ImageCLEFcoral task: Coral reef image annotation and localisation

Jon Chamberlain[1], Alba García Seco de Herrera[1], Antonio Campello[2] and Adrian Clark[1]

[1]University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK
[2]Wellcome Trust, UK

Abstract

This paper presents an overview of the ImageCLEFcoral 2022 task that was organised as part of the Conference and Labs of the Evaluation Forum - CLEF Labs 2022. The task addresses the problem of automatically segmenting and labelling a collection of underwater images that can be used in combination to create 3D models for the monitoring of coral reefs. The training data set contains images from four Worldwide geographical locations and the test data set contains images from only one of these locations. Therefore the participants could train on a subset of geographically similar images, which has been shown in previous editions of this task to be beneficial to performance. These images are grouped into image sets that can be used to create a 3D model of the environment using photogrammetry. The training dataset contained 1,374 images and 31,517 polygon objects. The test dataset comprises 200 images and 6,319 polygon objects. 6 teams registered to the ImageCLEFcoral 2022 task, of which 2 teams submitted 11 runs. Participants' entries showed that although automatic annotation of benthic substrates was possible, improving on the baselines set in previous years will be difficult.

Keywords

ImageCLEF, image annotation, image labelling, classification, segmentation, coral reef image annotation, 3D photogrammetry

CEUR-WS.org/Vol-3180/paper-97.pdf

1. Introduction

Marine ecosystem monitoring is a key priority for evaluating ecosystem conditions [1]. Despite a wide range of monitoring programs for tropical coral reefs, there is still a crucial need to establish an effective monitoring process. This process can be made by collecting 3D visual data using autonomous underwater vehicles. The ImageCLEFcoral task organisers have developed a novel multi-camera system that allows large amounts of imagery to be captured by a SCUBA diver or autonomous underwater vehicle in a single dive which will provide useful information for both annotation and further study of the coral. By releasing this data through an ImageCLEF lab [2], organised as part of the Conference and Labs of the Evaluation Forum – CLEF Labs 2022[1], advances can be made in the automatic processing at scale.

Previous editions of ImageCLEFcoral in 2019 [3] and 2020 [4] have shown improvements in task performance and promising results on cross-learning between images from geographical

CLEF 2022: Conference and Labs of the Evaluation Forum, September 5–8, 2022, Bologna, Italy
jchamb@essex.ac.uk (J. Chamberlain)
© 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)

[1]https://clef2022.clef-initiative.eu/

regions. The 3rd edition [5] increased the complexity of the task and size of data available to participants through supplemental data, resulting in lower performance than previous years. As with this 4rd edition, in 2022, the training and test data form a complete set of images required to form 3D reconstructions of the marine environment.

