Fully Automated Carbonate Petrography Using Deep Convolutional Neural Networks

Ardiansyah Koeshidayatullah[1], Michele Morsilli[2], Daniel J. Lehrmann[3], Khalid Al-Ramadan[4], Jonathan L. Payne[1]*

[1]Department of Geological Sciences, Stanford University, Stanford, CA 94305, USA

[2]Dipartimento di Fisica e Scienze della Terra, Università di Ferrara, 44122, Ferrara, Italy

[3]Department of Geosciences, Trinity University, San Antonio, TX, 78212, USA

[4]College of Petroleum Engineering & Geosciences, King Fahd University of Petroleum & Minerals, Dhahran, 31261, Saudi Arabia

*corresponding author: ardikoes@stanford.edu

Present Address: 450 Jane Stanford Way, Stanford University, Stanford, CA 94305, USA

ABSTRACT

Carbonate rocks are important archives of past ocean conditions as well as hosts of economic resources such as hydrocarbons, water, and minerals. Geologists typically perform compositional analysis of grain, matrix, cement and pore types in order to interpret depositional environments, diagenetic modification, and reservoir quality of carbonate strata. Such information can be obtained primarily from petrographic analysis, a task that is costly, labor-intensive, and requires in-depth knowledge of carbonate petrology and micropaleontology. Recent studies have leveraged machine learning-based image analysis, including Deep Convolutional Neural Networks (DCNN), to automate description, classification and interpretation of thin sections, subsurface core images and seismic facies, which would accelerate data acquisition and reproducibility for these tasks. In carbonate rocks, this approach has been applied primarily to recognize carbonate lithofacies, and no attempt has been made to individually identify and quantify various types of carbonate grains, matrix, and cement. In this study, the applicability and performance of DCNN-based object detection and image classification approaches are assessed with respect to carbonate compositional analysis. The training data comprised of more than 13,000 individually labeled objects from nearly 4000 carbonate petrographic images. The dataset is grouped into six and nine different classes for the image classification and object detection tasks, respectively. Even with a small and relatively imbalanced training set, the DCNN was able to achieve an F1 score of 92% for image classification and 84% mean precision for object detection by combining one-cycle policy, class weight, and label mixup-smoothing methods. This study highlights the inefficiency of image classification as an approach to replicating human description and classification of carbonate petrography. By contrast, DCNN-based object detection appears capable of approaching human speed and accuracy in the area of carbonate petrography because it is able to individually locate and identify different carbonate components with greater cost-efficiency, speed, and reproducibility than conventional (human) petrographic analysis.

Keywords: Carbonate; Petrography; Deep Learning; Object Detection; Automated analysis

1. INTRODUCTION

Recent advances in the field of artificial intelligence have been driven by the development of Deep Convolutional Neural Networks (DCNN), which can surpass human accuracy in computer vision tasks such as detailed image classification and object detection (e.g., Krizhevsky et al., 2012; Ren et al., 2015; Shin et al., 2016). DCNNs employ a non-linear function approximation that can perform better than shallow neural networks in the analysis of large data sets (LeCun et al., 2015). A high-speed graphics processing unit (GPU) allows vast arrays of tools/libraries to build DCNN architectures, which enable the machine to hierarchically learn and extract features within an image through a general-purpose learning process (e.g., Goodfellow et al., 2016).

Machine learning in general and deep learning in particular offer promising tools to build new, data-driven models Earth system sciences (Reichstein et al., 2019). In the field of petroleum exploration, these methods have been applied extensively to seismic phase inversion, multiphase flow, total organic carbon (TOC) estimation, reservoir characterization, fracture analysis, geophysical log correlation, drilling penetration, and porosity and permeability prediction with geophysical logs (Ashena and Thonhauser, 2015). Machine learning algorithms have also been applied successfully to other image-based, geoscience-related problems, such as such as satellite image mapping (e.g., Lynda, 2019), seismic facies classification (e.g., Qian et al., 2018), lithology classification (e.g., Saporetti et al., 2018), and mineral recognition and classification of igneous rocks (e.g., Izadi et al., 2017). Furthermore, Saporetti et al. (2018) and Silva et al. (2020) combined different machine learning algorithms (e.g., K-Nearest Neighbors, Decision Trees, and Support Vector Machine) to classify carbonate or siliciclastic rocks from rock physics- and well log-derived petrofacies, respectively.

Recently, the DCNN method has been applied to conduct various geoscience tasks, such as geological feature extraction from seismic attributes (Huang et al., 2017), history matching of geological facies models (Liu et al., 2019; Canchumuni et al., 2019), three dimensional porous media reconstruction (Mosser et al., 2017), reconstruction of relative geologic time from seismic image (Geng et al., 2020), and anomaly detection for geological carbon sequestration (Zhong et al., 2019). In addition, DCNN models have been applied widely to lithofacies classification in cores and thin section images (John and Kanagandran, 2019; Pires de Lima et al., 2019; Baraboshkin et al., 2020; Pires de Lima et al., 2020; Tang et al., 2020). To date, however, machine learning tools have been applied primarily to lithofacies classification. They have not yet been applied to the comprehensive identification, classification and quantification of individual grain and cement components in carbonate rocks.

Petrographic analysis is the most commonly used technique to identify and classify the components and textures of carbonate rocks. This approach is typically labor-intensive and requires substantial prior knowledge. This limitation is compounded by the potential for bias and subjectivity in the human interpretation and classification of carbonate rocks (e.g., Dunham Classification; Dunham, 1962), limiting reproducibility among researchers (Lokier and Al-Junaibi, 2016). Recently, automated approaches have been applied to perform quantitative petrographic analysis, such as point counting (Asmussen et al., 2015) and porosity measurement (Grove and Jerram, 2011; Amao et al., 2016), through image segmentation and thresholding. These methods were not fully automated, however, as the different types of carbonate constituents were still manually identified by geologists. Full automation would represent a major advance in the use of computerized tools for the petrographic description and classification of carbonate rocks.

Detailed and accurate characterization of carbonate constituents is important for applications to hydrocarbon reservoirs, freshwater aquifers, carbon capture and storage, the evolution of marine ecosystems, and paleoenvironmental reconstruction (e.g., Payne et al., 2006; Koeshidayatullah et al., 2016; Al-Ramadan et al., 2019; Koeshidayatullah et al., 2020).

To date, the majority of studies on geological image analysis have been focused in applying DCNN-based image classification to automate identification and interpretation processes (e.g., Baraboshkin et al., 2020). A typical petrographic thin section of a carbonate rock sample (and the image thereof) has various carbonate components; therefore, a DCNN-based image classification approach requires the image to be manually segmented into separate parts prior to classification if the image contains multiple components (Fig. 1). In addition, the classification task often returns only a single predicted class for each image, even when multiple object classes are present in the image (Fig. 1). Therefore, this study aims to explore the application and limitations of DCNN-based image classification on carbonate petrography images. Furthermore, this paper, for the first time, studies the application of coupled DCNN-based image classification and object identification tasks to simultaneously locate and identify multiple carbonate grain, matrix, and cement types within a single petrographic image (Fig. 1). Such a method more closely approximates the human approach to carbonate petrography. The performances of different DCNNs architectures or frameworks for image classification and object detection tasks are compared in order to identify the most effective approaches. In addition, the deep learning approach is extended to quantify the abundances of different components within a sample. Methods to improve network and detection performance in the presence of limited

training data are also explored. The results suggest that fully automated carbonate petrography is feasible given current technology.

2. DATA

For this study, nearly 4000 images of carbonate petrographic thin sections of different scales or magnification, both in plane- and cross-polarized light, were compiled from various sources. Samples were drawn primarily from the Upper Permian through Middle Triassic of south China, Turkey, and Saudi Arabia and supplemented by images from carbonate petrography textbooks (Adams and Mackenzie, 1998; Scholle and Scholle, 2003; Flügel, 2013). In addition, images from a publicly accessible site, www.carbonateworld.com (Della Porta and Wright, 2009) and other studies in the primary literature were compiled for both training and testing purposes. These images cover different geological periods and various type of carbonate grains and cements. Images in the database were resized to 512x512 pixels and 224x224 pixels in order to increase the efficiency of training processes while still preserving the original features of carbonate constituents. The image sizes were selected for the image classification task, the training dataset was segmented and split into smaller images that represent six different carbonate components (Fig. 2A): (i) coated grains (29%); (ii) bioclasts (27%); micrite (7%); (iv) calcite cement (9%); (v) replacement dolomite (18%); and (vi) porosity (10%). For the object detection task, more than 13,000 individual carbonate components were manually labelled from the training datasets and divided into nine classes (Fig. 2B): (i) ooids (37%); (ii) peloids (9%); (iii) foramifera (9%); (iv) molluscs (11%); (v) other skeletal grains (13%); (vii) micrite (5%); (vii) calcite cement (6%); (viii) replacement dolomite (7%); and (ix) porosity (3%).

3. METHODS

3.1 Computing Environment Setup

In this study, a publicly accessible cloud service, Google Colaboratory (Colab) environment with access to NVIDIA Tesla T4 (2560 CUDA® Cores) or P100 GPUs (3584 CUDA® Cores) and up to 25 GB of high-memory Virtual Machines to build the algorithms, was used to train the model and to make predictions. This Jupyter notebook-based environment comes with preinstalled Python libraries/dependencies for DCNN-based image classification and object detection, including *NumPy* (1.17) (Oliphant, 2006), *OpenCV* (4.1.2) (Bradski, 2000), *Scikit-learn* (0.22.2) (Pedregosa et al., 2011), *Imbalanced-learn* (Lemaître et al., 2017) and *Matplotlib* (2.2.2) (Hunter, 2007). Various back-end libraries (API): *TensorFlow* (1.13-1.15) (Abadi et al., 2016), *Keras* (2.2.5) (Chollet et al., 2015), *Pytorch* (1.5) (Paszke et al., 2019), and *fastai* (Howard and Gugger, 2020) were used to develop the DCNN architecture, preprocess the image dataset (e.g., image filtering and augmentation), and train the model.

3.2 Data Annotation and Pre-processing

ImageDataGenerator (*Keras* Library) and ImageDataBunch (*fastai* library) functions were used to create a labelled image dataset for the image classification task. For the object detection task, an open-source firmware, labelImg (Tzutalin, 2015), was used to generate object annotations in PASCAL visual object classes (VOC) and YOLO-accepted formats.

The training dataset has substantial class imbalance among carbonate components, which may negatively impact the network and prediction performance. To reduce the effects of bias and class imbalance in the training dataset, multiple mitigation approaches were performed, including: data augmentation, constructing cross-validation sets, oversampling-undersampling

the dataset, and pre-determining the class weights (*scikit-learn* and *Keras* libraries). For cross-validation sets StratifiedKFold by using *scikit-learn* library (Kohavi, 1995), the dataset was divided into 85% for training and 15% for testing and shuffled into five different combinations of training and testing datasets (K=5; Fig. 2C). This train-test split design was used based on optimum results during training and prediction processes compared to other test-split percentages (i.e. 10%, 20%, 25%). For training purposes, data augmentation techniques (e.g., transform, color transformation, add noise, pixel filling, and flip; *Keras* Library) were used to generate a pseudo training dataset and increase the diversity of the data without actually collecting additional images. In addition, a label mixup method ("Bag of Freebies"; Zhang et al., 2019) was used to minimize overfitting by giving the network two images from either similar or different classes and creating a linear combination of them.

Class imbalance is a common problem in machine learning, and it can be mitigated in several ways (e.g., data-level and algorithm-level methods; Buda et al., 2018; Johnson and Khoshgoftaar, 2019). In this paper, three different methods (data- and algorithm-level) were used and evaluated to address the effects of class imbalance. First, a combined oversampling-undersampling technique (SMOTEENN; *Imbalance-learn* library) was used following the Synthetic Minority Over-sampling Technique (SMOTE; Chawla et al., 2002) and Edited Nearest Neighbors (ENN; Alejo et al., 2010), respectively. In addition, the BalancedBatchGenerator (*Imbalance-learn* and *Keras* Libraries) was used to perform the SMOTEENN technique across the whole dataset. This method balances the dataset by duplicating samples in the minority classes while removing samples from majority classes. This method is possible with an assumption that the overall abundance of different carbonate grains, cement and micrite is roughly similar in the

study material. Second, an algorithm-level method was tested by assigning different class weights based on the amount of material for each class in the training data (*scikit-learn* and *Keras* libraries; King and Zen, 2001), which penalizes majority and minority classes equally (i.e. higher weights are given to minority classes while lower class weights are assigned to majority classes). For the image classification task, the calculated class weight for (i) bioclast (27%) is 0.65; (ii) coated grains (29%) is 0.51; (iii) calcite cement (9%) is 1.9; (iv) replacement dolomite (18%) is 1.1; (v) micrite (7%) is 3.6; and (vi) porosity (10%) is 1.4. Finally, a soft sampling approach was tested by applying a focal loss function (Lin et al., 2017a) instead of cross-entropy loss. This method works by put more weights on hard-to-classify samples while down-weighting the easy-to-classify examples during training. This method is applied primarily for the DCNN-based object detection task.

