

Overview of ImageCLEFcoral 2019 task

Jon Chamberlain¹, Antonio Campello^{1,2}, Jessica Wright¹,
Louis Clift¹, Adrian Clark¹ and Alba G. Seco de Herrera¹

¹ School of Computer Science and Electronic Engineering,
University of Essex, Colchester, UK

² Filament, Cargo Works, 1-2 Hatfields, London, UK
Corresponding author: jchamb@essex.ac.uk

Abstract. Understanding the composition of species in ecosystems on a large scale is key to developing effective solutions for marine conservation, hence there is a need to classify imagery automatically and rapidly. In 2019, ImageCLEF proposed for the first time the ImageCLEFcoral task. The task requires participants to automatically annotate and localize benthic substrate (such as hard coral, soft coral, algae and sponge) in a collection of images originating from a growing, large-scale dataset from coral reefs around the world as part of monitoring programmes. In its first edition, five groups participated submitting 20 runs using a variety of machine learning and deep learning approaches. Best runs achieved 0.24 in the annotation and localisation subtask and 0.04 on the pixel-wise parsing subtask in terms of MAP 0.5 IoU scores which measures the Mean Average Precision (MAP) when using the performance measure of Intersection over Union (IoU) bigger to 0.5 of the ground truth.

Keywords: ImageCLEF, image annotation, image labelling, classification, segmentation, coral reef image annotation, marine image annotation

1 Introduction

The ImageCLEFcoral task described in this paper is part of the ImageCLEF³ benchmarking campaign [1]. ImageCLEF is part of CLEF⁴ (Cross Language Evaluation Forum) and it provides a framework where researchers can share their expertise and compare their methods based on the exact same data and evaluation methodology in an annual rhythm. In 2019, there were four tasks in ImageCLEF: ImageCLEFlifelog; ImageCLEFmedical; ImageCLEFcoral; and ImageCLEFsecurity.

In its first edition, ImageCLEFcoral follows the successful ImageCLEF annotation tasks (2012-2016) [2–6] and requires participants to automatically annotate and localize a collection of images with types of benthic substrate, such as hard coral and sponge.

³ <http://www.imageclef.org/>

⁴ <http://www.clef-campaign.org/>

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0). CLEF 2019, 9-12 September 2019, Lugano, Switzerland.

The genesis of the ImageCLEFcoral task is from the ever-increasing amount of data being collected as part of marine conservation efforts. With the rise in popularity both of Scuba diving and underwater photography, high quality equipment to archive the underwater world has become inexpensive for marine biologists and conservation enthusiasts. The performance of so-called “action cameras” such as GoPros can yield high quality imagery with very little photographic expertise from the user. Such imagery can be used to assess the health of a coral reef, for measurements of individual species and wider benthic coverage of assemblages indicating phase shifts in the marine ecosystem [7]. However, this creates a volume of data that is too large to be annotated by human labellers, and this problem is substantially more difficult than the segmentation of other image types due to the complex morphologies of the objects [8].

A typical computer vision application involves image capture, then segmenting regions of interest from the surroundings, and finally classifying them. There has been a long-standing use of machine learning in the last of these stages in particular, because alternative strategies such as rule-based classification are generally less effective. Classifying man-made objects (such as in previous ImageCLEF annotations tasks), which may have simple, well-defined shapes, is more straightforward than classifying biological objects because the shape of the latter varies from one instance to another; correct classification will often involve a combination of shape and texture.

In the same way, segmenting a natural shape from its surroundings is not necessarily straightforward. For this task for example, colour might be enough to segment some parts of a single coral from a rock lying behind it but other parts of the same coral might lie in front of a similarly-coloured sponge or fish — this makes the segmenting task much more difficult.

The difficulty in both the segmentation and classification stages explains why there are two sub-tasks within the coral reef exercise. The first is to identify the type of substrate present within a region without having to identify its outline precisely, while the second involves segmenting the coral correctly as well as classifying it.

The rest of the paper is organised as follows: Section 2 presents the 2019 ImageCLEFcoral task; Section 3 provides an overview and analysis of the collection; Section 4 details the evaluation methodology; Sections 5 and 6 present and discuss the results of the participants; and finally, Section 7 concludes the paper indicating possible new directions for the challenge.