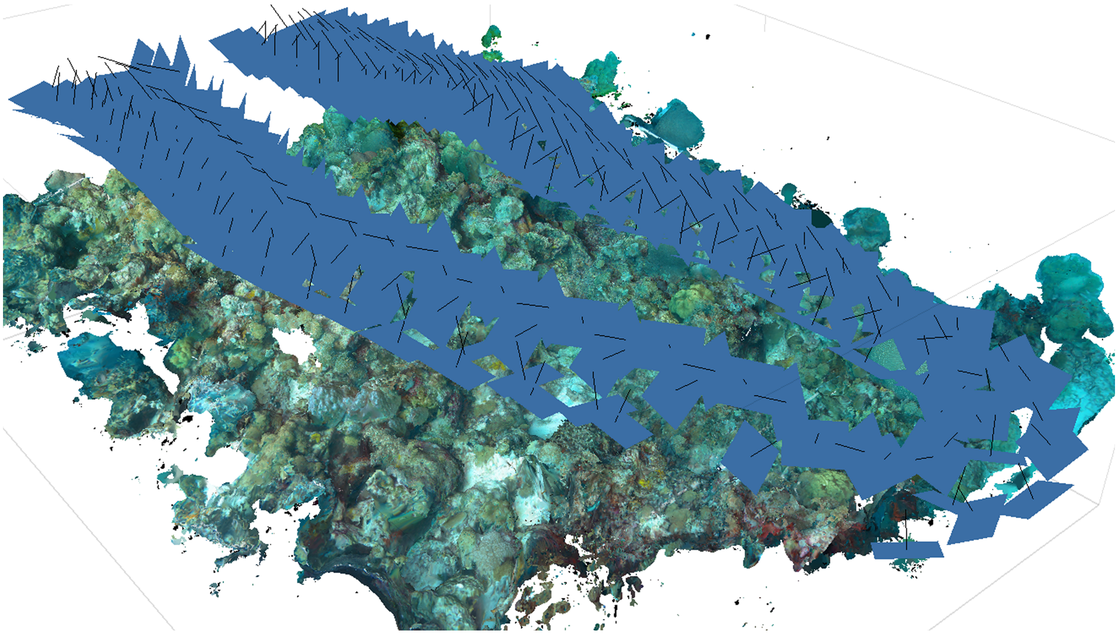

Figure 1: 3D reconstruction of a coral reef (approx. 4x6m). Each image in the subset to create this model is represented by a blue rectangle, with the track of multi-camera array clearly visible across the environment.

2. Task and Participation

In 2022, the ImageCLEFcoral task followed the format of previous editions [3, 4, 5]. Participants were again asked to devise and implement algorithms for automatically annotating regions in a collection of images containing several types of benthic substrate, such as hard coral or sponge.

As in previous editions, 2022 comprised two sub-tasks: T1-“Coral reef image annotation and localisation” and T2-“Coral reef image pixel-wise parsing” subtasks. The “Coral reef image annotation and localisation” subtask uses bounding boxes, with sides parallel to the edges of the image, for the annotation of regions in a collection of images containing several types of benthic substrates. The “Coral reef image pixel-wise parsing” subtask uses a series of boundary image coordinates which form a single polygon around each identified region in the coral reef images; this has been dubbed *pixel-wise parsing* (these polygons should not have self-intersections). Participants were invited to make submissions for either or both tasks with a limit of 10 runs per subtask.

In this 4th edition of the task 4 teams registered. Table 1 presents the two teams that submitted runs. They submitted a total of 11 valid runs. Unfortunately, this year were no participants to the “Coral reef image pixel-wise parsing” subtask.

Table 1
Participating groups of the ImageCLEF 2022 Coral task. Teams with previous participation are marked with an asterisk.

Team	Institution	Runs T1	Runs T2
HHU [6] *	Heinrich-Heine-Universität Düsseldorf, Germany	9	-
UTK [7]	University of Tennessee, Knoxville, UTK, USA	2	-

3. Data Set

The images used in the data set were captured using an underwater multi-camera system developed at the Marine Technology Research Unit at the University of Essex (MTRU), UK.

A complete set of images required to form a 3D reconstruction of the environment was provided with the training and test data. Figure 1 shows an example 3D reconstruction of one of the subsets of data (approx 4×6 m). Each image in the subset to create this model is represented by a blue rectangle, with the track of multi-camera array used for data collection clearly visible across the environment. The 3D models can be visualised online[2] and the corresponding `.obj` files were available to the participants.

The training set contains images from 5 locations (see Table 2). These images are grouped into image subsets that can be used to create a 3D model of the environment using photogrammetry and partially overlap. The test set contains images from a single location (K1, Kaledupa, Indonesia) so participants can choose which sets to train their systems with.

The ground truth annotations of the training and test sets were made by a combination of marine biology MSc students at the University of Essex and experienced marine researchers. All annotations were double checked by an experienced coral reef researcher. The annotations were performed using a web-based tool, designed to be simple to learn, quick to use and allows many people to work concurrently: full details are presented in the ImageCLEFcoral 2019 overview [3].

The data set used for the 2022 task includes data from previous versions of the task; however, all data underwent a review to improve the gold standard:

- a thumbnail for each polygon was generated and placed within a subfolder per class;
- the polygon thumbnail images for each class were reviewed at small sizing (approx. 50 per screen) to identify, remove and/or fix polygons that were very small, with an unusual shape, were a duplicate of another polygon, or had considerable overlap with another class;
- the polygon thumbnail images for each class were reviewed at medium resolution (approx. 20 per screen) to identify and correct classification errors.

The images contain annotations of the following 13 types of substrates: Hard Coral – Branching; Hard Coral – Submassive; Hard Coral – Boulder; Hard Coral – Encrusting; Hard Coral

[2] https://skfb.ly/oo6VZ

Table 2
Detatails of the ImageCLEFcoral training set.