3.3. Deep Convolutional Neural Networks

A full convolutional neural network (CNN) was first introduced nearly four decades ago to recognize handwriting and digits (LeNet; LeCun et al., 1989). Here, the convolution operation works as an element-wise matrix multiplication between the filter (a two-dimensional array of weights) and filter-sized patch of the input, which is followed by summing the elements of the produced matrix, resulting in a single number (scalar product). The rise of CNN occurred in 2012, when Krizhevsky et al. (2012) built a DCNN model (AlexNet) to recognize different image classes in an ImageNet dataset (Deng et al., 2009). This deep convolutional neural network perceives images as tensors and arranges its neurons as a 3D volume (height, weight and depth (Red-Green-Blue channel)) (Fig. 3). In general, the architecture of a DCNN is very similar to a CNN. It consists of input, output and multiple layers in between, including (Fig. 3): (i) convolutional layers; (ii)

activation layers; (iii) pooling layers; and (iv) fully connected layers (e.g., Goodfellow et al., 2016). Matrix operations between these layers and different filters allow the network to learn different features (high- to low-level feature) within an image. In the fully connected layer, a SoftMax function (Bridle, 1990) is used to calculate the prediction of different classes (Fig. 3). Furthermore, to optimize the learning processes (i.e., to find global minimum in order to minimize the loss), different hyper-parameters within the layers, such as learning rate, loss, and activation functions, can be fine-tuned (Goodfellow et al., 2016). In this paper, sparse categorical cross-entropy loss, rectified linear unit (ReLU) and adaptive learning rate optimization algorithm (Adam=Adaptive Momentum Estimation; Kingma and Bas, 2014) were used to increase the network performance.

DCNNs are typically trained on very large datasets (millions of images and thousands of classes; e.g., ImageNet). In geoscientific problems, however, such large-scale, annotated datasets are often difficult or impossible to obtain. The most commonly used technique employed when only limited training data are available is called "transfer learning" (e.g., Bengio, 2012; Shin et al., 2016; Tan et al., 2018). In transfer learning, the dataset does not need to be independent and identically distributed and the network is not trained from scratch. Instead, the networks or learning parameters are first trained on the domain model where very large datasets are available (e.g., CIFAR-10 and -100 (https://www.cs.toronto.edu/~kriz/cifar.html; Krizhevsky et al., 2009), ImageNet (http://www.image-net.org/; Deng et al., 2009) and Microsoft COCO (http://cocodataset.org/; Lin et al., 2014)). These pre-trained parameters (e.g., weights) and layers are transferred to perform certain computer vision tasks in the target model (e.g., DCNN-based carbonate image classification). In general, the training process of the transfer learning method requires less computation time and performs better because of the presence of pre-trained

weights and learning parameters to perform feature extraction from an image of interest. However, as the network was trained in a domain model where it contains images or objects that do not resemble carbonate components, networks were fine-tuned in this study by adding hidden layers with several different weight initializations (*Keras* Library; e.g., He initialization (He et al., 2015)). Furthermore, several studies have successfully used this transfer learning method to classify image-based carbonate-siliciclastic lithofacies (e.g., Pires de Lima et al., 2019; Baraboshkin et al., 2020).

3.4 Image Classification

In this study, the image classification task was conducted using two DCNN architectures: VGG16 (Virtual Geometry Group, Oxford University; Simonyan and Zisserman, 2014) and Inception-Resnet v2 (Szegedy et al., 2017) (Fig. 4A-B; Table 1). These two architectures have been successfully applied to large-scale image recognition tasks (e.g., ImageNet dataset). Therefore, this study applied transfer learning procedure by using these pre-existing architectures and pre-trained initialization weight.

The VGG16 architecture consists of 16 layers that have weights and approximately 138 million total trainable parameters (Simonyan and Zisserman, 2014); hence, it requires a high number of floating-point operations (FLOPS). The network consists of: (i) convolution layers of 3x3 filters with stride of 1; (ii) a maxpooling layer of 2x2 filter with a stride of 2; (iii) padding size; (iv) three fully connected layers at the end of the network output (Fig. 4A). In this paper, the VGG16 architecture was fine-tuned by adding batch normalization layers between convolution and pooling layers and by introducing dropout layers after the fully connected layers (Fig. 4A). The Inception-Resnet architecture is a combination of inception and residual architectures (Fig.

4B; Szegedy et al., 2017), with half as many parameters as the VGG16 (66 million), yet greater depth than VGG16 (up to 152 layers) and slightly lower FLOPS (Bianco et al., 2018). Here, the hybrid inception module consists of convolution layers with different size of filters (1x1, 3x3) and the different inception modules were combined with residual connection, which replaces the pooling layer in the naive inception module (Szegedy et al., 2017). The purpose of the 1x1 convolutions (bottleneck layers) in the inception module is to reduce feature depth of the output in order to match the dimension of input and output and to reduce the number of learning parameters. Furthermore, a dropout layer was added after the final fully connected layer to regularize the model and avoid overfitting (Fig. 4B).

During the training process, the network performance was monitored by comparing the training and validation/test loss as the training steps or epoch increased. Furthermore, several metrics (accuracy, precision, recall and F_1-score) (Fawcett, 2006) and the confusion matrix (Fawcett, 2006) were calculated to evaluate the performance of each network architecture on the prediction of image classes (Table 2). Here, TP is true positive, TN is true negative, FP is false positive, and FN is false negative. Accuracy measures the ratio between correct prediction and total observation whereas precision calculates the proportion of correct identification when the feature is present. Recall measures the sensitivity of the correct prediction and the F1-score is the weighted average of Precision and Recall. Overall, the F1-score is more useful than accuracy in evaluating the prediction performance.

3.5 Object Identification

The object identification task for carbonate rock components was performed to individually locate multiple objects within a single image. This study used one- and two-stage detectors and compared their performances (Fig. 5A-B; Table 1).

In general, a one-stage detection pipeline does not have a region proposal network module that creates a sophisticated end-to-end single detection framework (Liu et al., 2016) (Fig. 5A). For this one-stage detector, SSD (Single Shot Multibox Detector; Liu et al., 2016) and YOLOv3 (You Only Look Once version 3; Redmon and Farhadi, 2018) frameworks were applied. Both use a single convolutional neural network and a fixed set of default bounding boxes. The YOLO method uses Darknet architecture as the backbone and treats the object detection task as a regression problem in order to create different bounding boxes and predict their class probabilities (Redmon et al., 2016). For the SSD method, this study used and compared two different architectures as the backbone: ResNet (RetinaNet) and MobileNet (TensorFlow detection model zoo; https://github.com/tensorflow/models). The main differences between the two architectures are that the RetinaNet employs a feature pyramid network (FPN; Lin et al., 2017b), focal loss function and cosine annealing learning rate (Lin et al., 2017a).

A two-stage detection pipeline commonly consists of two modules: (i) a region proposal module that uses a DCNN to propose regions and the type of object to consider in the region; and (ii) a detection generator module that uses a DCNN for extracting features from the proposed regions and outputting the bounding box and class labels (Fig. 5B; Girshick et al., 2015). For the two-stage detector, this study compared the performance of R-FCN (Region-based Fully Convolutional Neural Networks; Dai et al., 2016) with a ResNet backbone and faster-RCNN

(faster-Region proposal Convolutional Neural Networks; Ren et al., 2015) with Inception-ResNet V2 backbone to classify and identify the different carbonate rock components. This framework uses a similar two-stage pipeline and input image size as faster R-CNN framework. The main differences between faster R-CNN and R-FCN are (i) the R-FCN framework uses ResNet101 to perform feature extraction and (ii) the R-FCN framework is computationally more efficient in proposing regions of interest by using a fully convolutional network.

Similar to the image classification problem, the performance of different object detection models was measured by calculating the pattern analysis, statistical modelling, and computational learning (PASCAL)-style average precision (AP) of each class (Fig. 5C) and the mean average precision (mAP) across all classes. In this detection task, the true positive can be predefined from the Intersection over Union (IoU) threshold value where the IoU measures the overlap between the ground truth bounding box and predicted bounding box (Fig. 5C). Here, the threshold value IoU of 0.5 was used to distinguish between true positive, false positive, and false negative (Fig. 5C). During the training process, network performance was evaluated by monitoring the classification and localization loss, instead of training and validation loss as in the image classification task.

4. RESULTS

4.1 Image Classification

VGG16 – The training process was conducted across 3500 training batches following a one-cycle policy method (Smith, 2018; *fastai* and *Keras* libraries) and the presented model represents the best accuracy from the cross-validation sets (Fig. 6A). Overall, the training loss decreased from > 3 to 0.09 as the number of processed training batches increases (Fig. 6A). While

the validation loss also dropped as the training progressed (from 2.3 to 0.8), it failed to converge with the training curve (Fig. 6A). Application of class weights produced a lower validation loss (~0.55) and improved prediction performance during the testing stage (Fig. 6B and Table 3). Across the six classes of carbonate components, the mean F_1-score for all classes is 0.84 and this architecture performed best at identifying bioclasts and coated grains, but struggled to accurately identify calcite cement (Table 3 and Fig. 6B). It achieved the highest precision for coated grains and dolomite classes (0.90) and the lowest precision for calcite cement (0.72) (Table 3 and Fig. 6B). Bioclasts and micrite have comparably high recall (0.89-0.90) whereas calcite cement has the lowest recall (0.74) (Table 3 and Fig. 6B).

Inception-Resnet v2 – A similar training process was performed for the Inception-Resnet v2 architecture (one-cycle policy and cross validation) but it involved less processed training batches (Fig 6C). In the Inception-Resnet v2 architecture, the training loss converged after 1800 training batches and the loss dropped from > 2 to 0.01 (Fig. 6C). Overall, the validation loss for this model is much lower than for the VGG16 (0.41) and the application of class weights further improved the validation loss (0.19) and prediction performance (Fig. 6C-D). The mean F_1-score of the six studied classes is 0.92 (Table 3) with four out of six classes have F_1-scores higher than 0.9 (Table 3; Fig. 6D). The highest precision and recall were obtained for bioclasts (0.97) and micrite (1.0), respectively (Table 3 and Fig. 6D). In contrast, precision (0.79) and recall (0.82) were lowest for calcite cement (Table 3 and Fig. 6D). Overall, the Inception-ResNet v2 architecture outperformed the VGG16 architecture during the prediction stage (Table 3). Performance across classes was similar between architectures, with coated grains showing the highest recall and calcite cement the lowest recall.

Although this method achieved satisfactory performance in distinguishing among carbonate components, it contains a fundamental weakness in that it can only recognize one class per image even when multiple components are present in the image (Fig. 7). It is therefore not well suited for interpreting petrographic images of complex rock samples. Application of this method alone is therefore unable to replicate human interpretation of carbonate petrography. In an attempt to overcome this fundamental limitation of the image classification technique, this study combined image classification with object identification in order to perform a more robust carbonate petrography description and interpretation.

4.2 Object Identification

Five different object detection frameworks were compared, including both one-stage (YOLOv3, SSD and RetinaNet) and two-stage (faster R-CNN and R-FCN) detectors. The training processes were conducted until the bounding box classification loss has reached less than 0.05 or plateaued. During the training processes, various features were extracted of different classes to help understand and improve the confidence during object recognition processes (Fig. 8). The performance for this object detection task is measured by the average precision on the IoU of 0.5 ($AP_{0.5}$) for each class, mean of average precision (mAP) and inference time required to perform a prediction (Fig. 9A-B and Table 4).

4.2.1. One-stage detectors

YOLOv3 – In this framework, the $AP_{0.5}$ varies from class to class, ranging from 0.59 (peloids) to 0.89 (ooids), with the average is 0.72 (Table 4). The average inference time is 18.4 milliseconds per image. Overall, the YOLOv3 framework has the highest $AP_{0.5}$ for the ooids and

molluscs classes among one-stage detectors (Fig. 9A). The YOLOv3 has the fastest detection speed and the second lowest detection accuracy out of the five different frameworks (Fig. 9B).

SSD– This framework shows the highest $AP_{0.5}$ in the ooids class (0.76) whereas the micrite class has the lowest $AP_{0.5}$ (0.53), with the mAP around 0.66 (Table 4; Fig. 9A). The average detection speed per image is 68.2 milliseconds (Fig. 9B). Compared to other one-stage detectors, this framework has the lowest $AP_{0.5}$ for all studied classes. This network displays the lowest detection accuracy (0.66) and second fastest detection speed (Fig. 9B).

