

Rip Current Segmentation: A Novel Benchmark and YOLOv8 Baseline Results

Andrei Dumitriu^{1,2}, Florin Tatui², Florin Miron², Radu Tudor Ionescu², Radu Timofte¹

¹Computer Vision Lab, CAIDAS & IFI, University of Würzburg, Germany

²University of Bucharest, Romania

andrei.dumitriu@uni-wuerzburg.de



Figure 1. An illustration of a visible rip current and its detection task, displayed from left to right: 1. The original photo, 2. The bounding box annotation from the YOLO-Rip dataset, 3. Our ground-truth annotation for instance segmentation, 4. Prediction using YOLOv8. This example highlights how bounding boxes may exclude relevant parts of the rip currents while also incorporating surrounding noise.

Abstract

Rip currents are the leading cause of fatal accidents and injuries on many beaches worldwide, emphasizing the importance of automatically detecting these hazardous surface water currents. In this paper, we address a novel task: rip current instance segmentation. We introduce a comprehensive dataset containing 2,466 images with newly created polygonal annotations for instance segmentation, used for training and validation. Additionally, we present a novel dataset comprising 17 drone videos (comprising about 24K frames) captured at 30 FPS, annotated with both polygons for instance segmentation and bounding boxes for object detection, employed for testing purposes. We train various versions of YOLOv8 for instance segmentation on static images and assess their performance on the test dataset (videos). The best results were achieved by the YOLOv8-nano model (runnable on a portable device), with an mAP50 of 88.94% on the validation dataset and 81.21% macro average on the test dataset. The results provide a baseline for future research in rip current segmentation. Our work contributes to the existing literature by introducing a detailed, annotated dataset, and training a deep learning model for instance segmentation of rip currents. The code, training details and the annotated dataset are made publicly available at https://github.com/Irikos/rip_currents.

1. Introduction

Rip currents are a common life-threatening hazard [6, 8, 17, 41] found at beaches bounded by large bodies of water, such as oceans, seas or large lakes. They can vary significantly in size and speed, being driven by changes in their controlling factors such as nearshore hydrodynamics and underwater morphology. While many are naturally occurring, some are the result of human activity, being related to the presence of coastal structures or other activities. The rip current’s main danger is its outward pulling power, some of them reaching up to 8.7km/h, which is faster than an Olympic swimmer [46]. Alongside the rip current’s drag, the actual danger is increased by the lack of knowledge on most people’s part, both in recognizing and properly reacting when caught in one. When trapped in a rip current, a common reaction is to panic and swim against it, resulting in fatigue and possible death. An advisable solution would be to swim perpendicular to it (parallel to the beach) and escape its grasp. An even better solution would be early identification of the rip current and not getting caught in it at all.

Rip currents are studied worldwide using traditional methods such as ecological dye experiments, drifter measurements and current meter surveys or video/drone footage with human analysis [3, 12, 33]. There have also been attempts of automatic rip current detection, using both traditional methods and deep learning.

The most relevant available dataset at the moment was started by de Silva *et al.* [18], with 1,740 images with rips and 700 images without rips, collected mostly from Google Earth imagery. Alongside these, there are several videos with and without rip currents, for which bounding box annotations have not been provided. Later, Zhu *et al.* [72] created a second dataset based on the images collected by de Silva, by discarding several of them and adding 1,352 high-resolution images (746 with rips and 606 without) from beaches along the coast of South China. The images are rescaled to a resolution of 640×640 pixels and are annotated with axis-aligned bounding boxes to label the rip currents.

These datasets - and the methods applied on them - only produce bounding box detections at best. While these are good, the amorphous property of rip currents and the limited format of bounding box detection imply many situations where either information is left out or background information is added to the detection. Creating a segmented dataset, while with its own challenges, opens the possibility of precisely detecting a rip current's position, shape and, possibly, direction. Furthermore, rip current detection and segmentation can increase the awareness, education and overall reaction of swimmers, surfers and other beach goers.

To this end, we propose a novel rip current segmentation benchmark, that contains both images and videos with polygonal annotations of rip currents. Our contribution can be summarized as:

- Annotating 2,466 images with polygonal annotations.
- Collecting 17 videos of rip currents (24,295 total frames) and annotating them with both bounding boxes and polygonal annotations.
- Training and comparing several versions of YOLOv8 [31] for rip current instance segmentation, serving as baselines for future research.

2. Related Works

2.1. Rip Current Detection Methods

Rip currents are widely studied in natural sciences [7, 22, 29, 33, 62, 63, 71]. Traditional observation methods include naked eye observations and camera monitoring systems [20, 27, 50]. While in-situ observations using floating objects or GPS drifters offer precision [11, 12, 61], these methods are costly, location-dependent, and impractical for detecting flash rips. Alternative methods like laser rangefinders and drones with tracer dye provide more flexibility and better views [16, 32, 51]. Our machine learning-based method surpasses these traditional approaches by being cost-effective, easily scalable, and capable of real-time



Figure 2. Rip current illustration showing the direction of the current and how to swim in order to escape it. Source: <https://www.noaa.gov/>.

detection, making it more accessible and practical for the general public.

Numerical models like SWAN [5, 19], FUNWAVE [28, 65], SWASH [13, 65], and XBEACH [19, 45, 59] can simulate rip currents based on environmental conditions [2, 21, 60]. These models help visualize rip current characteristics, but are not always precise representations due to the varied nature of rip currents. Our method offers real-time detection and could potentially be integrated with forecasting methods for improved accuracy and timely attention to high-risk periods.

Various methods analyze rip currents using collected video and image data [20, 26, 27, 50], including time exposure images or timex images. These images, averaging video frames over several minutes, enhance visual rip current detection [38, 47]. However, they are limited to wave breaking rip currents and shifting positions. Maryan *et al.* [42] compared detection algorithms on timex images, but their dataset was too small for deep learning. Pitman *et al.* [49] used synthetic rip current images, improving accuracy but underpredicting rip current channels. Liu *et al.* [39] employed a threshold and HSV-based segmentation method with 83% accuracy, but it is limited to rip currents with visible sediments. Our method uses a novel dataset and deep learning for more robust instance segmentation-based detection.