Image subset	Location	Similarity to test set	Images
K1-20180712-01	K1, Kaledupa, Indonesia	Same location	173
PK-20180714-01	PK, Hoga Indonesia	Similar location (within 10 miles)	244
PK-20180729-02	PK, Hoga Indonesia	Similar location (within 10 miles)	270
20180406-spermonde-keke	Keke, Spermonde, Indonesia	Geographically and ecologically similar	266
20190417-seychelles-BL	Curieuse Island, Seychelles	Geographically distinct but ecologically similar	120
20170803-dominica-cabrits	Cabrits, Dominica	Geographically and ecologically distinct	301
		Total images:	1,374

– Table; Hard Coral – Foliose; Hard Coral – Mushroom; Soft Coral; Soft Coral – Gorgonian; Sponge; Sponge – Barrel; Fire Coral – Millepora[3]; and Algae - Macro or Leaves. See Table 5 for description and example images of each class.

The training dataset contained 1,374 images from 6 subsets from 4 locations (see Table 2). All subsets were complete (containing all the images to build the 3D model), except K1-20180712-01 which was a partial collection. The test data (200 images) contained more images of the K1-20180712-01 dataset.

Participants were encouraged to use the publicly available NOAA NCEI data[4] and/or CoralNet[5] to train their approaches. The NOAA NCEI data typically contains 10 annotated pixels per image, with a considerably larger classification scheme than the classes used in ImageCLEFcoral. A NOAA Translation processor, used to capture the classification types within the data set and translate them via an expert-defined translation matrix into the ImageCLEFcoral classes, was provided. Furthermore, participants were encouraged to explore novel probabilistic computer vision techniques based around image overlap and transposition of data points.

Table 3 shows the distribution of polygons per class between the training and the test datasets. The training dataset had a higher proportion of algae, boulder coral, branching coral, submassive coral and sponge, compared to the test dataset which had much more soft coral. It was hoped the inclusion of additional large-scale public datasets from NOAA would allow the participants to address the lack of training examples for under-represented classes in the training data.

[3]After 2022 evaluation of the dataset, there were no examples of this class included in the training set.
[4]https://www.ncei.noaa.gov/
[5]https://coralnet.ucsd.edu/

Table 3
Distribution of polygons per class for training and test datasets.

Substrate	Training		Test	
algae_macro_or_leaves	1,870	5.93%	106	1.68%
fire_coral_millepora	0	0.00%	1	0.02%
hard_coral_boulder	7,373	23.39%	1,209	19.13%
hard_coral_branching	3,132	9.94%	183	2.90%
hard_coral_encrusting	380	1.21%	14	0.22%
hard_coral_foliose	233	0.74%	119	1.88%
hard_coral_mushroom	335	1.06%	55	0.87%
hard_coral_submassive	2,637	8.37%	150	2.37%
hard_coral_table	920	2.92%	37	0.59%
soft_coral	7,769	24.65%	3,349	53.00%
soft_coral_gorgonian	171	0.54%	222	3.51%
sponge	6,091	19.33%	815	12.90%
sponge_barrel	606	1.92%	59	0.93%
Total	31,517		6,319	

4. Evaluation Methodology

Algorithmic performance was evaluated on the unseen test data using the popular intersection over union metric from the PASCAL VOC[6] exercise. This computes the area of intersection of the output of an algorithm and the corresponding ground truth, normalising that by the area of their union to ensure its maximum value is bounded.

As in previous years we defined the following metric:

- *MAP 0.5 IoU*: the localised Mean Average Precision (MAP) for each submitted method using the performance measure of IoU >=0.5 of the ground truth.
- *MAP 0.0 IoU*: the localised Mean Average Precision (MAP) for each submitted method using the performance measure of IoU >=0.0 of the ground truth. It indicates whether the classes are detected in the image without any localisation.

5. Results

Table 1 presents the description of the teams who participated in this ImageCLEFcoral edition. To get a better overview of the submitted runs, the results for each team are presented in Table 4.

The training and testing datasets for the various editions of this coral annotation task have differed each year so direct comparisons have to be made with some caution. Nevertheless, for the "Coral reef image annotation and localisation" subtask, these results represent an improvement on the nearest comparable previous edition. Previous editions of this exercise have shown that the use of multiple locations in the training data impacts performance. There

[6]http://host.robots.ox.ac.uk/pascal/VOC/

Table 4
Coral reef image annotation and localisation performance in terms of *MAP 0.5 IoU* and *MAP 0.0 IoU*. The best run per team is selected.