2 Tasks

In its first edition, the ImageCLEFcoral task follows a similar format to 2015 and 2016 ImageCLEF annotation tasks [2, 3] and includes two subtasks:

Coral reef image annotation and localisation subtask For each image, participants produce a set of bounding boxes, predicting the benthic substrate for each bounding box in the images.

Coral reef image pixel-wise parsing subtask For each image, participants produce a set of polygons bounding each benthic substrate and predict the benthic substrate for each polygon in the image. This task aims to provide a more detailed segmentation of the substrate in the image.

3 Collection

The annotated dataset comprises several sets of overlapping images, each set taken in an area of underwater terrain (see Section 3.1). Figure 1 shows an example of an image in the collection. Each image was labelled by experts (see Section 3.2). The training set contains 240 images with 6670 substrate areas annotated and the test set contains 200 images with 5370 substrate areas annotated.



Fig. 1: An image from the ImageCLEFcoral 2019 task collection.

Some types of benthic substrates are easy to detect with obvious features, others are much more difficult. For example, hard corals are typified by having a rigid skeleton unlike soft corals which often have a soft, fluffy appearance. The classification and descriptions of organic benthic substrate are as follows:

Hard Coral - Branching: Morphologies grow like a tree, with branches coming out from a centre point and continually branching (secondary branching).

Hard Coral - Sub-Massive: Unlike branching colonies, these are columnar morphologies which do not have secondary branching. All branches come from the first column and typically they will only have columns.

- Hard Coral - Boulder:** Coral morphologies that grow in spherical shapes and are often called massive corals or brain corals (on some species polyp walls form ridges that look like a human brain).
- Hard Coral - Encrusting:** Morphologies that form a layer over hard surfaces, and they follow the contours of the surface. These are more difficult to see than more 3D corals, but some of those genera that form 3D morphologies also form encrusting morphologies.
- Hard Coral - Table:** Corals morphologies that look like a table, typically with a stalk connecting it with the substratum. Many of these colonies begin as branching colonies, but their branches form tight networks to create a flat surface with a table-like appearance.
- Hard Coral - Foliose:** These morphologies are named after their leaf-like structure. They are sometimes referred to as cabbage corals.
- Hard Coral - Mushroom:** These are single polyp corals. Their skeleton looks like that of an up-turned mushroom, with the ridges mimicking the “gills” of the underside of a mushroom. These can be found all over a reef, and are from only one family, the Fungidae. These single polyps are much larger than the polyps you find in colonies.
- Soft coral:** Soft corals cover a wide variety of species that are distinguished by their apparent flexibility, texture and colour. They may have clearly distinguishable open polyps. This category covers all soft corals, except Gorgonian sea fans (see below).
- Gorgonian:** Gorgonian sea fans are distinctive soft corals that grow as large branching planar structures that face into the prevailing current. They can resemble tree-like structures with a thick trunk and branches, as well as more uniformly branching structures that do not form complete plates.
- Sponge:** These can have a varied morphology and can also be highly cryptic. They differ from corals in that they do not have polyps but tiny holes, giving them a pitted appearance. This category covers all sponges, except barrel sponges (see below).
- Barrel sponge:** A sponge with a highly distinctive, dark orange/brown barrel shape with a wide opening at the top. Barrel sponges were classified morphologically, not taxonomically, so any with this growth form were grouped.
- Fire coral (Millepora):** This is not actually a coral but rather a colony of tiny animals called hydroids. They grow in a variety of shapes but typically a branching staghorn structure. They can also encrust rocks and other organic substrate such as gorgonian sea fans.
- Algae:** There are several classes of algae but we are interested only in the large-leaved foliose macro-algae and not turf algae, crustose coralline algae, encrusting algae or maerl which are difficult to classify from imagery. This type of algae can vary from large, vivid green leaves to paler, fluffy bushes.

3.1 Data acquisition

The images for the ImageCLEFcoral task are a subset from a growing, large-scale collection of images taken of coral reefs around the world as part of a coral

reef monitoring project with the Marine Technology Research Unit (MTRU) at the University of Essex. The subset used was collected from several locations in the Wakatobi Marine Reserve in Sulawesi, Indonesia in July 2018. The data was collected using SJCAM5000 Elite action cameras in underwater housings with a red filter attached, held at an oblique angle to the reef. Most images have a tape measure running through a portion of the image because they are part of a monitoring collection.

3.2 Image annotation

The images were manually annotated using a custom online polygon drawing tool (see Figure 2 for a screenshot of the tool).

