Article

Instance Segmentation of Microscopic Foraminifera

Thomas Haugland Johansen [1,*], Steffen Aagaard Sørensen [2], Kajsa Møllersen [3] and Fred Godtliebsen [1]

1   Department of Mathematics and Statistics, UiT The Arctic University of Norway, 9019 Tromsø, Norway;
    fred.godtliebsen@uit.no
2   Department of Geosciences, UiT The Arctic University of Norway, 9019 Tromsø, Norway;
    steffen.sorensen@uit.no
3   Department of Community Medicine, UiT The Arctic University of Norway, 9019 Tromsø, Norway;
    kajsa.mollersen@uit.no
*   Correspondence: thomas.h.johansen@uit.no

Abstract: Foraminifera are single-celled marine organisms that construct shells that remain as fossils in the marine sediments. Classifying and counting these fossils are important in paleo-oceanographic and -climatological research. However, the identification and counting process has been performed manually since the 1800s and is laborious and time-consuming. In this work, we present a deep learning-based instance segmentation model for classifying, detecting, and segmenting microscopic foraminifera. Our model is based on the Mask R-CNN architecture, using model weight parameters that have learned on the COCO detection dataset. We use a fine-tuning approach to adapt the parameters on a novel object detection dataset of more than 7000 microscopic foraminifera and sediment grains. The model achieves a (COCO-style) average precision of 0.78 on the classification and detection task, and 0.80 on the segmentation task. When the model is evaluated without challenging sediment grain images, the average precision for both tasks increases to 0.84 and 0.86, respectively. Prediction results are analyzed both quantitatively and qualitatively and discussed. Based on our findings we propose several directions for future work and conclude that our proposed model is an important step towards automating the identification and counting of microscopic foraminifera.

Keywords: foraminifera; instance segmentation; object detection; deep learning

1. Introduction

Foraminifera are microscopic (typically smaller than 1 mm) single-celled marine organisms (protists) that are ubiquitous in marine environments. During their life cycle, they construct shells from various materials that readily fossilize in sediments, which can be extracted and examined. Roughly 50,000 species have been recorded, of which approximately 9000 are living today [1]. Foraminiferal shells are abundant in both modern and ancient sediments. The foraminiferal faunal composition/abundance and chemical composition within shells reflect the ambient environment in which they live and, when fossilized, constitute records of past marine conditions. This makes the application of foraminifera valuable across many diverse fields of geoscientific inquiry within both industry and science. The applications include but are not limited to: (1) establishing of the overall foraminiferal concentration, abundance within groups (e.g., planktic, benthic calcareous and benthic agglutinated), and shell fragmentation to infer primary productivity of the water mass, calcium carbonate dissolution, and preservation state within sediments [2,3]; (2) usage of established modern foraminiferal distribution (e.g., Sejrup et al. [4]) for transfer function-based reconstruction of paleoenvironmental and paleowater mass property reconstruction [5,6]; (3) measurement of radiocarbon and stable isotopic- and trace element composition in shell material, used for dating, chemostratigraphy, and reconstruction of past climate and ocean conditions, including temperature, salinity, circulation, and ice volume [7]; (4) usage of foraminiferal fauna composition and index species as bio-indicators for timing, impact, and recovery related to anthropogenically introduced

stress to the marine environment [8] and interpreting the ages and paleoenvironments of sedimentary strata in oil wells via biostratigraphy [9].

After a sediment core has been retrieved from the seabed, a range of procedures are performed in the laboratory before the foraminiferal specimens can be identified and extracted under the microscope by a geoscientist using a brush or needle. From each core, several layers are extracted, and each layer is regarded as a sample. To establish a statistically robust representation of the fauna, 300–500 specimens are identified and extracted per sample. The time-consumption of this task is 2–8 h/sample, depending on the complexity of the sample and the experience level of the geoscientist. A typical study consists of 100–200 samples from one or several cores, and the overall time-consumption in just identifying the specimens is vast. Recently developed deep learning models show promising results towards automating parts of the identification and extraction process [10–14].

Figure 1 shows an example of a prepared foraminifera sample—microscopic objects spread out on a plate and photographed through a microscope. Of particular interest is the classification of each object into high-level foraminifera classes, which then serves as input for estimation of the environment in which sediment was produced. This task consists of identifying relevant objects, particularly separate sediment from foraminifera, and recognize foraminifera classes based on shapes and structures of each object.

First acquisition setup **Second acquisition setup**
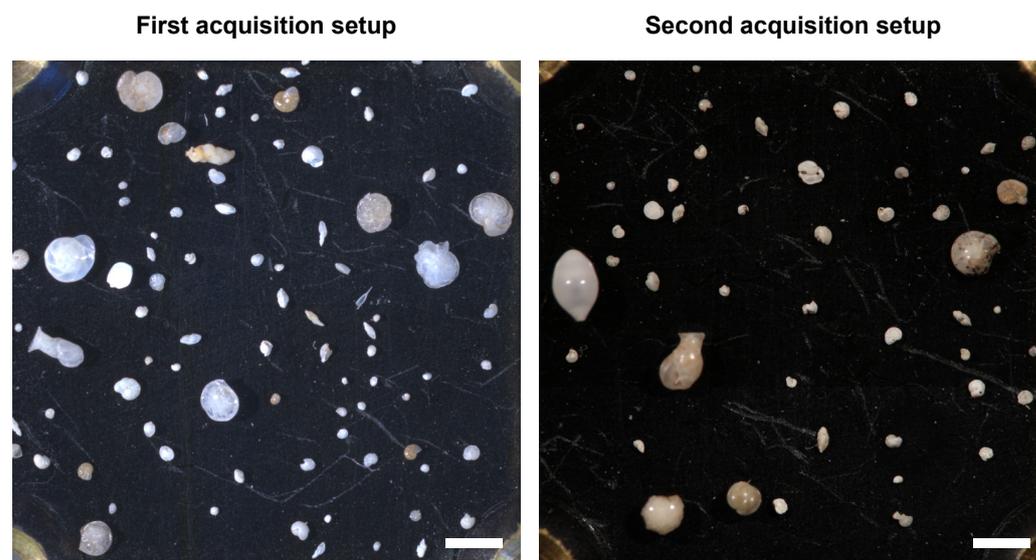
Figure 1. Examples of images from the two different image acquisition setups used during three data collection phases. The visible area in each image is approximately 10 mm × 10 mm and the scale bars correspond to 1000 μm. (**Left**): Calcareous benthics photographed with the first acquisition setup used during the first data collection phase. (**Right**): Calcareous benthics photographed with the second acquisition setup used during the second and third data collection phases.

Object classification in images is one of the great successes of deep learning, and conquer new applications as new methods are developed and high-quality data are made available for training and testing. In a deep learning context, a core task of object classification is instance segmentation. Not only must the objects be separated from the background, but the objects themselves must be separated from each other, so that adjacent objects are identified, and not treated as one single object.

Automatic foraminifera identification has great practical potential for three main reasons. Firstly, the time saved for highly qualified personnel is substantial, which can then be made available for other tasks. Secondly, it is the overall proportion of foraminifera classes that is the primary interest, and it is, therefore, robust to the occasional misclassification, as opposed to, e.g., cancer detection, where overlooking a small region of cancer cells can be fatal. Finally, the availability of deep learning algorithms that integrates object

detection, instance segmentation, and object classification makes the creation of such a method feasible.

The lack of publicly available datasets for this particular deep learning application has been an obstacle, but, with a curated private dataset, soon-to-be published, there is now an opportunity to investigate the potential of applying deep learning to foraminifera classification.

The manuscript is organized as follows:

In Section 2.1, we describe the acquisition and preparation of the dataset, and its final attributes. In Section 2.2, we give an overview of the Mask R-CNN model applied to foraminifera images. In Section 2.3, we give a detailed description of the experimental setup. To present the results, we have chosen to include training behavior (Section 3.1), since this is a first attempt for foraminifera application. Further, we give a detailed presentation of the performance from various aspects and different thresholds (Section 3.2) for a comprehensive understanding of strengths and weaknesses. Section 4 then emphasizes and discusses the most interesting findings, both in terms of promising performance and in terms of future work (Section 4.1). We round off with a conclusion (Section 5) to condense the discussion into three short statements.

2. Materials and Methods

The work presented in this article was performed in two distinct phases: first, a novel object detection dataset of microscopic foraminifera was created, and then a pre-trained Mask R-CNN model [15] was adapted and fine-tuned on the dataset.

2.1. Dataset Curation

All presented materials (foraminifera and sediment) were collected from sediment cores retrieved in the Arctic Barents Sea region. The specimens were picked from sediments influenced by Atlantic, Arctic, polar, and coastal waters representing different ecological environments. This was done to ensure good representation of the planktic and benthic foraminiferal fauna in the region. Foraminiferal specimens (planktics, benthics, agglutinated benthics) were picked from the 100 µm to 1000 µm size fraction of freeze dried and subsequently wet sieved sediments. Sediment grains representing a common sediment matrix were also sampled from the 100 µm to 1000 µm size range. Additional details about the sediment cores and materials can be found in Appendix B.