Run id	*Team*	*MAP 0.5 IoU*	*MAP 0.0 IoU*
183919	HHU	0.396	0.752
183914	HHU	0.371	0.726
183920	HHU	0.366	0.686
183911	HHU	0.365	0.721
183922	HHU	0.336	0.697
183912	HHU	0.318	0.646
183916	HHU	0.305	0.654
183913	HHU	0.297	0.661
183918	HHU	0.291	0.661
185373	UTK	0.003	0.327
184144	UTK	0.001	0.30

was no participation in the "Coral reef image pixel-wise parsing" this year: this is a more difficult task, albeit somewhat closer to the real-word problem.

Both submissions pointed out that the dataset is significantly unbalanced, reflecting the real distribution of the different types of coral in the regions in which the imagery was obtained. In particular, substrate type *c_soft_coral* accounts almost 25% of all annotations, while the least populous annotations provide under 1.5% of them. Moreover, both groups also mentioned that colours are not consistent across the dataset, a fact which again illustrates the kinds of variation that a real-world system would have to cope with. Finally, some minor problems with some of the annotations were noted by both groups.

The UTK submission [7] used a Convolutional Neural Network (CNN) architecture based around the popular VVG16 model. There was significant pre-processing in preparing the data for their model, and also some post-processing to produce the particular labels required for this exercise.

The submission from the HHU group [6] also used a CNN-based approach, though in this case centred around *Faster R-CNN* and *ResNet+FPN*. The colour cast alluded to above, the consequence of red wavelengths being extinguished more quickly with water depth than shorter wavelengths, was explicitly addressed while preparing the imagery for presentation to their system. A certain amount of hyperparameter tuning was performed. A non-maximum suppression phase was used to reduce overlapping predictions: when two bounding boxes of different classes had a IoU > 0.8, the one with the smaller confidence was discarded.

Submissions from this group use different depths of ResNet (-50, -101 and -150), and with or without colour balancing. In general, the deeper networks performed better, though the effects of colour cast removal were less clear.

6. Conclusion

The submissions to this task show that improvements in the research community's use of deep networks continues to improve performance in their ability to identify types of coral. This is an especially difficult task because, being a biological structure, coral types have characteristic features but are not necessarily similar in appearance. Hence, the best MAP 0.5 IoU score of about 0.4 represents very good performance on this extraordinarily difficult problem.

As with the previous edition of the task, the training and test data formed a complete set of images required to form 3D reconstructions of the marine environment. We believe this style of data can be explored in the future for probabilistic computer vision techniques based around image overlap and transposition of data points. A goal for the future is to collaborate with research groups to expand the training data and improve algorithms for benthic species identification.

Acknowledgments

The authors would like to thank the participants who have invested substantial amounts of time and effort in developing solutions to this task. The images used in this task were able to be gathered thanks to funding from the University of Essex and the ESRC Impact Acceleration Account, as well as logistical support from Operation Wallacea. We would also like to thank the MSc Tropical Marine Biology students who participated in the annotation of the images.

Table 5
Classes of benthic substrate, including an updated description and examples.

