colab-based GUI (Fig. 2c) extended from code in the Tensorflow Object Detection API (TensorFlow Developers, 2022) that allows users to view, modify, and save polygon-based grain segmentation masks. These polygon-masks are can either be loaded from a previous automated or GUI-based processing session or generated on-the-fly on a per-sample basis. After viewing or re-segmenting part or all of a dataset, users can send their grain segmentations for measurement and export (Sect. 4.3.2); grain dimension exports from the GUI will include additional tags indicating whether each grain was segmented by a human or by a deep learning model.

#### 4.3.4 Notebook-based exploratory data visualization interface

445 We do not provide any tools for assessing relationships between grain size or shape and age. Our processing notebooks do, however, include a simple interface that allows users to interactively load and filter (e.g., by scan name) colab\_zirc\_dims measurement data from their project folder before visualizing said data using parameterizable bar-whisker, histogram, and scatter plots (Fig. 2c).

## 5 Accuracy evaluations

450 We assessed the accuracy of our segmentation models by comparing a manually generated grain-dimension dataset (Leary et al., 2022) to automatically processing the Leary et al. (accepted) mosaic image dataset-generated grain dimensions from the same samples measured using all trained segmentation models with colab\_zirc\_dims. The test dataset from Leary et al. (2022) consists of samples collected from late Palaeozoic strata exposed across Arizona, USA. These samples were deposited in the same orogenic system—the Ancestral Rocky Mountains—as the Leary et al. (2020a) training dataset, and the grain ages and 455 depositional environments are largely similar. The test dataset is unrelated to the training dataset from UCSB (see above). The full dataset was automatically processed using model M-ST-C and pure Otsu thresholding via the colab\_zirc\_dims “Mosaic\_Process” notebook and comparing the resulting automated *best long axis length* and *best short axis length*

measurements ~~with~~ were compared to the manual (measured with the Chromium “Measure” tool) per-grain axial measurements from the same dataset. We picked “best” model checkpoints (Table 1) where models achieved apparent local maxima in validation accuracies (i.e., Fig. 4) and local minima in various For an n=301, sample-stratified random sub-sample of the Leary et al. (2022) dataset, colab\_zirc\_dims measurements of manual segmentation masks generated using the colab\_zirc\_dims semi-automated measurement error metrics (e.g., failure rate and absolute long axis error; Table 3). GUI were also evaluated.

**Table 32. Evaluation of error versus in colab\_zirc\_dims ‘best axis’ length measurements, with human measurements for automated in the Leary et al. (2022) dataset used as ‘ground truth’. For the full dataset (top), measurements produced by deep learning models on the Leary et al. (accepted) dataset, with fully automated segmentation via Otsu’s thresholding method as a baseline. The (using model M-ST-C) are compared against a baseline of Otsu thresholding. For the n=301 sample-stratified random subsample (bottom), measurements resulting from automated segmentation by model M-ST-C are compared to those resulting from new manual segmentations of the dataset using the colab\_zirc\_dims semi-automated processing GUI. Per-dataset best results on each metric are shown in bold type.**

Dataset	Model / method	n <sup>a</sup>	Failure rate <sup>b</sup> (%)	Average error <sup>c</sup> (μm)		Average absolute error <sup>d</sup> (μm)		Average error <sup>e</sup> (%)		Average absolute error <sup>f</sup> (%)		≥ 20% absolute error rate <sup>g</sup> (%)	Grain extent underestimate rate <sup>h</sup> (%)	Average segmentation time per image <sup>i</sup> (s)	
				Long axis	Short axis	Long axis	Short axis	Long axis	Short axis	Long axis	Short axis				
	M-ST-C	5003	<b>0.02</b>	<b>-2.3</b>	<b>-1.2</b>	<b>5.7</b>	<b>4.3</b>	<b>-2.06</b>	<b>-1.41</b>	<b>7.28</b>	<b>8.57</b>	<b>7.64</b>	<b>9.85</b>	<b>10.91</b>	<b>0.114</b>
Full <sup>†</sup>	Otsu thresholding	5003	<b>0.02</b>	-6.1	-5.7	10.1	7.8	-7.41	-10.94	13.03	15.82	18.19	25.78	27.00	<b>0.011</b>
Random sub-sample <sup>k</sup>	M-ST-C	301	<b>0.0</b>	<b>-2.9</b>	<b>-1.2</b>	<b>6.4</b>	<b>4.5</b>	<b>-2.71</b>	<b>-1.58</b>	<b>7.95</b>	<b>9.09</b>	<b>8.64</b>	<b>9.97</b>	10.30	<b>0.137</b>
	Manual segmentation <sup>l</sup>	301	<b>0.0</b>	3.6	3.9	6.4	4.8	5.84	8.73	8.62	10.65	8.64	14.29	1.66	~20

<sup>a</sup> Number of scan-images within dataset where a “central” grain mask could be identified with confidence ≥ 70%.

$$b \ 100 * \left( \frac{n_{total} - n}{n_{total}} \right)$$

$$c \ 1/n \sum_{i=1}^n (axis_{measured})_i - (axis_{Leary et al.})_i$$

$$d \ 1/n \sum_{i=1}^n |(axis_{measured})_i - (axis_{Leary et al.})_i|$$

$$e \ 100 * \frac{1}{n} \sum_{i=1}^n \frac{(axis_{measured})_i - (axis_{Leary et al.})_i}{(axis_{Leary et al.})_i}$$

$$f \ 100 * \frac{1}{n} \sum_{i=1}^n \left| \frac{(axis_{measured})_i - (axis_{Leary et al.})_i}{(axis_{Leary et al.})_i} \right|$$

$$g \ 100 * \frac{\text{number of grains with } \% \text{ error on either axis} \geq 20\%}{n}$$

$$h \ 100 * \frac{\text{number of grains with error on either axis} \leq -20\%}{n}$$

<sup>i</sup> Average time for model/method to successfully segment an image and return a measurable mask. Actual per-image processing times will be higher due to additional automated mask measurement and verification image saving time. Measured in Colab notebook with NVIDIA T4 GPU.

<sup>†</sup> The full Leary et al. (2022) dataset, with 5004 valid measurements.

<sup>k</sup> A sample-stratified random subsample of 301 measured grains from the Leary et al. (2022) dataset.

Formatted: Font: Times New Roman

Formatted: Font: Times New Roman

<sup>1</sup> By the first author, using the colab [zirc\\_dims](#) semi-automated segmentation GUI in Google Colab.

Model / method	n <sup>b</sup>	Failure rate (%)	Average <sup>a</sup> error (μm)				Average <sup>a</sup> absolute error (%)				Grain underestimate rate (%)	Average <sup>a</sup> extent segmentation time per spot <sup>d</sup> (s)		
			Long-axis	Short-axis	Long-axis	Short-axis	Long-axis	Short-axis	Long-axis	Short-axis				
M-R101-C	500 <sup>c</sup>	0.02	-1.2	-1.0	6.1	4.4	-0.45	-0.93	7.96	8.83	8.20	9.85	10.11	0.119 <sup>e</sup>
M-R50-C	4992	0.24	-2.0	-1.7	6.2	4.5	-1.47	-2.38	8.04	8.98	8.21	10.94	11.52	0.072 <sup>e</sup>
M-R101-S	4931	1.46	-2.2	-1.2	7.7	5.7	-1.31	-0.87	9.75	11.45	11.60	15.43	15.37	0.107 <sup>e</sup>
M-R50-S	4988	0.32	-3.6	-2.3	7.6	5.5	-3.13	-3.45	9.46	10.91	11.65	15.92	17.46	0.087 <sup>e</sup>
M-R50-S-NA	4749	5.10	-2.6	-0.2	13.1	8.6	1.01	17.1	18.16	23.23	23.08	21.75	21.75	0.108 <sup>e</sup>
M-ST-C	49930	2.2	-2.7	-2.4	6.0	4.5	2.79	-4.03	7.71	8.96	8.45	11.40	13.18	0.129 <sup>e</sup>
C-V-C	4998	0.12	-0.7	-0.7	6.0	4.4	0.12	-0.27	7.85	9.00	7.40	9.76	9.80	0.114 <sup>f</sup>
Otsu thresholding	5004	0.00	-3.4	-4.0	11.9	9.7	-3.27	-5.50	16.02	21.69	18.1	26.84	27.34	0.020

<sup>a</sup> Arithmetic mean.