The materials were prepared and photographed at three different points in time, with two slightly different image acquisition systems. During the first two rounds of acquisition, every image contains either pure benthic (agglutinated or calcareous), planktic assemblages, or sediment grains containing no foraminiferal specimens. In other words, each image contained only specimens belonging to one of four high-level classes: agglutinated benthic, calcareous benthic, planktic, and sediment grain. This approach greatly simplified the task of labeling each individual specimen with the correct class. Examples of specimens from each of the four high-level classes can be seen in Figure 2. In order to better mimic a real-world setting with mixed objects, the third acquisition only contained images where there was a realistic mixture of the four object classes. To get the necessary level of magnification and detail, four overlapping images were captured from the plates on which the specimens were placed, where each image corresponded to a distinct quadrant of the plate. The final images were produced by stitching together the mosaic of the four partially overlapping images.

Agglutinated benthics

Calcareous benthics

Planktics

Sediment grains
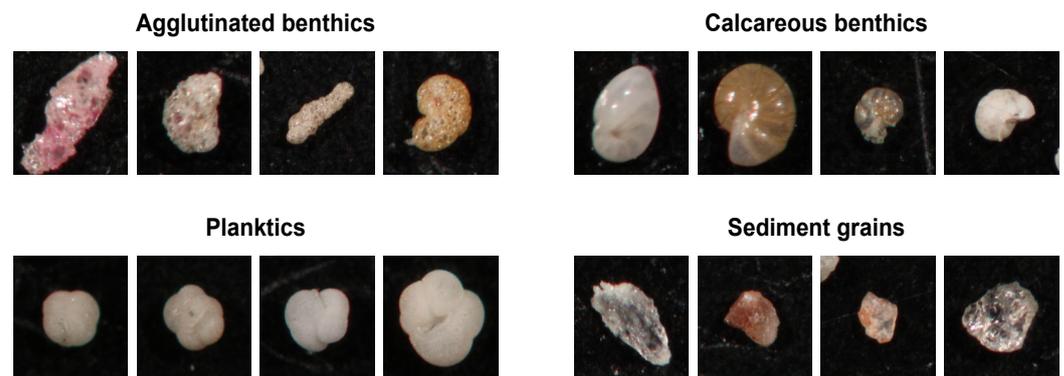
Figure 2. Examples of the four high-level classes of specimens found in the dataset.

All images from the first acquisition were captured with a 5 megapixel Leica Microsystems (Austria, Vienna) DFC450 digital camera mounted on a Leica Microsystems Z16 APO fully apochromatic zoom system. The remaining two acquisitions were captured using a 51 megapixel Canon (Tokyo, Japan) EOS 5DS R camera mounted on a Leica Microsystems M420 macroscope. The same Leica Microsystems CLS 150x (twin goose-neck combination light guide) was used for all acquisitions, but with slightly different settings. No illumination nor color correction was performed, in an attempt to mimic a real-world scenario of directly detecting, classifying, and segmenting foraminifera placed under a microscope. Examples of the differences in illumination settings can be seen in Figure 1.

To create the ground truth, a simple, yet effective, hand-crafted object detection pipeline [14] was run on each image, which produced initial segmentation mask candidates. The pipeline consisted of two steps of Gaussian smoothing, then grayscale thresholding followed by a connected components approach to detect individual specimens. Some parameters, such as the width of Gaussian filter kernel, as well as threshold levels, were hand-tuned to produce good results for each image in the dataset. A simple illustration of the preprocessing pipeline can be seen in Figure 3. For full details, see Johansen and Sørensen [14].

Input Image   **Binary Masks**   **Object Dataset**   **Mask Previews**
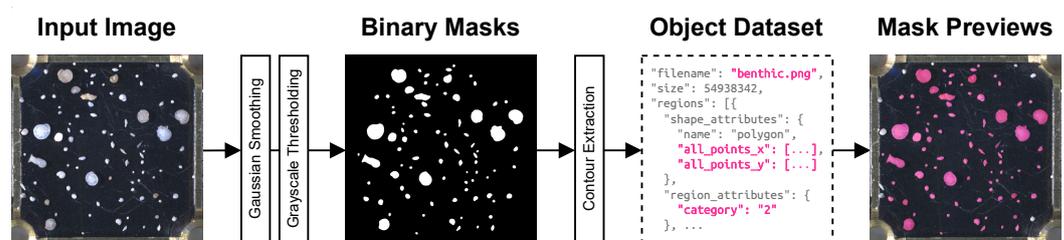
**Figure 3.** High-level summary of the dataset creation pipeline.

After obtaining the initial segmentation mask dataset, all masks were manually verified and adjusted using the VGG Image Annotator [16,17] software. Additionally, approximately 2000 segmentation masks were manually created (using the same software) for objects not detected by the detection pipeline. The end result is a novel object detection dataset consisting of 104 images containing over 7000 segmented objects. Full details on the final dataset can be found in Table 1.

Table 1. Detailed breakdown of the object dataset, where each row holds information for a specific microscope image acquisition phase. The first and second phases contain only "pure" images where every object is of a single class, whereas the third phase images contain only mixtures of several classes.

Phase	Images	Objects	Objects per Class			
			Agglutinated	Benthic	Planktic	Sediment
First	48	3775	172	897	726	1980
Second	41	2604	583	695	657	669
Third	15	633	154	156	155	168
Combined	104	7012	909	1748	1538	2817

2.2. Instance Segmentation Using Deep Learning

Mask R-CNN [15] is a proposal-based deep learning framework for instance segmentation, and it is an extension to Fast/Faster R-CNN [18,19]. In the Fast/Faster R-CNN framework the model has two output branches, one that performs bounding box regression and another that performs classification. The input to these two branches are pooled regions of interest (RoIs) produced from features extracted by a convolutional neural network (CNN) backbone. This is extended in Mask R-CNN by adding an extra (decoupled) output branch, which predicts segmentation masks on each RoI. Figure 4 shows a simple, high-level representation of the Mask R-CNN model architecture. Several alternatives to Mask R-CNN exist, such as PANet [20], TensorMask [21], CenterMask [22], and SOLOv2 [23]. We chose to use the Mask R-CNN framework for two key reasons: (1) it predicts bounding boxes, class labels, and segmentation masks at the same time in a single forward-pass, and (2) pre-trained model parameters are readily available, removing the need to train the model from scratch.
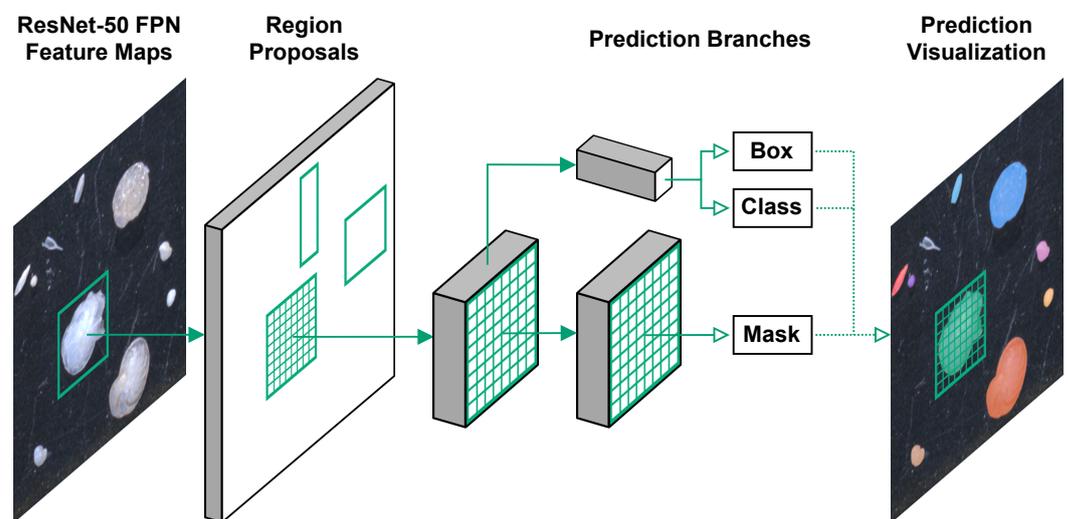
Figure 4. Simple, sketch-like depiction of the Mask R-CNN model architecture.

Due to its flexible architecture, there are numerous ways to design the feature extraction backbone of a Mask R-CNN model. We chose a model design based on a ResNet-50 [24] Feature Pyramid Network (FPN) [25] backbone for feature extraction and RoI proposals. To avoid having to train the model from scratch, we applied model parameters pre-trained on the COCO dataset [26]. The object detection model and all experiment code was implemented using Python 3.8, PyTorch 1.7.1 [27], and torchvision 0.8.2. The pre-trained model weights were downloaded via the torchvision library.

2.3. Experiment Setup and Training Details

The original Mask R-CNN model was trained using 8 GPUs and a batch size of 16 images, with 2 images per GPU. We did not have access to that kind of compute resources, and were instead limited to a single NVIDIA TITAN Xp GPU, which also meant our training batches only consisted of a single image. The end result of this was slightly more unstable loss terms and gradients, so we carefully tested many different optimization methods, learning rates, learning rate scheduling, and so on.