RetinaNet—The RetinaNet has the highest $AP_{0.5}$ for all classes except ooids and molluscs among one-stage detector pipelines (Table 4). Its highest $AP_{0.5}$ is for foraminifera (0.86) while its lowest $AP_{0.5}$ is for micrite (0.69) (Fig. 9A). In addition, this framework has the highest mAP (0.79), but slowest detection speed (101.8 milliseconds per image) when compared to other one-stage detection pipelines (Fig. 9B). Among other one- and two-stage detectors, the RetinaNet has the second highest accuracy and third fastest detection speed (Table 4 and Fig. 9B).

4.2.2. Two-stage detectors

Faster R-CNN – In this framework, the highest $AP_{0.5}$ is observed for ooids (0.95) (Fig. 9 A; Table 4) while the lowest $AP_{0.5}$ is for calcite cement (0.75) (Fig. 9A; Table 4). Different examples of the performance of this framework on the validation or test dataset, specifically the high accuracy in detecting different types of carbonate grains (ooids, molluscs, foraminifera) were evaluated by using petrographic images from published literature (Figs. 10 and 11). Overall, this framework outperformed the other one-stage and two-stage detection frameworks, where it shows the highest $AP_{0.5}$ for all studied classes and mAP (0.84) (Fig. 9B). However, this framework has the

slowest inference time compared to other one-s and two-stage detectors when predicting the different objects in a single image (814 milliseconds per image; Fig. 9B; Table 4).

R-FCN – In this framework, the $AP_{0.5}$ varies substantially among classes, from 0.54 (calcite cement) to 0.92 (ooids) (Fig. 9A; Table 4). Among the five detection frameworks, the mAP of this framework is the third highest (0.76) while the detection speed is the second slowest (277.8 milliseconds per image) (Fig. 9B; Table 4).

4.3 Object Counting and Porosity Calculation

For these tasks, the number of predicted bounding boxes for each class in an image were calculated to perform automatic object counting (Figs. 10 and 11). A total of 15 tests were conducted on samples from an oolitic grainstone facies and this automated object counting method never overcounts the number of grains, with the level of agreement varies from 0.56 to 0.95 (Table 5). In some analysis, the deep learning method can match the manual grain counts, in particular when the grain counts are less than 20 (Table 5). Overall, a good agreement between manual grain counting and deep learning-guided object counting was observed (av. 87% accuracy; Table 5). Furthermore, this study applied a combined adaptive Gaussian threshold and HSV filtering on the detected bounding box for the porosity class (Fig. 12) to isolate the blue-dyed pore space with the surrounding grains/matrix in the samples. Here, the proposed method provides a very close porosity estimation (av. 90% accuracy; Table 5). There are two examples where the estimated porosity values are substantially higher than the actual core plug porosity.

5. DISCUSSION

5.1. Fully Automated Carbonate Petrography

This study presents an evaluation of two different DCNN-based image analysis techniques, image classification and object detection, in describing and classifying carbonate petrographic images. A few prior studies have attempted DCNN-based image classification in order to identify carbonate lithofacies (e.g., Pires de Lima et al., 2019; John and Kanagandran, 2019), but no previous studies have attempted an object detection approach. In general, DCNN-based object detection is more efficient than image classification because: (i) there is no need to split the image into several smaller images to label the image class (Fig. 3) and (ii) it can provide more direct information about the identified features as well as their relative abundances and volumetric contributions. Although the image classification task can be used to predict multiple classes within a single image (Fig. 3), the different predicted classes may not necessarily represent different components, but may just occur because of a misclassified prediction of a single component (e.g., predicted output for micritized coated grains can be as either ooids, peloids or micrite; Fig. 3). In contrast, the object detection task can provide different predicted bounding boxes for an image that individually allocate different grains and cement types with various confidence levels.

Based on the results of the tests run in this study, it appears that the object detection approach is the more direct and effective route toward fully automated carbonate petrographic description. This method also more closely mimics the process of human-based petrographic description where individual carbonate components are identified and characterized simultaneously for a single image. In addition, this method can be used to estimate the number,

size and volume of carbonate grains from the locations and sizes of the detected bounding boxes. This information is crucial for performing textural identification of carbonate lithofacies based on the Dunham textural classification scheme (Dunham, 1962). One of the main reasons for inconsistencies in manual carbonate lithofacies classification is the incorrect identification of carbonate grain types and incorrect estimation of the proportions of different carbonate grain types, cement, and micrite (i.e. misclassifications between matrix and grain supported facies, and fine and coarse-grained facies) (Lokier and Junaibi, 2016). This weakness of manual classification persists in DCNN-based automatic identification of carbonate facies using the Dunham classification scheme (John and Kanagandran, 2019). The object detection approach using DCNN, proposed in this study, could lay a foundation towards a more robust method for classifying lithofacies based on the detection and quantification of individual carbonate components. Similar conclusions have been reached in other fields, such as medical imaging, object detection outperforms image classification for characterizing medical microscopic images (Wang et al., 2019).

In both image classification and object detection tasks, calcite cement and micrite are recognized less accurately than carbonate grains. There are two contributing factors for this difference: (i) calcite cement and micrite often have diffuse object boundaries and (ii) carbonate grains can be micritized or recrystallized, blurring the conceptual distinctions between micrite, cement, and grains. For example, micritized foraminifera can be labelled as either foraminifera or micrite and the more appropriate label will depend upon the study objective. Furthermore, the training dataset used in this study consists predominantly of Permo-Triassic samples; hence, the

ability to recognize certain types of carbonate grains which are not present in this interval may be less accurate.

The object detection approach entails a tradeoff between detection speed and accuracy (Fig. 9B). In this study, the highest accuracy is obtained from the faster R-CNN method, but it has the slowest detection speed. This issue is not significant for the analysis of a small number of static images, where the overall computational time is still sufficiently short for research purposes. However, if one were to conduct the object detection task in real time, for example on streaming video, the YOLOv3 or RetinaNet may combine speed and accuracy in more useful ways (Fig. 9B; Table 4) (see supplementary material Video_1_Petrography streaming or access the file from https://github.com/ardikoes/Carbonate_ML.git). Furthermore, object detection tools can increase the accuracy and efficiency of lithological description and classification from both subsurface cores and wellbore cuttings. In such cases, carbonate reservoir characterization for hydrocarbon exploration processes could benefit from the application of this approach.

5.2. Performance Evaluation

This study compared different deep learning models for both image classification and object identification tasks. For classifying different carbonate components through image classification, the inception-ResNet architecture has clearly outperformed the VGG16 model as the VGG16 model fails to converge, hence giving lower accuracy when performing new predictions due to probably overfitting problem. Furthermore, the VGG16 model require more training time and computing resources than the inception-ResNet model, although the latter has a much deeper layer. The creator of VGG16 model (Simonyan and Zisserman, 2014) encountered a similar issue and suggested pre-training of the model in smaller networks to mitigate this issue.

Pre-training is time consuming, however, and may not be worth the sacrifice of speed. Similar optimization processes, such as cross validation and one-cycle policy (Smith, 2018), did not significantly improve the VGG16 model in this study relative to the inception-ResNet model. One reason why the VGG16 model performs poorly could be the limited size of the dataset (even after data augmentation). The total number of trainable parameters in VGG16 is twice that of the inception-ResNet model, which makes the VGG16 mode data "hungry" and also prone to overfitting. In addition, there is the vanishing or exploding gradient problem for very deep networks, as highlighted by Goodfellow et al. (2016), which makes the network unstable and learning process harder to progress. This problem is not encountered in the ResNet architecture because the residual blocks allow the gradient to flow continuously or uninterrupted. Although it is possible to increase the performance of VGG16 by fine tuning the hyper-parameters, the high accuracy achieved by the inception-ResNet architecture made this approach unnecessary.

For the object detection task, this study utilized either ResNet or Inception-ResNet architecture as the feature extractor, hence the issues encountered in the image classification task were not present. The comparison of five distinct object detection frameworks show that there is a significant tradeoff between detection accuracy and speed (Fig. 9B), and this tradeoff appears to be general to the method (Huang et al., 2017). In our study, the faster R-CNN (two-stage detectors) show the highest accuracy but slowest detection speed. The main reason is the one-stage detectors run the object detection task directly over a dense sampling of proposed bounding boxes without first performing a region proposal step (Fig. 5). Furthermore, the one stage detectors struggle when the image contains multiple objects of different classes. Soviany and Ionescu (2018) highlighted this issue and utilized PASCAL VOC 2007 dataset to propose an image

difficulty predictor (i.e. differentiate easy versus hard image) to split the test image instead of random split as a pre-treatment stage before the detection task. By feeding the difficult image to two-stage detectors and easy image to one-stage detectors, the tradeoff between detection speed and accuracy can be reduced.

5.3. Improving Accuracy with Limited and Imbalanced Datasets

DCNN-based methods are generally data hungry when performing a narrow task. These datasets typically comprise tens of millions of images, divided into hundreds of classes. Depending on the applications, collecting this quantity of data for geological applications, including petrographic images, can be very difficult to achieve. Therefore, different algorithm optimizations are necessary to better train the networks. Network optimization is often achieved by applying various hyperparameters fine-tuning approaches such as weight initialization, batch normalization, and regularization (Goodfellow et al., 2016). Here, in addition to applying these methods, a one-cycle policy (Smith, 2018) by using the *fastai* library was used to optimize the learning rate. One of the important methods under the one-cycle policy is the use of a cosine scheduler for learning rate decay that allows the network to train better and converge faster. Furthermore, this study applied different data augmentation procedures, including: (i) transforming the image (reshape, filter, transform, and introduce noise; *Keras* library) and (ii) label Mixup (class mixing in the training dataset) and smoothing (decrease the label confidence to 0.9 to replace one hot encoded labelling) using the *fastai* library. The first method, which is commonly used in DCNN-based image classification analysis (e.g., Perez and Wang, 2017), achieved little improvement in the classification and detection accuracies (+ 0.2 to 1.5%). This failure occurred because such data augmentation does not increase the variability of dataset

significantly. The second method shows a more significant impact on the overall prediction performance and improved the average precision of all classes (AP0.5: + 2 to 5%; Table 4). This approach was more successful because it makes the neural networks more robust to noise and model overfitting.

Another common and significant challenge is class imbalance in the training dataset, which has been observed in other research fields (Buda et al., 2018), such as medical analysis (Grzymala-Busse et al., 2004) and remote sensing (Johnson et al., 2013). Many studies have discussed the significant detrimental impact of this issue to the learning processes and proposed different approaches to attenuate its effect (Buda et al., 2018 and references therein). Here, two methods were compared and evaluated. The class weight method outperformed the oversampling-undersampling method in this study (Fig. 6A and C). The class weight method allows the networks to put more emphasis on the minority classes and less emphasis on the majority classes by equally penalizing their weights. While oversampling has been suggested to outperform other mitigation methods and does not cause overfitting in DCNN (Buda et al., 2018), the main disadvantages of oversampling-undersampling method are the possibility to remove important features in the original dataset and synthetically creating dataset which may ignore the actual distribution of data abundance and availability. In addition, the comparison conducted herein shows that the sampling method tends to cause overfitting (Fig. 6A and C). Furthermore, this paper used a cost-sensitive learning by applying a focal loss function (Lin et al., 2017a) instead of a categorical cross-entropy loss function to attenuate the class imbalance issue. Several studies in different fields have shown the advantage of applying a focal loss function to improve the network and prediction performance (Chatterjee et al., 2019; Sun et al., 2019; Celik et al., 2020). In

this study, these augmentation and class imbalance mitigation methods, coupled with one-cycle policy, were the most successful approach to mitigating the challenges of a limited, imbalanced training dataset.

6. FUTURE RECOMMENDATIONS

Deep Convolution Neural Networks show substantial promise in the tasks of image classification and object identification in carbonate rocks. The main advantage of DCNN-based object detection is the ability to accurately locate and identify different components individually. In addition, this method enables real-time carbonate petrography with optimum accuracy and speed.

There remain several key areas for improvement. First, classes may be mislabeled even when the model loss is very low (< 0.1). This phenomenon may occur because of the complex nature and diagenetic alteration of carbonate components (e.g., completely recrystallized skeletal grains often mislabeled as calcite cement or micritized grain can be mistakenly classified as micrite). In general, grains tend to be bounded objects that can be fit well with bounding boxes whereas cements tend to be interstitial and so are not easily isolated by bounding boxes if they are continuous in the pore space between grains. An expanded annotated training dataset for these classes and other carbonate grain types across geologic timescales may improve the overall learning processes, performance and prediction. In addition, a higher-level skeletal grain classification (e.g., phylum or class/genera level) is necessary for applications in carbonate paleoecology and biostratigraphy as well as for refined interpretation of depositional environment. Furthermore, certain carbonate grains, especially skeletal grains, have limited

stratigraphic distributions and so a tool that can be applied to any particular time interval would require at least some training images of the same age.