<sup>b</sup> Number of scan images out of a total of 5004 where a "central" grain mask could be identified with confidence  $\geq 70\%$  for Mask-RCNN models or  $\geq 50\%$  for Centermask2 (C-V-C) model.

<sup>c</sup> Proportion of grains where new measurements differ by  $\leq 20\%$  along long and/or short axes from published (Leary et al., accepted) measurements.

<sup>d</sup> Time for model/method to segment (or fail to segment) an image and return a measurable mask. Actual per-spot processing times will be higher due to additional automated mask measurement and image saving time.

<sup>e</sup> Measured in Colab Notebook with NVIDIA T4 GPU.

<sup>f</sup> Measured in Colab Notebook with NVIDIA P100 GPU.

### 5.1 Impact Machine error

Otsu thresholding as implemented in colab [zirc\\_dims](#) is a reasonably performant baseline segmentation method and apparently produces dimensionally accurate masks for the majority of grains in the Leary et al. (2022) dataset (Table 2). Our default

Formatted: Font: Calibri

Inserted Cells

Formatted: Font: Calibri

Formatted: Font: Calibri

Inserted Cells

Deleted Cells

Deleted Cells

Deleted Cells

Merged Cells

Formatted: Caption, Left

Formatted: Font: Calibri

Formatted: Font: Calibri

Inserted Cells

Inserted Cells

Inserted Cells

Inserted Cells

Formatted: Font: Calibri

Inserted Cells

Deleted Cells

Merged Cells

Formatted: Caption, Left

Inserted Cells

Inserted Cells

Formatted: Default Paragraph Font, Font: Times New Roman

Formatted: Font: Times New Roman

Formatted: Normal

Formatted: Font: Times New Roman

475 model architecture and training parameters, however, significantly outperforms the baseline method of Otsu thresholding in every metric except for speed (Table 2). Given that segmentation time for M-ST-C is still a fraction of a second (Table 2) when run on a GPU-equipped computer, deep-learning-based instance segmentation appears to be superior for producing high-quality segmentation masks from reflected light images. The Leary et al. (2022) image dataset is also mostly free of artefacts (e.g., Fig. A1), and we expect that the gulf in accuracy between the two methods would widen if evaluated on a lower-quality dataset.

Formatted: Font: Times New Roman

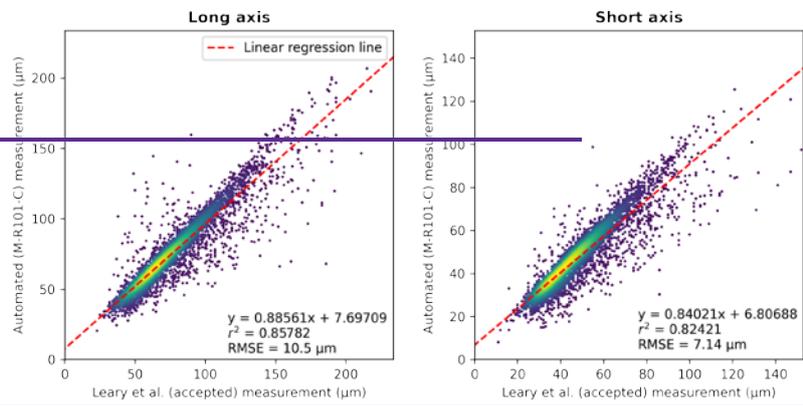
480 Some of our results are unsurprising: models trained from scratch with training image augmentation fail to match pre-trained models in validation set segmentation (Fig. 4) and full dataset (Leary et al., accepted) measurement accuracies (Table 3) despite converging on the same training loss values (Fig. 4). Our model trained from scratch without training image augmentation (M-R50-S-NA) additionally underperforms pure Otsu's method thresholding in segmentation measurement accuracy in several metrics with a concurrently high rate of failure to recognize grains in images (Table 3), likely due to poor generalization ability. 485 Pre-training and training image augmentation are essential to training models that accurately and consistently segment mineral grains.

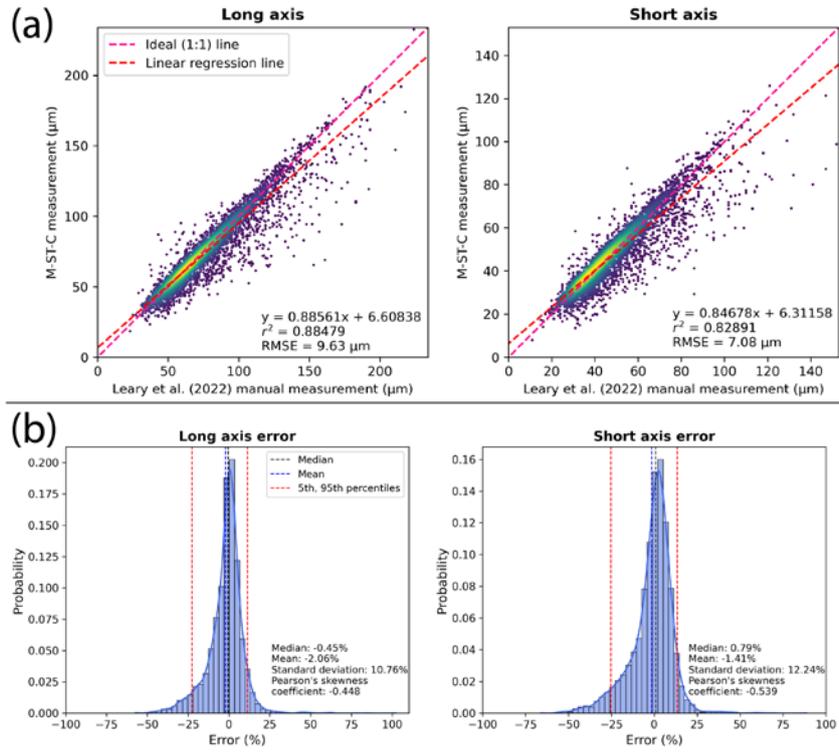
Formatted: Font: Times New Roman

Our three best performing models on the full dataset are M-R101-C, C-V-C, and M-ST-C, and we find that M-R50-C produces faster but somewhat less accurate measurements (Table 3). C-V-C achieves the highest measurement accuracy in most absolute error metrics as well as impressive, though likely not particularly robust (e.g., M-R50-S-NA; Table 3),  $< 0.5\%$  average long and short axis errors. In combination with its high validation bounding box AP (Fig. 4), this suggests that Centermask2's anchor free detection algorithm allows C-V-C to predict grain extents prior to segmentation with greater accuracy than Mask RCNN models (Table 3). M-ST-C notably fails to outperform our other models by most metrics (Table 3) despite its state-of-the-art Swin-T backbone, and we attribute this to our small training dataset; the basic Swin architecture tends to perform best when given large amounts of training data (Lee et al., 2021). We expect that all our models could improve 495 with more training data and that Centermasks2 and Mask RCNN models with Swin backbones could significantly outperform Mask RCNN ResNet models given a much larger training dataset or one that requires models to distinguish between grains of different minerals. Although it is slightly less accurate overall than C-V-C, M-R101-C has the lowest grain identification failure rate (Table 3) and we observe that it can accurately segment grains from artefactitious (e.g., very blurry) images where other models fail. This suggests that model M-R101-C has the best generalization ability of our trained models, and, though 500 users are encouraged to experiment with different models, we have made M-R101-C the default colab\_zirc\_dims segmentation model to maximize potential applicability to images from LA-ICP-MS facilities not represented in our training or test datasets.

## 5.2 Machine error

505 Figure 6. Automated (M-R101-C) and manual (Leary et al., accepted) Figure 3. Plots displaying error distributions when comparing measurements produced by automated (M-ST-C) colab\_zirc\_dims segmentation against manual measurements (i.e., Leary et al. 2022). (a) Automated (y-axis) versus manual (x-axis; Leary et al., 2022) measurement plots for long and short grain axes with linear regression lines plotted and gaussian KDE density shown via heatmap. Root mean squared error (RMSE) is shown at the bottom-right of each plot. (b) Histogram-KDE plots showing error distributions along long and short axes. Statistical information is shown at the bottom right of each plot.