The dataset was split (with class-level stratification) into separate training and test sets, using a 2.47:1 ratio, which produced 74 training images and 30 test images. The training and test sets remained the same for all experiments. During training, images were randomly augmented, which included horizontal and vertical flipping, brightness, contrast, saturation, hue, and gamma adjustments. (The validity of flipping in this context w.r.t. biological handedness is discussed in Section 4.1.2.) Both the horizontal and vertical flips were applied independently, with a flip probability of 50% for both cases. Brightness and contrast factors were randomly sampled from $[0.9, 1.1]$, the saturation factor from $[0.99, 1.01]$, and hue from $[-0.01, 0.01]$. For the random gamma augmentation, the gamma exponent was randomly sampled from $[0.8, 1.2]$.

We ran the initial experiments using the Stochastic Gradient Descent (SGD) optimization method with Nesterov momentum [28] and weight decay. The learning rates tested were $\{10^{-3}, 5 \times 10^{-3}, 10^{-5}\}$, and the momentum parameter was set 0.9. For weight decay, we tested the values $\{0, 10^{-4}, 10^{-5}, 5 \times 10^{-5}\}$. In some experiments, the learning rate was reduced by a factor of 10 after either 15 or 25 epochs. Training was stopped after 50 epochs. After the initial experiments with SGD, we tested the Adam [29] optimization method. We tested the learning rates $\{10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}\}$. The weight decay parameter values were $\{0, 10^{-4}, 10^{-6}\}$. We used the same scheduled learning rate decay as with SGD for the initial experiment, and training was stopped after 50 epochs.

From on our initial experiments with SGD and Adam, we saw that the latter gave more stable loss terms during training. Based on this discovery, we also experimented with a recent variant of the Adam optimizer with decoupled weight decay, referred to as AdamW [30], since we believed it would further stabilize and improve model training. We implemented a slightly adjusted scheduled learning rate decay, with a factor of 10 reduction after both 25 and 45 epochs of training. Because we used model parameters pre-trained on the COCO dataset, we also ran experiments with fine-tuning the backbone model to adapt it to our target domain. For the fine-tuning experiments, we tested when to "freeze" and "unfreeze" the backbone model parameters, i.e., when to fine-tune the backbone, as well as which layers of the backbone to fine-tune.

Based on all probing experiments and the hyperparameter tuning, our final model was trained using AdamW for 50 epochs. During the first 25 epochs of training, the last three ResNet-50 backbone layers were fine-tuned, and then they were subsequently frozen. The initial learning rate was set to 10^{-5} and was reduced to 10^{-6} after 25 epochs, and further reduced to 10^{-7} after 45 epochs. We set the weight decay parameter to 10^{-4}. Using this configuration, we trained the model 10 times using different random number generator states to ensure valid results and to measure the robustness of the model.

3. Results

Model performance is evaluated using the standard COCO metrics for object detection and instance segmentation [31]. Specifically, we are using the average precision (AP) and average recall (AR) metrics averaged over 10 intersection-over-union (IoU) thresholds and all classes, where the IoU thresholds range from 0.5 to 0.95 in increments of 0.05. We also use the more traditional definition of AP, which is evaluated at a specific IoU, e.g., AP_{50} denotes the AP evaluated with an IoU of 0.5. Additionally, we present conventional precision-recall curves with different evaluation configurations, e.g., per-class, per-IoU, and so forth. All presented precision and recall results were produced by evaluating models on the test split of the dataset. The mean precision and recall scores for repeated training runs are

presented without the standard deviation since this fell below the significance level of all presented results.

The precision and recall are calculated as a combination of detection and classification. The ground truth consists of a segmentation mask, its minimum enveloping rectangle, and the object class. The corresponding output is a segmentation mask, a bounding box, and the object label. Note that the bounding box is not necessarily the minimum rectangle of the segmentation mask, and we report the results separately.

In the bounding box setting, a detection is defined as an IoU between a bounding box and a minimum rectangle above a pre-specified threshold. A true positive (TP) is a detected object with label identical to the object class. A false positive (FP) has either no IoU above the specified threshold, or its IoU is with an object with a different class. A false negative (FN) is an undetected object. TP, FP, and FN are defined correspondingly in the segmentation mask setting, with IoU calculated w.r.t. the ground truth mask and the output mask.

3.1. Model Training

During training, all training losses were carefully monitored and reported both per-batch and per-epoch. Four of the key loss terms for the Mask R-CNN model can be seen in Figure 5, where each curve represents one of the 10 repeated training runs, with different initial random state. At the end of every training epoch, we evaluated the model performance in terms of the AP metric for both the detection and segmentation task on the test images. The per-epoch results for all 10 runs can be seen in Figure 6.
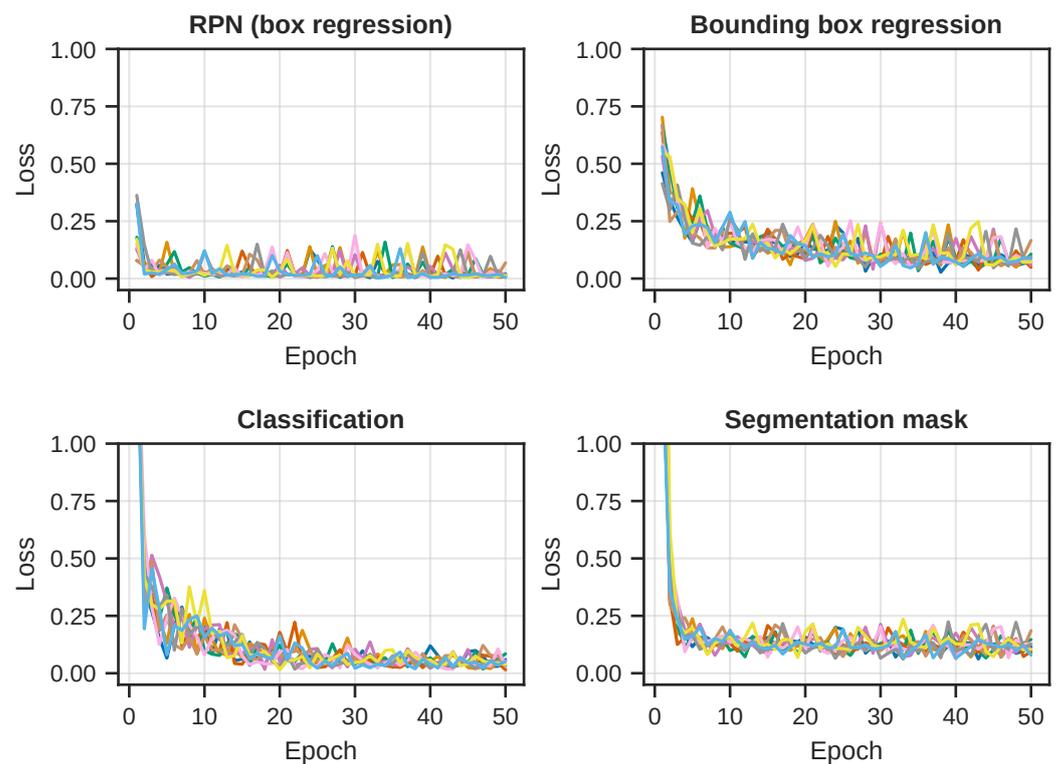
Figure 5. Evolution of the individual loss terms for each of the 10 training runs. (**Top left**): The loss term for the RPN box regression sub-task. Fairly rapid convergence, but we can see the effect of the single-image batches in the curves. (**Top right**): Bounding box regression loss for the detection branch. The convergence is slower when compared to the RPN loss, and perhaps slightly less stable. (**Bottom left**): Object classification loss for the detection branch. A similar story can be seen here as with the bounding box regression, which suggests a possible challenge with the detection branch. (**Bottom right**): The segmentation mask loss for the segmentation branch. Fast convergence, but, again, we see the effect of the single-image training batches.
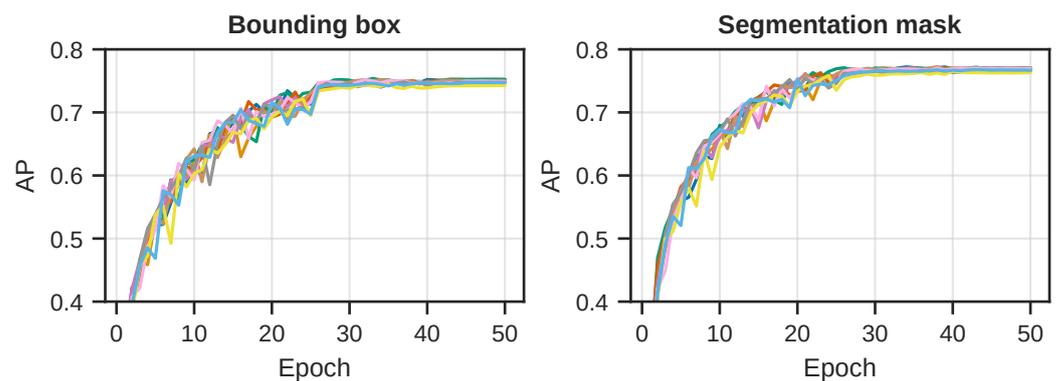
Figure 6. Evolution of the AP for both the detection and segmentation task, for each of the 10 runs, per epoch of training. Note that these results are based on evaluations using a maximum of 100 detections per image. (**Left**): The AP for the detection (bounding box) task. A plateau is reached after about 30 epochs. (**Right**): The AP for the segmentation mask task. We observe that the same type of plateau is reached here as with the detection task.

