

RipViz: Finding Rip Currents by Learning Pathline Behavior

Akila de Silva, Mona Zhao, Donald Stewart, Fahim Hasan Khan, Gregory Dusek, James Davis, and Alex Pang

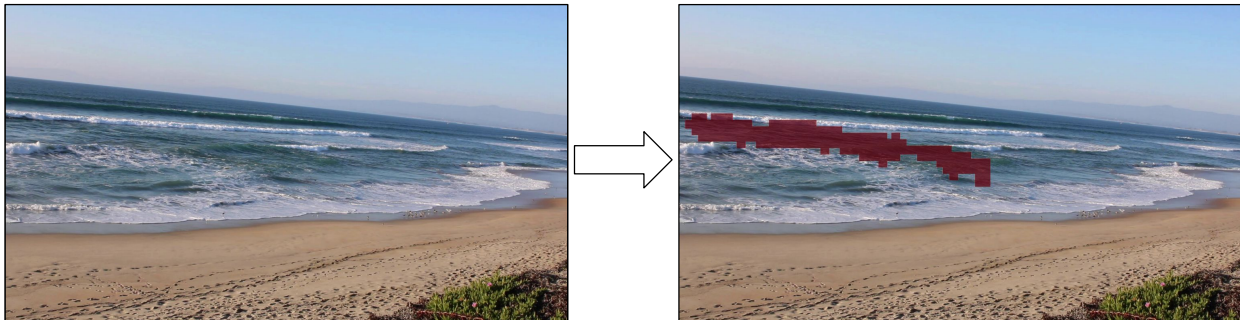


Fig. 1: **Rip currents are deadly but remain invisible to many:** We propose a feature detection method, RipViz, to make these invisible rip currents visible. RipViz highlights locations of rip currents as a red region. This region is determined by finding seed points that produce pathline sequences deviating from normal ocean flow. Non-experts and experts alike can use RipViz to visualize rip currents.

Abstract—We present a hybrid machine learning and flow analysis feature detection method, RipViz, to extract rip currents from stationary videos. Rip currents are dangerous strong currents that can drag beachgoers out to sea. Most people are either unaware of them or do not know what they look like. In some instances, even trained personnel such as lifeguards have difficulty identifying them. RipViz produces a simple, easy to understand visualization of rip location overlaid on the source video. With RipViz, we first obtain an unsteady 2D vector field from the stationary video using optical flow. Movement at each pixel is analyzed over time. At each seed point, sequences of short pathlines, rather than a single long pathline, are traced across the frames of the video to better capture the quasi-periodic flow behavior of wave activity. Because of the motion on the beach, the surf zone, and the surrounding areas, these pathlines may still appear very cluttered and incomprehensible. Furthermore, lay audiences are not familiar with pathlines and may not know how to interpret them. To address this, we treat rip currents as a flow anomaly in an otherwise normal flow. To learn about the normal flow behavior, we train an LSTM autoencoder with pathline sequences from normal ocean, foreground, and background movements. During test time, we use the trained LSTM autoencoder to detect anomalous pathlines (i.e., those in the rip zone). The origination points of such anomalous pathlines, over the course of the video, are then presented as points within the rip zone. RipViz is fully automated and does not require user input. Feedback from domain expert suggests that RipViz has the potential for wider use.

Index Terms—Flow visualization, 2D unsteady flow fields, pathlines, LSTM autoencoders, anomaly detection

1 INTRODUCTION

Rip currents are powerful, narrow channels of fast-moving water flowing towards the sea from the nearshore [3, 6, 20, 27]. The speed of seaward rips can be very strong, reaching two meters per second, faster than an Olympic swimmer. They are a dangerous beach hazard that most people do not recognize. As a result, there are thousands of drownings each year due to rip currents globally [10, 25]. The goal of this work is to improve public safety by helping the beachgoers *see* the rip currents when they are present.