Second, smaller carbonate grains (< 50 microns) are either difficult to detect or detected with low accuracy (< 30 % confidence) during the object detection task. The small object detection problem has also been encountered in recognizing real world objects using DCNN (Hu et al., 2018). This problem occurs because of the iterative combination of low-level features during the convolution process which make small object difficult to detect and also the saliency DCNN model where it tends to focus on bigger objects in an image (Hu et al., 2018) because small objects are represented by fewer pixels. To mitigate this issue, Hu et al. (2018) suggested multiscale feature extractions from different convolution layers and combined these features to a 1-D vector. Another method that may be applied is remote sensing region-based convolutional neural network (R^2-CNN; Pang et al., 2019) by using Tiny-Net and intermediate global attention block as the backbone.

Furthermore, future work should focus to apply either supervised or unsupervised semantic and instance segmentation (e.g., mask R-CNN; He et al., 2017 or U-Net; Ronneberger et al., 2015) as the next step after object detection in order to provide a more reliable interpretation of carbonate petrography and lithofacies classification. This method will also allow a better characterization of pore types in carbonate rocks. In addition, future work could explore the application of LSTM model (Long Short-Term Memory) with Recurrent Neural Networks (RNN) (Hossain et al., 2019) to provide automatic carbonate microfacies description and the application of AttnGAN (Xu et al., 2018) to create carbonate thin section image from image description. When the size of the training dataset cannot be significantly increased, an alternative approach to create

significantly larger and statistically meaningful synthetic dataset is to apply Generative Adversarial Networks (GANs) (Goodfellow et al., 2014; Goodfellow et al., 2016; Choi et al., 2018). Several survey studies have shown how image augmentation from GANs method can significantly improve the performance of the algorithms and provide a more close-to-real world examples for the model to learn (Wu et al., 2017; Pan et al., 2019). Nanjo and Tanaka (2020) have successfully applied this method to reconstruct and interpret carbonate thin section image from a sketch of labelled image using GAN method. Hence, further development of this method maybe worth exploring in the future.

7. CONCLUSIONS

This study has successfully performed a fully automated identification of different carbonate constituents by coupling DCNN-based image classification and object detection approaches. The results show a promising potential of deep learning application to automate carbonate petrography description and interpretation where high level performance was achieved from both image classification and object detection tasks (> 80% accuracy) even with limited training data. Furthermore, this method can perform grain identification and quantification simultaneously at much greater speed than human-based analysis.

For the image classification task, the Inception-ResNet architecture shows a better performance than the VGG16 architecture in terms of training speed and classification accuracy. The presence of residual connections in Inception-ResNet network helps the network to train better even with limited dataset. However, this study determined that performing image classification alone would not suffice to mimic how a human describes or interprets a carbonate petrographic image. Therefore, this study, for the first time, used DCNN-based object detection

task to compliment image classification analysis in order to perform a more robust and closer to human approach of carbonate petrography interpretation. In the object detection task, the highest detection accuracy is obtained from the faster R-CNN framework, although its detection speed is the slowest among other frameworks. Depending on the type of analysis, the other four object detection frameworks may also be useful to perform fully automated carbonate petrography. While for image analysis the detection speed is not really an issue, if to apply this system for real-time petrography analysis the YOLOv3 or RetinaNet will perform better.

In addition, this study explored different approaches to improve network performance and accuracy in the availability of a limited dataset and proposed that data augmentation (label mixup and smoothing) generally works better than other data augmentation procedures (resize, flip and introduce noise to the image). Coupling the proposed data augmentation with a focal loss function and cosine learning rate scheduler can improve the accuracy ($AP_{0.5}$: +2 to 5%). Furthermore, this study discussed the various future pathways that can be considered to further advance the application of deep learning to solve carbonate petrography and geoscience-related problems, in general. Compiling more datasets across geologic timescales and different grain components is a necessary step in order to improve the prediction capability and the ability for the model to generalize and perform a robust fully automated carbonate petrography. Ultimately, application and development of this advanced method could aid to understand variation in carbonate factories across geologic timescales and optimize the characterization of carbonate reservoir.

ACKNOWLEDGMENTS

REFERENCES

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M. and Kudlur, M., 2016. Tensorflow: A system for large-scale machine learning. In 12th {USENIX} Symposium on Operating Systems Design and Implementation (OSDI), pp. 265-283.

Adams, A. and Mackenzie, I., 1998. Carbonate sediments and rocks under the microscope: a colour atlas. CRC Press, p. 181.

Al-Ramadan, K., Koeshidayatullah, A., Cantrell, D. and Swart, P.K., 2019. Impact of basin architecture on diagenesis and dolomitization in a fault-bounded carbonate platform: outcrop analogue of a pre-salt carbonate reservoir, Red Sea rift, NW Saudi Arabia. Petroleum Geoscience, 26 (3), pp. 448-461.

Alejo, R., Sotoca, J.M., Valdovinos, R.M. and Toribio, P., 2010, June. Edited nearest neighbor rule for improving neural networks classifications. In International Symposium on Neural Networks (pp. 303-310). Springer, Berlin, Heidelberg.

Amao, A.O., Al-Ramadan, K. and Koeshidayatullah, A., 2016. Automated mineralogical methodology to study carbonate grain microstructure: an example from oncoids. Environmental earth sciences, 75(8), p.666.

Asmussen, P., Conrad, O., Günther, A., Kirsch, M. and Riller, U., 2015. Semi-automatic segmentation of petrographic thin section images using a "seeded-region growing algorithm" with an application to characterize wheathered subarkose sandstone. Computers & Geosciences, 83, pp.89-99.

Ashena, R., and Thonhauser, G., 2015. Application of Artificial Neural Networks in Geoscience and Petroleum Industry, in: Artificial Intelligent Approaches in Petroleum Geosciences. Springer International Publishing, Cham, pp. 127–166.

Baraboshkin, E.E., Ismailova, L.S., Orlov, D.M., Zhukovskaya, E.A., Kalmykov, G.A., Khotylev, O.V., Baraboshkin, E.Y. and Koroteev, D.A., 2020. Deep convolutions for in-depth automated rock typing. Computers & Geosciences, 135, p.104330.

Bengio, Y., 2012, June. Deep learning of representations for unsupervised and transfer learning. In Proceedings of ICML workshop on unsupervised and transfer learning, pp. 17-36.

Bianco, S., Cadene, R., Celona, L. and Napoletano, P., 2018. Benchmark analysis of representative deep neural network architectures. IEEE Access, 6, pp. 64270-64277.

Bradski, G., 2000. The OpenCV library. Dr. Dobb's J. Softw. Tools 120, pp. 122–125.

Bridle, J.S., 1990. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In Neurocomputing. Springer, Berlin, Heidelberg, pp. 227-236.

Burchette, T.P., 2019. Are carbonate reservoirs 'difficult'? GEOExPro, v.16 (5).

Celik, Y., Talo, M., Yildirim, O., Karabatak, M. and Acharya, U.R., 2020. Automated invasive ductal carcinoma detection based using deep transfer learning with whole-slide images. Pattern Recognition Letters.

Canchumuni, S.W., Emerick, A.A. and Pacheco, M.A.C., 2019. History matching geological facies models based on ensemble smoother and deep generative models. Journal of Petroleum Science and Engineering, 177, pp.941-958.

Chatterjee, S., Dutta, R.K., Ganguly, D., Chatterjee, K. and Roy, S., 2019. Bengali handwritten character classification using transfer learning on deep convolutional neural network. arXiv preprint arXiv:1902.11133.

Chawla, N.V., Bowyer, K.W., Hall, L.O. and Kegelmeyer, W.P., 2002. SMOTE: synthetic minority over-sampling technique. Journal of artificial intelligence research, 16, pp.321-357.

Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S. and Choo, J., 2018. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 8789-8797.

Chollet, F. Keras. https://github.com/fchollet/keras, 2015.

Dai, J., Li, Y., He, K. and Sun, J., 2016. R-FCN: Object detection via region-based fully convolutional networks. In Advances in neural information processing systems, pp. 379-387.

de Lima, R.P., Suriamin, F., Marfurt, K.J. and Pranter, M.J., 2019. Convolutional neural networks as aid in core lithofacies classification. Interpretation, 7(3), pp. SF27-SF40.

de Lima, R.P., Duarte, D., Nicholson, C., Slatt, R. and Marfurt, K.J., 2020. Petrographic microfacies classification with deep convolutional neural networks. Computers & Geosciences, p.104481.

Della Porta, G. and Wright, P.V., 2009. www.carbonateworld.com. Carbonate World: a web tutorial for the petrographic analysis of carbonate rocks.

Deng, J., Dong, W., Socher, R., Li, L.J., Li, K. and Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 248-255.

Dunham, R.J., 1962. Classification of carbonate rocks according to depositional textures. AAPG Special Volumes: Classification of Carbonate Rocks—A Symposium, pp. 108-121.

Eberli, G.P., Baechle, G.T., Anselmetti, F.S. and Incze, M.L., 2003. Factors controlling elastic properties in carbonate sediments and rocks. The Leading Edge, 22(7), pp. 654-660.

Eltom, H.A., Abdullatif, O.M., Babalola, L.O., Bashari, M.A., Yassin, M., Osman, M.S. and Abdulraziq, A.M., 2017. Integration of facies architecture, ooid granulometry and morphology for prediction of reservoir quality, Lower Triassic Khuff Formation, Saudi Arabia. Petroleum Geoscience, 23(2), pp. 177-189.

Fawcett, T., 2006. An introduction to ROC analysis. Pattern Recognition Letters. 27, pp. 861–874.

Flügel, E., 2013. Microfacies of carbonate rocks: analysis, interpretation and application. Springer Science & Business Media, p. 1006.

Geng, Z., Wu, X., Shi, Y. and Fomel, S., 2020. Deep learning for relative geologic time and seismic horizons. Geophysics, 85(4), pp. 87-100.

Girshick, R., 2015. Fast R-CNN. In Proceedings of the IEEE international conference on computer vision, pp. 1440-1448.

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y., 2014. Generative adversarial nets. In Advances in neural information processing systems (pp. 2672-2680).

Goodfellow, I., Bengio, Y. and Courville, A., 2016. Deep learning. MIT press.

Grove, C. and Jerram, D.A., 2011. jPOR: An ImageJ macro to quantify total optical porosity from blue-stained thin sections. Computers & Geosciences, 37(11), pp.1850-1859.

Grzymala-Busse, J. W., Goodwin, L. K., Grzymala-Busse, W. J., and Zheng, X., 2004. An approach to imbalanced data sets based on changing rule strength. In Rough-neural computing (pp. 543–553). Springer.

He, K., Zhang, X., Ren, S. and Sun, J., 2015. Spatial pyramid pooling in deep convolutional networks for visual recognition. IEEE transactions on pattern analysis and machine intelligence, 37(9), pp.1904-1916.

He, K., Gkioxari, G., Dollár, P. and Girshick, R., 2017. Mask R-CNN. In Proceedings of the IEEE international conference on computer vision and pattern recognition, pp. 2961-2969.

Hossain, M.Z., Sohel, F., Shiratuddin, M.F. and Laga, H., 2019. A comprehensive survey of deep learning for image captioning. ACM Computing Surveys (CSUR), 51(6), pp.1-36.

Howard, J. and Gugger, S., 2020. Fastai: A layered API for deep learning. Information, 11(2), p.108.

Hu, G.X., Yang, Z., Hu, L., Huang, L. and Han, J.M., 2018. Small object detection with multiscale features. International Journal of Digital Multimedia Broadcasting, 2018, Article ID 4546896, pp. 1-10.

Huang, J., Rathod, V., Sun, C., Zhu, M., Korattikara, A., Fathi, A., Fischer, I., Wojna, Z., Song, Y., Guadarrama, S. and Murphy, K., 2017. Speed/accuracy trade-offs for modern convolutional object detectors. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7310-7311.

Huang, L., Dong, X. and Clee, T.E., 2017. A scalable deep learning platform for identifying geologic features from seismic attributes. The Leading Edge, 36(3), pp.249-256.

Hunter, J.D., 2007. Matplotlib: A 2D graphics environment. Computing in Science & Engineering, 9(3), pp. 90-95.