These results indicate that, even though we reached some kind of plateau during training, we did not end up overfitting or otherwise hurt the performance on the test dataset. The AP for both tasks also reach a plateau, which is almost identical for all of the learned model parameters. This suggests that the training runs reached an upper limit on performance given the dataset, model design, and hyperparameters.

3.2. Evaluating the Model Performance

After the 10 training runs had concluded, we evaluated each model on the test data using their respective parameters from the final training epoch. Note that all precision and recall evaluations presented from this point onward are based on a maximum of 256 detections per image, which is an increase from the maximum of 100 detections per image used during training due to lower computational costs and faster evaluation. The mean AP across all of the 10 run was evaluated as 0.78 for the detection task, and 0.80 for the segmentation task. Table 2 shows a summary of the AP and AR metrics for both tasks, where each result is the mean and standard deviation of all training runs. Note that this table shows results averaged over all four classes, as well as with the "sediment" class omitted from each respective evaluation, which will be discussed later.

Table 2. AP and AR scores for different IoU thresholds, evaluated with all object classes being considered and with the "sediment" class excluded.

	All Classes		Sans "Sediment" Class	
	Bound. Box	Segm. Mask	Bound. Box	Segm. Mask
AP_{50}	0.90	0.90	0.94	0.95
AP_{75}	0.88	0.90	0.94	0.94
AP	0.78	0.80	0.84	0.86
AR	0.83	0.84	0.89	0.90

The precision-recall curves computed by averaging over all 10 training runs can be seen in Figure 7. Note the sharp and sudden drop in the curve around the recall threshold of 0.75, for both tasks.
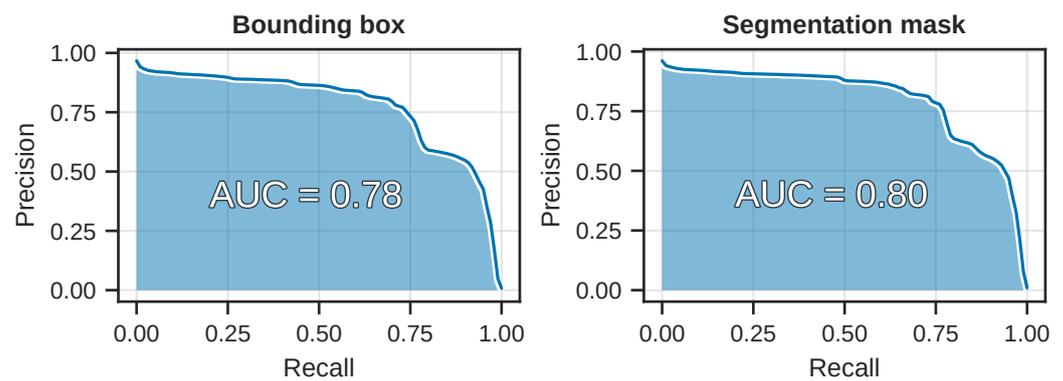
Figure 7. The mean average precision-recall curves for the 10 training runs, for both the detection and segmentation tasks. The area-under-the-curve (AUC) shown here is the same as our definition of AP, which can be seen in Table 2.

In order to investigate the sharp drop in precision and recall, we computed per-class precision and recall; the results can be seen in Figure 8. From the curves in the figure, it is clear that the model is finding the "sediment" class particularly challenging. Notice how the precision rapidly goes towards zero slightly after the recall threshold of 0.75.
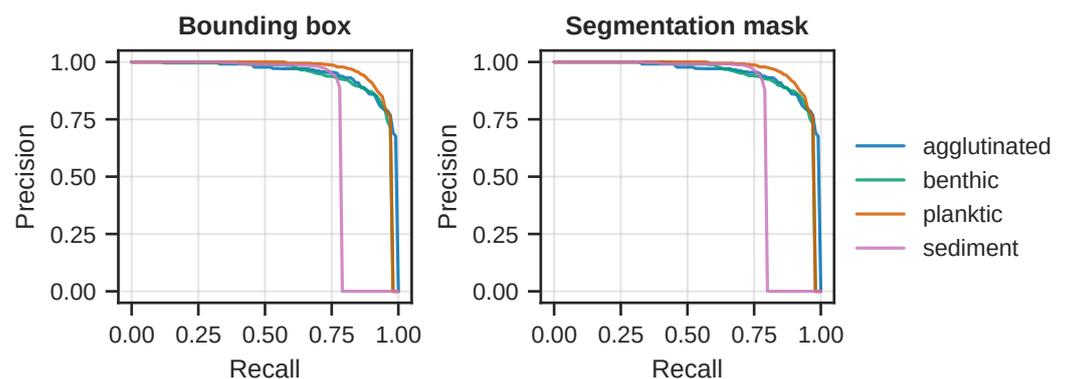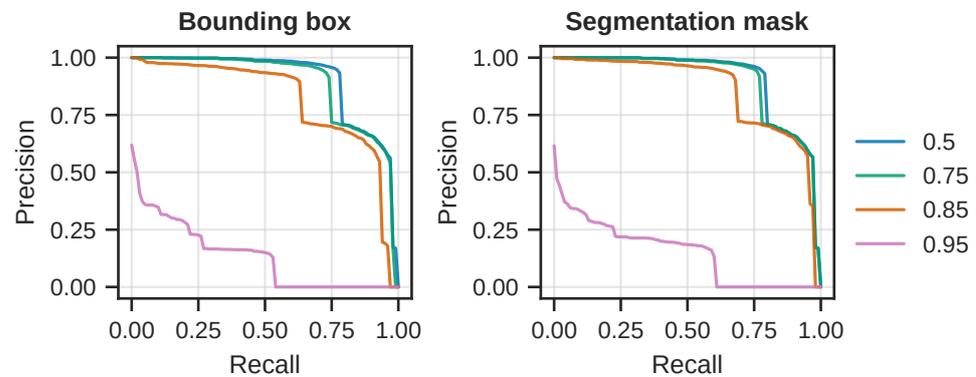
Figure 8. Precision-recall curves for each of the four object classes. (**Left**): Per-class curves for the detection (bounding box) task. Performance is approximately the same for the "agglutinated", "benthic", and "planktic" classes, but it is significantly worse for the "sediment" class. (**Right**): The per-class curves for the segmentation task, which tells the same story as for the detection task.

We also wanted to determine how well the model performed at different IoU thresholds, so precision and recall were evaluated for the IoU thresholds $\{0.5, 0.75, 0.85, 0.95\}$. Figure 9 shows the precision-recall curves for all object classes, and Figure 10 shows the sans "sediment" class curves. From these results, it is clear that the model performs quite well at IoU thresholds up to and including 0.85, but, at 0.95, the model does not perform well.

Based on the per-class and per-IoU results, it became evident that some test images containing only "sediment" class objects were particularly challenging. This can in part be explained by the object density in these images, with multiple objects sometimes overlapping or casting shadows on each other. In the COCO context, these types of object clusters are referred to as a "crowd", and they receive special treatment during evaluation. Importantly, none of the objects in our dataset have been annotated as being part of a "crowd" due to the resources required to annotate more than 7000 objects based on their proximity to other objects with sufficient precision and recall. Some examples of these dense object clusters can be seen in Figures A2 and A3. By removing the "sediment" class from the evaluation, the AP score for the bounding box increased to 0.84, and for instance segmentation it increased to 0.86. Recall also increased significantly, which means that more target objects were correctly detected and segmented. This increase can also be seen

by comparing the per-IoU curves shown in Figure 10 with those in Figure 9, as well as the results presented in Table 2.

Bounding box   **Segmentation mask**

Figure 9. Precision-recall curves at different IoU thresholds, where each curve is based on the average for all four object classes. (**Left**): PR curves for the detection (bounding box) task. There is a sharp drop in precision at the approximate recall thresholds $\{0.65, 0.74, 0.8\}$, which corresponds to the lower precision of the "sediment" class. (**Right**): The same drop in precision is observer for the predicted masks, which can again be explained by the performance on the "sediment" class.

Bounding box   **Segmentation mask**
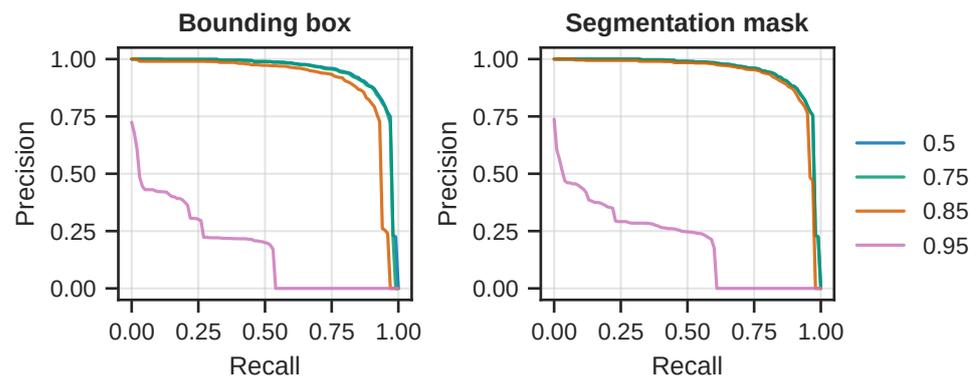
Figure 10. The average PR curves without the "sediment" class, at different IoU thresholds for both tasks. (**Left**): Without the "sediment" class, the curves for thresholds $\{0.5, 0.75, 0.85\}$ are almost identical, whereas there is still a major decrease for the 0.95 IoU threshold. (**Right**): The PR curves for the segmentation task paint the same picture as for the detection task, indicating that few predictions are correct above 0.95 IoU, and that very many targets are not being predicted.