Optical flow has been utilized for rip current detection in various studies. Philip *et al.* [48] applied the Lukas-Kanade optical flow algorithm to determine water flow direction and isolate rip currents. Despite increasing resilience to small camera movements, their method requires a stable platform and only detects the predominant rip current direction. Mori *et al.* [44] modified flow visualization fields, improving rip current detection and clarity. However, it also demands a stable camera platform. McGill *et al.* [43] used dense optical flows with the Farneback method and timex images, achieving 67.3% and 96.2% accuracy in detecting

rip currents and channels, respectively. Their method is slow, taking up to 30 minutes to process a 10-minute video, and depends on camera placement and beach morphology. Optical flow methods offer certain advantages, such as rip current detection in the absence of wave-breaking patterns and providing flow information for comparison with in-situ observations. However, they are best suited for specific fixed camera placements and orientations. Our method is more noise-resistant and adaptable to various camera placements.

Rip current detection has been analyzed by a Rashid *et al.* [53] in their proposed RipNet architecture. The authors approach rip current detection as an anomaly detection problem rather than an object detection or classification problem, improving their performance through an autoencoder-based anomaly detection, reducing the need of additional negative samples during training, and achieving a considerable increase in accuracy, specificity and sensitivity metrics on the same dataset as Maryan *et al.* [42].

The same authors proposed RipDet [54], a rip current detector based on Tiny-YOLOv3, pre-trained on the COCO-2017 dataset [37] and fine-tuned on the same rip current dataset, expanded by augmentations. They obtain a significant 98.13% detection rate on the proposed benchmark, with reduced model dimensions and increased inference speed. They continued with a proposal of RipDet+ [52], which improves the previous architecture by adding residual blocks. Similar to RipDet, RipDet+ uses the weights of Tiny-YOLOv3 pre-trained on the COCO-2017 dataset for initializing the weights of all layers, except in residual and penultimate layers. Afterwards, the entire model is fine-tuned on the rip current dataset of 4,192 images. This strategy increased their detection accuracy to 98.55%.

De Silva *et al.* [18] introduced a new dataset, mainly collected from Google Earth, that includes images of rip currents and similar beaches without rip currents. They used this dataset to train a Faster R-CNN [58] model, which was tested on video data. Their approach also involves a custom temporal aggregation stage that leverages the continuity of video data to improve detection accuracy and reduce bounding box variance. The authors demonstrated the superiority of their method compared to previous approaches. Although the exact results are credible, they are not entirely reproducible, since the ground-truth annotations of the test dataset is not known. Moreover, the validation results are not provided.

Zhu *et al.* [72] expanded the training dataset collected by de Silva *et al.* [18] with additional images from the internet and South China coastal beach sites. They developed a YOLO-Rip network based on the YOLOv5s architecture, adding modifications to improve rip current detection. Their network showed increased accuracy on their dataset compared to YOLOv5, as well as faster training speeds than

previous implementations.

Our method advances existing approaches by employing instance segmentation for more precise rip current detection, utilizing a comprehensive dataset to enhance generalization, and leveraging the YOLOv8 architecture for faster and more accurate results. These improvements enable detailed analysis and facilitate real-time applications, ultimately contributing to enhanced beach safety measures.

2.2. Rip Current Detection Datasets

The development of rip current detection methods, particularly those based on machine learning and computer vision, relies heavily on the availability of high-quality, annotated datasets. In this section, we review some of the notable datasets that have been assembled and utilized for rip current detection research.

Maryan *et al.* dataset [42]: this dataset contains 514 rip channel examples, including the test dataset. The rip channels are images of 24×24 pixels extracted from images of 1334×1334 pixels. These larger images are timex images downloaded from the backlog of beach imagery on the Oregon State University website [1]. The images have been orthorectified and are time-averaged over 1,200 frames collected at 2 Hz over 10 minutes. The rip channel samples are extracted from these larger images using the GIMP image editor and normalized to 24×24 pixels. The samples are also converted to grayscale to minimize the effects of different lighting conditions on the detection models. To create a larger dataset suitable for training convolutional neural networks, data augmentation is applied, resulting in a dataset of 4,000 rip channel images. The dataset is used for training and evaluating various rip current detection algorithms in several studies, such as those conducted by Rashid *et al.* [52–54].

De Silva *et al.* dataset [18]: the primary source for the dataset was Google Earth, which provided high-resolution aerial images of rip currents and non-rip current beach scenes. In total, the dataset consists of 1,740 rip current images and 700 non-rip current images, with sizes ranging from 1086×916 to 234×234 pixels. Ground-truth annotations with axis-aligned bounding boxes were added to the rip current images. This dataset was used for training the models described in the paper. In addition to the static image dataset, De Silva *et al.* collected a test dataset consisting of 23 video clips, totaling 18,042 frames. Among these, 9,053 frames contain rip currents, and 8,989 frames do not. The image size varies from 1280×720 to 1080×920 pixels. A co-author, who is also a rip current expert at NOAA, verified the ground-truth annotations for this dataset. The frames of this video dataset were used for testing the models. It is important to note that the static images in the training set were taken from high elevation, while the test videos were taken from a lower perspective. Bounding box anno-