The mechanism for rip currents is an increase in the mean water level, referred to as setup, which occurs when waves break against the shore. This setup can vary along a shoreline depending on the amount of water or height of breaking waves. Rip currents form as water tends to flow from regions of high setup (larger waves) to regions of lower setup (smaller waves), where currents converge to form a seaward flowing rip.

Detecting rip currents with machine learning (ML) is challenging

because there are different types of rips, each with a different appearance. The three major factors that lead to different types of rips are the shape of the shoreline, the bathymetry, and hydrodynamic factors (e.g., wave height and direction, tides). The combination of these lead to rip currents with different visual signatures. Rips may also either be transient or persistent in space and time. Detecting rip currents pose unique challenges compared to detecting other objects such as cars, or people, etc. Rips are amorphous without a well defined shape or boundary, and are ephemeral without well defined temporal bounds.

Using an object detector, such as those included in a recent survey [22], would require a substantial training dataset for each type of rip current. To date, there are only two publicly available training datasets for rip detection: individually labeled frames [11] and time averaged images [29], both of the same type of rip current. There are no existing training dataset for other types of rip.

By nature, rips always eventually flow seaward regardless of their visual appearance. Therefore, we propose a novel hybrid approach that incorporates flow analysis with ML for rip detection, rather than relying on image colors. We also propose encoding rips as locations with anomalous flow behavior, transforming the detection task into differentiating normal from anomalous behavior. This greatly simplifies the collection of training data since only a single dataset is needed.

There are two existing approaches to detecting rip currents from stationary videos. The first approach uses the appearance of rip currents

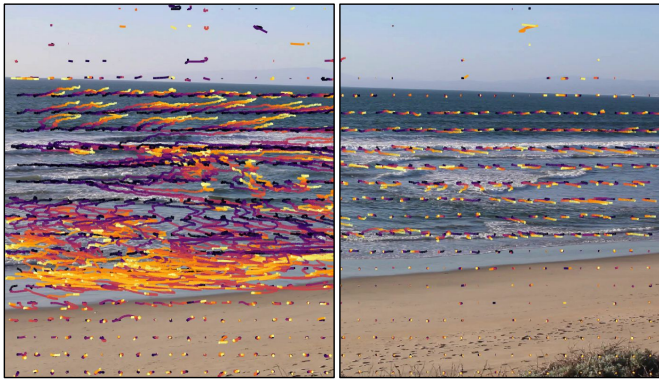


Fig. 2: **Long pathlines get cluttered making it difficult for ML to learn:** Left pane shows longer pathlines integrated over the entire length of the video. Shorter pathlines, integrated over 900 time steps are shown in the right pane. Longer pathlines are much more cluttered and noisy, making it difficult for ML algorithms to learn its behavior. Notice the relatively large amount of noisy pathlines in the sky and the beach of the left pane. The pathlines are colored by age with yellow representing most recent.

to detect them. This includes human experts analyzing time-averaged (Timex) video or running automated object detectors [11, 29, 35, 36] to detect rip currents. The second approach relies on flow analysis, where the flow behavior of ocean waves is analyzed to detect rip currents. Direction-based clustering of flow vectors [30, 34] and timelines [30] placed parallel to the beach are used in this approach. While able to detect weaker rip currents or those where the appearance is not obvious, timelines require user input to specify their initial placement.

In this paper, we introduce RipViz, a fully automated, hybrid of deep learning and flow analysis, feature detection method to find rip currents, as shown in Figure 1. We first obtain the time varying flow field from a stationary video using optical flow. We use pathline sequences to capture the flow behavior. One could simply seed pathlines at every point and trace each of them for the duration of the video. However, due to the quasi-periodic nature of wave activity, this simple approach produces too much clutter to be of much use as shown in Figure 2. Furthermore, pathlines integrated over the full length of the video accumulate more error, especially in noisy real world datasets. Instead, we generate *sequence of shorter pathlines* for each seed point. By staggering the initiation of pathlines, our expectation is that pathline sequences outside the rip zone will behave differently. For example, pathlines seeded in the surf zone where the waves are breaking will have large variations in their trajectories. Pathlines seeded further out to sea, on the beach, and sky would be fairly still. Pathlines seeded within the rip zone will have less variations in their trajectories yet will be different from the stationary ones.