3.3. Qualitative Analysis of Predictions

When evaluating the predictions manually, it became apparent that the overall accuracy and quality of the segmentation masks produced by the model are good. The boundary of the masks quite precisely delineates the foraminifera and sediment grains from the background. For the most part, the predicted bounding boxes correspond well with the masks. One of the biggest challenges seems to lie in the classification of object labels; there are (for trained observers) many obvious misclassifications. The exact cause is somewhat uncertain, but, in many cases, the objects are relatively small and feature-less. It is not hard to imagine how a feature-less planktic foraminifera can be misclassified as benthic, especially if the object is small. Other cases of misclassifications are likely caused by a lack of training examples; many seem like out-of-distribution examples due to the high confidence score. Examples of predictions can be seen in Figures 11 and 12.

Predictions for benthic specimens   **Predictions for mixed specimens**
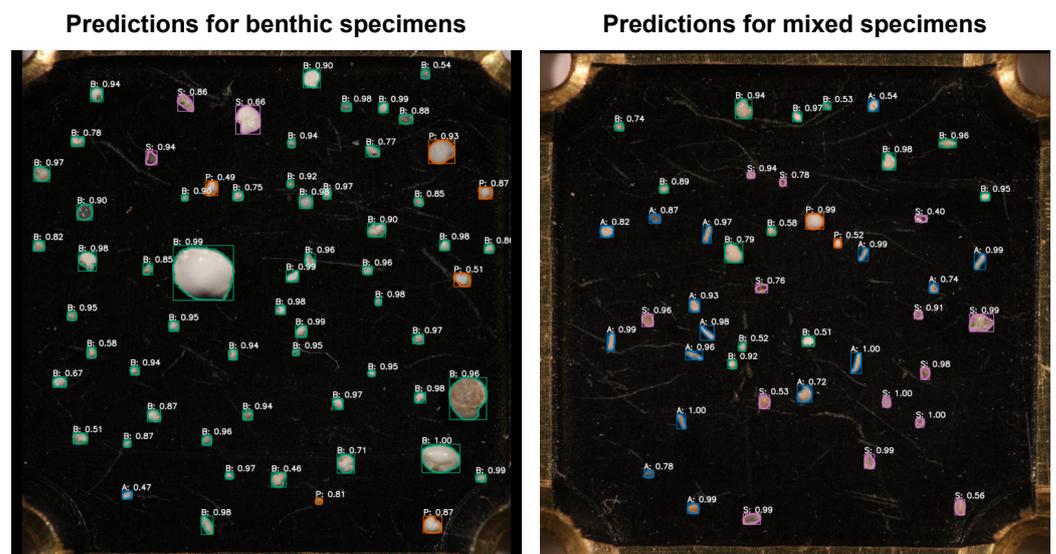
Figure 11. Examples of predictions for two images from the test dataset. (**Left**): Predictions for purely calcareous benthic specimens. The accuracy and quality of the predicted masks and bounding boxes are good, but there are several misclassified objects. (**Right**): Predictions for a mixture of specimen types. The accuracy and quality of the predicted masks and bounding boxes are good. However, there are misclassified detections for this image, as well.

Predictions for planktic specimens   **Predictions for agglutinated specimens**
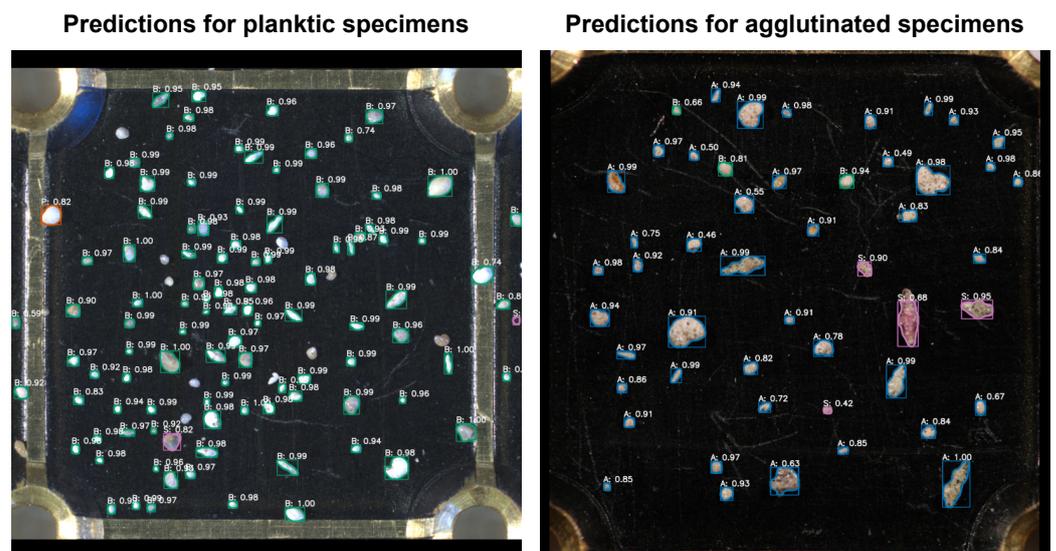
Figure 12. Additional examples of predictions for two images from the test dataset. (**Left**): Predictions for purely planktic specimens. There are a few false positives, but the accuracy and quality of true positive detections are good. Note that some objects that have been misclassified. (**Right**): Predictions for purely agglutinated benthic specimens. Good accuracy and quality of predicted masks and bounding boxes the majority of detections. However, low quality masks and misclassified detections are visible.

4. Discussion

The results presented clearly show that a model built on the Mask R-CNN architecture is capable of performing instance segmentation of microscopic foraminifera. Using model parameters pre-trained on the COCO dataset, we adapted and fine-tuned the model for our novel dataset and achieved AP scores of 0.78 and 0.80 on the bounding box and instance segmentation tasks, respectively. There were significant increases in precision and recall when going from averaging over all IoU thresholds (i.e., AP and AR) to specific IoU thresholds. When evaluated with an IoU of 0.5, precision increased to 0.90 for both

tasks, and, with an IoU of 0.75, the precision was 0.88 for the detection task and 0.90 for the segmentation task. This means that predicting bounding boxes and segmentation masks that almost perfectly overlap with their respective ground-truth is challenging for the given dataset, and possibly for the model architecture or hyperparameters. In all manual annotated datasets, there is the occasional error, both in terms of inaccuracies at the pixel-level, but also potential false positives or false negatives, meaning that achieving perfect predictions are very unlikely. Importantly, depending upon the specific application of an instance segmentation model, pixel-perfect predictions might not be a necessity.

Omitting the "sediment" class also lead to significant increases in model performance, which can be explained by the challenging nature of some test images that contained very dense clusters of sediment grains. This can in part be mitigated in practical applications by ensuring objects are not clustered, but, ideally, this also should be addressed at the model-level. It is possible that this can, to some extent, be overcome by introducing much more training examples with crowded scenes, as well as correctly annotating all objects as being in a crowd. Additionally, it is possible that the issue can also be reduced by tuning the hyperparameters of the Mask R-CNN architecture.

Both quantitative and qualitative analysis of the predicted detections and segmentation masks suggest that the model is performing well. However, the results also show that there are some challenges that should to be investigated further and addressed in future work.

4.1. Future Research

Based on the experiments and results, we propose a few research ideas worth investigating in future efforts.

4.1.1. Expanding and Revisiting the Dataset

Expanding the dataset is perhaps the most natural extension of the presented work. If carefully curated, a more exhaustive dataset should help improve some of the corner cases where the model is struggling to produce accurate predictions. Additionally, with the appropriate resources, it would be valuable to ensure every object in the existing dataset is appropriately labeled as part of a "crowd" or not. Improving the accuracy of the e.g., densely packed "sediment" objects, will improve model performance, as well as make the model more applicable to real-world situations. Another important aspect of expanding the dataset is introduce species-level object classes, as opposed to the high-level categories used today. Accurately detecting microscopic foraminiferal species is vital to most downstream geoscience applications.

4.1.2. Revising the Random Image Augmentations

Applying random image augmentations, such as rotation, flipping, and cropping, is a common practice in deep learning, often leading to significant improvements in task performance. It is essential to ensure that the applied random augmentations are valid in a given context, e.g., ensuring that applying random flips does not interfere with biological handedness. We suggest that artificially generated coiling directions due to the random flipping of training images pose little to no interference when distinguishing between overall groups of foraminifera. However, in future endeavors where the foraminiferal fauna will be identified at the species level, random flipping of the specimen images should not be done. This is critical for planktic foraminifera, especially in the Arctic region, where the coiling direction has profound implications for paleoenvironmental research and interpretation. For example, the species Neogloboquadrina pachyderma ("cold" species) and Neogloboquadrina incompta ("warm" species) are visually distinguished primarily via their coiling direction. Regarding benthic foraminifera, some studies have also been conducted elucidating the underlying mechanisms behind the preferred coiling direction [32,33]. In the Arctic realm, varying coiling direction in benthic foraminifera is not a well-studied phenomenon; little is known as to why this happens and what implications it

might have for paleoreconstructions. We expect that novel information related to the coiling direction of benthic foraminifera might be uncovered with carefully crafted experiments.