510

Per-grain automated (M-R101ST-C) measurements for the full Leary et al. (accepted2022) dataset generally hew close to ground-truth measurements but skew slightly negative with a significant number of datapoints plotting well below the 1:1 measured versus ground truth (i.e., Leary et al., 2022) line (Fig. 63a). The apparent dominant cause of this negative skew (i.e., Equation 2; Fig. 3b) is under-segmentation of grains that are incompletely exposed at the surface of epoxy mounts but whose full grain areas are interpretable by humans from “shadows” visible in the (mostly) reflected light images (Fig. 74). We did not train our modelsmodel to interpret beyond clearly visible grain boundaries and they consequently fail to reproduce human measurements for these grains, but models could likely might be trainable to do so without diminished accuracy on “normal” grains given training on a larger, more interpretively segmented training dataset. Positive measurement errors are relatively rare (Fig. 6Figs. 3A, 3B) but are probably mainly attributable to segmentation masks that merge different grains (Fig. 74). Failure to identify the correct “central grain” in images (Fig. 74) is likewise rare but may cause positive, negative, or negligible measurement error depending on the respective sizes of the target and mistakenly identified grains. Cases where

515

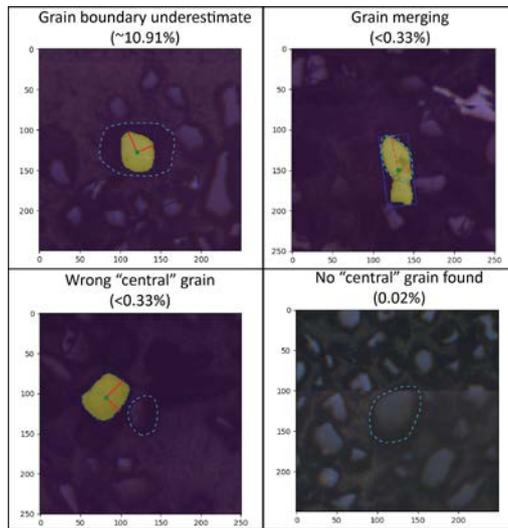
520

no grain could be identified are exceedingly rare (Table 32, Fig. 74) and do not contribute directly to measurement error but, like all identified errors, necessitate manual re-segmentation of grains for production of accurate measurements.

**Equation 2:**

$$\text{Pearson's skewness coefficient} = \frac{3(\text{mean} - \text{median})}{\text{standard deviation}}$$

Figure 74. Examples of automated (M-R191ST-C) segmentation mask error modes with estimated occurrence rates, with axes scaled in  $\mu\text{m}$  and correct grain segmentations outlined in light blue. Rates for “grain boundary underestimate” and “no central grain found” errors are estimated from analysis of the entire Leary et al. (accepted2022) dataset (i.e., Table 3), and “2). No “grain merging” and “wrong central grain” rates are based on errors were identified in a manual error identification and classification in an review of the n = 301 sample of the full dataset (i.e., Table 4), and their occurrence rates are estimated from their non-appearance therein.

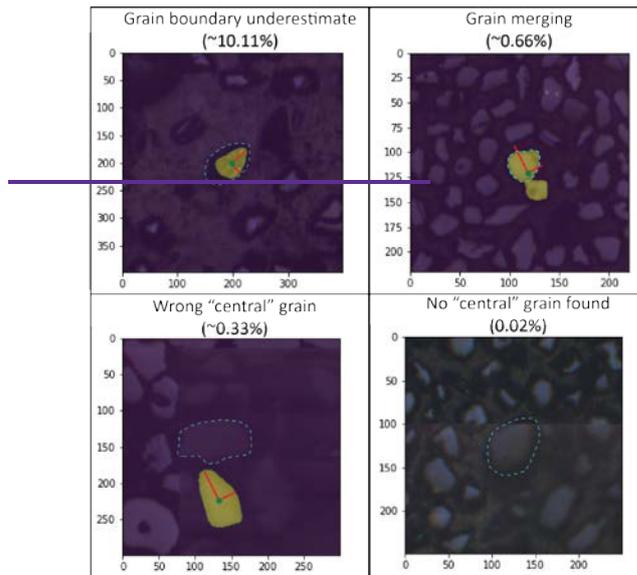


Formatted: Centered

## 5.2 Human error

Automated measurement error metrics (e.g., Table 2) likely encompass some error that would be present even if grains were manually segmented, due to differential interpretations of grain areas between researchers, especially in samples with underexposed grains. In the n = 301, randomly picked sample-stratified grain subsample from the Leary et al. (2022) dataset, we find that our default automated segmentation model (M-ST-C) achieves similar axial measurement absolute error metrics to the first author (M.S.) of this manuscript (Table 2). Though apparently mostly free of interpretive grain extent underestimates, the first author’s measurements skew higher than dataset measurements (Table 2). Apparent over-interpretations of grain extents by the first author likely reflect different image display conditions (e.g., higher zoom and different contrast) during

manual re-segmentation versus those present during collection of dataset measurements. Various features of colab\_zirc\_dims, namely automated segmentation of most grains and uniform image display conditions during manual segmentation of other grains, may enhance grain measurement dataset reproducibility in addition to collection speed.



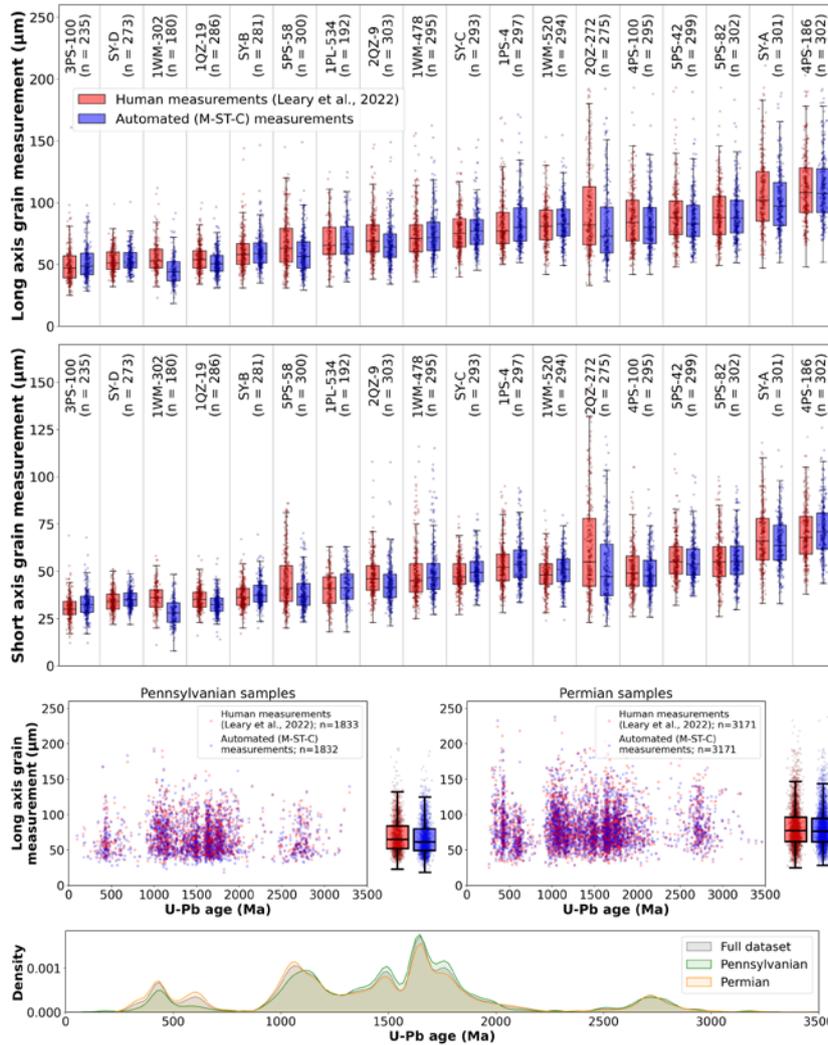
545 5.3 Impact of grain exposure

We find that automated processing using colab\_zirc\_dims and our default model (M-R101-ST-C) can approximately reproduce aggregate long and short grain axis length distributions for most samples in the Leary et al. (aeecepted2022) mosaic image and measurement dataset (Fig. 85). Systemic negative errors along both grain axes are concentrated within four samples (1WM-302, 5PS-58, 2QZ-9, and 2QZ-272; Fig. 85). We found that grains in these samples were consistently underexposed above mount surfaces and that "grain extent underestimate" (Table 3; Fig. 7) segmentation errors were as a result sufficiently common to negatively impact sample axis length distributions. Based on these results, we believe that fully automated measurement using colab\_zirc\_dims is a viable method for rapid approximation of grain size distributions in optimal (e.g., with well-exposed grains) samples. Due to low but significant segmentation error rates (Fig. 7), manual segmentation verification and correction (i.e., semi-automated measurement) is, however, 2; Fig. 4) segmentation errors were as a result common enough to negatively impact sample axis length distributions. Because these images are of sufficiently high quality that subsurface grain extents were interpretable by Leary et al. (2022), and because model M-ST-C generally only segments grain areas above resin surfaces, errors in these samples can also be used as a proxy for dimensional data loss from using reflected light versus transmitted light