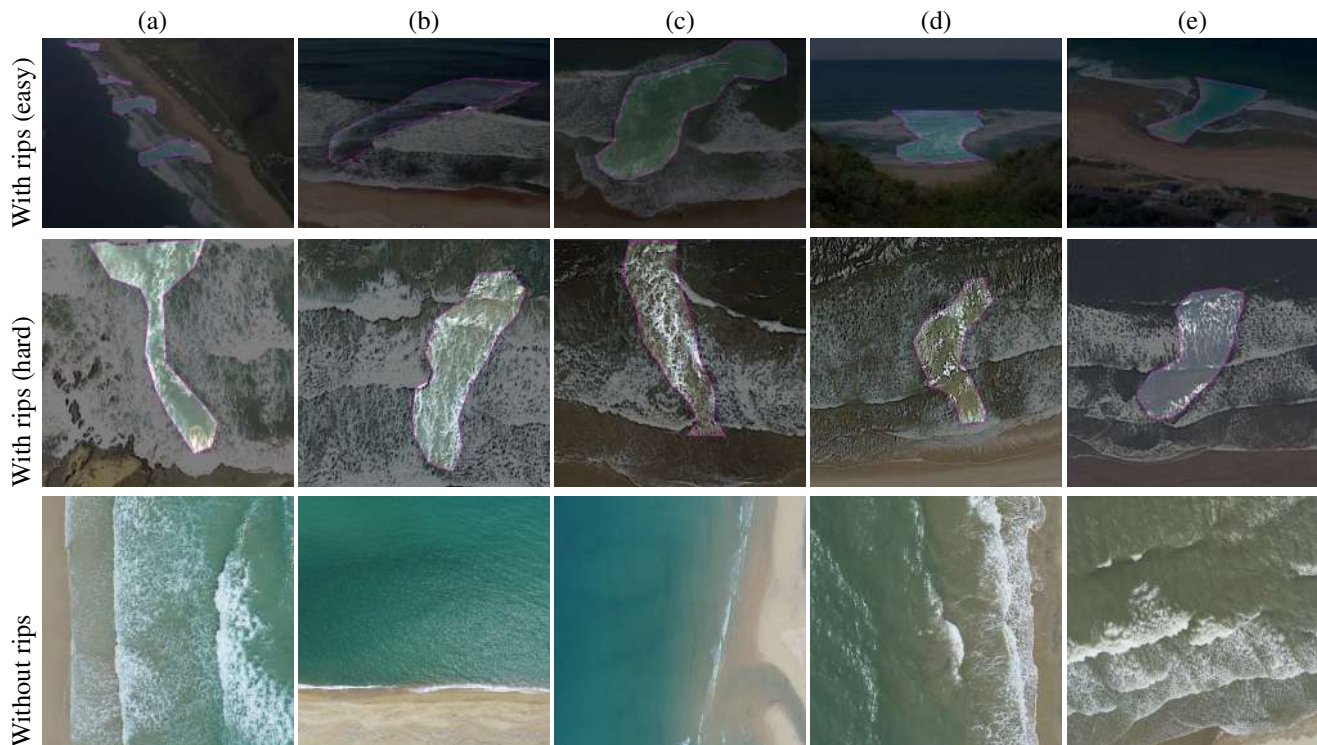


Figure 3. Selected images from the training dataset with corresponding ground-truth annotations. Rip currents exhibit an amorphous nature, resulting in varying appearances even for the same rip current, depending on the snapshot’s timing and surrounding conditions. The top and middle rows display rip currents that are easy and difficult to spot, respectively. The bottom row showcases images without rip currents in various settings.

tations are provided for the 1,740 rip current images but not for the videos.

YOLO-Rip dataset [72]: the YOLO-Rip authors expanded the dataset provided by de Silva *et al.* by collecting several sets of real beach scene photographs along the coast of South China. They selected 1,352 high-resolution images, 746 containing rip currents and 606 without, ranging in size from 4000×2250 to 480×360 pixels. Axis-aligned bounding boxes were used to label the rip boundaries in the images containing rip currents. This expanded dataset aimed to improve the model’s ability to accurately identify rip currents from diverse types of images and enhance its real-world applicability.

3. Proposed Benchmark Dataset

Rip currents exhibit amorphous characteristics, lacking a definite shape. Existing datasets predominantly consist of bounding box annotations, which, despite being more convenient to create, provide only a rough estimation of the rip current channel location. This often results in either the inclusion of noise or the exclusion of relevant information. Our approach builds upon the datasets provided by Zhu *et al.* and De Silva *et al.*. Similar to De Silva *et al.*, we divide

the dataset into training images and test videos, supplying files and annotations for both.

3.1. Training images

We have annotated 2,466 images from the YOLO-Rip dataset for instance segmentation. As rip currents are amorphous, we assumed an approximate shape based on expert knowledge in cases where the entire shape was not visibly apparent. Additionally, we included 1,307 images of similar beaches without rip currents from the same dataset.

Scenes: numerous images in the dataset are satellite views of beaches, captured from Google Earth. While most images display a clear wave-breaking pattern, rip currents occasionally reveal their shape through distinct water colors. Rip currents, defined by their outward-flowing sea channels, exhibit periodic apparent shapes due to continuous water movement and waves. In videos, the rip current channel’s shape is more easily discernible, thanks to extended observation. However, in images, the rip current channel’s visibility depends on the snapshot’s timing, which can vary from obvious to obstructed. In the latter case, informed assumptions are made based on the shape.

Annotations: annotations were generated using

Model	Last epoch	Best epoch	Train time (h)	FPS	Precision	Recall	mAP50	mAP50:95	Stddev (mAP50)
YOLOv8n-seg	180.8	130.8	1.58	73.33	88.27	84.67	88.94	45.75	2.09
YOLOv8s-seg	167.3	117.3	1.83	33.74	87.62	84.67	89.25	46.66	1.95
YOLOv8m-seg	147.2	97.20	2.23	15.54	88.04	85.04	89.92	47.48	2.23
YOLOv8l-seg	135.6	85.6	2.91	9.664	89.04	84.97	89.72	46.85	1.90
YOLOv8x-seg	139.6	89.6	4.86	6.324	88.74	84.26	90.19	47.45	1.45

Table 1. Validation results on the testing dataset. Each metric value is the average over all 10 folds, except for stddev. The difference in mAP50 (our most relevant metric) is marginal between the models, even when comparing the smallest to the largest. However, as the model size increases, the FPS rate drops drastically.

Roboflow’s polygon tool [23]. One of the challenges in rip current segmentation is the annotation stage. As previously mentioned, images showcase rip currents at different moments and with varying apparent shapes. Lacking context, assumptions must be made. Depending on the image and the rip current’s apparent shape, our approach varied between annotating the obvious shape, the assumed channel, or a combination of both to provide as much information as possible to the model. Some images contain a significant amount of calm offshore water, devoid of easily distinguishable characteristics. In these instances, decisions were made regarding where to terminate the rip current segmentation based on the image itself. Occasionally, the rip current’s outer boundaries are segmented by waves before breaking, while at other times, waves are included in the segmentation due to the rip current’s apparent shape in front of and behind the waves. The annotations were checked by a co-author who is an expert in the field of rip currents.