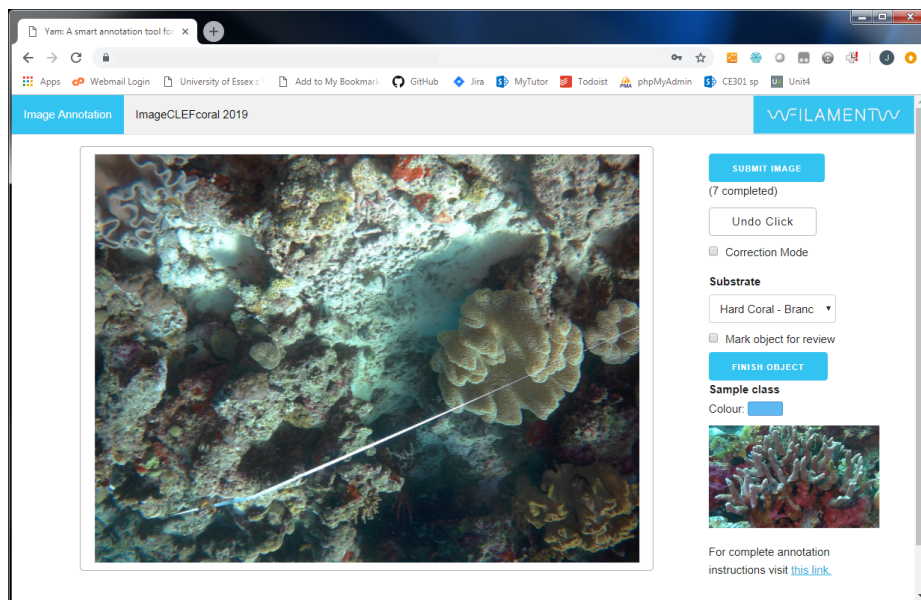


Fig. 2: A screenshot from the polygon annotation tool used for the manual labelling.

Each image was hand-annotated by postgraduate coral biology students to identify benthic substrate and validated by an administrator. Using the online tool, areas of organic benthic substrate were identified. First a polygon area was created by clicking points on the image (see Figure 3) and then a substrate type was chosen and the label completed (see Figure 4).




Fig. 3: A cropped screenshot of the annotation interface showing how dots are drawn around the object.

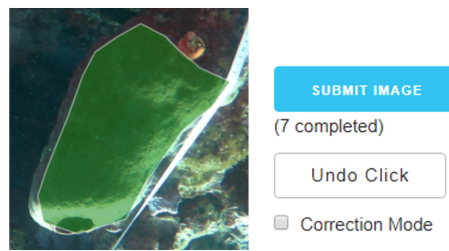


Fig. 4: A cropped screenshot of the annotation interface showing a completed polygon classified by colour.

Annotators were instructed to ensure the following rules:

- Click points must be in sequence around the object to avoid overlapping bounding boxes;
- Click points and eventual bounding box should be inside the annotated object;
- One bounding box can be used for multiple individuals (substrates) so long as they are all the same benthic substrate type;
- Bounding boxes must not overlap.




Fig. 5: An example of a training image used by ImageCLEFcoral annotators to ensure they were annotating correctly.

As a quality control to rank and train annotators, they were provided with five training images with a rich diversity of substrates, annotated previously to pixel level (see Figure 5). These training images are considered as “golden annotations” and used for the per-class agreement of the users and the ground-truth.

The agreement per class has been calculated, on a pixel level, with the intersection over union (IoU) metric:

$$\text{agreement} = \frac{\# \text{ of true positives}}{\# \text{ of false positives} + \# \text{ of false negatives} + \# \text{ of true positives}}.$$

After discarding the annotators with an agreement level smaller than 15% for quality control, the average annotation agreement and standard deviation was calculated per category of the benthic substrate type.

3.3 Collection analysis

For the training set, the proportion of “background” pixels (i.e., pixels not annotated with any of substrates) was 76.44%. Excluding the background, the benthic substrate type pixel-distribution can be found in Figure 6. The figure shows a large imbalance towards soft-corals, whereas seven benthic substrate types have been underrepresented in less than 3% of the total non-background pixels (“Hard Coral - Submassive”, “Hard Coral - Table”, “Hard Coral - Foliose”, “Hard Coral - Mushroom”, “Soft Coral - Gorgonian”, “Fire Coral - Millepora”, “Algae - Macro or Leaves”).