555

560 images to measure shapes of very poorly exposed grains in cases where reflected light images do not reveal any information  
about subsurface grain extents (Sect. 1; Leary et al., 2020a). In the worst-evaluated sample, 1WM-302 (n=180), M-ST-C  
565 produces axial measurements that underestimated manually measured grain axes by at least 20% 66.6% of the time, with  
average grain measurement errors of -18.0% and -22.0% along long and short axes, respectively. Treating these automatically  
generated axial measurements as ground truth data could result in significantly flawed analysis of relationships between grain  
size and age. Such shape parameter underestimates present only a minor (though potentially time-consuming) problem for  
570 colab\_zirc\_dims users with poorly exposed grains whose actual areas are still interpretable by humans (e.g., in the case of  
1WM-302); erroneous segmentation masks can simply be corrected manually using the GUI. Users who observe that their  
mounted crystals are both very poorly exposed and invisible below the resin surface in their reflected light images, though,  
may consider re-imaging their samples using transmitted light and then measuring grains using a different program (e.g.,  
AnalyZr) to avoid collecting flawed data. Researchers should consider excluding grain mounts that appear heavily over-  
polished from their datasets, as accurate two-dimensional grain dimensions will not be resolvable under any lighting  
conditions.

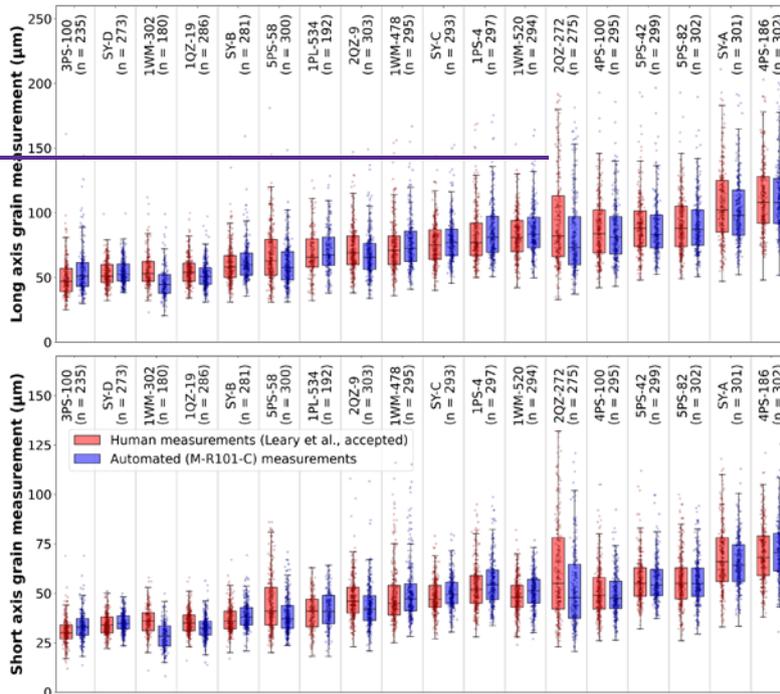
575 **Figure 5. Top: sample-by-sample boxplot comparison of human (Leary et al., 2022) and automated (M-ST-C) measurements along  
long and short grain axes. Below: additional scatter and bar-whisker plots showing relationships between human and automated  
grain long axis length measurements and U-Pb age, with samples binned by depositional period. Bottom: a KDE plot of detrital  
zircon U-Pb ages in the Leary et al. (2022) dataset. Boxplot boxes extend from Q1 to Q3, and whiskers extend from  $Q1 - 1.5 * (Q3 - Q1)$   
to  $Q3 + 1.5 * (Q3 + Q1)$ ; sample medians are indicated by black horizontal lines within each box.**



## 6 Viability of fully automated measurement

Due to low but significant segmentation error rates (Fig. 4) stemming almost entirely from poor grain exposure, we believe that manual segmentation verification and correction (i.e., semi-automated measurement) is necessary for production of publication-quality grain measurement datasets. Assuming time requirements of 2635 minutes total to automatically generate segmentation masks, one second per grain to manually check masks, and 20 seconds to correct each mis-segmentation, and, conservatively (Fig. 74), that 15% of grains must be re-segmented via GUI, we estimate that it would take about six hours to semi-automatically collect zircon grain measurements for the full (n = 5,004) Leary et al. (accepted2022) dataset using colab\_zirc\_dims.

Figure 8. A sample-by-sample boxplot comparison of human (Leary et al., accepted) and automated (M-R101-C) measurements along long and short grain axes. Boxes extend from Q1 to Q3, and whiskers extend from  $Q1 - 1.5 * (Q3 - Q1)$  to  $Q3 + 1.5 * (Q3 - Q1)$ ; sample medians are indicated by black horizontal lines.



### 5.3 Human error

Automated measurement error metrics (e.g., Table 3) likely encompass some error that would be present even if grains were manually segmented, due to differential interpretations of grain areas between researchers, especially in samples with under-exposed grains. In an  $n = 301$ , randomly picked sample proportionate grain subsample from the Leary et al. (accepted) dataset, we find that our default automated segmentation model (M-R101-C) achieves similar axial measurement absolute error metrics to the first author (M.S.) of this manuscript (Table 4). Though free of clear mis-segmentations (e.g., grain merging; Fig. 7) and mostly free of interpretive grain extent underestimates, the first author's measurements skew higher than dataset measurements (Table 4). Apparent over-interpretations of grain extents by the first author likely reflect different image display conditions (e.g., higher zoom and different contrast) during manual re-segmentation versus those present during collection of dataset measurements. Various features of colab\_zirc\_dims, namely automated segmentation of most grains and uniform image display conditions during manual segmentation of other grains, may enhance grain measurement dataset reproducibility in addition to collection speed.

Table 4. A comparison between human (first author) and automated (model M-R101-C) measurement errors in an  $n = 301$ , sample-proportionate random subsample of the Leary et al. (accepted) image-measurement dataset. The best results for each metric are shown in bold type.

Segmented by	n	Average* error ( $\mu\text{m}$ )		Average* absolute error ( $\mu\text{m}$ )		Average* error (%)		Average* absolute error (%)		RMSE (%)		$\geq 20\%$ absolute error rate (%)		Grain extent underestimate rate <sup>b</sup> (%)
		Long axis	Short axis	Long axis	Short axis	Long axis	Short axis	Long axis	Short axis	Long axis	Short axis	Long axis	Short axis	
Human <sup>c</sup>	304	4.0	3.9	<b>6.6</b>	4.8	6.19	8.70	8.83	10.65	<b>12.55</b>	14.06	<b>8.64</b>	14.29	<b>1.66</b>
M-R101-C	304	<b>-1.8</b>	<b>-1.2</b>	6.9	<b>4.6</b>	<b>-1.13</b>	<b>-1.37</b>	<b>8.56</b>	<b>9.10</b>	13.33	<b>12.87</b>	9.30	<b>9.30</b>	10.30

\*Arithmetic mean.

<sup>b</sup>Proportion of grains where new measurements differ by  $\leq 20\%$  along long and/or short axes from published (Leary et al., accepted) measurements.

<sup>c</sup>Manual grain segmentation by the first author using the colab\_zirc\_dims GUI; measurements derived from the resulting masks.