Izadi, H., Sadri, J. and Bayati, M., 2017. An intelligent system for mineral identification in thin sections based on a cascade approach. Computers & Geosciences, 99, pp. 37-49.

John, C. and Kanagandran, S., 2019. AI to Improve the Reliability and Reproducibility of Descriptive Data: A Case Study Using Convolutional Neural Networks to Recognize Carbonate Facies in Cores. In AAPG Annual Convention and Exhibition, San Antonio, TX.

Johnson, B. A., Tateishi, R., and Hoan, N. T., 2013. A hybrid pansharpening approach and multiscale object-based image analysis for mapping diseased pine and oak trees. International Journal of Remote Sensing, 34(20), 6969–6982.

Johnson, J.M. and Khoshgoftaar, T.M., 2019. Survey on deep learning with class imbalance. Journal of Big Data, 6(1), p.27.

King, G. and Zeng, L., 2001. Logistic regression in rare events data. Political analysis, 9(2), pp.137-163.

Kingma, D.P. and Ba, J., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Koeshidayatullah, A., Al-Ramadan, K., Collier, R. and Hughes, G.W., 2016a. Variations in architecture and cyclicity in fault-bounded carbonate platforms: Early Miocene Red Sea Rift, NW Saudi Arabia. Marine and Petroleum Geology, 70, pp.77-92.

Kohavi, R., 1995. A study of cross-validation and bootstrap for accuracy estimation and model selection. Proc., 14th Int. Joint Conf. on Artificial Intelligence, Morgan Kaufmann, San Francisco, pp. 1137-1143.

Krizhevsky, A., Nair, V. and Hinton, G., 2009. Cifar-10 and cifar-100 datasets. URl: https://www. cs. toronto. edu/kriz/cifar. html, 6.

Krizhevsky, A., Sutskever, I. and Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pp. 1097-1105.

LeCun, Y., Boser, B., Denker, J.S., Henderson, D., Howard, R.E., Hubbard, W., Jackel, L.D., 1989. Backpropagation applied to handwritten zip code recognition. Neural Computation, 1, 541–551.

LeCun, Y., Bengio, Y. and Hinton, G., 2015. Deep learning. Nature, 521, 436–444.

Lemaître, G., Nogueira, F. and Aridas, C.K., 2017. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. The Journal of Machine Learning Research, 18(1), pp.559-563.

Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P. and Zitnick, C.L., 2014, September. Microsoft coco: Common objects in context. In European conference on computer vision, Springer, Cham, pp. 740-755.

Lin, T.Y., Goyal, P., Girshick, R., He, K. and Dollár, P., 2017a. Focal loss for dense object detection. In Proceedings of the IEEE international conference on computer vision and pattern recognition, pp. 2980-2988.

Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B. and Belongie, S., 2017b. Feature pyramid networks for object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 2117-2125.

Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y. and Berg, A.C., 2016. SSD: Single shot multibox detector. In European conference on computer vision, Springer, Cham, pp. 21-37.

Liu, Y., Sun, W. and Durlofsky, L.J., 2019. A deep-learning-based geological parameterization for history matching complex models. Mathematical Geosciences, 51(6), pp.725-766.

Lokier, S.W. and Al Junaibi, M., 2016. The petrographic description of carbonate facies: are we all speaking the same language?. Sedimentology, 63(7), pp. 1843-1885.

Lynda, N.O., 2019. Systematic survey of convolutional neural network in satellite image classification for geological mapping. In 15th International Conference on Electronics, Computer and Computation (ICECCO), pp. 1-6.

Mosser, L., Dubrule, O. and Blunt, M.J., 2017. Reconstruction of three-dimensional porous media using generative adversarial neural networks. Physical Review E, 96(4), pp.043309.

Nanjo, T. and Tanaka, S., 2020, January. Carbonate Lithology Identification with Generative Adversarial Networks. In International Petroleum Technology Conference. International Petroleum Technology Conference.

Oliphant, T.E., 2006. A guide to NumPy. Trelgol Publishing. USA, Vol. 1, p. 85.

Pan, Z., Yu, W., Yi, X., Khan, A., Yuan, F. and Zheng, Y., 2019. Recent progress on generative adversarial networks (GANs): A survey. IEEE Access, 7, pp. 36322-36333.

Pang, J., Li, C., Shi, J., Xu, Z. and Feng, H., 2019. R^2-CNN: Fast Tiny Object Detection in Large-Scale Remote Sensing Images. IEEE Transactions on Geoscience and Remote Sensing, 57(8), pp. 5512-5524.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L. and Desmaison, A., 2019. PyTorch: An imperative style, high-performance deep learning library. In Advances in Neural Information Processing Systems, pp. 8024-8035.

Payne, J.L., Lehrmann, D.J., Wei, J. and Knoll, A.H., 2006. The pattern and timing of biotic recovery from the end-Permian extinction on the Great Bank of Guizhou, Guizhou Province, China. Palaios, 21(1), pp.63-85.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E., 2011. Scikit-learn: machine learning in Python. J. Mach. Learn. Res. 12, pp. 2825–2830.

Perez, L. and Wang, J., 2017. The effectiveness of data augmentation in image classification using deep learning. arXiv preprint arXiv:1712.04621.

Qian, F., Yin, M., Liu, X.Y., Wang, Y.J., Lu, C. and Hu, G.M., 2018. Unsupervised seismic facies analysis via deep convolutional autoencoders. Geophysics, 83(3), pp. A39-A43.

Redmon, J., Divvala, S., Girshick, R. and Farhadi, A., 2016. You only look once: Unified, real-time object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 779-788.

Redmon, J. and Farhadi, A., 2018. Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767.

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., Prabhat, 2019. Deep learning and process understanding for data-driven Earth system science. Nature, 566, pp. 195–204.

Ren, S., He, K., Girshick, R. and Sun, J., 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In Advances in neural information processing systems, pp. 91-99.

Ronneberger, O., Fischer, P. and Brox, T., 2015, October. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical image computing and computer-assisted intervention (pp. 234-241). Springer, Cham.

Samek, W., Binder, A., Montavon, G., Lapuschkin, S. and Müller, K.R., 2016. Evaluating the visualization of what a deep neural network has learned. IEEE transactions on neural networks and learning systems, 28(11), pp. 2660-2673.

Saporetti, C.M., da Fonseca, L.G., Pereira, E. and de Oliveira, L.C., 2018. Machine learning approaches for petrographic classification of carbonate-siliciclastic rocks using well logs and textural information. Journal of Applied Geophysics, 155, pp. 217-225.

Scholle, P.A. and Ulmer-Scholle, D.S., 2003. A Color Guide to the Petrography of Carbonate Rocks: Grains, Textures, Porosity, Diagenesis, AAPG Memoir 77. Tulsa, OK, USA.

Shin, H.C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D. and Summers, R.M., 2016. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. IEEE transactions on medical imaging, 35(5), pp. 1285-1298.

Silva, A.A., Tavares, M.W., Carrasquilla, A., Misságia, R. and Ceia, M., 2020. Petrofacies classification using machine learning algorithms. Geophysics, 85(4), pp. WA101-WA113.

Simonyan, K. and Zisserman, A., 2014. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.

Smith, L.N., 2018. A disciplined approach to neural network hyper-parameters: Part 1--learning rate, batch size, momentum, and weight decay. arXiv preprint arXiv:1803.09820.

Soviany, P. and Ionescu, R.T., 2018. Optimizing the trade-off between single-stage and two-stage deep object detectors using image difficulty prediction. In 20th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC), IEEE, pp. 209-214.

Sun, X., Dong, K., Ma, L., Sutcliffe, R., He, F., Chen, S. and Feng, J., 2019. Drug-drug interaction extraction via recurrent hybrid convolutional neural networks with an improved focal loss. Entropy, 21(1), p.37.

Szegedy, C., Ioffe, S., Vanhoucke, V. and Alemi, A.A., 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In Thirty-first AAAI conference on artificial intelligence.

Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C. and Liu, C., 2018, October. A survey on deep transfer learning. In International conference on artificial neural networks, Springer, Cham, pp. 270-279.

Tzutalin, LabelImg, Git code (2015). https://github.com/tzutalin/label Img

Wang Q, Bi S, Sun M, Wang Y, Wang D, Yang S., 2019. Deep learning approach to peripheral leukocyte recognition. PLoS ONE 14(6): e0218808.

Wu, X., Xu, K. and Hall, P., 2017. A survey of image synthesis and editing with generative adversarial networks. Tsinghua Science and Technology, 22(6), pp. 660-674.

Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X. and He, X., 2018. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 1316-1324.

Zhang, Z., He, T., Zhang, H., Zhang, Z., Xie, J. and Li, M., 2019. Bag of freebies for training object detection neural networks. arXiv preprint arXiv:1902.04103.

Zhong, Z., Sun, A.Y., Yang, Q. and Ouyang, Q., 2019. A deep learning approach to anomaly detection in geological carbon sequestration sites using pressure measurements. Journal of hydrology, 573, pp.885-894.