In RipViz, we frame detecting rip currents as a flow anomaly detection problem. An LSTM autoencoder with a custom weighted loss function is used to learn the spatiotemporal features of pathline sequences for normal ocean flow (i.e. not rip currents). The trained LSTM autoencoder can predict anomalous pathline sequences (i.e. rip currents) during test time. The origination points of anomalous pathlines are identified and highlighted as a means of visualizing the rip zone. Our target users are general public who are not familiar with rip current dynamics nor visual analytic systems. Hence, our design for the visualization output is to make it as simple and unambiguous as possible.

The main contribution of this paper is:

- A hybrid feature detection method that combines machine learning and flow analysis techniques to automatically find and visualize dangerous rip currents.

In order to realize this, the following innovations are necessary:

- For flow fields with a quasi-periodic behavior, such as ocean flow, working with a *sequence of shorter pathlines* is better than a single long pathline.
- A weighted binary cross entropy, unlike the unweighted binary cross entropy used in previous works, is more effective when learning from sparsely distributed pathlines generated from ocean flow.

2 RELATED WORK

Streamline Selection: Streamline selection and seed placement are well studied in flow visualization. Sane et al. [39] provide a survey of the body of work over the last two decades. These works share overlapping goals with research on streamline clustering [5], and flow simplification [23]. Each aims to produce an uncluttered presentation of the flow field while still capturing the essential flow features and behaviors. For example, Marchesin et al. [28] proposed dynamically selecting a set of streamlines that leads to intelligible and uncluttered streamline selection. Ma et al. [26] used an importance-driven approach to view-dependent streamline selection that guarantees coherent streamline update when the view changes gradually. Yu et al. [44] proposed hierarchical streamline bundles by producing streamlines near critical points without enforcing dense seeding throughout the volume. They grouped the streamlines to form a hierarchy from which they extracted streamline bundles at different levels of detail. Tao et al. [41] proposed two interrelated channels between candidate streamlines and sample viewpoints. They selected streamlines by taking into account their contribution to all sample viewpoints. However, the methods mentioned above use handcrafted features to define the feature representation of streamlines. Handcrafting features to represent pathlines in complex unsteady flow fields, such as ocean flow fields, is a challenging task. In contrast, the method presented in this paper learns complex feature representations without the need for handcrafted features.

Deep Learning for Flow Visualization: In recent years, the visualization community has worked with ML in two ways: visualization to understand the ML model, and use of ML in visualization tasks. On the latter, particularly for flow visualization tasks, Berenjkoub et al. [2] used U-net, a deep learning neural network, to identify vortex boundaries. Kim and Günther [19] used a neural network to extract a steady reference frame from an unsteady vector field. Han et al. [17] used an autoencoder-based deep learning model, FlowNet, to learn feature representations of streamlines in 3D steady flow fields which are then used to cluster the streamlines. They used longer streamlines that were integrated through the entire extent of the data. Furthermore, they used one streamline per seed point. The work presented in this paper uses a sequence of pathlines per seed point to learn the flow behavior from noisy unsteady 2D flow fields. This work uses an LSTM autoencoder to detect anomalous pathline sequences from noisy unsteady 2D flow fields.