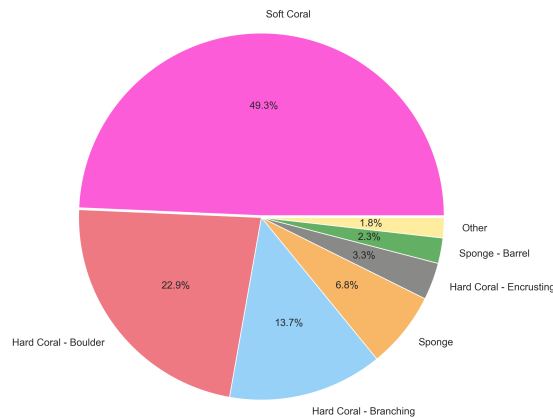


Fig. 6: Benthic substrate types distribution (per-pixel), excluding background, in the ImageCLEFcoral 2019 training set.

Figure 7 presents the histograms showing the number of objects (i.e., distinct polygons) per image and Figure 8 presents the number of points per object completed by the annotators in the training set. This shows the substrate type imbalance in the data making the task more challenging.




Fig. 7: Histogram representing the number objects per image.




Fig. 8: Histogram representing the number points (or clicks) per object.

4 Evaluation Methodology

The task was evaluated using the methodology of previous ImageCLEF annotation tasks [2, 3], which follows a PASCAL style metric of IoU. We used the following three measures:

- MAP 0.5 IoU*: the localised Mean Average Precision (MAP) for each submitted method using the performance measure of $\text{IoU} \geq 0.5$ of the ground truth;
- R 0.5 IoU*: the localised mean recall for each submitted method using the performance measure of $\text{IoU} \geq 0.5$ of the ground truth;
- MAP 0 IoU*: the image annotation average for each method in which the concept is detected in the image without any localisation.

In addition, to further analyse the results per types of benthic substrate, the measure *accuracy per substrate* was used, in which the segmentation accuracy for a substrate was assessed using the number of correctly labelled pixels of that substrate, divided by the number of pixels labelled with that class (in either the ground truth labelling or the inferred labelling).

5 Results

In 2019, 13 teams registered for the first edition of the ImageCLEFcoral task. Five individual teams submitted 20 runs. Table 1 gives an overview of all par-

ticipants and their runs. There was a limit of at most 10 runs per team and subtask.

Table 1: Participating groups

Team	Institution	# Runs T1	# Runs T2
ISEC [9]	Coimbra Institute of Engineering, Portugal	1	0
VIT [10]	Vellore Institute of Technology, India	5	0
HHUD [11]	Heinrich-Heine-Universität Duesseldorf, Germany	9	1
SOTON	University of Southampton, UK	0	3
MTRU [12]	Marine Technology Research Unit, University of Essex, UK	0	1

5.1 Coral reef image annotation and localisation (subtask 1)

Tables 2 and 3 present the performance of the participants on the coral reef image annotation and localisation subtask. 15 runs were submitted in this subtask by 3 teams.

Table 2: Coral reef image annotation and localisation performance in terms of the following scores: $MAP\ 0.5\ IoU$; $R\ 0.5\ IoU$; and $MAP\ 0\ IoU$.

Run id	team	$MAP\ 0.5\ IoU$	$R\ 0.5\ IoU$	$MAP\ 0\ IoU$
27417	HHUD	0.2427	0.1309	0.4877
27416	HHUD	0.2294	0.1307	0.5010
27419	HHUD	0.2199	0.1216	0.4421
27418	HHUD	0.2100	0.1216	0.4547
27349	VIT	0.1400	0.0682	0.4310
27348	VIT	0.1344	0.0723	0.4240
27115	VIT	0.0849	0.0456	0.4240
27350	VIT	0.0483	0.0287	0.2871
27347	VIT	0.0410	0.0274	0.2716
27421	HHUD	0.0026	0.0037	0.2051
27414	HHUD	0.0029	0.0043	0.2284
27415	HHUD	0.0027	0.0045	0.2910
27398	HHUD	0.0026	0.0043	0.2715
27413	HHUD	0.0021	0.0030	0.2028
27497	ISEC	0.0006	0.0006	0.0006