FIGURES AND FIGURE CAPTIONS
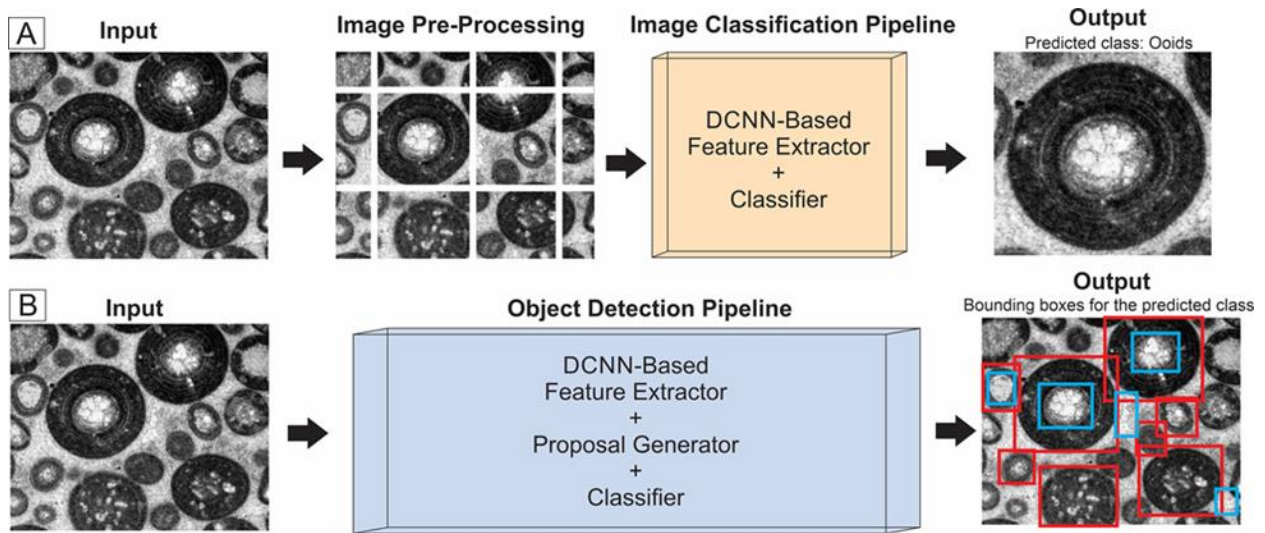
Figure 1. Comparison of two DCNN-based image analysis protocols (modified from Wang et al., 2019)**.** (A) Image classification pipeline. (B) Object detection pipeline. The object detection pipeline differs from the image classification pipeline in that it does not require image segmentation and is capable of returning multiple classes from a single image.
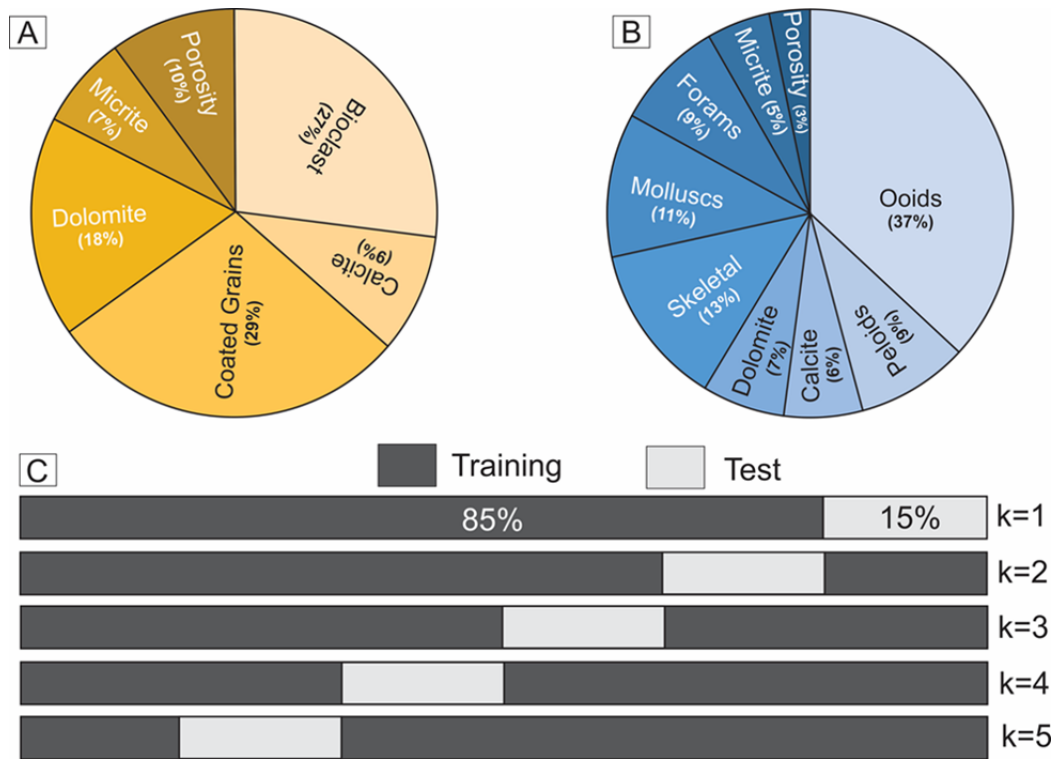
Figure 2. Distribution of the dataset for each DCNN-based image analysis. (A) Image classification dataset. (B) Object detection dataset. (C) Cross validation set (k=5) used in this study, where the dataset is divided into 85% training data and 15% validation/test data.
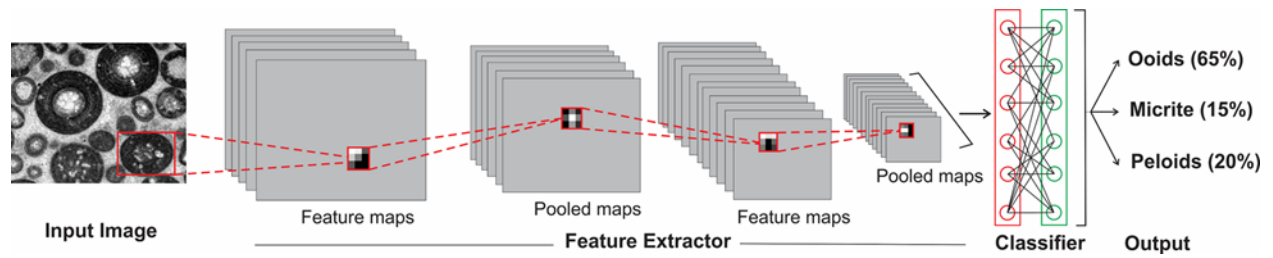
Figure 3. A typical architecture of CNN with multiple convolutional and activation layers as feature maps and pooling layers as pooled maps (modified from Krizhevsky et al., 2012). In addition, there is a fully connected layer at the end of the model as the classifier.
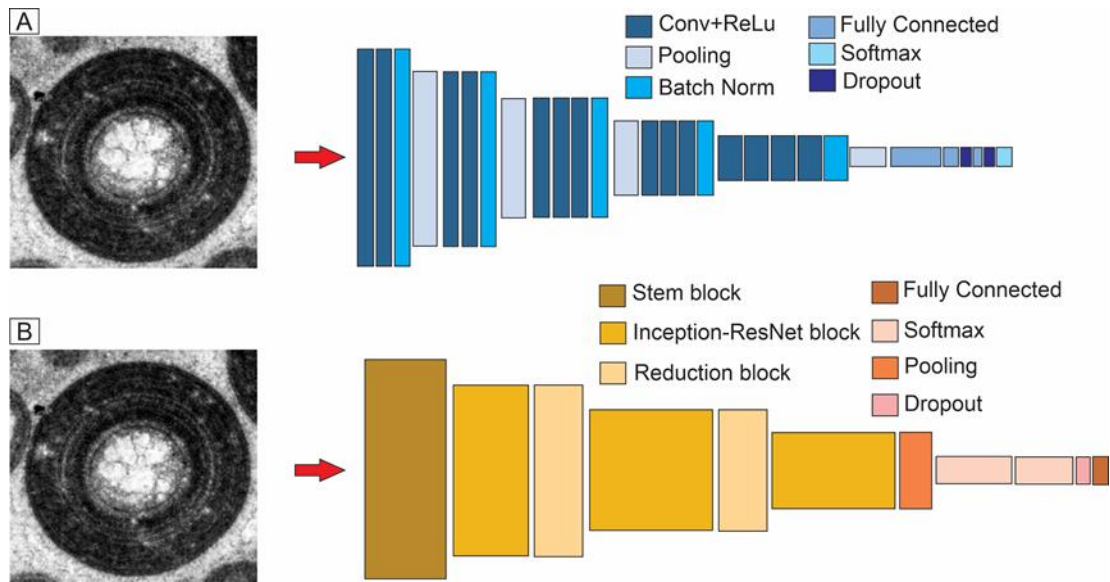
Figure 4. Schematic representation of the two different DCNN-based image classification architectures used for this work. (A) VGG16 architecture. (B) Inception-ResNet v2 architecture.
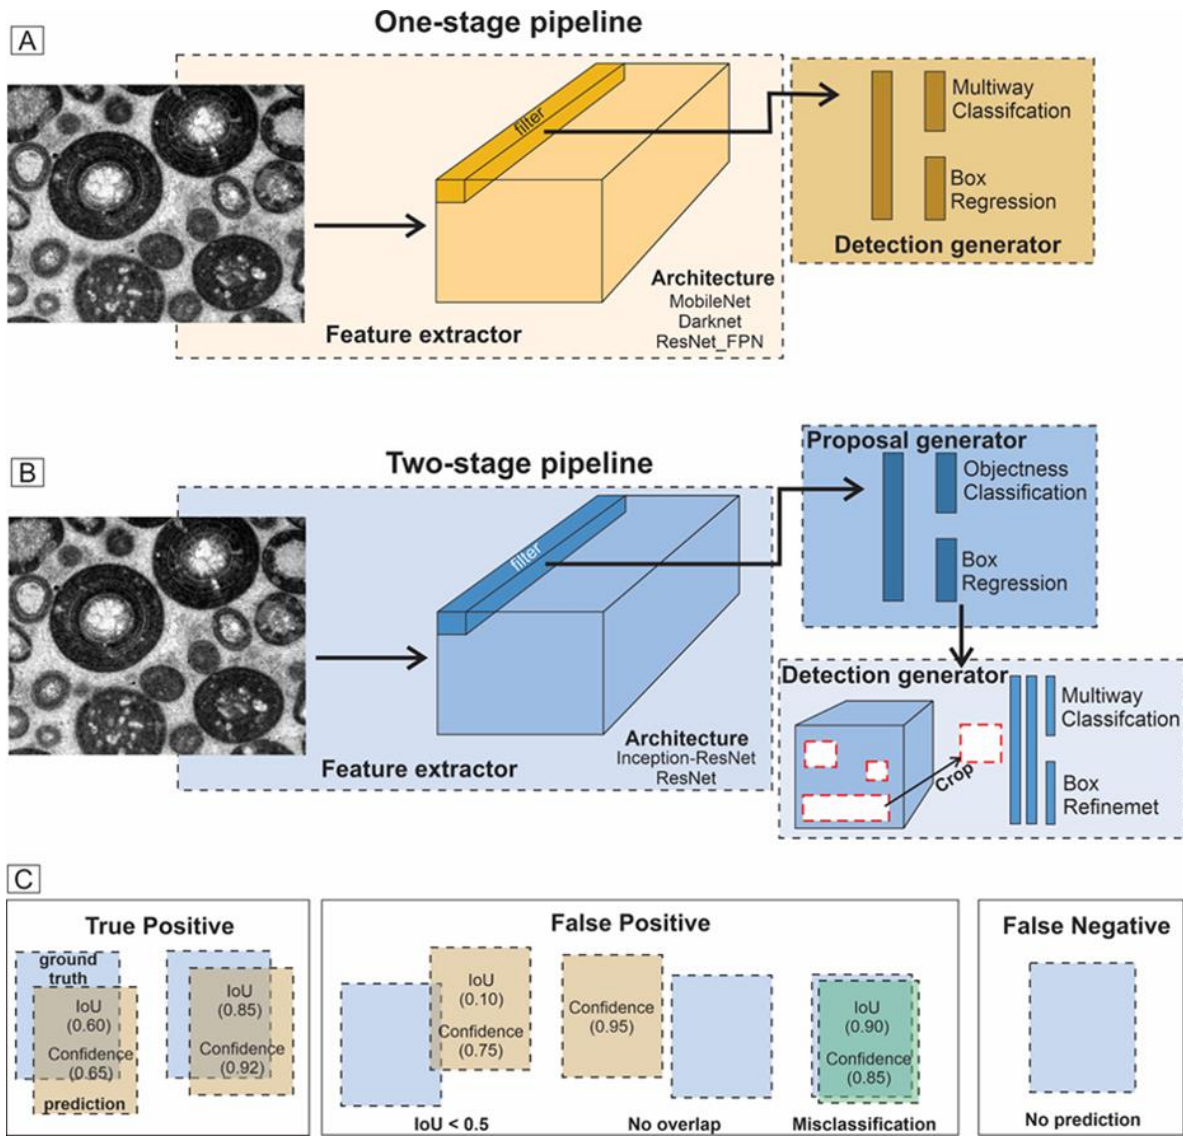
Figure 5. Visualization of object detection pipelines (after Huang et al., 2017). (A) One-stage detector framework showing an end-to-end detection pipeline from feature extraction to detection. (B) Two-stage detector framework with an addition of region proposal network (proposal generator) after the feature extraction and before the detection processes. (C) Illustration of true positive, false positive and false negative with IoU threshold value of 0.5.
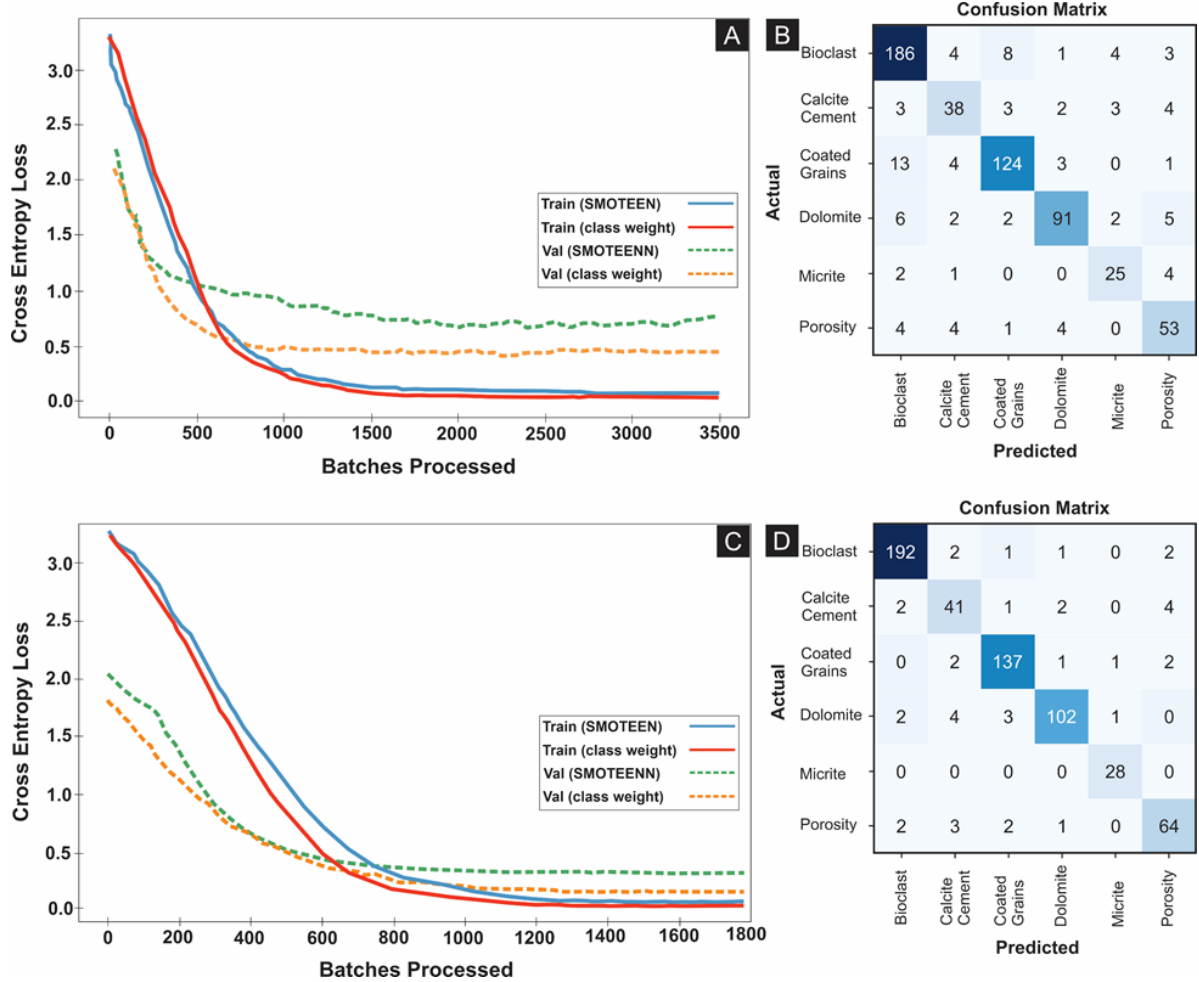
Figure 6. Evaluation of network performance during the training and prediction processes. (A-B) Model performance (loss) during training process for VGG16 architecture and the resulted prediction presented as a confusion matrix. (C-D) Model performance (loss) during training process for Inception-Resnet v2 and the resulted prediction presented as a confusion matrix. Overall, the Inception-Resnet network outperforms the VGG16 network.
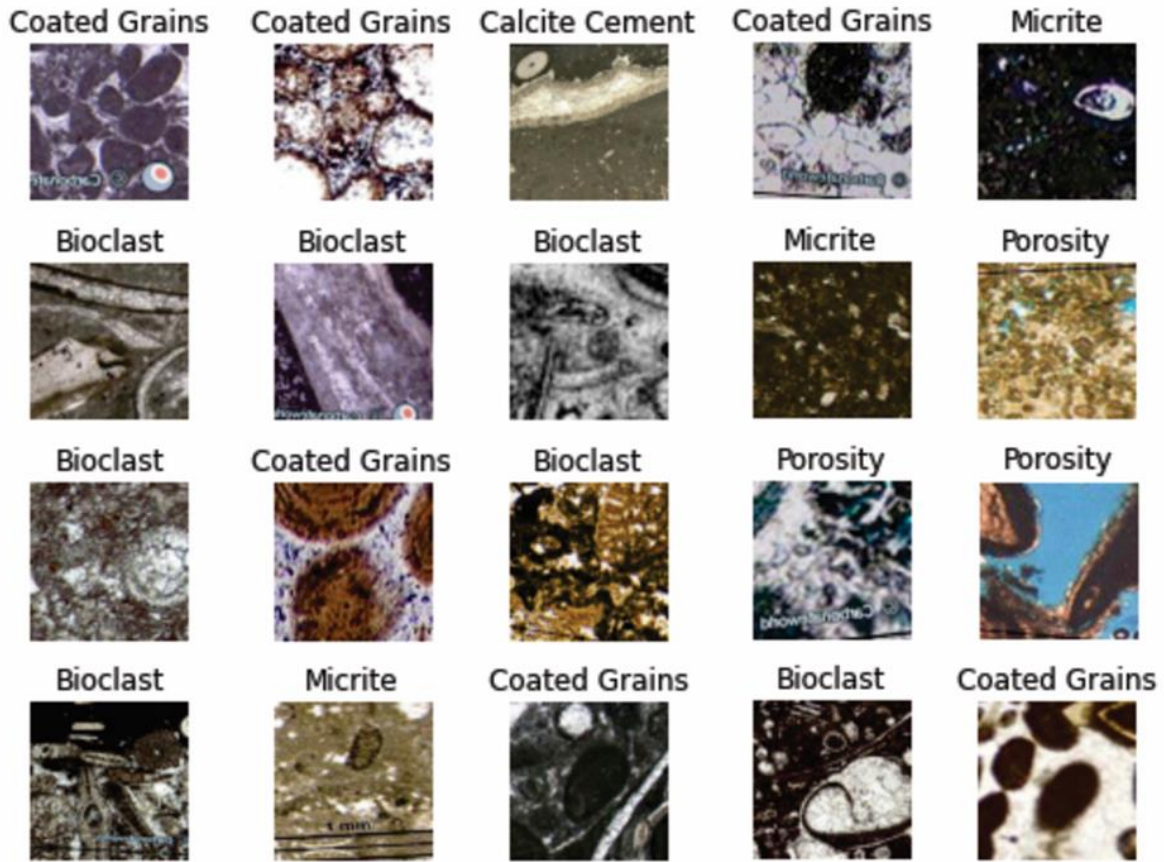
Figure 7. Visualization of the correctly predicted image classes. Note that although the predicted class is present in the image, there are other unidentified objects that should belong to other classes.