Rip Current Detection: Traditional rip current detection generally involves in-situ instrumentation such as GPS-equipped drifters and current meters [21, 27]. Remote sensing of rip currents is also possible with aerial imaging of the spread of fluorescein dye in the rip zone captured by drones, and the use of marine radar [16, 21]. These require expensive equipment or a team of observers which makes them impractical for rip current monitoring or detection purposes. However, with simple optical video capture such as from surf webcams, it is possible to detect certain rip currents using a time-space (timestack) display or a simulated long time exposure images (timex) display [18]. Timex images reveal rip locations as darker regions in the image that correspond to deeper channels in the bathymetry where water may flow seaward. In contrast, incoming water associated with the breaking waves in the surf zone appear much brighter in timex images. This type of rip is referred to as bathymetry controlled rips and are characterized by a quiet region in the rip channel that is flanked by breaking waves on either side. Interpreting timex images require some domain expertise. Maryan et al. [29] used a Viola-Jones framework to train their model on timex images, and indicate the location of the rips via bounding boxes. Likewise, Rashid et al. [35, 36] also used timex images but utilized

a modified version of the Tiny-Yolo V3 architecture. Similarly, Ellis and McGill [12] used timex images in conjunction with environmental information such as tides, wave height, and period to cluster offshore movements to rip currents. However, their method produced false positives in situations where non-water objects such as surfers, paddle boarders, etc., are also moving offshore. Rather than working with timex images, de Silva et al. [11] trained a Faster R-CNN model with an accumulation buffer to detect bathymetry rips using individual frames of the video. All of these methods detect bathymetry rips based on the appearance of the sea state. However, in instances where appearance is different or weak, these methods fail to detect rip currents.

An alternative approach is to detect rips based on the observed behavior. Philip and Pang [34] obtained an unsteady flow field from the video and hypothesized that the rip current is directly opposite the dominant flow in the vector field due to the incoming wave motion. They grouped vectors based on the direction and magnitude to visualize the rip current. Mori et al. [30] used a similar approach to group vectors and improved the visualization of the rip current by mapping direction to color and magnitude to hue. In the same work, Mori et al. [30] also used timelines to visualize rip currents. They placed timelines parallel to the beach and observed its shape as it gets dragged by the rip current. The work presented in this paper also uses optical video of the sea state to detect rip currents. The underlying approach is also based on the flow behavior, and hence not constrained to detecting bathymetry rips. However, the presented methodology in this paper is the first to propose combining ML and flow analysis to detect different types of rip currents. Because detection is based on treating flow behavior in rip currents as anomalous, the task of collecting and labeling training data for an ML model is unified and simplified. In contrast, a standard ML rip detection model would require a training data set for each type of rip current – a costly and time consuming process.

Autoencoders: Autoencoders are a type of neural network that can learn object representations without any supervision [7, 45]. Autoencoders are trained with a loss function that compares the input and reconstructed output by using a reconstruction error. Long short term memory (LSTM) autoencoders are a specialized type of autoencoder that can learn from sequential data. These types of autoencoders are equipped with LSTM layers that can learn how data is temporally related. Furthermore, autoencoders are also used as anomaly detectors [38]. For anomaly detection, autoencoders are first trained with *normal* data. When presented with *anomalous* data, the trained autoencoder will produce a high reconstruction error. We exploit this property to detect anomalous flow behavior.

The flow visualization community has used autoencoders to find feature representations of objects that then can be used for clustering. In their work, Han et al. [17] proposed to learn feature representation of streamlines and stream surfaces of 3D steady flow fields by using an autoencoder with a binary cross-entropy function. They then further reduced the dimensionality of the features and used those to generate clusters. However, their method does not learn how pathlines are temporally related. Additionally, their unweighted binary cross entropy loss function does not account for sparsely distributed pathlines. In our work, we use an LSTM autoencoder with a custom weighted binary cross-entropy loss function, to learn spatiotemporal behavior of sparsely distributed pathline sequences from ocean scenes.

3 RIPVIZ

Identifying rip currents in complex and chaotic ocean flow is challenging. We first reconstructed a 2D unsteady flow field using optical flow from ocean videos. Our method captured the ocean flow characteristics by using short pathline sequences which are then represented as stacks of binary images. We used an LSTM autoencoder to learn spatiotemporal features of these pathline sequences. We trained our model with pathline sequences of normal ocean flow using a custom weighted binary cross-entropy function, suited for learning pathline sequences of ocean flow. We then used this trained LSTM autoencoder to identify pathline sequences of abnormal ocean flow or rip currents. We visualized the rip currents by projecting the seed points of these abnormal pathline sequences back onto the video frame and growing a

transparent red region around these points.