We also believe, however, that fully automated measurement using colab\_zirc\_dims is a viable method for rapid approximation of grain dimensions in optimal samples (i.e., with well-exposed grains) as well as in larger datasets where the majority of samples have well-exposed grains. Meaningful relationships between grain dimensions and age appear to be resolvable solely based on fully automated measurement of such datasets. Leary et al. (2022) used zircon grain-dimension data to reinterpret the provenance and transport mechanism of 500-800 Ma zircons within the Pennsylvanian-Permian Ancestral Rocky Mountains system in southwest Laurentia. Based on the arrival of dominantly small ( $< 60 \mu\text{m}$ ), 500-800 Ma zircons in that study area at the Pennsylvanian-Permian boundary, Leary et al. (2022) interpreted these grains as having been transported into the study area principally by wind and reinterpreted their provenance as Gondwanan (as opposed to Arctic and/or northern Appalachian as previously interpreted by Leary et al., 2020b). We find (Fig. 5) that this relationship observable in fully automated (i.e., M-ST-C) measurement results from the dataset. Our hope is that the increased ability to explore such age-grain-dimension relationships and to generate large grain-dimension datasets from toolsets such as those presented here and

615 by Scharf et al. (2022) will improve future provenance interpretations, specifically as they relate to grain transport processes  
(e.g. Lawrence et al., 2011; Ibanez et al., 2018; Leary et al., 2020a; Cantine et al., 2021).

## 7 Limitations

620 Though our models (e.g., M-ST-C) evidentially generalize well to our test set, and we believe that they will most likely  
generalize well to other datasets, they are still untested on data from facilities not represented in their training dataset (i.e.,  
besides ALC and UCSB). And, though they have been exposed to some relatively euhedral detrital zircon grains in the UCSB  
training images, our models are notably also untested on crystals derived from primary igneous and volcanic rocks. Some  
uncertainty remains in how well our models will work when applied to more diverse data by colab\_zirc\_dims users. We hope  
that any users who find that colab\_zirc\_dims struggles with their image data will share said data with us so that we can use it  
to expand on our training dataset and so improve our models' utility.

625 Measurements produced using colab\_zirc\_dims will persist all uncertainties that are innate to the methodology of  
measuring grain dimensions from reflected light images. Although most facilities aspire to polish their laser ablation zircon  
mounts to half the thickness of the zircons, it is possible that differences in sample preparation methods could produce  
significant systematic interfacility or even intra-facility (i.e., between different analysts) biases in measurable two-dimensional  
grain dimensions; it remains somewhat unclear whether data derived from sample preparation and imaging at different facilities  
630 can be compared. Additionally, because there is some variability in quality of polish this was achieved at ALC in the test  
dataset (Leary et al., 2020a; see above discussion of samples 1WM-302, 5PS-58, 2QZ-9, and 2QZ-272), careful manual  
checking of polish quality will always be required in any dataset as described above. Ultimately, a study in which pre- (e.g.  
Finzel, 2017) and post-mount (Leary et al., 2020a; Scharf et al., 2022; current study) grain dimension measurements can be  
collected on the same samples, or one in which differential preparation methods are simulated (e.g., through slicing of three-  
635 dimensional micro-CT data), will be the best way to quantify the bias introduced by polishing and/or by different facilities.  
However, such a test is well beyond the scope of the current study.

## 8 Future developments

The colab\_zirc\_dims package and ~~Google Colab~~Jupyter-style notebooks make it significantly faster and easier to augment an  
appropriate LA-ICP-MS dataset with grain measurements. We will continue to maintain and update colab\_zirc\_dims, and in  
640 the future hope to test and, if necessary, modify our code to extend full support to datasets from facilities beyond ALC and  
UCSB, possibly including those using targeting software other than Chromium. Though individual researchers are our intended  
userbase for colab\_zirc\_dims, we also believe that deep learning models ~~also~~ hold great potential utility for LA-ICP-MS  
facilities. Such facilities are well-resourced to create ~~larger~~large, customized training datasets ~~than ours~~ and could implement  
trained models in a variety of applications including provision of per-spot grain measurements as a standard data product, fully

645 automated spot picking, and possibly automated phase identification. Our training-validation dataset and pre-trained models (Sitar, 2022) may lower the barrier to entry for researchers and/or facilities hoping to apply machine- or deep-learning-based methods to similar problems.

## **79 Conclusions**

650 We ~~have trained~~created a set of deep learning models to segment ~~new, large dataset for instance segmentation of~~ detrital zircon grain instances from reflected light images saved during LA-ICP-MS analysis. Using this dataset, we trained a suite of deep learning models and developed code that uses the models to rapidly extract per-grain dimensional measurements from ~~images saved during LA-ICP-MS analyses~~images collected at facilities using Chromium targeting software. We present this code as the colab\_zirc\_dims Python package, and we implement it in a collection of interactive, ~~ready-to-run Google Colab Jupyter~~ notebooks. These notebooks ~~are highly accessible and~~ allow users to automatically or semi-automatically process datasets utilizing freely available high-end GPUs and can be run locally after installation of code dependencies or online in Google Colab with zero setup, hardware requirements, or installation.

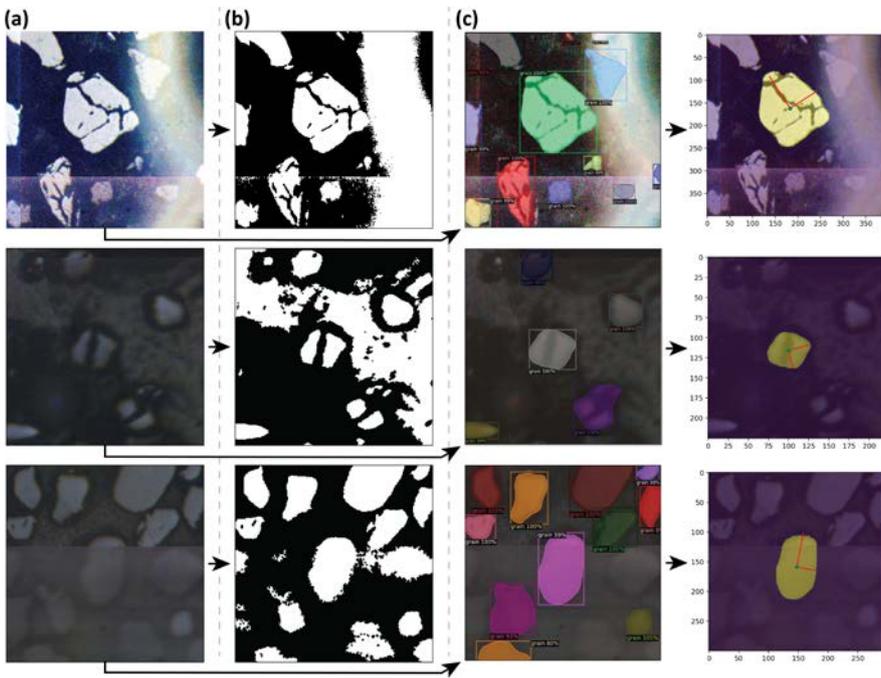
660 ~~Based on evaluations of our trained models against manual grain measurements from an in press, large n detrital zircon dataset, we conclude that transfer learning and training image augmentation are essential to training deep learning models that can accurately segment and identify mineral grains given a relatively small training dataset. We find that Centermask2-VoVNetV2 and Mask RCNN models with ResNet-101 and Swin-T backbones achieve the greatest measurement accuracies on the evaluated dataset and expect that greater still accuracies could be reached with more training data. All our trained models are available to colab\_zirc\_dims users for application to their datasets.~~

665 The colab\_zirc\_dims deep-learning-based automated measurement algorithm approaches human measurement accuracy on a sample-by-sample basis and can be used to rapidly approximate grain size distributions for samples with well-exposed zircon grains, without any human involvement. Our semi-automated segmentation workflow allows researchers to create manually reviewed and corrected grain size measurements for large-n datasets in under a day-, though data collected through this process inherit all uncertainties related to the methodology of measuring mounted, polished grains in reflected light images.

670 ~~We~~ believe that colab\_zirc\_dims makes it drastically easier to augment applicable LA-ICP-MS datasets with grain measurements, and ~~we~~ hope that allowing more researchers to do so will expand our understanding of the relationships between zircon dimensions and age in varied environments. We also hope to extend full colab\_zirc\_dims support to datasets that do not currently work with its processing notebooks in the future and encourage users to share samples of such datasets with the first author.