The HHUD team [11] achieved the best results in terms of $MAP\ 0.5\ IoU$ by applying a state-of-the-art deep learning approach, YOLO. Unlike the regional

convolutional neural network (R-CNN) approach adopted by [10], YOLO works on the whole image at the same time, by dividing the entire image into square cells which are predicted to contain bounding boxes of substrate. This should mean that there are fewer background errors compared to R-CNNs because more context is taken into account. In addition, the authors devised an approach of their own, first locating and then classifying, and then utilising machine learning. The strategy for locating possible substrate is that they differ from background regions, so the authors partition the image into small “tiles” and construct a feature vector containing colour, shape and texture measures: normalized histograms, grey-level co-occurrence matrices and Hu moments, respectively. From these features, a binary classifier is trained to distinguish coral and non-coral tiles. For classification of the substrate types, both k-nearest neighbour and convolutional neural network (CNN) approaches were examined, the latter contrasting shallow and deep networks and utilised the popular pre-trained VGG19 with transfer learning.

The contribution from the ISEC [9] team took a traditional computer vision approach. The authors established that colour alone is not a suitable measure for classification and established a feature vector that encapsulated colour and texture information: the mean, standard deviation, entropy of a grey-scale version of the original colour image, plus a *hue ratio* which measures colour content. All of these were calculated in a 5×5 region around each pixel. Using this feature vector, the authors explored a variety of machine learning algorithms (e.g., nearest neighbour, decision trees, discriminant analysis and support vector machines). They found that the most effective machine learning algorithm was random forests.

The VIT submission [10] approaches the annotation task using CNNs, which have been applied to a wide range of image classification tasks in recent years. A standard CNN would fare badly on this task because it is tuned to the entire image containing a single instance of the object to be classified, so they have used a variant known as Faster R-CNN. It overcomes the issue by extracting “region proposals” from the image, then using an algorithm to combine them into those passed on to the CNN. They have explored layering this on top of three existing CNNs: NASnet, Inception V2, and Resnet101. All these models are provided with the Tensorflow object detection API and are provided pre-trained using the COCO dataset, which contains 300,000 images in 80 categories of natural and man-made objects. Using the ImageCLEF training images, the authors fine-tuned the ability of the various networks to perform the classification task. They also explored the effect of augmenting the networks’ training dataset by adding noise, changing the brightness and contrast, and performing geometrical distortions such as shears, shifts, rotations and mirror imaging. CNNs currently represent the state-of-the-art in machine learning for image classification, and this use of well-known techniques provides a good benchmark.

Table 3: Coral reef image annotation and localisation performance in terms of the Intersection over Union (IoU) per benthic substrate.

Run id	Team	hard-coral-branching	hard-coral-submassive	hard-coral-boulder	hard-coral-encrusting	hard-coral-table	hard-coral-foliose	hard-coral-mushroom	soft-coral	soft-coral-gorgonian	sponge	sponge-barrel	fire-coral-millepora	algae-macro-or-leaves
27497	ISEC	0.0198	0.0	0.0007	0.0121	0	0.0	0.0	0.0079	0.0	0.0277	0.0	0	0.0
27421	HHUD	0.0007	0.0	0.0167	0.0048	0	0.0006	0.0106	0.0571	0.0	0.0072	0.0	0	0.0
27419	HHUD	0.0261	0.0	0.0315	0.0042	0	0.0023	0.0228	0.0423	0.0005	0.0088	0.0304	0	0.0012
27418	HHUD	0.0246	0.0	0.0321	0.0038	0	0.0	0.0269	0.0447	0.0005	0.0086	0.0318	0	0.0012
27417	HHUD	0.0356	0.0	0.033	0.0069	0	0.0	0.0406	0.0505	0.0008	0.0094	0.0213	0	0.0157
27416	HHUD	0.0346	0.0	0.0343	0.0064	0	0.0	0.0437	0.055	0.0008	0.0094	0.0222	0	0.0158
27415	HHUD	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0731	0.0	0.0	0.0	0	0.0
27414	HHUD	0.0089	0.0	0.016	0.0015	0	0.0	0.0	0.0562	0.0	0.0104	0.0054	0	0.0
27413	HHUD	0.0068	0.0021	0.0063	0.0014	0	0.0022	0.0	0.0523	0.0	0.0063	0.0	0	0.0
27398	HHUD	0.0013	0.0	0.0116	0.0	0	0.0	0.0	0.0702	0.0	0.002	0.0	0	0.0
27350	VIT	0.0597	0.0	0.0305	0.0141	0	0.0	0.0422	0.0808	0.0	0.0299	0.0598	0	0.0
27349	VIT	0.0637	0.0	0.1012	0.0195	0	0.0028	0.0758	0.0804	0.0	0.0329	0.0619	0	0.0004
27348	VIT	0.0548	0.0	0.0956	0.0171	0	0.0129	0.119	0.0782	0.0	0.0365	0.0579	0	0.0
27347	VIT	0.0456	0.0	0.0374	0.0055	0	0.0	0.0204	0.0918	0.0	0.0239	0.0498	0	0.0
27115	VIT	0.0436	0.0	0.0809	0.0168	0	0.0128	0.0664	0.0722	0.0	0.0349	0.0526	0	0.0