3.1 Flow Field Reconstruction

An unsteady 2D flow field is obtained from the stationary video using optical flow. Many optical flow algorithms use the relative motion of neighboring pixels between consecutive frames in the video to calculate the local flow. We use Lucas-Kanade [24] sparse optical flow function in the OpenCV library [14] to trace pathlines. We verified each integration step by comparing the forward and backward integration of the flow field.

3.2 Sequence of Pathlines

Each pathline is represented by a 1D vector, $\mathbf{p} = \{x_1, y_1, \dots, x_n, y_n\}$, where (x_i, y_i) is a point in the coordinate system of the video frame and n is the length of the pathline. For the LSTM autoencoder to learn about the pathlines, we transformed these pathlines into their own 2D pathline-centric coordinate system. To do this, each pathline is represented by an $L \times L$ binary image \mathbf{I} , where each point in \mathbf{p} is translated to center the seed point in the binary image. For longer pathlines that extend beyond $L \times L$, we increased the size of the binary image to accommodate these longer pathlines. We then resized all of these larger binary images down to $L \times L$. The reason why all pathlines were not simply centered then resized to $L \times L$ is that we need to differentiate pathlines that barely moved from their initial seed position versus those that actually travelled beyond $L \times L$. Also note that by centering, the 2D representation of the pathlines are now location agnostic. This combination allows us to compare pathlines from different parts of the video to identify those with similar behavior. As noted earlier, tracing pathlines over the entire length of a video of quasi-periodic motion derived from noisy optical flow calculations result in unusable cluttered flow representations. Instead, we reseeded pathlines at the same seed point over regular time intervals to generate pathline series of length S . We represented each pathline sequence as a stack of binary images of size $S \times L \times L$ associated with each seed point.

3.3 LSTM Autoencoder

We used an LSTM autoencoder to learn from the pathline sequences. The LSTM autoencoder consisted of two components, an encoder, and a decoder, as shown in Figure 3. The encoder learns the spatiotemporal features of the pathline sequence. The decoder reconstructs the pathline sequence by using the learned spatiotemporal feature representation.

We specify the input layer of the LSTM autoencoder, to expect pathline sequences of shape $S \times L \times L$. The first and second layers of the encoder are time-distributed 2D convolutional layers. These layers consist of 128 and 64 convolutional filters, respectively. These layers process each pathline of the sequence separately and learn the spatial features of each pathline. The third and fourth layers are 2D convolutional LSTM layers. Each layer consists of 64 and 32 convolutional filters, respectively. These convolutional LSTM layers process the sequence of pathlines together and learn the temporal features of the pathline sequences. The first layer of the decoder is a convolutional LSTM layer with 64 convolutional filters. The second and third layers of the decoder are time-distributed deconvolutional layers with 64, and 128 respectively. The last layer of the LSTM autoencoder is a time-distributed convolutional layer with 1 convolutional filter. Each layer, except for the input and the output layers, is followed by a normalization layer [1]. The input volume has no padding; therefore, we set the stride of all layers to 1. The autoencoder outputs a sequence of pathlines of shape $S \times L \times L$. For the hidden layers, we use the ReLU activation [31]. For the output layer, we use the sigmoid activation.

In comparison to recent autoencoder based stream line selection methods, our method not only learns the spatial features of each pathline but also learns how each pathline is temporally related to other pathlines in the same pathline sequence. Learning these spatio-temporal features allows our method to better learn about the quasi-periodic nature of the chaotic ocean flow.