**Appendix A: Additional examples of segmentation results**

675 **Figure A1. Comparison between Otsu thresholding and CNN-based instance segmentation results in the presence of diverse grain morphologies and image artefacts, including anomalous bright spots (top row), heavily fractured grains (middle row), and tiling artefacts (bottom row). (a) Original grain-centred images clipped from ALC mosaics. (b) Segmentation masks produced via Otsu's thresholding method (Otsu, 1979). (c) Instance segmentation results produced by a Mask RCNN model (M-ST-C) (at left) and resulting colab\_zirc\_dims verification image plots (at right).**



680

**Appendix B: Glossary of deep learning terminology**

**Table B1: A glossary of deep learning terminology used in this study.**

<u>Term</u>	<u>Explanation</u>	<u>Reference(s)</u>
<u>Weights</u>	<u>Training-optimizable parameters that are applied to data at various points within a neural network.</u>	-
<u>Convolutional neural network (CNN)</u>	<u>A neural network wherein convolutional layers (roughly, these pass sliding filters over inputs) are used to abstract data. This allows processing of larger data (e.g., images) with fewer weight parameters.</u>	-

<u>Backbone network</u>	<u>A module within a larger model that abstracts input data into an intermediate 'feature map' that is passed to other modules to produce the final model outputs. Larger model architectures are commonly referred to using the syntax "[model architecture name]-[backbone network name]".</u>	-
<u>FPN</u>	<u>Feature Pyramid Network' -- a network that enhances feature maps via convolutional upsampling. Can be attached to a backbone network within a larger model to improve resolution of small objects.</u>	<u>Lin et al. (2016)</u>
<u>Mask-RCNN</u>	<u>A CNN-based model architecture developed by He et al. (2018). Internal modules use the feature map returned by a backbone network to propose regions which may contain objects. Later, independent modules fit bounding boxes to and create masks for each detected object. The most commonly used Mask-RCNN backbone is the ResNet network (He et al., 2015).</u>	<u>He et al. (2015, 2018)</u>
<u>Swin-T</u>	<u>Swin-'Tiny': the smallest variant of the 'Swin' model architecture (Liu et al., 2021), which is based on 'transformer' networks (Vaswani et al., 2017). In transformer networks, inputs are respectively translated to and from a higher dimensional space by 'encoder' and 'decoder' modules. These are impractical for direct application to images, as computational complexity scales exponentially with pixel count. The Swin architecture deals with this by splitting up image data using smaller, shifting windows.</u>	<u>Liu et al., (2021), Vaswani et al. (2017)</u>
<u>Centermask</u>	<u>A CNN-based model architecture developed by Lee and Park (2020). Similar to Mask-RCNN, except objects are detected and fit with bounding boxes by a single module, without an intermediate region proposal stage, prior to mask generation for each object. The standard backbone is VoVNet (Lee et al., 2019).</u>	<u>Lee et al. (2019), Lee and Park (2020)</u>

### Code availability

The colab\_zirc\_dims source code, small example datasets, and links to pre-formatted template project folders and the latest versions of colab\_zirc\_dims Google Colab notebooks are available at the colab\_zirc\_dims GitHub page (Sitar, 2022):(Sitar, 2022): [https://github.com/MCSitar/colab\\_zirc\\_dims](https://github.com/MCSitar/colab_zirc_dims). Additional code for reproducing error evaluations and figures presented in this manuscript using new or included automatically generated measurements is included in the supplementary data repository (Sitar and Leary, 2022): [\(Sitar and Leary, 2022\): https://doi.org/10.5281/zenodo.6412303](https://doi.org/10.5281/zenodo.6412303).  
<https://doi.org/10.5281/zenodo.7434851>.

### 690 Data availability

The full Leary et al. (accepted2022) dataset of images and measurements that we used for model evaluation, our training dataset, model training scripts with copies of our trained model weights, and full measurement and evaluation data supporting the results presented in our manuscript can be found in the supplementary data repository (Sitar and Leary, 2022):(Sitar and Leary, 2022): <https://doi.org/10.5281/zenodo.7434851>.

695 **Author contribution**

M.S. wrote the first draft of the manuscript and both authors contributed to subsequent drafts. M.S. segmented the training dataset, trained the models, developed the code, and evaluated model-derived measurements. R. L. provided contextualized image and measurement datasets for model training and evaluation and feedback for improvement of the code and processing notebooks.

700 **Competing interests**

The authors declare that they have no conflict of interest.

**Acknowledgements**

We would like to thank Kurt Sundell for insights into imaging systems at the Arizona LaserChron Center. Michael Sitar is additionally grateful to John Singleton for the leeway to finish this project. Previously unpublished UCSB training image data were collected in collaboration with Alaina Rosenthal-Guillot, with assistance from Andrew Kylander-Clark. Work by Michael Sitar was completed concurrently with work supported by grants from the National Science Foundation (award # 2115719) and USGS EDMAP Program (award # G21AC10493).

**References**

- Alzubaidi, L., Zhang, J., Humaidi, A. J., Al Dujaili, A., Duan, Y., Al Shamma, O., Santamaría, J., Fadhel, M. A., Al Amidie, M., and Farhan, L.: Review of deep learning: concepts, CNN architectures, challenges, applications, future directions, *J. Big Data*, 8, 53, <https://doi.org/10.1186/s40537-021-00444-8>, 2021.
- Python Package Index—PyPI: <https://pypi.org/>, last access: 13 April 2022.
- Augustsson, C., Voigt, T., Bernhart, K., Kreißler, M., Gaupp, R., Gärtner, A., Hofmann, M., and Linnemann, U.: Zircon size-age sorting and source area effect: The German Triassic Buntsandstein Group, *Sediment. Geol.*, 375, 218–231, <https://doi.org/10.1016/j.sedgeo.2017.11.004>, 2018.
- Bradski, G.: The OpenCV Library, Dr Dobbs J. Softw. Tools, 2000.
- Bukharev, A., Budenny, S., Lokhanova, O., Belozarov, B., and Zhukovskaya, E.: The Task of Instance Segmentation of Mineral Grains in Digital Images of Rock Samples (Thin Sections), in: 2018 International Conference on Artificial Intelligence Applications and Innovations (IC-AIAI), 2018 International Conference on Artificial Intelligence Applications and Innovations (IC-AIAI), 18–23, <https://doi.org/10.1109/IC-AIAI.2018.8674449>, 2018.
- Cantine, M. D., Setera, J. B., Vantongeren, J. A., Mwinde, C., and Bergmann, K. D.: Grain size and transport biases in an Ediacaran detrital zircon record, *J. Sediment. Res.*, 91, 913–928, <https://doi.org/10.2110/jsr.2020.153>, 2021.

- TensorFlow Developers.: TensorFlow, Zenodo, <https://doi.org/10.5281/zenodo.5949169>, 2022.
- 725 Dutta, A. and Zisserman, A.: The VIA Annotation Software for Images, Audio and Video, Proc. 27th ACM Int. Conf. Multimed., 2276–2279, <https://doi.org/10.1145/3343031.3350535>, 2019.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Fernández del Río, J., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E.: Array programming with NumPy, *Nature*, 585, 357–362, <https://doi.org/10.1038/s41586-020-2649-2>, 2020.
- 730 He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, ArXiv151203385 Cs, 2015.
- He, K., Gkioxari, G., Dollár, P., and Girshick, R.: Mask R-CNN, ArXiv170306870 Cs, 2018.
- Ibañez Mejía, M., Pullen, A., Pepper, M., Urbani, F., Ghoshal, G., and Ibañez Mejía, J. C.: Use and abuse of detrital zircon U-Pb geochronology: A case from the Río Orinoco delta, eastern Venezuela, *Geology*, 46, 1019–1022, <https://doi.org/10.1130/G45596.1>, 2018.
- 735 J. D. Hunter: Matplotlib: A 2D Graphics Environment, *Comput. Sci. Eng.*, 9, 90–95, <https://doi.org/10.1109/MCSE.2007.55>, 2007.
- Jiang, F., Li, N., and Zhou, L.: Grain segmentation of sandstone images based on convolutional neural networks and weighted fuzzy clustering, *IET Image Process.*, 14, 3499–3507, <https://doi.org/10.1049/iet-ipr.2019.1761>, 2020.
- Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C.: Jupyter Notebooks – a publishing format for reproducible computational workflows, in: *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, 87–90, 2016.
- 740 Latif, G., Bouchard, K., Maitre, J., Back, A., and Bédard, L. P.: Deep Learning-Based Automatic Mineral Grain Segmentation and Recognition, *Minerals*, 12, 455, <https://doi.org/10.3390/min12040455>, 2022.
- Lawrence, R. L., Cox, R., Mapes, R. W., and Coleman, D. S.: Hydrodynamic fractionation of zircon age populations, *GSA Bull.*, 123, 295–305, <https://doi.org/10.1130/B30151.1>, 2011.
- 745 Leary, R., Smith, M. E., and Umhoefer, P.: Mixed eolian longshore sediment transport in the Late Paleozoic Arizona-Pedregosa basin, USA: a case study in grain size analysis of detrital zircon datasets, *J. Sediment. Res.*, accepted for publication.
- Leary, R. J., Smith, M. E., and Umhoefer, P.: Grain Size Control on Detrital Zircon Cycloprovenance in the Late Paleozoic Paradox and Eagle Basins, USA, *J. Geophys. Res. Solid Earth*, 125, e2019JB019226, <https://doi.org/10.1029/2019JB019226>, 2020.
- 750 Lee, S. H., Lee, S., and Song, B. C.: Vision Transformer for Small Size Datasets, ArXiv211213492 Cs, 2021.
- Lee, Y. and Park, J.: CenterMask: Real-Time Anchor-Free Instance Segmentation, ArXiv191106667 Cs, 2020.
- Lee, Y., Hwang, J., Lee, S., Bae, Y., and Park, J.: An Energy and GPU Computation Efficient Backbone Network for Real-Time Object Detection, ArXiv190409730 Cs, 2019.
- 755 Lin, T. Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C. L., and Dollár, P.: Microsoft COCO: Common Objects in Context, ArXiv14050312 Cs, 2015.