3.2. Test videos

We have compiled 17 videos over three years (2021, 2022, and 2023), captured at 30 FPS, resulting in a total of 28 minutes and 58 seconds of aerial footage of rip currents and 24,295 frames. Among these videos, 2 are recorded at a resolution of 3840×2160 , while the rest at a resolution of 1920×1080 . We recorded the videos at various dates and times from a Black Sea beach site in Eforie Nord, Romania.

The duration of the videos range from as short as 2 seconds to as long as 10 minutes and 46 seconds. This variation is intentional to evaluate the model’s performance on both short and long instances, since it was only trained on images. The average duration of the videos is 1 minute and 42 seconds, with a median value of 1 minute. Most videos feature a hovering fixed standpoint overlooking the rip current, with only minor adjustments in position and angle due to wind. Some footage captures either the camera or the drone moving, taken from different elevations. The 1920×1080 videos were recorded at f2.8 and ISO 100. We also included 5 beachfront videos without rip currents from De Silva *et al.*

test data in order to evaluate for false positives, as well.

We encourage other researchers and organizations to contribute to the collection of diverse rip current samples, fostering collaboration and expansion of the dataset for improved model performance.

4. Baseline method: YOLOv8

YOLOv8 [31] is the latest version in the YOLO (You Only Look Once) series of object detection algorithms, known for their real-time detection capabilities and accuracy [4, 30, 34, 35, 55–57, 64]. YOLOv8 employs an anchor-free detection mechanism, an enhanced feature pyramid network, and a modified loss function to achieve improved performance over its predecessors. It is designed to segment objects in images by predicting bounding boxes and associated class probabilities using a single convolutional neural network (CNN). Before choosing YOLOv8, we considered other models that are in the top ranking of COCO instance segmentation test-dev, such as EVA [24], FD-SwinV2-G [68], BEiT-3 [67], MasK Dino [36], ViT-Adapter-L [15] and SwinV2-G [40], as well as older methods, such as RDSnet [66], TensorMask [14], PolarMask [69], SipMask [9], D2Det [10] and Mask R-CNN [25]. However, we selected YOLOv8 as our baseline method for several reasons. First, YOLO algorithms have consistently demonstrated strong performance in object detection tasks, making them competitive and suitable for our rip current instance segmentation problem. The real-time detection capabilities of YOLOv8, in particular, are essential for applications involving public safety and emergency response, as demonstrated by the success of previous YOLO versions in various object detection tasks. Second, YOLOv8 is a well-established method with an extensive user community, providing accessible implementation resources and making it an appropriate choice for a benchmark study. Third, it is easy to deploy and use in multiple version sizes.

In our study, we opted to use YOLOv8 out-of-the-box without any modifications. This decision ensures a clear and unbiased evaluation of our dataset, and provides a straightforward starting point for future research. By using

YOLOv8 unmodified, we establish a performance baseline that can be easily replicated and compared with other methods applied to our dataset.

To determine the most suitable model size for our experiments, we conducted a comparative analysis of the various YOLOv8 models. This comparison helps to identify the trade-offs between detection accuracy and computational efficiency, ensuring that the selected model is appropriate for the specific hardware and application scenario. We present a detailed analysis of the model sizes and their performance on our dataset in the experiments section.

5. Experiments

5.1. Environment and parameters

The training and inference were run on a 24GB NVIDIA GeForce RTX 3090 GPU, using YOLOv8.0.53, Python 3.10.4, PyTorch 1.12.1 and CUDA 11.7. We trained on images of 640×640 pixels for 300 epochs, on mini-batches ranging from 64 to 12, depending on model size, with an early stopping patience parameter of 50 epochs. The models stopped before overfitting somewhere between epoch 104 and 214. The detailed model parameters and logs can be accessed in the github repository.

5.2. Evaluation metrics

For evaluation on the validation data, we take a look at several metrics: last epoch, to see where the model training stopped due to overfitting (with a patience of 50 epochs), best epoch, *i.e.* the epoch with the best results before starting to overfit, training time (average per epoch), frames per second (FPS), precision (2), recall (3), mAP50, map50:95 (5), and standard deviation over the 10 folds for the mAP50 values. All values, except for the standard deviation, are the average over all 10 folds. We use the mean Average Precision (mAP) as our primary evaluation metric, specifically mAP50. We also look at frames per second, considering real-world applications. The mAP is based on the Intersection over Union (IoU) (1), which is derived from the Jaccard index, and represents the average of the precision values at different recall levels for a given IoU threshold.

The IoU, or Jaccard index, measures the similarity between two sets and is defined as the ratio of their intersection to their union:

$$IoU(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (1)$$

Precision and Recall are calculated as follows:

$$Precision = \frac{True\ Positives}{True\ Positives + False\ Positives}, \quad (2)$$

$$Recall = \frac{True\ Positives}{True\ Positives + False\ Negatives}. \quad (3)$$

In general, the Average Precision (AP) for a single class is calculated by ranking the model's predictions by their confidence scores, and then computing the area under the precision-recall curve:

$$AP = \sum_n (Recall_n - Recall_{n-1}) \cdot Precision_n. \quad (4)$$

Since our model detects only one class (rip currents), the mAP is equal to the AP of that class. For mAP50, we compute the AP at an IoU threshold of 0.5. For mAP50:95, we calculate the AP for each IoU threshold from 0.5 to 0.95 in increments of 0.05, and then average those values:

$$mAP = \frac{AP_{IoU=0.5} + AP_{IoU=0.55} + \dots + AP_{IoU=0.95}}{k}, \quad (5)$$

where k is the number of IoU thresholds considered.

By using mAP50 and mAP50:95, we evaluate our model's ability to accurately segment rip currents at varying IoU thresholds. We also assess the model's efficiency by considering training speed and inference time, ensuring that our solution is practical for real-world applications in rip current detection and analysis.

For evaluation on the test videos, we examine the number of frames with accurately segmented rip currents using the mAP50 metric. We assess the accuracy for each video and analyze the failure cases. The final results are presented as both macro average, considering the average accuracy across all videos, and micro average, taking into account the number of true positives across all frames.