5.2 Coral reef image pixel-wise parsing (subtask 2)

Tables 4 and 5 present the performance of the participants on the coral reef image pixel-wise parsing subtask. Five runs were submitted in this subtask by three teams. In this subtask, a team comprising of researchers at Filament working

Table 4: Coral reef image pixel-wise parsing performance in terms of the following scores: $MAP\ 0.5\ IoU$; $R\ 0.5\ IoU$; and $MAP\ 0\ IoU$.

Run id	team	$MAP\ 0.5\ IoU$	$R\ 0.5\ IoU$	$MAP\ 0\ IoU$
27505	HHUD	0.0	0.0	0.0
27500	MTRU	0.0419	0.049	0.2398
27343	SOTON	0.0004	0.0015	0.0484
27324	SOTON	0.0	0.0	0.0899
27212	SOTON	0.0	0.0	0.0712

with the University of Essex [12] achieved best results in terms of $MAP\ 0.5\ IoU$. They developed a classification system based around Deeplab V3, a deep CNN that reduces the amount of post-processing necessary to deliver a final, semantic segmentation and classification. Their post-processing involved connected-component labelling, morphological opening and closing to delete small regions

Table 5: Coral reef image pixel-wise parsing performance in terms of the Intersection over Union (IoU) per benthic substrate type.

Run id	team	hard-coral-branching	hard-coral-submassive	hard-coral-boulder	hard-coral-encrusting	hard-coral-table	hard-coral-foliose	hard-coral-mushroom	soft-coral	soft-coral-gorgonian	sponge	sponge-barrel	fire-coral-millepora	algae-macro-or-leaves
27505	HHUD	0.0003	0.0	0.0	0.0	0.0	0.0	0.0	0.0432	0.0	0.0	0.0	0.0	0.0
27500	MTRU	0.0958	0.0	0.1659	0.0446	0.0	0.0065	0.219	0.13	0.0186	0.0573	0.0889	0.0	0.0007
27343	SOTON	0.0235	0.0	0.023	0.0124	0.0	0.0012	0.048	0.1145	0.0	0.0108	0.0	0.0038	0.0
27324	SOTON	0.0262	0.0296	0.0183	0.0188	0.0138	0.0135	0.0025	0.0851	0.0254	0.0283	0.0135	0.0212	0.0167
27212	SOTON	0.0121	0.0047	0.0153	0.0188	0.0	0.0042	0.0021	0.0851	0.0	0.0	0.0156	0.0	0.0008

and polygon approximation, all reasonably conventional image processing functions.

The SOTON team used a Keras implementation of Deeplab V3+, pre-trained on the well-known Pascal VOC and fine-tuned for each class using a one-versus-all, pixel-wise classifier which was trained on the ImageCLEF dataset, with the loss weighted by the ratio of pixels belonging to the class versus not. As is common with deep networks, the training data were augmented with rotation, flipping, shearing and elastic distortions. The result was then passed through a conditional random field (CRF), which allows groups of similar entities to be assigned the same label.

The HHUD [11] team also participated in this subtask and used a similar approach from subtask 1 (see Section 5.1).

HHUD and SOTON submitted self-intersecting polygons, like the one in Figure 9, which were excluded from the evaluation.



Fig. 9: An example of a self-intersecting polygon submitted to the ImageCLEF coral reef image pixel-wise parsing subtask that were excluded from the evaluation.