- Lin, T. Y., Dollár, P., Girshick, R., He, K., Hariharan, B., and Belongie, S.: Feature Pyramid Networks for Object Detection, <https://doi.org/10.48550/arXiv.1612.03144>, 2016.
- 760 Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B.: Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, <https://doi.org/10.48550/ARXIV.2103.14030>, 2021.
- McKinney, W.: Data Structures for Statistical Computing in Python, in: Proceedings of the 9th Python in Science Conference, 56–61, <https://doi.org/10.25080/Majora-92bf1922-00a>, 2010.
- Muhlbauer, J. G., Fedo, C. M., and Farmer, G. L.: Influence of textural parameters on detrital zircon age spectra with application to provenance and paleogeography during the Ediacaran–Terreneuvian of southwestern Laurentia, *GSA Bull.*, 129, 1585–1601, <https://doi.org/10.1130/B31611-1>, 2017.
- 765 Nachtergaele, S. and De Grave, J.: AI-Track-tive: open source software for automated recognition and counting of surface semi tracks using computer vision (artificial intelligence), *Geochronology*, 3, 383–394, <https://doi.org/10.5194/gechron-3-383-2021>, 2021.
- Otsu, N.: A Threshold Selection Method from Gray-Level Histograms, *IEEE Trans. Syst. Man Cybern.*, 9, 62–66, <https://doi.org/10.1109/TSMC.1979.4310076>, 1979.
- 770 Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S.: PyTorch: An Imperative Style, High Performance Deep Learning Library, *ArXiv1912.01703 Cs Stat*, 2019.
- Pérez, F. and Granger, B. E.: IPython: a System for Interactive Scientific Computing, *Comput. Sci. Eng.*, 9, 21–29, <https://doi.org/10.1109/MCSE.2007.53>, 2007.
- 775 Ren, S., He, K., Girshick, R., and Sun, J.: Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks, *ArXiv1506.01497 Cs*, 2016.
- Resentini, A., Andò, S., and Garzanti, E.: Quantifying Roundness of Detrital Minerals By Image Analysis: Sediment Transport, Shape Effects, and Provenance Implications, *J. Sediment. Res.*, 88, 276–289, <https://doi.org/10.2110/jsr.2018.12>, 2018.
- 780 Scharf, T., Kirkland, C. L., Daggitt, M. L., Barham, M., and Puzryev, V.: AnalyZr: A Python application for zircon grain image segmentation and shape analysis, *Comput. Geosci.*, 162, 105057, <https://doi.org/10.1016/j.cageo.2022.105057>, 2022.
- Sitar, M. C.: MCSitar/colab\_zirc\_dims: colab\_zirc\_dims v1.0.8, Zenodo, <https://doi.org/10.5281/ZENODO.6410745>, 2022.
- Sitar, M. C. and Leary, R. J.: colab\_zirc\_dims: full results, datasets, and replication code repository, <https://doi.org/10.5281/zenodo.6412304>, 2022.
- 785 Soloy, A., Turki, I., Fournier, M., Costa, S., Peuziat, B., and Lecoq, N.: A Deep Learning-Based Method for Quantifying and Mapping the Grain Size on Pebble Beaches, *Remote Sens.*, 12, 3659, <https://doi.org/10.3390/rs12213659>, 2020.
- Sundell, K., Gehrels, G. E., Quinn, D. P., Pecha, M., Giesler, D., Pepper, M., George, S., and White, A.: AGEALCML: AN OPEN SOURCE MATLAB-BASED DATA REDUCTION PLATFORM FOR LA ICP-MS GEOCHRONOLOGY AND

- 790 GEOCHEMISTRY DATA FROM THE ARIZONA LASERCHRON CENTER, GSA 2020 Connects Online, 358944, <https://doi.org/10.1130/abs/2020AM-358944>, 2020.  
Teledyne Photon Machines: Chromium 2.4, Teledyne Photon Machines, 2020.  
Tian, Z., Shen, C., Chen, H., and He, T.: FCOS: Fully Convolutional One Stage Object Detection, ArXiv190401355 Cs, 2019.  
Umesh, P.: Image Processing in Python, CSI Commun., 23, 2012.
- 795 Van Rossum, G.: The Python Library Reference, release 3.8.2, Python Software Foundation, 2020.  
van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., Yu, T., and contributors., the scikit-image: scikit-image: Image processing in Python, PeerJ, 2, e453, <https://doi.org/10.7717/peerj.453>, 2014.  
Detectron2: <https://github.com/facebookresearch/detectron2>.
- 800 Ye, H., Yang, Y., and L3str4nge: SwinT\_detectron2: v1.2, Zenodo, <https://doi.org/10.5281/ZENODO.6468976>, 2021.  
Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M. A., Al-Amidie, M., and Farhan, L.: Review of deep learning: concepts, CNN architectures, challenges, applications, future directions, J. Big Data, 8, 53, <https://doi.org/10.1186/s40537-021-00444-8>, 2021.  
Python Package Index - PyPI: <https://pypi.org/>, last access: 13 April 2022.
- 805 Augustsson, C., Voigt, T., Bernhart, K., Kreißler, M., Gaupp, R., Gärtner, A., Hofmann, M., and Linnemann, U.: Zircon size-age sorting and source-area effect: The German Triassic Buntsandstein Group, Sediment. Geol., 375, 218–231, <https://doi.org/10.1016/j.sedgeo.2017.11.004>, 2018.  
Bradski, G.: The OpenCV Library, Dr Dobbs J. Softw. Tools, 2000.  
Bukharev, A., Budennyi, S., Lokhanova, O., Belozarov, B., and Zhukovskaya, E.: The Task of Instance Segmentation of Mineral Grains in Digital Images of Rock Samples (Thin Sections), in: 2018 International Conference on Artificial Intelligence Applications and Innovations (IC-AIAD), 2018 International Conference on Artificial Intelligence Applications and Innovations (IC-AIAD), 18–23, <https://doi.org/10.1109/IC-AIAD.2018.8674449>, 2018.  
Cantine, M. D., Setera, J. B., Vantongeren, J. A., Mwinde, C., and Bergmann, K. D.: Grain size and transport biases in an Ediacaran detrital zircon record, J. Sediment. Res., 91, 913–928, <https://doi.org/10.2110/jsr.2020.153>, 2021.
- 815 Detectron2: <https://github.com/facebookresearch/detectron2>.  
Dutta, A. and Zisserman, A.: The VIA Annotation Software for Images, Audio and Video, Proc. 27th ACM Int. Conf. Multimed., 2276–2279, <https://doi.org/10.1145/3343031.3350535>, 2019.  
Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., Fernández del Río, J., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E.: Array programming with NumPy, Nature, 585, 357–362, <https://doi.org/10.1038/s41586-020-2649-2>, 2020.
- 820 He, K., Zhang, X., Ren, S., and Sun, J.: Deep Residual Learning for Image Recognition, ArXiv151203385 Cs, 2015.  
He, K., Gkioxari, G., Dollár, P., and Girshick, R.: Mask R-CNN, ArXiv170306870 Cs, 2018.