5.3. Training Method

Given the diverse nature of our dataset, which includes satellite, aerial, and beachfront images, we employed a k-fold cross-validation technique with 10 folds and a 90% train and 10% validation split. The dataset comprises 2,466 images of rip currents and 1,307 images without rip currents. To avoid an imbalanced fold, we separately applied a 10-fold cross-validation on each category (rip and non-rip), and then concatenated the file lists before training. This approach, known as stratified sampling, ensures a consistent class distribution in each fold.

For full reproducibility, the code and files used in each fold are provided in the github repository. We trained all model sizes (n, s, m, l, and x) on the same 10-fold split, comparing all the results, training and inference times, but focusing on mAP50 and FPS. The metrics compared for each model represent the average over all 10 folds, except for standard deviation, which is the metric for the standard deviation of the mAP50 values between all 10 folds.

5.4. Results

We compared all the models via cross-validation on images and testing on videos. The validation results and

Video	Frames	YOLOv8-n (%)	YOLOv8-s (%)	YOLOv8-m (%)	YOLOv8-l (%)	YOLOv8-x (%)
DJI_0092.mp4	278	100.0	75.90	10.43	100.0	0.720
DJI_0093.mp4	646	96.90	72.45	97.52	100.0	96.75
DJI_0094.mp4	5248	50.84	71.55	91.54	89.81	84.81
DJI_0095.mp4	2516	61.29	30.45	76.27	56.68	18.76
DJI_0384.mp4	4304	93.66	76.65	66.94	72.10	85.04
DJI_0393.mp4	84	100.0	100.0	100.0	100.0	100.0
DJI_0394.mp4	3623	71.93	86.28	93.43	85.04	97.90
DJI_0395.mp4	4583	75.98	98.19	95.99	87.67	92.84
DJI_0410.mp4	1397	94.56	47.24	95.49	59.06	85.25
DJI_0423.mp4	1825	65.04	81.04	47.18	14.85	46.79
DJI_0438.mp4	1238	81.83	70.19	92.08	52.18	57.27
DJI_0448.mp4	10	100.0	10.00	10.00	100.0	100.0
DJI_0670.mp4	1812	76.49	46.52	46.30	56.73	74.78
DJI_0678.mp4	1238	89.74	98.87	91.68	22.29	99.35
DJI_0679.mp4	634	98.58	100.0	100.0	97.16	99.37
DJI_0680.mp4	1359	83.08	93.97	100.0	91.39	89.04
DJI_0808.mp4	2194	54.51	0.270	0.000	8.200	0.730
no_rip_01.mp4	939	98.94	89.03	97.23	93.29	95.63
no_rip_02.mp4	924	100.0	100.0	100.0	100.0	100.0
no_rip_03.mp4	921	12.49	27.25	92.40	31.38	18.13
no_rip_04.mp4	926	82.07	50.76	89.09	90.71	85.85
no_rip_05.mp4	922	82.21	80.91	73.21	88.61	81.24
no_rip_06.mp4	451	93.57	49.22	72.73	73.61	49.67
no_rip_07.mp4	900	83.67	93.44	89.00	48.67	54.00
no_rip_11.mp4	2576	82.92	52.60	78.22	76.36	74.18
All videos	(macro avg)	81.21	68.11	76.27	71.83	71.52
All videos	(micro avg)	74.83	69.43	78.86	69.62	73.30

Table 2. Results on the test dataset (videos) for rip current segmentation. Values represent the percentage of successfully segmented frames using mAP50. Macro and micro averages are shown in the last two rows. Best results are highlighted in blue, extreme failure cases (<20% accuracy) in red, and 20%-50% accuracy in orange.

speeds can be seen in Table 1. On the training and validation data, as expected, the difference in FPS or inference time between models is significant and correlates with size. Contrary to initial expectations, the larger models do not yield a drastic improvement in segmentation performance on the validation data.

The test results can be seen in Table 2, and manually selected frames have been included in Figure 4. While the results vary, we report a maximum macro average accuracy of 81.21% on all videos and a maximum micro average accuracy of 78.86% on all frames. These results could be improved with several post-processing methods, such as taking into consideration the temporal dimension of the videos. It is interesting that the best performing model in most situations is the smallest one. Zeiler and Fergus [70] have shown that deeper layers of CNNs tend to learn more and more complex patterns. Seeing that we do not detect an exact shape, but more of a general pattern (due to their amorphous property), it is possible that this problem is easier to

solve with low-level features.

The nano model successfully detects the rip currents in many situations (more than 50% of the frames of each video), leading to a clear segmentation with minimal inclusion of noise and minimal exclusion of relevant information, even in situations where multiple rip currents are visible. For the tested videos, the false negatives are less than 20% of the frames, depending on the video. To all these, there is only one clear exception: the video no_rip_03.mp4. This video is taken at around head-level elevation. The false positive is detected in the horizontal breaking pattern of the waves. This is an important error, as the high number of false positives makes it harder to correct even by post-processing the information. An in-depth analysis of false negatives reveals that certain factors, such as wave patterns and lighting conditions, could be responsible for these errors. Future work should focus on addressing these factors by fine-tuning model parameters and incorporating additional context information.