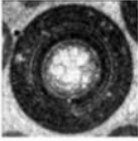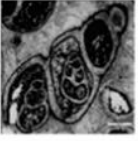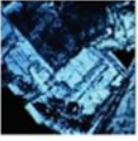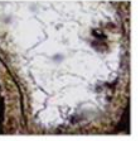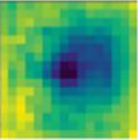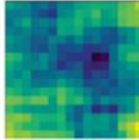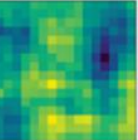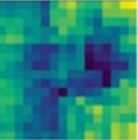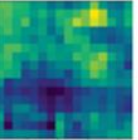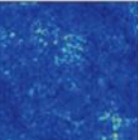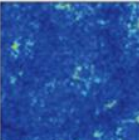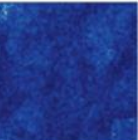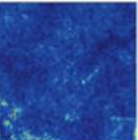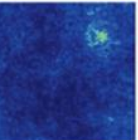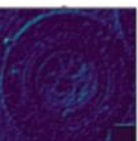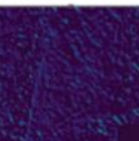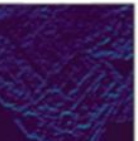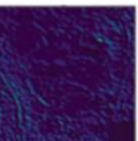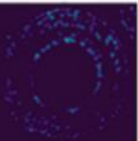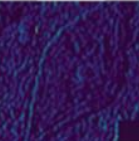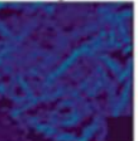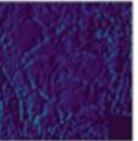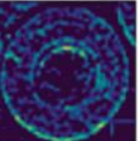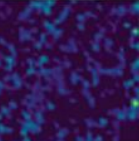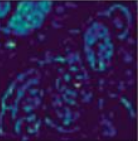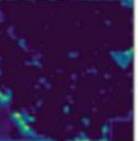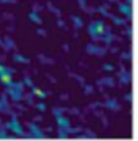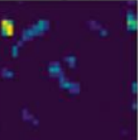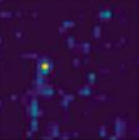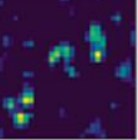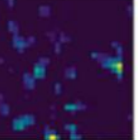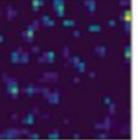
Figure 8. Different visualization maps, including heat map, saliency map and feature maps that were generated during the training process (feature extraction stage) using *Keras* and *Pytorch* libraries (Samek et al., 2016). These maps help to illustrate how the network performed and what features were extracted during different convolutional operation blocks.
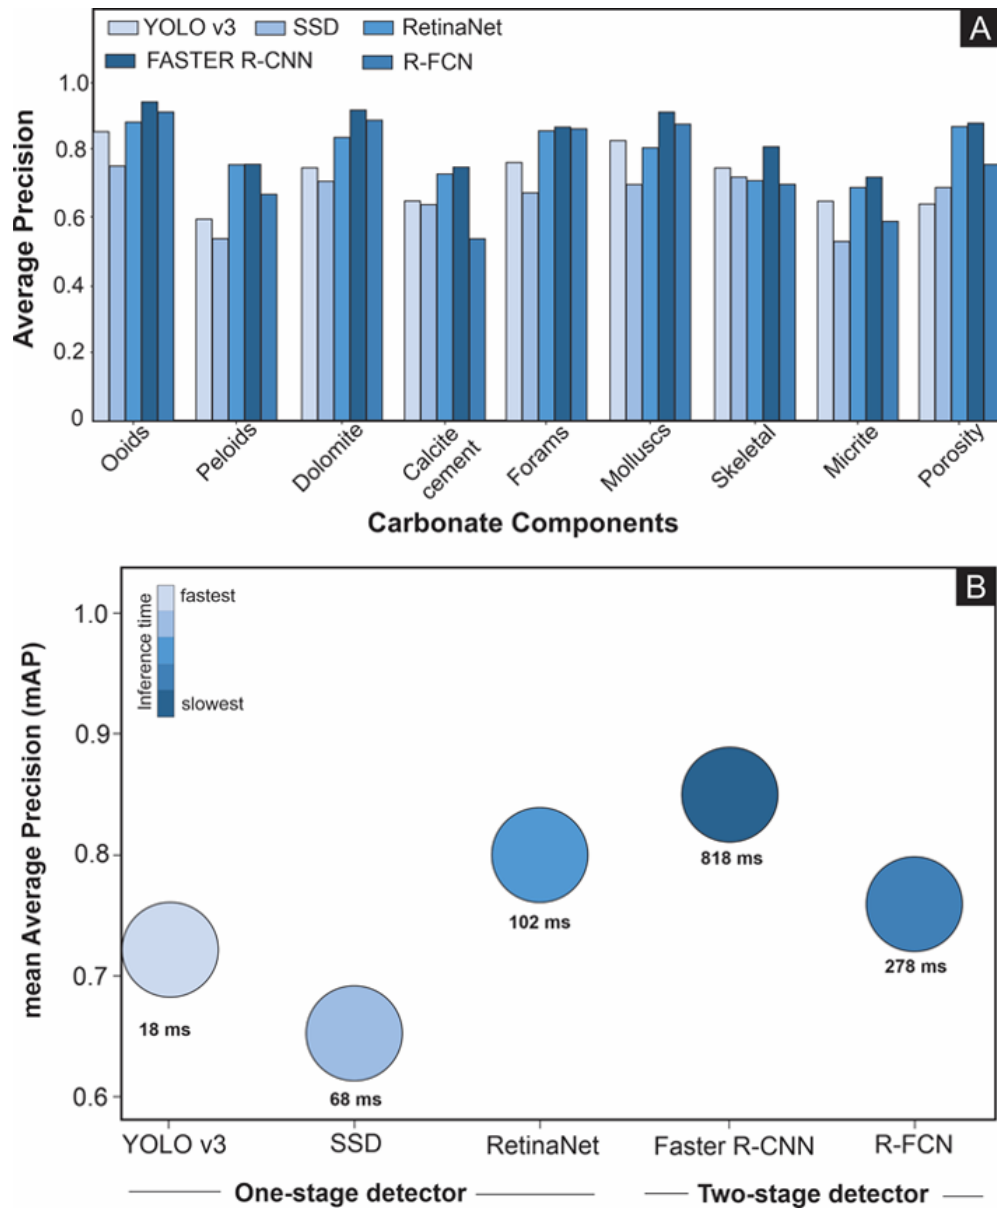
Figure 9. Evaluation performance of the different object detector frameworks on the validation or test set. (A) Average precision (AP) of different carbonate components. Note that the Faster R-CNN method achieved the highest AP for all class. (B) Cross plot between detection accuracy (mAP) and speed (inference time). The fastest detection speed is achieved with the YOLO v3 framework whereas the highest accuracy is observed when using Faster R-CNN framework. The best tradeoff between detection speed and accuracy is obtained from the RetinaNet framework.
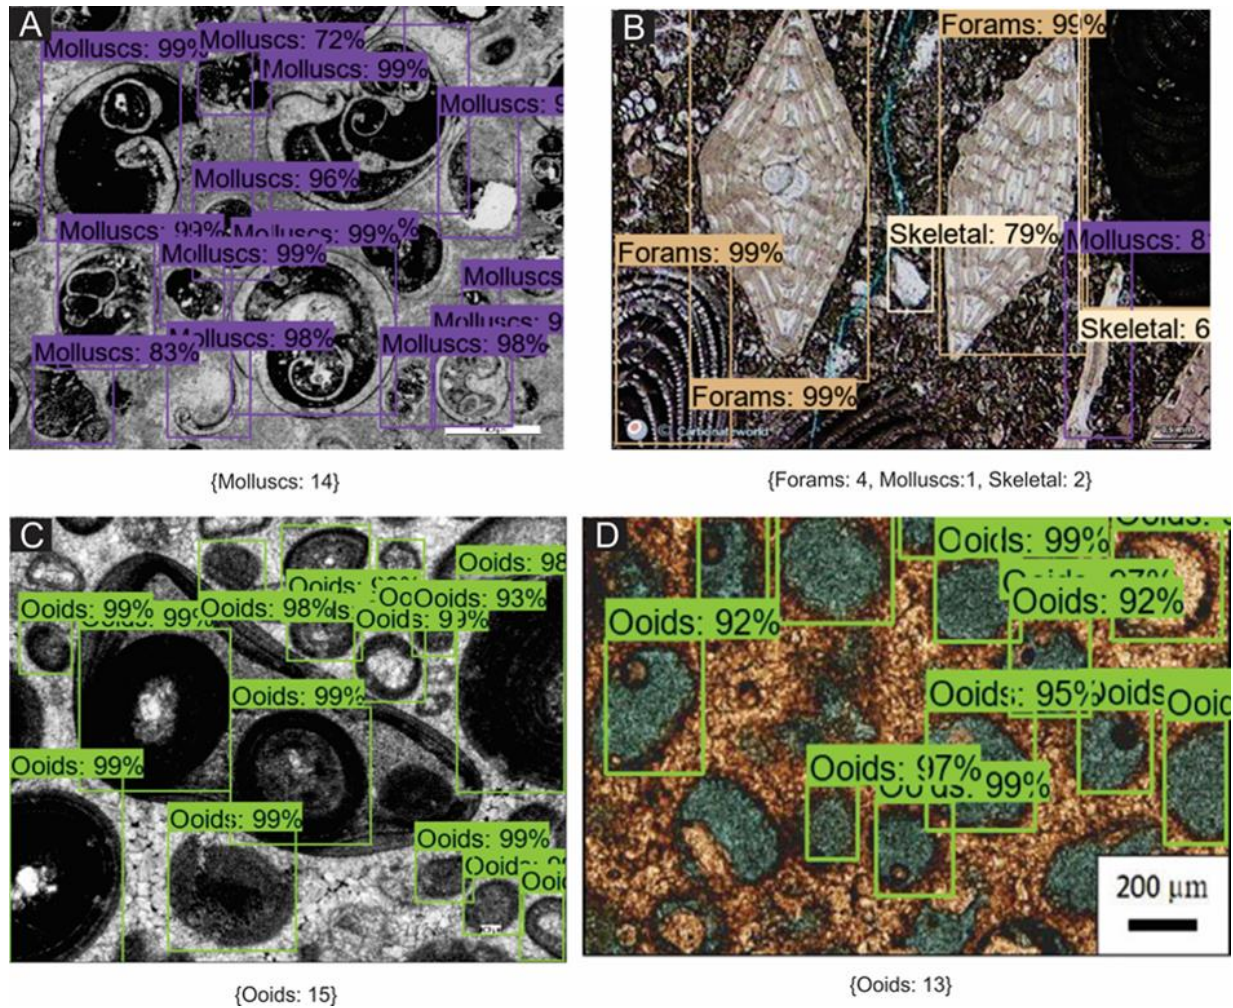
Figure 10 Examples of petrographic images with predicted carbonate components for each identified object. (A) High confidence detection of molluscs grain (sample name: TK-61.5). (B) Foraminifera and skeletal grains (carbonateworld.com). (C) Ooids grain with concentric cortices (sample name: TK-45.5). Note that the detection can separately detect two individual ooids within a compound grain. (D) Correctly classified ooids in a case where the internal structure is completely obliterated or dissolved (image from Eltom et al., 2017).
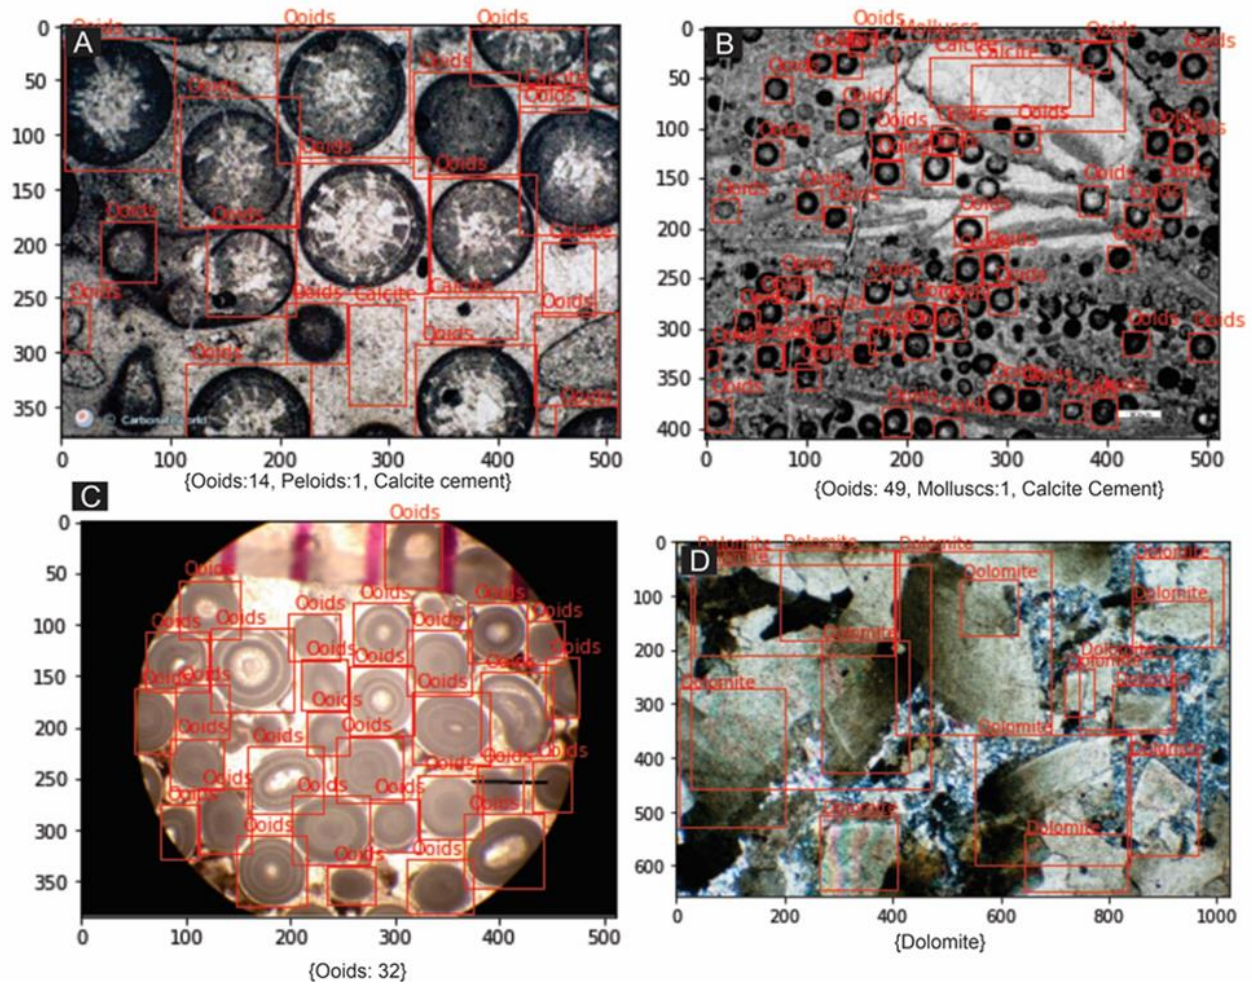
Figure 11. Examples of petrographic images where both grain and cement were detected and with high level accuracy. (A) Ooids with radial cortices and calcite cement (carbonateworld.com). (B) Small ooids grains and recrystallized molluscs grain by calcite cement (sample name: TK-18.1). (C) Ooids grain with concentric cortices (sample name: DJ12-2). (D) Dolomite minerals (image from: https://www.flickr.com/photos/zeesstof/1439328806/in/photostream/).