6 Discussion

The ultimate goal of this ImageCLEF task is the reconstruction of 3D models of coral reefs from images and using measurements of complexity, surface area and volume from the reconstructed models for advances in marine conservation monitoring. The 3D reconstruction process is known as *visual structure from motion* (SfM) and requires only a set of uncalibrated images of the object to work from. This has been used for a wide variety of reconstructions ranging from small archaeological finds to large-scale environments with some success. In the context of this work, reconstructions range from individual coral specimens (from about 100 images) up to entire reefs, the latter using an innovative multi-camera capture system. To assess the health of a coral reef, marine biologists need to measure individual specimens and it has been established that reconstructions are accurate enough for this to be possible. However, this cannot be automated because there is no easy way to establish which reconstructed points belong to which substrate type; indeed, this problem is substantially more difficult than the segmentation of images. Hence, we have conjectured that, given annotated image regions in 2D, it should be possible to carry the assigned class labels forward through the SfM processing pipeline, resulting in a 3D model in which individual substrates are labelled [13]. The 2D labelling task is what this ImageCLEF task addresses.

As such the task is different from similar image classification and marine substrate classification tasks [14–16]. Firstly, the images were collected using low-cost action cameras (approx. \$200 per camera) with a fixed lens and firing on a three second lapse. The effect of this on the imagery is that there is some blurring (in some images this is quite bad), the colour balance is not always correct (as the camera adjusts the white balance automatically based on changing environmental variables) and final image quality is lower than what could be achieved using high-end action cameras or DSLRs which are more typically used in this type of research. However, all of the images used in the task are used for building the 3D model and therefore have useful information in the pipeline. Low cost cameras were used to show this approach could be replicated affordably for marine conservation projects.

Additionally, the distance and angle the camera was facing the reef was unpredictable due to how they were placed on the multicamera array. This meant that some images were close-range and downward facing whilst other images were oblique across the reef. This has a big impact on the number of objects in the field of view and the ability of annotators to label them.

Tables 3 and 5 shows how the accuracy of the proposed approach varied between benthic substrate types. In fact, a marine dataset is difficult to annotate due to the wide variety of growth forms of benthic organisms that causes errors in image hand-annotation. Sponges are widely varied in their morphology and can often look like both hard and soft coral species, particularly encrusting varieties. Branching and submassive hard coral growth forms can be easily mistaken for each other due to their similar morphologies and the difficulty of noting secondary branching from images alone. Branching hard coral and lobed soft coral

are also difficult to distinguish. However, some morphologies are easier to detect, such as gorgonian sea fans and barrel sponges, hence these were annotated as separate categories.

Despite the difficulties, the participants applied a variety of machine learning and deep learning approaches with promising results. We noticed that in the coral reef image pixel-wise parsing subtask many self-intersecting polygons were submitted and the evaluation approach excluded this type of polygon which could cause the low performance in the subtask.

7 Conclusions

The first edition of the ImageCLEFcoral task required participants to automatically annotate and localize benthic substrate (such as hard coral, soft coral, algae and sponge) in a collection of images used for marine conservation monitoring. Five groups participated in the task with a variety of machine learning and deep learning approaches.

We hope that future editions of this task will include images from different geographical areas, meaning that the visual features of the substrate classes will be different; however, it may be possible to employ a cross-learning technique from training data from other regions. Additionally, we hope to develop methods for evaluating the subtasks based on insitu evaluation and photogrammetric evaluation, giving participants a richer set of metadata to use within their computer vision approaches.

Acknowledgments

The work of Antonio Campello was supported by Innovate UK, Knowledge Transfer Partnership project KTP010993, and hosted at Filament Consultancy Group Limited. The data collection was funded by a University of Essex Impact Acceleration Account grant ES/M500537/1 with support from Professor David Smith and Operation Wallacea. We would also like to thank the annotators Edward Longford, Ekin Yagis, Nida Sae Yong, Abigail Wink, Gareth Naylor, Hollie Hubbard, Nicholas Adamson, Laura Macrina, James Burford, Duncan O'Brien, Deanna Atkins, Hollie Sams and Olivia Beatty.