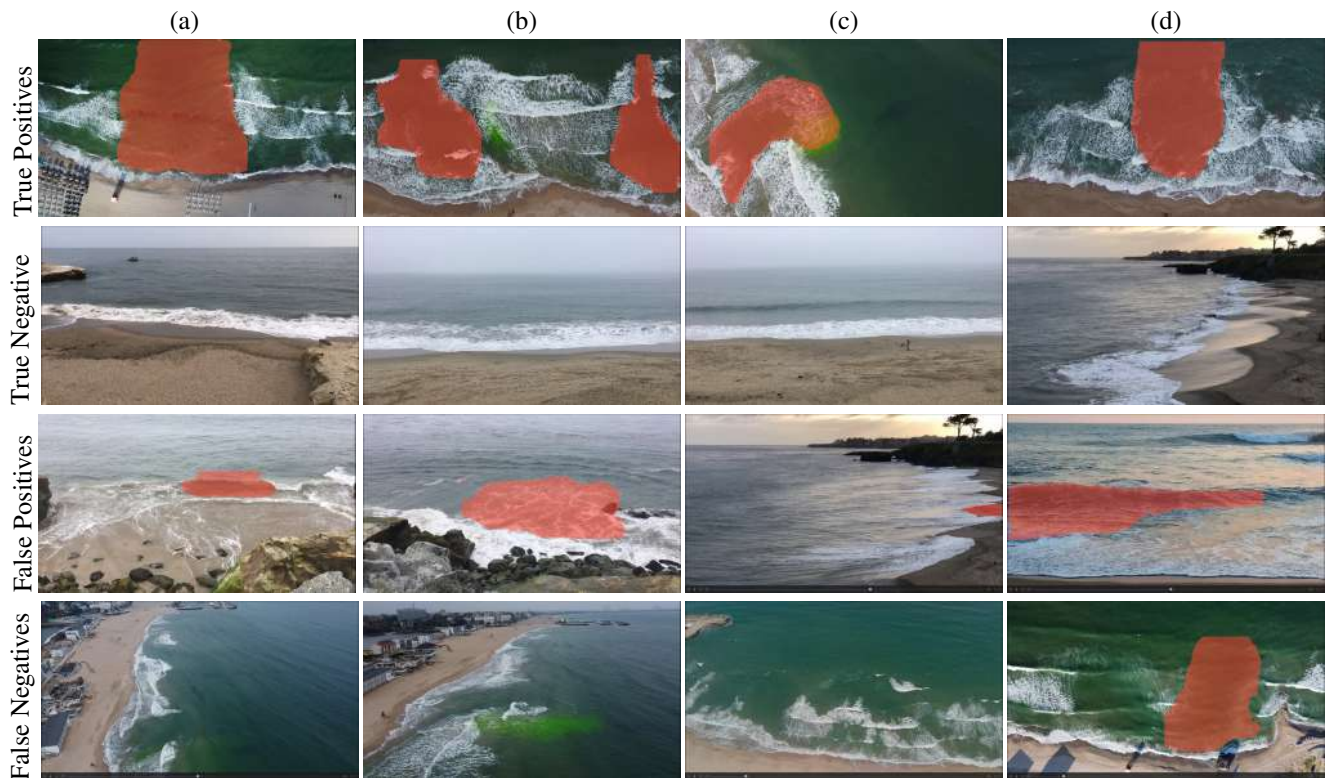


Figure 4. Example of results on testing dataset (using the nano model). On the first two rows we have the correct predictions, the true positives and the true negatives. On the last row we have the incorrect predictions, the false positives and the false negatives. Notice column (d) on the last row, where the model manages to correctly predict one rip current, but misses the second one.

6. Conclusion

In this study, we introduced a novel training dataset for rip current instance segmentation, a task not previously approached using deep learning. We also presented 17 videos with both bounding box and instance segmentation annotations, paving the way for methods leveraging the temporal continuity of videos. This dataset is the first of its kind. We employed the YOLOv8 model in all available sizes (n, s, m, l, and x) and compared their training speed, inference speed, mAP50, and mAP50:95 performance. The best performing model is the smallest one (nano), likely due to the amorphous nature of rip currents. The nano model performs well on all videos except one, where it yields a strong false negative. Addressing false negatives is crucial for user safety. This study demonstrates the potential of deep learning models, specifically YOLOv8, for rip current segmentation. However, we acknowledge the limitations of our approach due to the dataset distribution, and the fact that the model can only generalize based on the available data.

In future work, we aim to focus on expanding the dataset, investigating alternative architectures, fine-tuning model parameters, developing real-time detection systems for videos, including temporal information, and adding

multiple classes such as swimmers, surfers, and boats in the segmentation task. The developed model holds promise for real-world applications, such as integration into beach safety systems or mobile devices for user alerts. Further work is required to optimize the model for deployment in these scenarios. Moreover, we plan to explore collaboration opportunities with domain experts, organizations, and beach authorities to improve data collection and model deployment in real-world environments.

The presented dataset, models, and evaluation methods aim to advance the field and improve beach safety by preventing rip current-related incidents. To facilitate future research, we have made the code, logs, exact k-fold splits, and datasets publicly available.

Acknowledgments

This work was partly supported by the Humboldt Foundation.

We thank previous authors for their research and dataset contributions. Through continued effort, a larger, diversified dataset is being developed. Rip current detection has life-saving potential, and each study contributes to enhancing global beach safety.