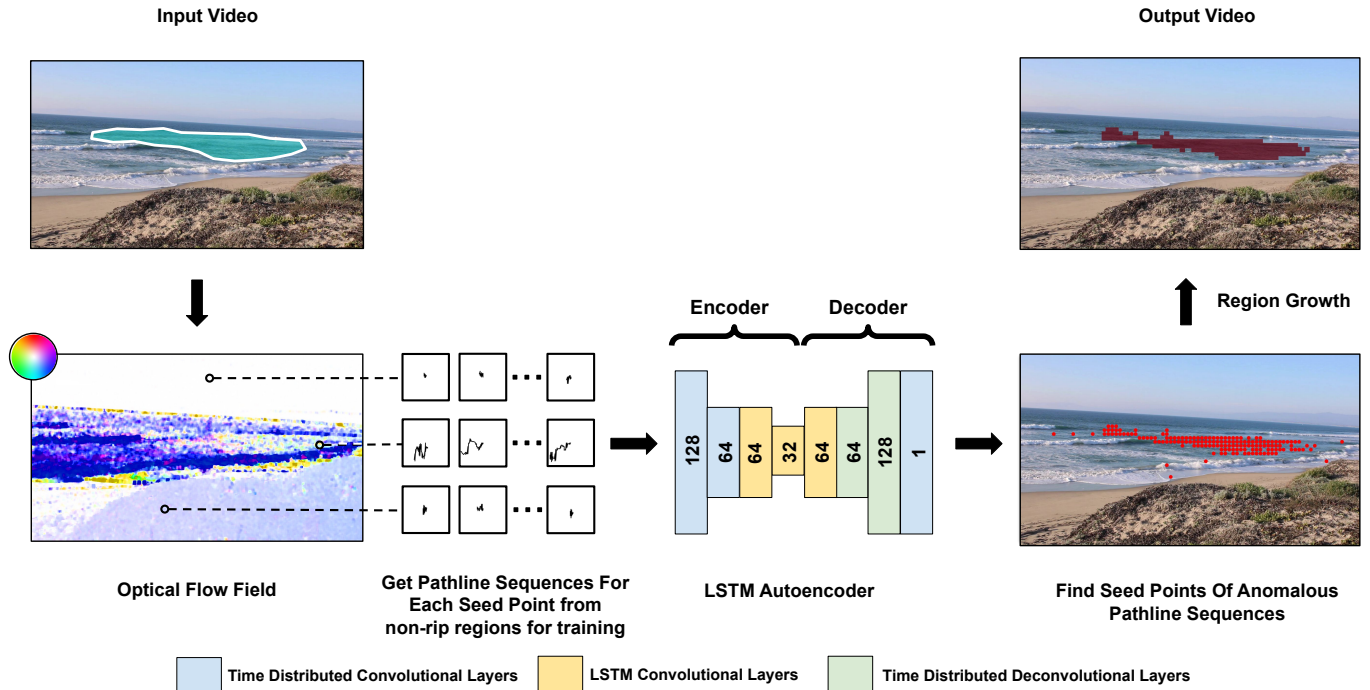


Fig. 3: **RipViz pipeline:** First, an optical flow field was generated from the input video (magnitude and direction of velocities are mapped to value and hue, respectively.) Then pathline sequences were generated for each seed point in non-rip regions (only three seed points are shown) where a new pathline is seeded at every frame. Finally, an LSTM autoencoder was trained with these pathline sequences generated from non-rip regions. For rip detection, the trained LSTM autoencoder is used to detect anomalous pathline sequences by regularly seeding the video frame. We then applied a region growing algorithm to the anomalous points to find the anomalous region, while filtering out singleton seeds.

3.4 Weighted Loss Function

We trained the LSTM autoencoder by minimizing the difference between each training sample against the corresponding model prediction. The loss function used to compute this difference is a weighted binary cross-entropy loss function. Since each training sample is a pathline sequence represented as a stack of binary images of size $S \times L \times L$, we treat each sample as a binary volume where p_i is 1 when the pathline crosses that voxel and 0 otherwise. \hat{p}_i is the corresponding value from the predicted volume. The loss is calculated over all voxels (i.e., $N = S \times L \times L$).