4.1.3. Additional Hyperparameter Tuning

If sufficient computational resources are available, performing more exhaustive hyperparameter tuning should be pursued. While this should include experiments with optimizers, learning rates, and so forth, it should more crucially be focused on the numerous hyperparameters of the Mask R-CNN model components. Specifically, the parameters of the regional proposal network, and the fully-convolutional network (for mask prediction) should be validated and experimented with. It is entirely possible some number of these parameters are sub-optimal for the given dataset.

4.1.4. Improved GPU Training

While training on multiple GPUs might not lead to big improvements in model performance, the increased effective batch size will help stabilize and speed up training. Additionally, given the small size of the most objects relative to the image dimensions, training without having to resize the images to fit in GPU memory will increase model performance. This could be solved directly by using GPUs with more memory, or possibly by partitioning each image across multiple GPUs, predicting on a sub-region per GPU.

4.1.5. Other Segmentation Models

We chose to use Mask R-CNN primarily because of its capabilities, but also because proven, pre-trained weights were readily available. Recently, numerous models have been published that surpass Mask R-CNN in several performance metrics, and importantly also seem to have much faster inference times (which is important for real-world applications.) Examples of alternative models that should be tested include PANet [20], TensorMask [21], CenterMask [22], and SOLOv2 [23].

4.1.6. Uncertainty Estimation

We have shown that the model is robust to training runs with different random seeds, and the next natural step is to investigate robustness with regards to different training/test data splits, and to estimate the uncertainty of the model predictions. Some work has been published on estimating model predictive uncertainty of Mask R-CNN models [34–36]. However, it should be possible to avoid the need for introducing Monte Carlo dropout sampling [37], which requires making changes to existing models, by leveraging the more recent Monte Carlo batch normalization sampling [38] technique instead.

5. Conclusions

The proposed model achieved an AP of 0.78 on the bounding box (detection) task and 0.80 on the segmentation task, based on 10 training runs with different random seeds. We also evaluated the model without the challenging sediment grain images, and the AP for both tasks increased to 0.84 and 0.86, respectively.

When evaluating predictions both qualitatively and quantitatively, we saw the predicted bounding boxes and segmentation masks were good for the majority of test cases. However, there were many cases of incorrect class label predictions, mostly for small objects, or objects that we hypothesize can be considered out-of-distribution.
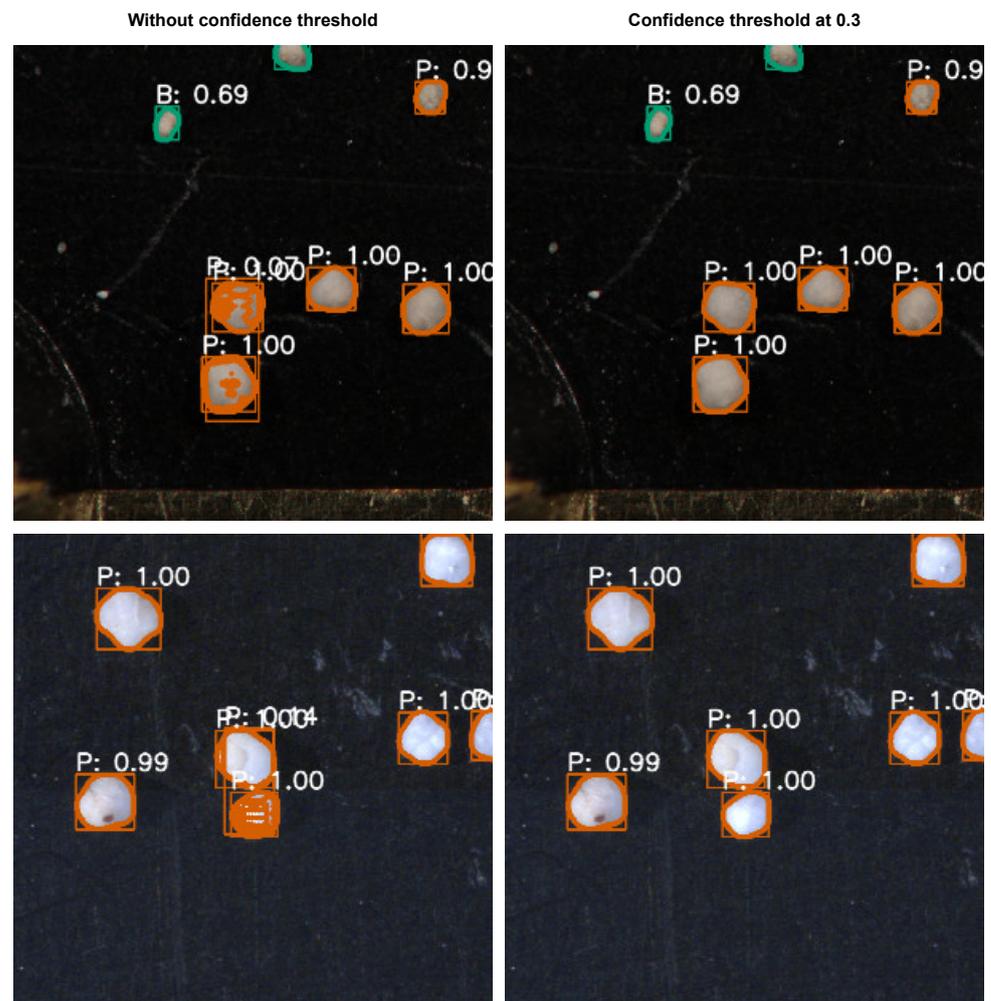
Based on the presented results, our proposed model for semantic segmentation of microscopic foraminifera is a step towards automating the process of identifying, counting, and picking of microscopic foraminifera. However, work remains to be done, such as expanding the dataset to improve the model accuracy, experimenting with other architectures, and implementing uncertainty estimation techniques.

