Mineral Grain Localization and Classification using Deep Neural Networks

Raymond Donelick[1], Andrew J Donelick[1], and Raymond A Donelick[1]

[1]Affiliation not available

September 11, 2023

Mineral Grain Localization and Classification using Deep Neural Networks

Apatite.com Partners, LLC
Viola, Idaho USA

Andrew J. Donelick[1], Raymond A. Donelick[1]
adonelick@g.hmc.edu donelick@apatite.com

Introduction

Locating high quality apatite and zircon mineral grains on a microscope mount is one of the first steps in a thermochronology analysis, and it is often a tedious and time-consuming process. Fortunately, this problem is well suited to be solved with deep-learning based image analysis techniques.

Traditional deep-learning based object recognition pipelines require massive amounts of expert-labeled data (both object class and position). The time consuming nature of collecting such data can be an insurmountable hurdle to creating a useable object recognition system.

This system leverages a modern pre-trained image segmentation model to solve the object localization problem without expert input. An expert can then quickly label the localized objects by class, which serves as the basis for a mineral & quality classifier. Together these components form a unified object recognition and localization system that is faster to train and easier to maintain than a traditional pipeline.

Step 1: Image Capture

Start with reflected light images of a mineral mount collected from an optical microscope. We use 1200 x 1600 pixel images with a spatial size of approximately 250 µm x 330 µm, respectively.

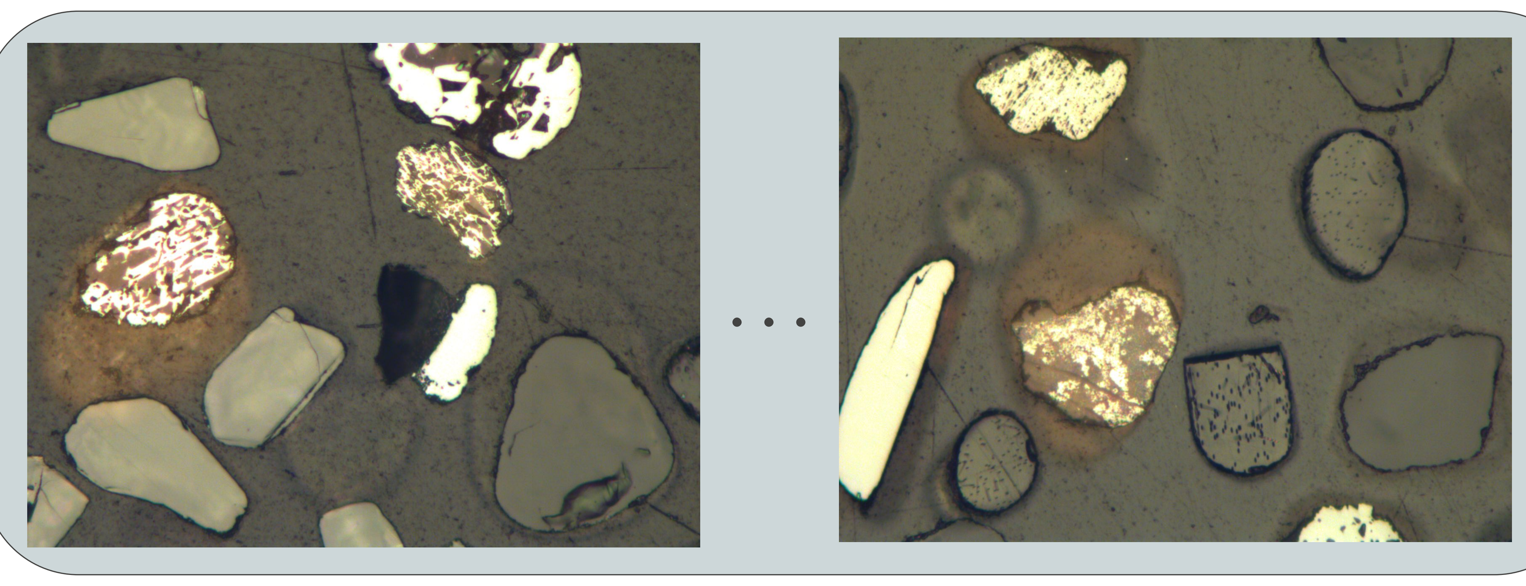

Step 2: Image Segmentation

Segment the image to distinguish grain-like objects from epoxy mount background. The output from the segmentation can be masks or bounding boxes locating the grain within the image.

This project uses the Segment Anything Model (SAM) from Facebook Research. No fine tuning or training is required to obtain high quality grain-like object masks.

However, not all the masks identified are grain objects. Filtering is required.

Step 3: Grain Candidate Filter

Extract each candidate grain from the 20X image, and then filter the grain-like objects to remove incomplete grains or objects of no analytical value. This step uses a neural network trained to identify objects likely to be grains, regardless of quality or mineral type.

Incomplete grains and objects of no analytical value are discarded and not considered for further analysis.

True Grain Candidates

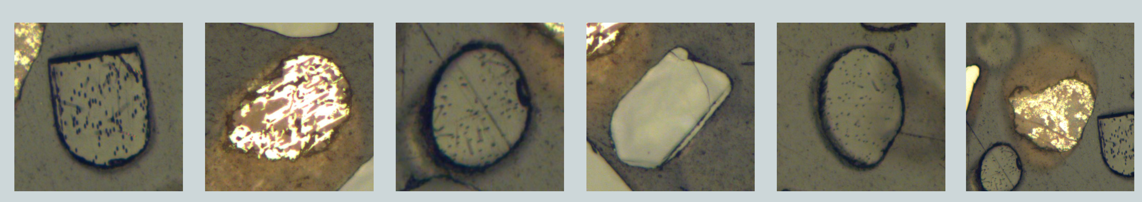

False Positives

If we wish to analyze these grains, they can be sent on to a grain classifier for mineral and quality labeling.