- 825 [Ibañez-Mejia, M., Pullen, A., Pepper, M., Urbani, F., Ghoshal, G., and Ibañez-Mejia, J. C.: Use and abuse of detrital zircon U-Pb geochronology—A case from the Río Orinoco delta, eastern Venezuela, \*Geology\*, 46, 1019–1022, <https://doi.org/10.1130/G45596.1>, 2018.](#)
- [J. D. Hunter: Matplotlib: A 2D Graphics Environment, \*Comput. Sci. Eng.\*, 9, 90–95, <https://doi.org/10.1109/MCSE.2007.55.2007>.](#)
- 830 [Jiang, F., Li, N., and Zhou, L.: Grain segmentation of sandstone images based on convolutional neural networks and weighted fuzzy clustering, \*IET Image Process.\*, 14, 3499–3507, <https://doi.org/10.1049/iet-ipr.2019.1761>, 2020.](#)
- [Kluyver, T., Ragan-Kelley, B., Pérez, F., Granger, B., Bussonnier, M., Frederic, J., Kelley, K., Hamrick, J., Grout, J., Corlay, S., Ivanov, P., Avila, D., Abdalla, S., and Willing, C.: Jupyter Notebooks – a publishing format for reproducible computational workflows, in: \*Positioning and Power in Academic Publishing: Players, Agents and Agendas\*, 87–90, 2016.](#)
- 835 [Latif, G., Bouchard, K., Maitre, J., Back, A., and Bédard, L. P.: Deep-Learning-Based Automatic Mineral Grain Segmentation and Recognition, \*Minerals\*, 12, 455, <https://doi.org/10.3390/min12040455>, 2022.](#)
- [Lawrence, R. L., Cox, R., Mapes, R. W., and Coleman, D. S.: Hydrodynamic fractionation of zircon age populations, \*GSA Bull.\*, 123, 295–305, <https://doi.org/10.1130/B30151.1>, 2011.](#)
- 840 [Leary, R. J., Smith, M. E., and Umhoefer, P.: Grain-Size Control on Detrital Zircon Cycloprovenance in the Late Paleozoic Paradox and Eagle Basins, USA, \*J. Geophys. Res. Solid Earth\*, 125, e2019JB019226, <https://doi.org/10.1029/2019JB019226>, 2020a.](#)
- [Leary, R. J., Umhoefer, P., Smith, M. E., Smith, T. M., Saylor, J. E., Riggs, N., Burr, G., Lodes, E., Foley, D., Licht, A., Mueller, M. A., and Baird, C.: Provenance of Pennsylvanian–Permian sedimentary rocks associated with the Ancestral Rocky Mountains orogeny in southwestern Laurentia: Implications for continental-scale Laurentian sediment transport systems, \*Lithosphere\*, 12, 88–121, <https://doi.org/10.1130/L1115.1>, 2020b.](#)
- 845 [Leary, R. J., Smith, M. E., and Umhoefer, P.: Mixed eolian–longshore sediment transport in the late Paleozoic Arizona shelf and Pedregosa basin, U.S.A.: A case study in grain-size analysis of detrital-zircon datasets, \*J. Sediment. Res.\*, 92, 676–694, <https://doi.org/10.2110/jsr.2021.101>, 2022.](#)
- [Lee, Y. and Park, J.: CenterMask: Real-Time Anchor-Free Instance Segmentation, \*ArXiv191106667 Cs\*, 2020.](#)
- 850 [Lee, Y., Hwang, J., Lee, S., Bae, Y., and Park, J.: An Energy and GPU-Computation Efficient Backbone Network for Real-Time Object Detection, \*ArXiv190409730 Cs\*, 2019.](#)
- [Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., and Belongie, S.: Feature Pyramid Networks for Object Detection, <https://doi.org/10.48550/arXiv.1612.03144>, 2016.](#)
- [Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B.: Swin Transformer: Hierarchical Vision Transformer using Shifted Windows, <https://doi.org/10.48550/ARXIV.2103.14030>, 2021.](#)
- 855 [McKinney, W.: Data Structures for Statistical Computing in Python, in: \*Proceedings of the 9th Python in Science Conference\*, 56–61, <https://doi.org/10.25080/Majora-92bf1922-00a>, 2010.](#)
- [Muhlbauer, J. G., Fedo, C. M., and Farmer, G. L.: Influence of textural parameters on detrital-zircon age spectra with application to provenance and paleogeography during the Ediacaran–Terreneuvian of southwestern Laurentia, \*GSA Bull.\*, 129, 1585–1601, <https://doi.org/10.1130/B31611.1>, 2017.](#)

- 860 [Nachtergaele, S. and De Grave, J.: AI-Track-tive: open-source software for automated recognition and counting of surface semi-tracks using computer vision \(artificial intelligence\). \*Geochronology\*, 3, 383–394, <https://doi.org/10.5194/gchron-3-383-2021>, 2021.](#)
- [Otsu, N.: A Threshold Selection Method from Gray-Level Histograms. \*IEEE Trans. Syst. Man Cybern.\*, 9, 62–66, <https://doi.org/10.1109/TSMC.1979.4310076>, 1979.](#)
- 865 [Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library, \*ArXiv191201703 Cs Stat\*, 2019.](#)
- [Pérez, F. and Granger, B. E.: IPython: a System for Interactive Scientific Computing, \*Comput. Sci. Eng.\*, 9, 21–29, <https://doi.org/10.1109/MCSE.2007.53>, 2007.](#)
- 870 [Resentini, A., Andò, S., and Garzanti, E.: Quantifying Roundness of Detrital Minerals By Image Analysis: Sediment Transport, Shape Effects, and Provenance Implications, \*J. Sediment. Res.\*, 88, 276–289, <https://doi.org/10.2110/jsr.2018.12.2018>.](#)
- [Scharf, T., Kirkland, C. L., Daggitt, M. L., Barham, M., and Puzyrev, V.: AnalyZr: A Python application for zircon grain image segmentation and shape analysis, \*Comput. Geosci.\*, 162, 105057, <https://doi.org/10.1016/j.cageo.2022.105057>, 2022.](#)
- 875 [Sitar, M. C.: MCSitar/colab\\_zirc\\_dims: v1.0.10, <https://doi.org/10.5281/zenodo.7425633>, 2022.](#)
- [Sitar, M. C. and Leary, R. J.: colab\\_zirc\\_dims: full results, datasets, and replication code repository, <https://doi.org/10.5281/zenodo.7434851>, 2022.](#)
- [Soloy, A., Turki, I., Fournier, M., Costa, S., Peuziat, B., and Lecoq, N.: A Deep Learning-Based Method for Quantifying and Mapping the Grain Size on Pebble Beaches, \*Remote Sens.\*, 12, 3659, <https://doi.org/10.3390/rs12213659>, 2020.](#)
- 880 [Sundell, K., Gehrels, G. E., Quinn, D. P., Pecha, M., Giesler, D., Pepper, M., George, S., and White, A.: AGEALCML: AN OPEN-SOURCE MATLAB-BASED DATA REDUCTION PLATFORM FOR LA-ICP-MS GEOCHRONOLOGY AND GEOCHEMISTRY DATA FROM THE ARIZONA LASERCHRON CENTER, \*GSA 2020 Connects Online\*, 358944, <https://doi.org/10.1130/abs/2020AM-358944>, 2020.](#)
- [Teledyne Photon Machines: Chromium 2.4, 2020.](#)
- 885 [TensorFlow Developers: TensorFlow, <https://doi.org/10.5281/zenodo.5949169>, 2022.](#)
- [Umesh, P.: Image Processing in Python, \*CSI Commun.\*, 23, 2012.](#)
- [Van Rossum, G.: The Python Library Reference, release 3.8.2, Python Software Foundation, 2020.](#)
- [van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Gouillart, E., Yu, T., and contributors, the scikit-image: scikit-image: Image processing in Python, \*PeerJ\*, 2, e453, <https://doi.org/10.7717/peerj.453>, 2014.](#)
- 890 [Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I.: Attention Is All You Need, <https://doi.org/10.48550/arXiv.1706.03762>, 5 December 2017.](#)
- [Ye, H., Yang, Y., and L3str4nge: SwinT\\_detectron2: v1.2, <https://doi.org/10.5281/ZENODO.6468976>, 2021.](#)