Appendix A. Prediction Examples

Figure A1. Examples of predicted bounding boxes and segmentation masks for the "planktic" class. (**Left column**): Overlapping predictions can be seen near the middle of both images. The confidence score for the overlapping predictions with low-quality masks is significantly lower than the high-quality predictions. We can also see that some smaller objects in the top image have been misclassified as the "benthic" class. (**Right column**): The overlapping predictions have been removed by thresholding the confidence score at 0.3.
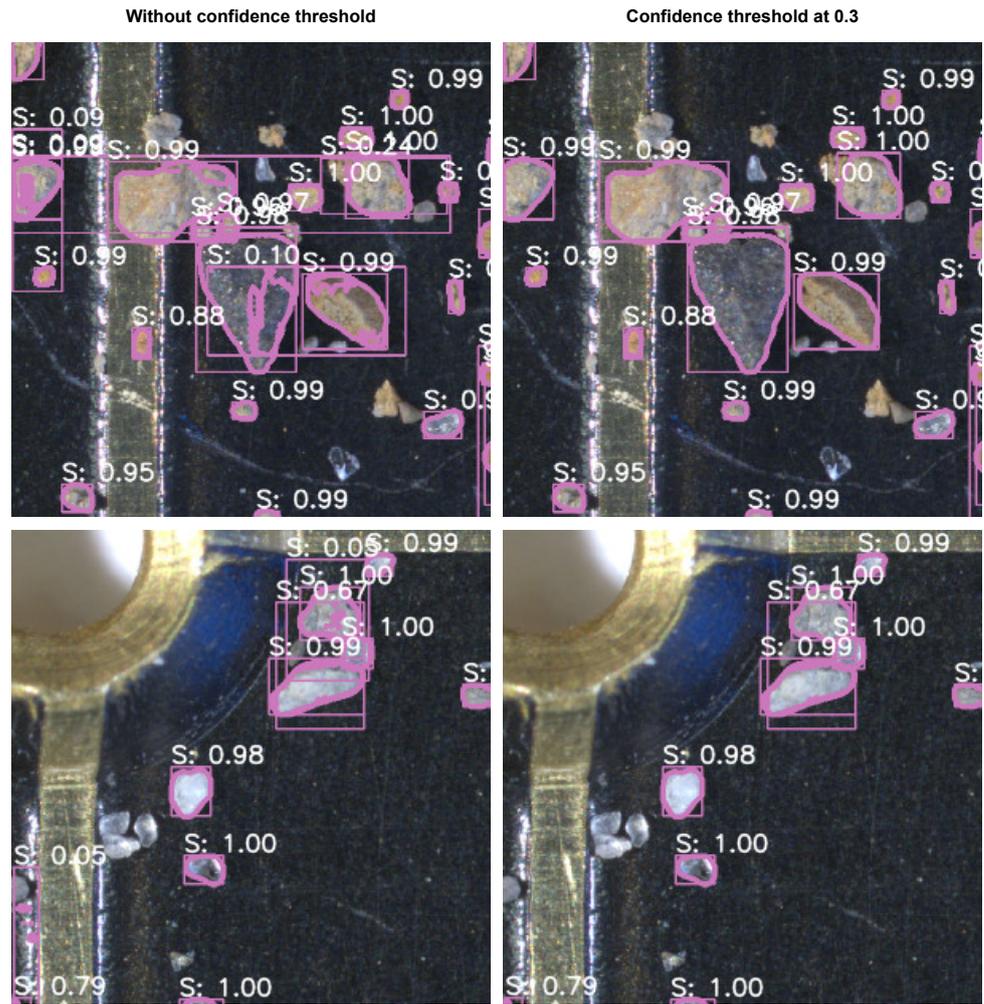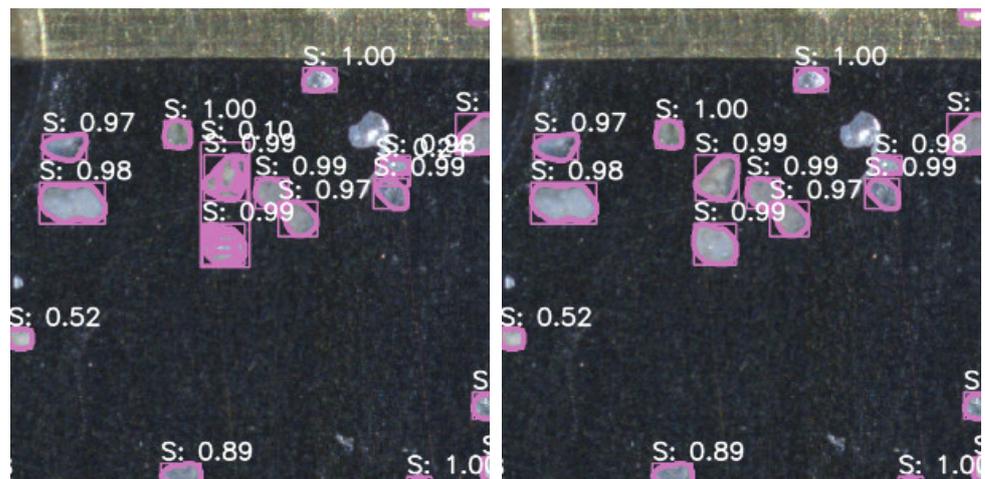
Figure A2. Examples of predicted bounding boxes and segmentation masks for the "sediment" class. (**Left column**): Overlapping predictions can be seen near the middle of both images. In addition, notice that several objects have been missed entirely. (**Right column**): Most of the overlapping predictions have been removed by thresholding the confidence score at 0.3.
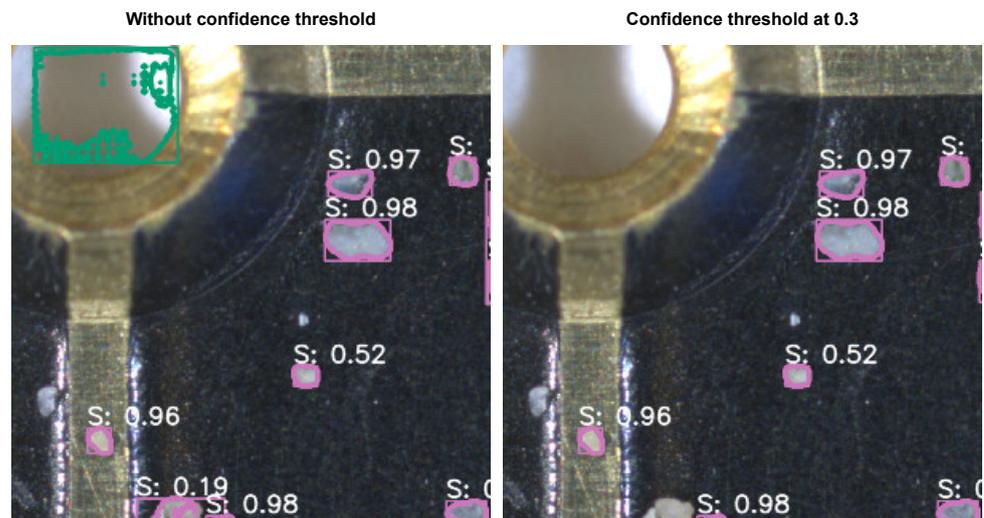
Figure A3. *Cont.*

Figure A3. Additional examples of predicted bounding boxes and segmentation masks for the "sediment" class. (**Left column**): A very obvious false positive detection of the "benthic" class can be seen near the top-left corner of the first image. For the second image, several overlapping predictions can be seen. (**Right column**): The false positive detection and the overlapping predictions have been removed by thresholding the confidence score at 0.3.

Appendix B. Additional Dataset Details
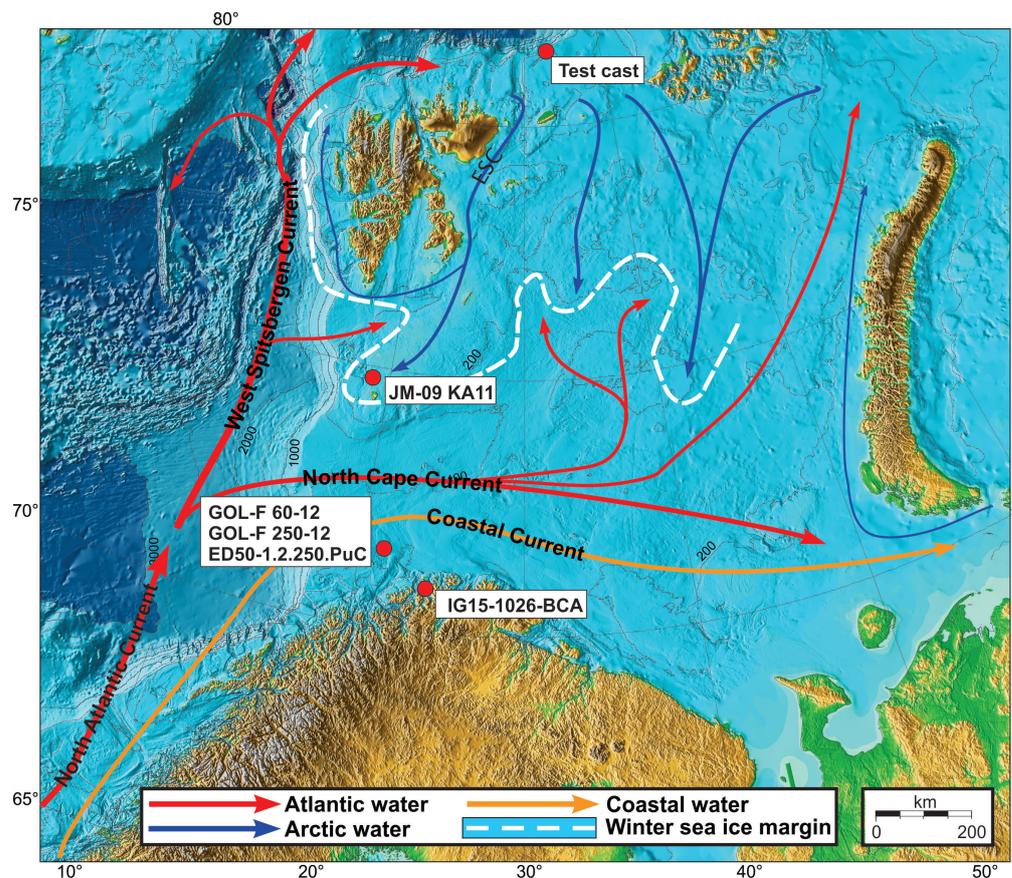
Figure A4. Topographic map of the Barents Sea region with main surface currents and winter sea ice margin. Coring sites marked with red dots. *ESC = East Spitsbergen Current.*

Table A1. Details about collected materials (object type, approximate amount, approximate age) and corresponding core information (location, water depth, core depth, coring method) used to create the presented dataset.

	Materials			Core				
Object Type	Amount	Age [Years]	Location	Name	Water Depth [m]	Depth [cm]	Type
Calcareous benthic	30	Modern 200–0	Revsbotn	IG15-1026-BCA	40	3–4	Multi-corer
	400	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	0–1	Push-core
	600	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	7–8	Push-core
	400	Modern 200–0	SW Barents Sea	GOL-F 250-12	390	2–3	Push-core
	345	13,000	West Barents Sea	JM09 KA11	345	208–209	Gravity core
	100	Modern 200–0	North slope Barents Sea	Test cast	810	0–1	Multi-corer
Planktic	500	Modern 200–0	SW Barents Sea	ED50-1.2.250.PuC	350	8–9	Push-core
	650	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	7–8	Push-core
	400	Modern 200–0	North slope Barents Sea	Test cast	810	0–1	Multi-corer
Agglutinated benthic	300	Modern 200–0	Revsbotn	IG15-1026-BCA	40	1–2	Multi-corer
	450	Modern 200–0	Revsbotn	IG15-1026-BCA	40	3–4	Multi-corer
	100	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	0–1	Push-core
	50	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	1–2	Push-core
	10	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	7–8	Push-core
Sediments	600	Modern 200–0	Revsbotn	IG15-1026-BCA	40	3–4	Multi-corer
	700	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	0–1	Push-core
	700	Modern 200–0	SW Barents Sea	GOL-F 60-12	390	7–8	Push-core
	800	13,000	West Barents Sea	JM09 KA11	345	208–209	Gravity core