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N [w_1 \cdot p_i \log \hat{p}_i + w_0 \cdot (1 - p_i) \log(1 - \hat{p}_i)] \quad (1)$$

w_1 is weight assigned to voxels when $p_i = 1$. w_0 is weight assigned to voxels when $p_i = 0$. In contrast, previous autoencoder based neural network architectures used in flow visualization tasks used an unweighted version of the same loss function (i.e., $w_1 = w_0 = 1$). We found that a weighted binary cross-entropy function is better suited for our application domain for reasons discussed in Section 4.2.

3.5 Detecting and Visualizing Rip Currents

During inference/test time, we calculated the reconstruction error for each pathline sequence by using the loss function defined earlier. If the error is larger than threshold T , we labeled those pathline sequences as anomalous (i.e., rip currents). Otherwise, we labeled the sequences as belonging to normal flow. We discuss how T was selected in Section 4.2. Once the anomalous pathline sequences were found, the corresponding seed points were connected by using a region growth algorithm to generate a region. Isolated seeds without neighbors are discarded. To visualize the rip zone, we projected this region back onto the frames in the video as shown in Figure 1.

3.6 Network Training

We implemented our neural network in Tensorflow by using the Keras application programming interface (API) [9]. We trained the network using an NVIDIA Tesla A100 graphical processing unit (GPU). In the training process, we initialized parameters in all layers of the neural network using a normal distribution $\mathcal{N}(\mu, \sigma^2)$, where mean $\mu = 0$ and variance $\sigma^2 = 0.01$. We applied the Adam optimizer [38] to update the parameters with a learning rate of 10^{-6} . We used the minibatch size of 10 and trained our model with 100 epochs.

4 RESULTS AND DISCUSSION

We describe our dataset, provide an analysis of the components and parameters used in RipViz, and present comparisons with other methods.

4.1 Dataset

Our dataset consists of 55 stationary videos. Each video is 1.5 to 4 minutes long and of size 1920×1080 pixels. We collected the data from Salinas, Marina, and Davenport beaches in California in winter of 2022. We used either a tripod-mounted Canon EOS Rebel T7 DSLR camera or a Samsung Android phone to acquire the data. We chose 4 videos to generate 16000 non-rip pathline sequences for training our model by using the seeding strategy discussed in Section 4.2. Our validation dataset consists of 12000 pathlines generated from three videos. The remaining videos were used for testing. Note that our training data does not include examples of each of the many types of rip currents, only examples of normal ocean flow. This greatly simplifies collection of training data. Each video was labeled under the guidance of the rip current expert. Pathlines seeded in the rip current area were labeled as “rip” and others as “non-rip”.

The framing of these videos, which includes the distance of camera to the water and the focal length or zoom factor of the lens, is meant to be representative of surf webcams (e.g. surfline.com) that might be

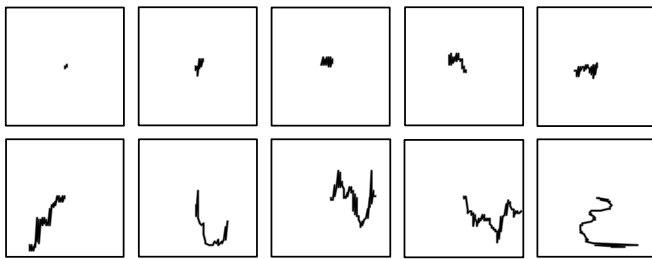


Fig. 4: **Sparsely distributed pathlines in binary image:** This figure shows a collection of pathlines from our data. Notice the imbalance in 0s (white) and 1s (black) in the binary images. The weighted loss function, as described in Equation 1 allows the autoencoder to learn with imbalanced data.

useful for rip current detection. Parameters such as integration length n and binary image size L are based on such framing. Sensitivity of these parameters on different framings are discussed in Section 4.2.