References

- [1] Oregon state university: Coastal imaging lab. <http://cil-www.coas.oregonstate.edu>. Accessed: March 2023. **3**
- [2] SVV Arun Kumar and KVS Prasad. Rip current-related fatalities in India: a new predictive risk scale for forecasting rip currents. *Natural hazards*, 70:313–335, 2014. **2**
- [3] Guido Benassai, Pietro Aucelli, Giorgio Budillon, Massimo De Stefano, Diana Di Luccio, Gianluigi Di Paola, Raffaele Montella, Luigi Mucerino, Mario Sica, and Miela Pennetta. Rip current evidence by hydrodynamic simulations, bathymetric surveys and UAV observation. *Natural Hazards and Earth System Sciences*, 17(9):1493–1503, 2017. **1**
- [4] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020. **5**
- [5] N Booij, LH Holthuijsen, and RC Ris. The “swan” wave model for shallow water. In *Coastal Engineering*, pages 668–676, 1996. **2**
- [6] R Brander, Dale Dominey-Howes, C Champion, O Del Vecchio, and B Brighton. Brief Communication: A new perspective on the Australian rip current hazard. *Natural hazards and earth system sciences*, 13(6):1687–1690, 2013. **1**
- [7] Robert W Brander and AD Short. Morphodynamics of a large-scale rip current system at Muriwai Beach, New Zealand. *Marine Geology*, 165(1-4):27–39, 2000. **2**
- [8] B Chris Brewster, Richard E Gould, and Robert W Brander. Estimations of rip current rescues and drowning in the United States. *Natural Hazards and Earth System Sciences*, 19(2):389–397, 2019. **1**
- [9] Jiale Cao, Rao Muhammad Anwer, Hisham Cholakkal, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. Sipmask: Spatial information preservation for fast image and video instance segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 1–18. Springer, 2020. **5**
- [10] Jiale Cao, Hisham Cholakkal, Rao Muhammad Anwer, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. D2det: Towards high quality object detection and instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11485–11494, 2020. **5**
- [11] Bruno Castelle, Rafael Almar, Matthieu Dorel, Jean-Pierre Lefebvre, Nadia Senechal, Edward J Anthony, Raoul Laibi, Rémy Chuchla, and Yves du Penhoat. Rip currents and circulation on a high-energy low-tide-terraced beach (Grand Popo, Benin, West Africa). *Journal of Coastal Research*, 70(10070):633–638, 2014. **2**
- [12] B Castelle, T Scott, RW Brander, and RJ McCarroll. Rip current types, circulation and hazard. *Earth-Science Reviews*, 163:1–21, 2016. **1, 2**
- [13] Sung-Yeol Chang, Jin-Hoon Kim, Hyung-Seok Lee, and Inho Kim. Mechanism of rip current generation at Deoksan beach, South Korea. *Journal of Coastal Research*, 114(SI):295–299, 2021. **2**
- [14] Xinlei Chen, Ross Girshick, Kaiming He, and Piotr Dollár. Tensormask: A foundation for dense object segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2061–2069, 2019. **5**
- [15] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. *arXiv preprint arXiv:2205.08534*, 2022. **5**
- [16] David B Clark, Luc Lenain, Falk Feddersen, Emmanuel Boss, and RT Guza. Aerial imaging of fluorescent dye in the near shore. *Journal of Atmospheric and Oceanic Technology*, 31(6):1410–1421, 2014. **2**
- [17] AH Da F. Klein, GG Santana, FL Diehl, and JT De Menezes. Analysis of hazards associated with sea bathing: results of five years work in oceanic beaches of Santa Catarina state, southern Brazil. *Journal of Coastal Research*, pages 107–116, 2003. **1**
- [18] Akila de Silva, Issei Mori, Gregory Dusek, James Davis, and Alex Pang. Automated rip current detection with region based convolutional neural networks. *Coastal Engineering*, 166:103859, 2021. **2, 3**
- [19] Aleksandra Dudkowska, Aleksandra Boruń, Jakub Malicki, Jan Schönhofer, and Gabriela Gic-Grusza. Rip currents in the non-tidal surf zone with sandbars: numerical analysis versus field measurements. *Oceanologia*, 62(3):291–308, 2020. **2**
- [20] Gregory Dusek, Debra Hernandez, Mark Willis, Jenna A Brown, Joseph W Long, Dwayne E Porter, and Tiffany C Vance. Webcat: Piloting the development of a web camera coastal observing network for diverse applications. *Frontiers in Marine Science*, 6:353, 2019. **2**
- [21] G Dusek and H Seim. A probabilistic rip current forecast model. *Journal of Coastal Research*, 29(4):909–925, 2013. **2**
- [22] G Dusek and H Seim. Rip current intensity estimates from lifeguard observations. *Journal of coastal research*, 29(3):505–518, 2013. **2**
- [23] B. Dwyer, J. Nelson, J. Solawetz, et al. Roboflow (version 1.0) [software]. <https://roboflow.com>, 2022. Computer vision. **5**
- [24] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. *arXiv preprint arXiv:2211.07636*, 2022. **5**
- [25] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. **5**
- [26] K Todd Holland, Robert A Holman, Thomas C Lippmann, John Stanley, and Nathaniel Plant. Practical use of video imagery in nearshore oceanographic field studies. *IEEE Journal of oceanic engineering*, 22(1):81–92, 1997. **2**
- [27] Rob A Holman and John Stanley. The history and technical capabilities of Argus. *Coastal engineering*, 54(6-7):477–491, 2007. **2**
- [28] Xiao Hong, Yao Zhang, Bin Wang, Shuihua Zhou, Shengbin Yu, and Juan Zhang. Numerical study of rip currents interlaced with multichannel sandbars. *Natural hazards*, 108(1):593–605, 2021. **2**