References

1. Hayward, B.W.; Le Coze, F.; Vachard, D.; Gross, O. World Foraminifera Database. 2021. Available online: http://www.marinespecies.org/foraminifera (accessed on 29 April 2021)
2. Steinsund, P.I.; Hald, M. Recent calcium carbonate dissolution in the Barents Sea: Paleoceanographic applications. *Mar. Geol.* **1994**, *117*, 303–316. [CrossRef]
3. Conan, S.M.H.; Ivanova, E.M.; Brummer, G.J.A. Quantifying carbonate dissolution and calibration of foraminiferal dissolution indices in the Somali Basin. *Mar. Geol.* **2002**, *182*, 325–349. [CrossRef]
4. Sejrup, H.P.; Birks, H.J.B.; Klitgaard Kristensen, D.; Madsen, H.B. Benthonic foraminiferal distributions and quantitative transfer functions for the northwest European continental margin. *Mar. Micropaleontol.* **2004**, *53*, 197–226. [CrossRef]
5. Hald, M.; Andersson, C.; Ebbesen, H.; Jansen, E.; Klitgaard-Kristensen, D.; Risebrobakken, B.; Salomonsen, G.R.; Sarnthein, M.; Sejrup, H.P.; Telford, R.J. Variations in temperature and extent of Atlantic Water in the northern North Atlantic during the Holocene. *Quat. Sci. Rev.* **2007**, *26*, 3423–3440. [CrossRef]
6. Rasmussen, T.L.; Thomsen, E.; Nielsen, T. Water mass exchange between the Nordic seas and the Arctic Ocean on millennial timescale during MIS 4–MIS 2. *Geochem. Geophys. Geosystems* **2014**, *15*, 530–544. [CrossRef]
7. Katz, M.E.; Cramer, B.S.; Franzese, A.; Hönisch, B.; Miller, K.G.; Rosenthal, Y.; Wright, J.D. Traditional and emerging geochemical proxies in foraminifera. *J. Foraminifer. Res.* **2010**, *40*, 165–192. [CrossRef]
8. Suokhrie, T.; Saraswat, R.; Nigam, R.; Kathal, P.; Talib, A. Foraminifera as bio-indicators of pollution: A review of research over the last decade. In *Micropaleontology and Its Applications*; Scientific Publishers: Jodhpur, India, 2017; pp. 265–284.
9. Sharma, C.S.B.; Malarkodi, N.; Kumar, S.C. Role of Foraminifera in Hydrocarbon Exploration. *Int. J. Geol. Earth Sci.* **2018**, *4*, 63–65.
10. Zhong, B.; Ge, Q.; Kanakiya, B.; Marchitto, R.M.T.; Lobaton, E. A comparative study of image classification algorithms for Foraminifera identification. In Proceedings of the 2017 IEEE Symposium Series on Computational Intelligence (SSCI), Honolulu, HI, USA, 27 November–1 December 2017; pp. 1–8.
11. Ge, Q.; Zhong, B.; Kanakiya, B.; Mitra, R.; Marchitto, T.; Lobaton, E. Coarse-to-fine foraminifera image segmentation through 3D and deep features. In Proceedings of the 2017 IEEE Symposium Series on Computational Intelligence (SSCI), Honolulu, HI, USA, 27 November–1 December 2017; pp. 1–8.
12. De Garidel-Thoron, T.; Marchant, R.; Soto, E.; Gally, Y.; Beaufort, L.; Bolton, C.T.; Bouslama, M.; Licari, L.; Mazur, J.C.; Brutti, J.M.; et al. Automatic Picking of Foraminifera: Design of the Foraminifera Image Recognition and Sorting Tool (FIRST) Prototype and Results of the Image Classification Scheme. In Proceedings of the AGU Fall Meeting Abstracts, New Orleans, LA, USA, 11–15 December 2017.
13. Mitra, R.; Marchitto, T.; Ge, Q.; Zhong, B.; Kanakiya, B.; Cook, M.; Fehrenbacher, J.; Ortiz, J.; Tripati, A.; Lobaton, E. Automated species-level identification of planktic foraminifera using convolutional neural networks, with comparison to human performance. *Mar. Micropaleontol.* **2019**, *147*, 16–24. [CrossRef]
14. Johansen, T.H.; Sørensen, S.A. Towards detection and classification of microscopic foraminifera using transfer learning. In Proceedings of the Northern Lights Deep Learning Workshop, Tromsø, Norway, 20–21 January 2020.
15. He, K.; Gkioxari, G.; Dollár, P.; Girshick, R. Mask R-CNN. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *42*, 386–397.

16. Dutta, A.; Gupta, A.; Zissermann, A. VGG Image Annotator (VIA). Version: 2.0.9. 2016. Available online: http://www.robots.ox.ac.uk/~vgg/software/via/ (accessed on 10 March 2020).

17. Dutta, A.; Zisserman, A. The VIA Annotation Software for Images, Audio and Video. In Proceedings of the 27th ACM International Conference on Multimedia, MM'19, Nice, France, 21–25 October 2019. [CrossRef]

18. Girshick, R. Fast R-CNN. In Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV), Santiago, Chile, 7–13 December 2015; pp. 1440–1448.

19. Ren, S.; He, K.; Girshick, R.; Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *39*, 1137–1149.

20. Liu, S.; Qi, L.; Qin, H.; Shi, J.; Jia, J. Path Aggregation Network for Instance Segmentation. In Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–22 June 2018; pp. 8759–8768.

21. Chen, X.; Girshick, R.; He, K.; Dollar, P. TensorMask: A Foundation for Dense Object Segmentation. In Proceedings of the 2019 IEEE/CVF International Conference on Computer Vision (ICCV), Seoul, Korea, 27 October–2 November 2019; pp. 2061–2069.

22. Wang, Y.; Xu, Z.; Shen, H.; Cheng, B.; Yang, L. CenterMask: Single shot instance segmentation with point representation. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Virtual, 14–19 June 2020; pp. 12190–12199.

23. Wang, X.; Zhang, R.; Kong, T.; Li, L.; Shen, C. SOLOv2: Dynamic and Fast Instance Segmentation. *NeurIPS* **2020**, *33*, 17721–17732.

24. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Boston, MA, USA, 7–12 June 2015; pp. 770–778.

25. Lin, T.Y.; Dollar, P.; Girshick, R.; He, K.; Hariharan, B.; Belongie, S. Feature Pyramid Networks for Object Detection. In Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, USA, 21–26 July 2017; pp. 936–944.

26. Lin, T.Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Zitnick, C.L.; Dollár, P. Microsoft COCO: Common Objects in Context. 2015. Available online: https://arxiv.org/abs/1405.0312 (accessed on 2 May 2021).

27. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems*; Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2019; Volume 32, pp. 8024–8035.

28. Sutskever, I.; Martens, J.; Dahl, G.; Hinton, G. On the importance of initialization and momentum in deep learning. In Proceedings of the 30th International Conference on Machine Learning, Atlanta, GA, USA, 16–21 June 2013; McAllester, S.D., Ed.; PMLR: Atlanta, GA, USA, 2013; Volume 28, pp. 1139–1147.

29. Kingma, D.P.; Ba, J. Adam: A Method for Stochastic Optimization. International Conference on Learning Representations. 2015. Available online: https://arxiv.org/abs/1412.6980 (accessed on 14 April 2021).

30. Loshchilov, I.; Hutter, F. Decoupled Weight Decay Regularization. International Conference on Learning Representations. 2019. Available online: https://arxiv.org/abs/1711.05101 (accessed on 14 April 2021).

31. COCO: Common Object in Context—Detection Evaluation. Available online: https://cocodataset.org/#detection-eval (accessed on 17 April 2021).

32. Nigam, R.; Rao, A. The Intriguing Relationship between Coiling Direction and Reproductive Mode in Benthic Foraminifera. *J. Palaeontol. Soc. India* **1989**, *34*, 79–82.

33. Collins, L.S. The correspondence between water temperature and coiling direction in Bulimina. *Paleoceanography* **1990**, *5*, 289–294. [CrossRef]

34. Miller, D.; Nicholson, L.; Dayoub, F.; Sünderhauf, N. Dropout Sampling for Robust Object Detection in Open-Set Conditions. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–26 May 2018; pp. 3243–3249. [CrossRef]

35. Miller, D.; Dayoub, F.; Milford, M.; Sünderhauf, N. Evaluating Merging Strategies for Sampling-based Uncertainty Techniques in Object Detection. In Proceedings of the 2019 International Conference on Robotics and Automation (ICRA), Montreal, QC, Canada, 20–24 May 2019; pp. 2348–2354. [CrossRef]

36. Morrison, D.; Milan, A.; Antonakos, N. Uncertainty-aware Instance Segmentation using Dropout Sampling. In Proceedings of the Robotic Vision Probabilistic Object Detection Challenge (CVPR 2019 Workshop), Long Beach, CA, USA, 16–20 June 2019.

37. Gal, Y.; Ghahramani, Z. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In Proceeding of the 33rd International Conference on Machine Learning, New York, NY, USA, 19–24 June 2016; Volume 48, pp. 1050–1059.

38. Teye, M.; Azizpour, H.; Smith, K. Bayesian Uncertainty Estimation for Batch Normalized Deep Networks. In Proceedings of the Machine Learning Research, Stockholm, Sweden, 10–15 July 2018; Volume 11, pp. 7824–7833.