Figure 12. Petrographic images showing various porosity types in carbonate rocks. (A) Intraparticle porosity (carbonateworld.com). (B) Moldic porosity in crystalline dolomite (image from: Burchette, 2019). (C) Inter- and intracrystalline porosity (carbonateworld.com). (D). Oomoldic porosity (image from: Eberli et al., 2003).

TABLES

Task	Architecture
DCNN-based Image Classification	VGG16
	Inception-Resnet v2
DCNN-based Object Detection	
One-stage detector	YOLO v3_Darknet53
	SSD_Mobilenet v2
	SSD_Resnet FPN (RetinaNet)
Two-stage detector	Faster R-CNN_Inception-Resnet v2 (FPN+Cosine scheduler)
	RFCN_Resnet101

Table 1. Summary of different architectures or frameworks used in this study to evaluate DCNN-based image classification and object detection tasks for carbonate petrography images.

Evaluation Metrics	
Accuracy	$\dfrac{TP + TN}{TP + TN + FP + FN}$
Precision	$\dfrac{TP}{TP + FP}$
Recall	$\dfrac{TP}{TP + FN}$
F1 − score	$\dfrac{2*(Precision*Recall)}{(Precision + Recall)}$

Table 2. Calculation of evaluation metrics based on the percentage between correctly classified and misclassified classes.

VGG16			
	Precision	**Recall**	**F$_1$-score**	**Support**
Bioclast	0.87	0.90	0.89	206
Calcite Cement	0.72	0.78	0.75	49
Coated Grains	0.90	0.86	0.88	145
Dolomite	0.90	0.84	0.87	108
Micrite	0.74	0.89	0.81	28
Porosity	0.85	0.80	0.83	66
mean/total	0.83	0.85	0.84	602

Inception-ResNet-V2			
	Precision	**Recall**	**F$_1$-score**	**Support**
Bioclast	0.97	0.97	0.97	198
Calcite Cement	0.79	0.82	0.80	49
Coated Grains	0.95	0.96	0.95	143
Dolomite	0.95	0.91	0.93	112
Micrite	0.93	1.00	0.97	28
Porosity	0.89	0.89	0.89	72
mean/total	0.91	0.92	0.92	602

Table 3. Performance evaluation (precision, recall and F-1 score) during the prediction stage.

Detection Pipeline	Backbone	Ooids	Peloids	Dolomite	Calcite Cement	Forams	Molluscs	Skeletal	Micrite	Porosity	mAP (%)	Average Inference Time (ms/image)
						AP$_{0.5}$(%)						
One-stage detector												
YOLOv3 (512x512)	Darknet-53	**0.89**	0.60	0.75	0.65	0.77	**0.83**	0.71	0.65	0.64	0.72	18.4
SSD (512x512)	MobileNet	0.76	0.54	0.71	0.64	0.68	0.70	0.72	0.53	0.69	0.66	68.2
RetinaNet (324x324)	ResNet50-FPN	0.87	**0.76**	**0.84**	**0.73**	**0.86**	0.82	**0.77**	**0.69**	**0.83**	0.80	101.8
Two-stage detector												
Faster R-CNN (512x512)	Inception-ResNet-v2	**0.95**	**0.76**	**0.92**	**0.75**	**0.88**	**0.91**	**0.81**	**0.72**	**0.88**	0.84	818.4
R-FCN (512x512)	ResNet 101	0.92	0.67	0.89	0.54	0.87	0.88	0.70	0.59	0.76	0.76	277.8

Table 4. Performance evaluation of different object detection frameworks from both one-stage and two-stage detectors. Bold numbers represent the best precision for each class when compared to other object detector frameworks.

	Image Analysis		Deep Learning		Level of Agreement	
	Point Counting	Porosity (%)	Object Counting	Porosity (%)	Grain Counting	Porosity (%)
Sample 1	14	11.4	11	12.4	0.79	1.09
Sample 2	84	8.6	71	7.8	0.85	0.91
Sample 3	42	15.9	24	14.6	0.57	0.92
Sample 4	154	3.6	122	3	0.79	0.83
Sample 5	29	10.5	29	10.1	1.00	0.96
Sample 6	32	22.9	27	21.5	0.84	0.94
Sample 7	6	2.6	6	2.1	1.00	0.81
Sample 8	101	1.4	88	1.2	0.87	0.86
Sample 9	12	3.4	12	3.2	1.00	0.94
Sample 10	16	2.4	15	2.4	0.94	1.00
Sample 11	39	2.1	22	3.5	0.56	1.67
Sample 12	44	1.5	40	1.3	0.91	0.87
Sample 13	11	14.8	11	14.2	1.00	0.96
Sample 14	5	12.3	5	9.7	1.00	0.79
Sample 15	22	11.1	21	10.5	0.95	0.95

Table 5. Comparison between traditional image analysis and deep learning-based image analysis in conducting point counting and porosity measurement.