Class	Description	Examples
Algae - Macro or Leaves	Leafy or bulbous structures that can also overgrow other benthic substrates. Fine (grass-like) turf algae is not included. Typically vibrant green.	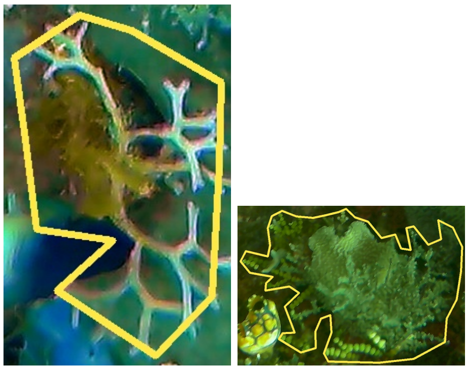
Sponge	Includes encrusting, leafy, tubular, boulder-like, vase and chimney morphologies that can appear in a variety of colours. Often have a "rough" looking surface from spicules and small holes.	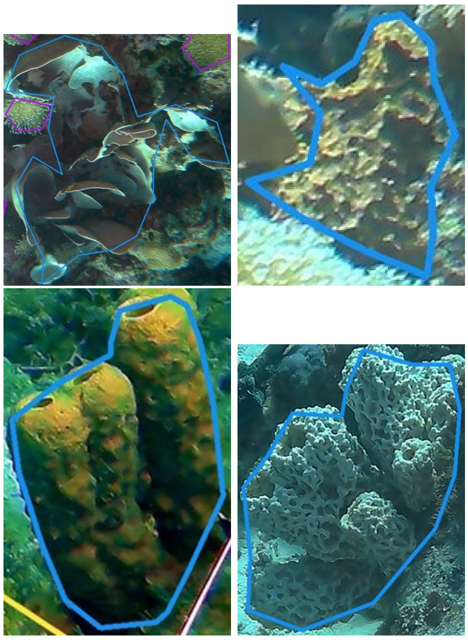
Sponge – Barrel	Includes all large barrel-sponge shaped species such as *Xestospongia muta*, but also includes young, small barrel sponges.	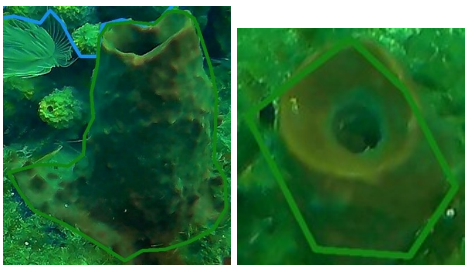

Hard Coral – Foliose	Leaf-like or cabbage-like leaf structures	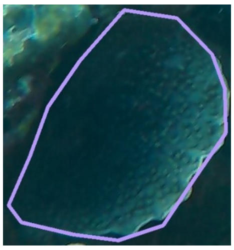 
Hard Coral – Table	Circular, broad horizontal forms originating from a single, thick stem. Polyps on the edge appear lighter.	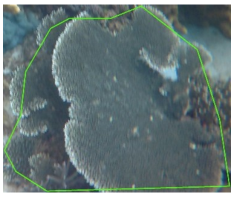 
Hard Coral – Branching	Numerous branches with secondary branching. Includes plate corals such as Elk Horn coral. Can grow in bushes similar to Table Coral but rounded at the top (not flat).	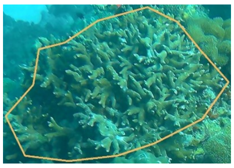 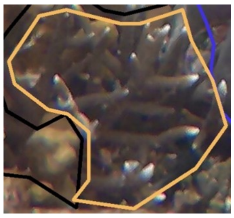
Hard Coral – Submassive	Digitate or pillar forms growing upwards from a thick stem. Includes small, packed finger-like structures and thick branching structures without secondary branching.	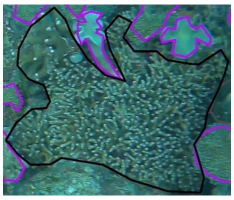 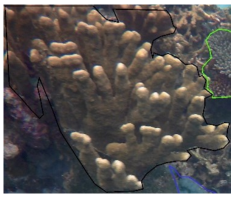
Hard Coral – Boulder	Boulder-like corals with polyps arranged evenly across the surface. Includes thin, hard encrusting type corals.	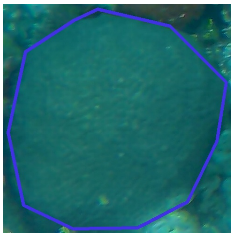 