Alternatively, the good candidate grains can be used to build the training dataset for the mineral and quality classifier.

Step 4A: Grain Classification

Classify the grains by mineral and quality using a Convolutional Neural Network (CNN) model trained on the dataset collected and labeled by the expert.

High, Medium Quality Apatite
High, Medium Quality Zircon

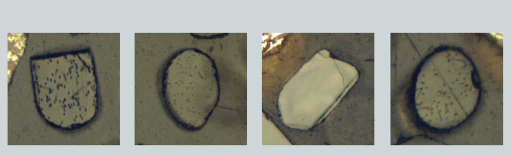

Low Quality Apatite, Zircon
Holes, Barite, Quartz
Oxides and Sulfides

Currently these classes are not used for analysis, but may support future analytical techniques.

Step 4B: Grain Classifier Dataset and Training

Each grain candidate is labeled by an expert with its mineral type and quality. The labeled image dataset is used to train a convolutional neural network (CNN) model to classify images by mineral and quality.

High Quality Apatite

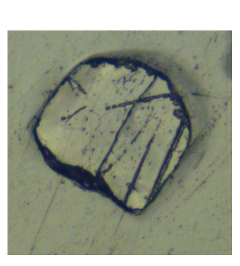

Medium Quality Apatite

Low Quality Apatite

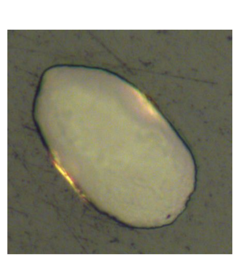

High Quality Zircon

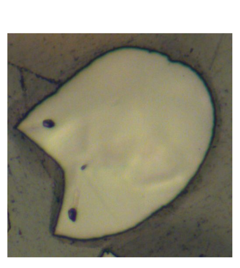

Medium Quality Zircon

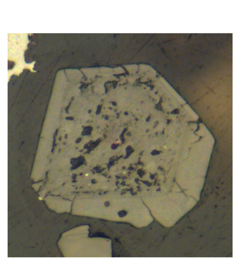

Low Quality Zircon

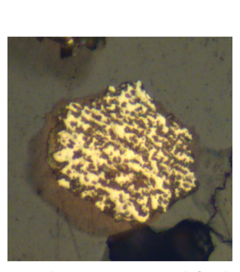

Oxides or Sulfides

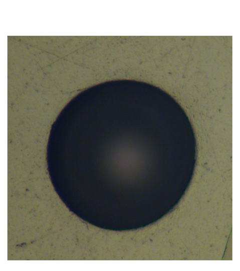

Holes

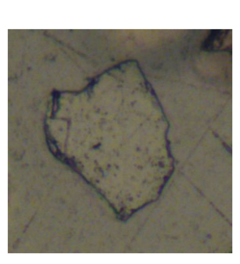

Barite

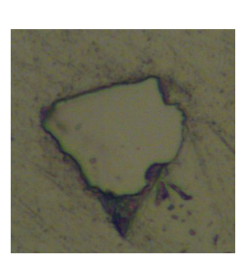

Quartz

Step 5: Thermochronology Analysis

Finally, we combine the candidate grain mask with image position metadata to provide the grain's location on the epoxy mount. This information allows the grain(s) to be located and analyzed.

Status of Grain Classification Dataset

Total Number of Samples: 57,202

Significant imbalances are present in the dataset. Data augmentation during training helps, but significant work remains in building out and balancing this dataset to improve classification accuracy results.

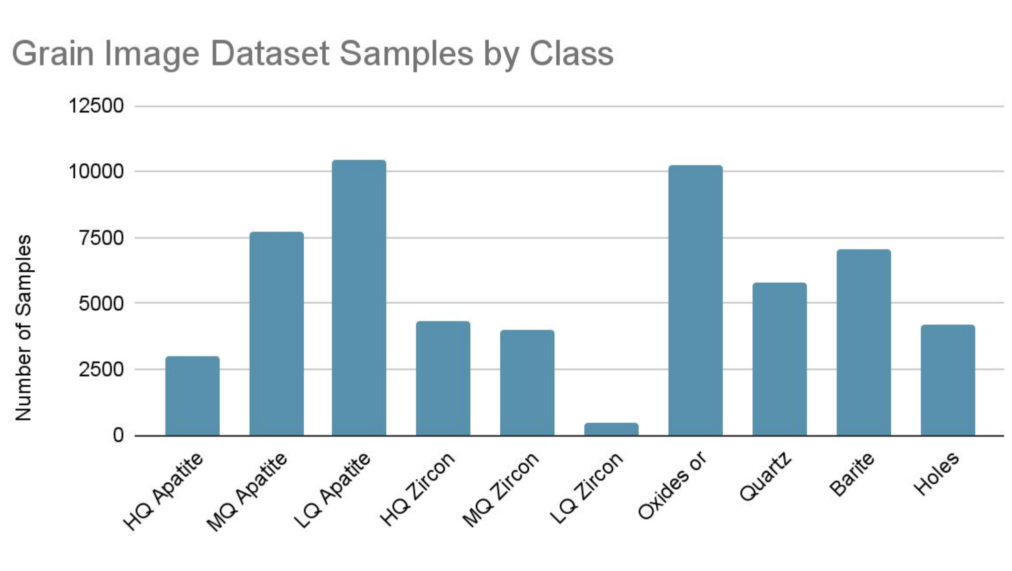

Grain Classification Results

Validation accuracy results for each of the classes in the grain dataset. These results are for a classifier built with a ResNet-34 model architecture.