References

1. Ionescu, B., Müller, H., Péteri, R., Dicente Cid, Y., Liauchuk, V., Kovalev, V., Klimuk, D., Tarasau, A., Ben Abacha, A., Hasan, S.A., Datla, V., Liu, J., Demner-Fushman, D., Dang-Nguyen, D.T., Piras, L., Riegler, M., Tran, M.T., Lux, M., Gurrin, C., Pelka, O., Friedrich, C.M., García Seco de Herrera, A., Garcia, N., Kavalieratou, E., del Blanco, C.R., Cuevas Rodríguez, C., Vasilopoulos, N., Karampidis, K., Chamberlain, J., Clark, A., Campello, A.: ImageCLEF 2019: Multimedia retrieval in medicine, lifelogging, security and nature. In: Experimental IR Meets

- Multilinguality, Multimodality, and Interaction. Proceedings of the 10th International Conference of the CLEF Association (CLEF 2019), Lugano, Switzerland, LNCS Lecture Notes in Computer Science, Springer (September 9-12 2019)
2. Gilbert, A., Piras, L., Wang, J., Yan, F., Ramisa, A., Dellandrea, E., Gaizauskas, R.J., Villegas, M., Mikolajczyk, K.: Overview of the ImageCLEF 2016 scalable concept image annotation task. In: CLEF Working Notes. (2016) 254–278
 3. Gilbert, A., Piras, L., Wang, J., Yan, F., Dellandrea, E., Gaizauskas, R.J., Villegas, M., Mikolajczyk, K.: Overview of the ImageCLEF 2015 scalable image annotation, localization and sentence generation task. In: CLEF Working Notes. (2015)
 4. Villegas, M., Paredes, R.: Overview of the ImageCLEF 2014 scalable concept image annotation task. In: CLEF Working Notes, Citeseer (2014) 308–328
 5. Villegas, M., Paredes, R., Thomee, B.: Overview of the ImageCLEF 2013 scalable concept image annotation subtask. In: CLEF Working Notes. (2012)
 6. Villegas, M., Paredes, R.: Overview of the ImageCLEF 2012 scalable web image annotation task. In: CLEF Working Notes. (2012)
 7. Young, G.C., Dey, S., Rogers, A.D., Exton, D.: Cost and time-effective method for multi-scale measures of rugosity, fractal dimension, and vector dispersion from coral reef 3d models. *PLOS ONE* **12**(4) (04 2017) 1–18
 8. Howell, K.L., Bullimore, R.D., Foster, N.L.: Quality assurance in the identification of deep-sea taxa from video and image analysis: response to henry and roberts. *ICES Journal of Marine Science: Journal du Conseil* **71**(4) (2014) 899–906
 9. Caridade, C.M.R., Marcal, A.R.S.: Automatic classification of coral images using color and textures. Volume 2380., Lugano, Switzerland, CEUR-WS.org <<http://ceur-ws.org>> (September 9-12 2019)
 10. Jaisakthi, S.M., Mirunalini, P., Aravindan, C.: Coral reef annotation and localization using faster r-cnns. Volume 2380., Lugano, Switzerland, CEUR-WS.org <<http://ceur-ws.org>> (September 9-12 2019)
 11. Bogomasov, K., Grawe, P., Conrad, S.: A two-stage approach for localization and classification of coral reef structures. Volume 2380., Lugano, Switzerland, CEUR-WS.org <<http://ceur-ws.org>> (September 9-12 2019)
 12. Steffens, A., Campello, A., Ravenscroft, J., Clark, A., Hagrais, H.: Deep segmentation: Using deep convolutional networks for coral reef pixel-wise parsing. Volume 2380., Lugano, Switzerland, CEUR-WS.org <<http://ceur-ws.org>> (September 9-12 2019)
 13. Stathopoulou, E., Remondino, F.: Semantic photogrammetry boosting image-based 3d reconstruction with semantic labeling. *ISPRS - International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* **XLII-2/W9** (01 2019) 685–690
 14. Schoening, T., Bergmann, M., Purser, A., Dannheim, J., Gutt, J., Nattkemper, T.W.: Semi-automated image analysis for the assessment of megafaunal densities at the Arctic deep-sea observatory HAUSGARTEN. *PLoS ONE* **7**(6) (2012)
 15. Culverhouse, P., Williams, R., Reguera, B., Herry, V., González-Gil, S.: Do experts make mistakes? A comparison of human and machine identification of dinoflagellates. *Marine Ecology Progress Series* **247** (2003) 17–25
 16. Beijbom, O., Edmunds, P.J., Kline, D.I., Mitchell, B.G., Kriegman, D.: Automated annotation of coral reef survey images. In: Proceedings of the 25th IEEE Conference on Computer Vision and Pattern Recognition (CVPR'12), Providence, Rhode Island (June 2012)