Hard Coral – Encrusting	Fleshy or boulder-like structures with polyps arranged in channels rather than individually. Includes brain corals, rose corals and bubble corals.	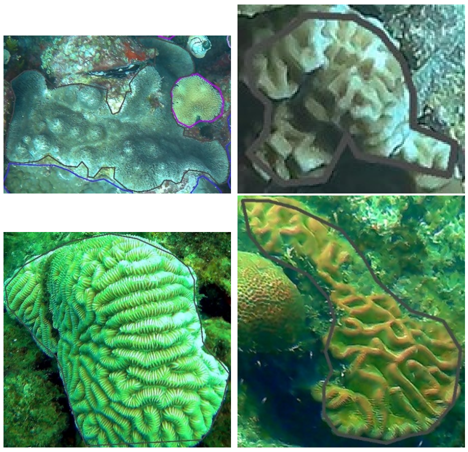
Soft Coral	A wide range of morphologies from clumped, branching types (that can be confused with branching coral) to lobed structures. Can have a fleshy, soft appearance.	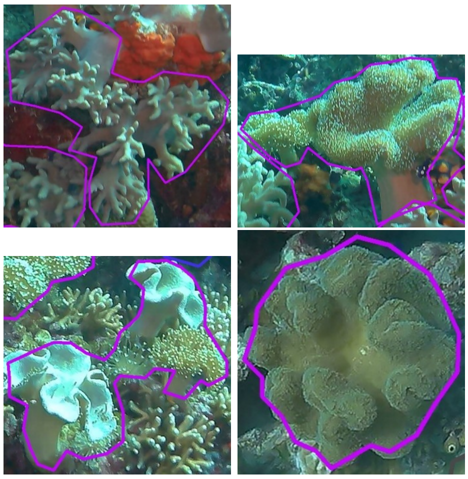
Soft Coral – Gorgonian	Sea fans (thin vertical branching plates from a single stem) and sea whips (long, thin soft coral from a single stem).	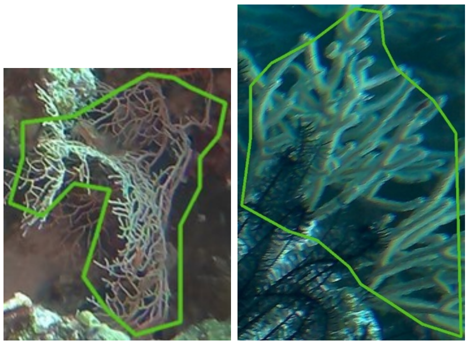
Fire Coral – Millepora	Fine branching structures similar to branching coral. Very few substrates were in the dataset and were hard to distinguish from Hard Coral - Branching so this category is not used.	

References

[1] D. M. Carrillo-García, M. Kolb, Indicator framework for monitoring ecosystem integrity of coral reefs in the western caribbean, Ocean Science Journal (2022) 1–24.

[2] B. Ionescu, H. Müller, R. Peteri, J. Rückert, A. Ben Abacha, A. G. S. de Herrera, C. M. Friedrich, L. Bloch, R. Brüngel, A. Idrissi-Yaghir, H. Schäfer, S. Kozlovski, Y. D. Cid, V. Kovalev, L.-D. Ştefan, M. G. Constantin, M. Dogariu, A. Popescu, J. Deshayes-Chossart, H. Schindler, J. Chamberlain, A. Campello, A. Clark, Overview of the ImageCLEF 2022: Multimedia retrieval in medical, social media and nature applications, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction, Proceedings of the 13th International Conference of the CLEF Association (CLEF 2022), LNCS Lecture Notes in Computer Science, Springer, Bologna, Italy, 2022.

[3] J. Chamberlain, A. Campello, J. P. Wright, L. G. Clift, A. Clark, A. García Seco de Herrera, Overview of ImageCLEFcoral 2019 task, in: CLEF2019 Working Notes, CEUR Workshop Proceedings, CEUR-WS.org, 2019.

[4] J. Chamberlain, A. Campello, J. P. Wright, L. G. Clift, A. Clark, A. García Seco de Herrera, Overview of the ImageCLEFcoral 2020 task: Automated coral reef image annotation, in: CLEF2020 Working Notes, CEUR Workshop Proceedings, CEUR-WS.org, 2020.

[5] J. Chamberlain, A. García Seco de Herrera, A. Campello, A. Clark, T. A. Oliver, H. Moustahfid, Overview of the ImageCLEFcoral 2021 task: Coral reef image annotation of a 3d environment, in: CLEF2021 Working Notes, CEUR Workshop Proceedings, CEUR-WS.org, Bucharest, Romania, 2021.

[6] F. Kerlin, K. Bogomasov, S. Conrad, Monitoring coral reefs using faster R-CNN, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction, Proceedings of the 13th International Conference of the CLEF Association (CLEF 2022), LNCS Lecture Notes in Computer Science, Springer, Bologna, Italy, 2022.

[7] R. R. Gunti, A. Rorissa, A dual convolutional neural networks and regression model based coral reef annotation and localization, in: Experimental IR Meets Multilinguality, Multimodality, and Interaction, Proceedings of the 13th International Conference of the CLEF Association (CLEF 2022), LNCS Lecture Notes in Computer Science, Springer, Bologna, Italy, 2022.