- [29] Douglas L Inman, James A Zampol, Thomas E White, Daniel M Hanes, Walton B Waldorf, and Kim A Kastens. Field measurements of sand motion in the surf zone. *In Coastal Engineering*, pages 1215–1234, 1980. [2](#)
- [30] Glenn Jocher. YOLOv5 by Ultralytics, 5 2020. [5](#)
- [31] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. YOLO by Ultralytics, 1 2023. [2](#), [5](#)
- [32] Hyun Dong Kim and Kyu-Han Kim. Analysis of rip current characteristics using dye tracking method. *Atmosphere*, 12(6):719, 2021. [2](#)
- [33] SB Leatherman and SP Leatherman. Techniques for detecting and measuring rip currents. *Int J Earth Sci Geophys*, 3:014, 2017. [1](#), [2](#)
- [34] Chuyi Li, Lulu Li, Yifei Geng, Hongliang Jiang, Meng Cheng, Bo Zhang, Zaidan Ke, Xiaoming Xu, and Xiangxiang Chu. Yolov6 v3. 0: A full-scale reloading. *arXiv preprint arXiv:2301.05586*, 2023. [5](#)
- [35] Chuyi Li, Lulu Li, Hongliang Jiang, Kaiheng Weng, Yifei Geng, Liang Li, Zaidan Ke, Qingyuan Li, Meng Cheng, Weiqiang Nie, et al. Yolov6: A single-stage object detection framework for industrial applications. *arXiv preprint arXiv:2209.02976*, 2022. [5](#)
- [36] Feng Li, Hao Zhang, Shilong Liu, Lei Zhang, Lionel M Ni, Heung-Yeung Shum, et al. Mask DINO: Towards a unified transformer-based framework for object detection and segmentation. *arXiv preprint arXiv:2206.02777*, 2022. [5](#)
- [37] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. [3](#)
- [38] Tom C Lippmann and Rob A Holman. Quantification of sand bar morphology: A video technique based on wave dissipation. *Journal of Geophysical Research: Oceans*, 94(C1):995–1011, 1989. [2](#)
- [39] Yuli Liu and Chin H Wu. Lifeguarding operational camera kiosk system (locks) for flash rip warning: Development and application. *Coastal engineering*, 152:103537, 2019. [2](#)
- [40] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12009–12019, 2022. [5](#)
- [41] James B Lushine. A study of rip current drownings and related weather factors. *National weather digest*, 16(3):13–19, 1991. [1](#)
- [42] Corey Maryan, Md Tamjidul Hoque, Christopher Michael, Elias Ioup, and Mahdi Abdelguerfi. Machine learning applications in detecting rip channels from images. *Applied Soft Computing*, 78:84–93, 2019. [2](#), [3](#)
- [43] Sean P McGill and Jean T Ellis. Rip current and channel detection using surfcams and optical flow. *Shore & Beach*, 90(1):50, 2022. [2](#)
- [44] Issei Mori, Akila de Silva, Gregory Dusek, James Davis, and Alex Pang. Flow-based rip current detection and visualization. *IEEE Access*, 10:6483–6495, 2022. [2](#)
- [45] Arthur Mouragues, Philippe Bonneton, Bruno Castelle, and Kévin Martins. Headland rip modelling at a natural beach under high-energy wave conditions. *Journal of Marine Science and Engineering*, 9(11):1161, 2021. [2](#)
- [46] National Oceanic and Atmospheric Administration (NOAA). What is a rip current? <https://oceanservice.noaa.gov/facts/ripcurrent.html>. Accessed: March, 2023. [1](#)
- [47] Varjola Nelko and R Dalrymple. Rip current prediction in ocean city Maryland. In *Rip Currents: Beach Safety, Physical Oceanography and Wave Modeling*, pages 45–58. CRC Press, 2011. [2](#)
- [48] Shweta Philip and Alex Pang. Detecting and Visualizing Rip Current Using Optical Flow. In Enrico Bertini, Niklas Elmqvist, and Thomas Wischgoll, editors, *EuroVis 2016 - Short Papers*. The Eurographics Association, 2016. [2](#)
- [49] Sebastian Pitman, Shari L Gallop, Ivan D Haigh, Sasan Mahmoodi, Gerd Masselink, and Roshanka Ranasinghe. Synthetic imagery for the automated detection of rip currents. *Journal of coastal research*, 75(10075):912–916, 2016. [2](#)
- [50] S Prodger. Argus observations of rip current variability along a macro-tidal beach. *Unpublished Master thesis, Plymouth University*, 2012. [2](#)
- [51] NN Pujianiki, ING Antara, IGRM Temaja, IGDY Partama, and T Osawa. Application of UAV in rip current investigations. *International Journal on Advanced Science Engineering Information Technology*, 10(6):2337–2343, 2020. [2](#)
- [52] Ashraf Haroon Rashid, Imran Razzak, M Tanveer, and Michael Hobbs. Reducing rip current drowning: An improved residual based lightweight deep architecture for rip detection. *ISA transactions*, 132:199–207, 2023. [3](#)
- [53] Ashraf Haroon Rashid, Imran Razzak, Muhammad Tanveer, and Antonio Robles-Kelly. Ripnet: A lightweight one-class deep neural network for the identification of rip currents. In *Neural Information Processing: 27th International Conference, ICONIP 2020, Bangkok, Thailand, November 18–22, 2020, Proceedings, Part V 27*, pages 172–179. Springer, 2020. [3](#)
- [54] Ashraf Haroon Rashid, Imran Razzak, Muhammad Tanveer, and Antonio Robles-Kelly. RipDet: A fast and lightweight deep neural network for rip currents detection. In *2021 international joint conference on neural networks (IJCNN)*, pages 1–6. IEEE, 2021. [3](#)
- [55] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. [5](#)
- [56] Joseph Redmon and Ali Farhadi. Yolo9000: Better, faster, stronger, 2016. [5](#)
- [57] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. [5](#)
- [58] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [3](#)
- [59] Dano Roelvink, AJHM Reniers, A Van Dongeren, J Van Thiel de Vries, Jamie Lescinski, and Robert McCall. Xbeach

- model description and manual. *Unesco-IHE Institute for Water Education, Deltares and Delft University of Technology. Report June*, 21:2010, 2010. [2](#)
- [60] L.E. Sembiring. *Rip Current Prediction System for Swimmer Safety: Towards operational forecasting using a process based model and nearshore bathymetry from video*. PhD thesis, TU Delft, 2015. [2](#)
- [61] AD Short and CL Hogan. Rip currents and beach hazards: their impact on public safety and implications for coastal management. *Journal of Coastal Research*, pages 197–209, 1994. [2](#)
- [62] Choule J Sonu. Field observation of nearshore circulation and meandering currents. *Journal of Geophysical Research*, 77(18):3232–3247, 1972. [2](#)
- [63] A Valipour and AA Bidokhti. An analytical model for the prediction of rip spacing in intermediate beaches. *Journal of Earth system Science*, 127:1–11, 2018. [2](#)
- [64] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. *arXiv preprint arXiv:2207.02696*, 2022. [5](#)
- [65] Hong Wang, Shouxian Zhu, Xunqiang Li, Wenjing Zhang, and Yu Nie. Numerical simulations of rip currents off arc-shaped coastlines. *Acta Oceanologica Sinica*, 37:21–30, 2018. [2](#)
- [66] Shaoru Wang, Yongchao Gong, Junliang Xing, Lichao Huang, Chang Huang, and Weiming Hu. Rdsnet: A new deep architecture for reciprocal object detection and instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12208–12215, 2020. [5](#)
- [67] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhi-liang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: BEiT pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. [5](#)
- [68] Yixuan Wei, Han Hu, Zhenda Xie, Zheng Zhang, Yue Cao, Jianmin Bao, Dong Chen, and Baining Guo. Contrastive learning rivals masked image modeling in fine-tuning via feature distillation. *arXiv preprint arXiv:2205.14141*, 2022. [5](#)
- [69] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12193–12202, 2020. [5](#)
- [70] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part I 13*, pages 818–833. Springer, 2014. [7](#)
- [71] Yao Zhang, Wanru Huang, Xunan Liu, Chi Zhang, Guodong Xu, and Bin Wang. Rip current hazard at coastal recreational beaches in china. *Ocean & Coastal Management*, 210:105734, 2021. [2](#)
- [72] Daoheng Zhu, Rui Qi, Pengpeng Hu, Qianxin Su, Xue Qin, and Zhiqiang Li. Yolo-rip: A modified lightweight network for rip currents detection. *Frontiers in Marine Science*, 9:930478, 2022. [2](#), [3](#), [4](#)