Note that we assume a stable video source and hence did not perform any video stabilization. Wind can cause slight movements of the camera, but also movements of grass, clouds, etc. Also, video processing seldom work with raw video but rather on compressed video. Deriving the optical flow field on compressed video may also introduce some motion artifacts especially in highly compressed regions such as the sky or empty sandy beaches. RipViz handles these type of motion artifacts better than existing methods for our dataset.

4.2 Analysis of RipViz components

The RipViz method contains two important changes from prior methods: the use of a weighted loss function, and the use of a sequences of pathlines rather than a single pathline. Without these changes we found the ML fails to learn ocean flow. In addition, the method has parameters like detection threshold, T , and binary image size L which are likely dependent on our specific application domain. In this section we analyze each of these factors, showing that our changes are necessary, and providing the method by which we determined parameters.

Threshold T : Threshold T is used to filter anomalous pathlines (i.e., rip currents) based on the reconstruction error as discussed in Section 3.5. In order to find the optimal threshold T we calculated F_1 score at varying threshold values in a subset of our data. F_1 score is defined as,

$$F_1 = \frac{2}{\frac{1}{\text{recall}} + \frac{1}{\text{precision}}} = \frac{2}{\frac{FN_s+TP_s}{TP_s} + \frac{FP_s+TP_s}{TP_s}} \quad (2)$$

If a detected pathline falls within the expert annotated boundary of the rip current, then it's counted as a TPs (true positive). Otherwise, it's considered an FP (false positive). Suppose a pathline originating within the rip current is not detected, then it's an FN (false negative). The range of the F_1 score is between 0 and 1. If most pathlines fall within the rip current boundary, the score will be closer to 1, otherwise closer to 0. We found that threshold, $T = 1.0$ produced the highest F_1 score, and was used as T for the experiments in this paper.

Weights of the loss function w_0 and w_1 : We reconstructed the flow field from videos using optical flow. As discussed in Section 3.2, all pathlines are represented in a binary image I . We observed that some pathlines are short and do not extend far from the seed point, while other pathlines are longer and travel farther from the seed point, as shown in Figure 4. We needed the model to learn the distribution of these shorter and longer pathlines because it allowed the model to learn the *normal* flow behavior in an ocean scene. However, this leads to an imbalance in the number of voxels with 1s and 0s, especially for shorter pathlines. For the deep learning model to learn effectively about these small features/pathlines, we needed to increase the weight of the loss function when the target probability (p_i) is 1. The loss function, as described in Section 3.4, penalizes the model more for

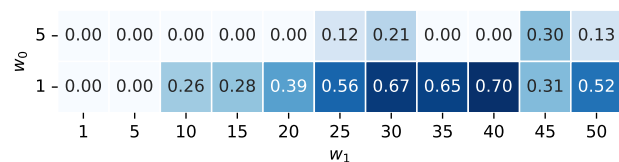


Fig. 5: **F_1 Scores for different combinations of w_0 and w_1 :** Notice that the method does not converge when w_0 and w_1 are balanced. The highest F_1 score of 0.70 is achieved when w_0 is set to 1 and w_1 is set to 40.

making mistakes when predicting pixels with the target probability 1 in the training process. Using the weighted binary cross-entropy function made the model better learn the distribution of shorter from longer pathlines. In contrast, in FlowNet [17], the streamlines were well spread out across their input volumes, making the use of a weighted loss function unnecessary.

To find the optimal w_0 and w_1 for our application domain, we randomly choose a small subset of pathline sequences, some short and some long, to train models at different combinations of w_0 and w_1 . We trained each model for 100 epochs, using the same training conditions specified in Section 3.6. For each trained model, we observed how many voxels were correctly recreated. If the predicted probability for a voxel with target probability of 1 is greater than 0.7, we marked that voxel as TP; otherwise, it is considered to be FN. If the model predicted pathline voxels in areas with no pathline, we marked that as FP. We calculated the F_1 score for each model with different weight combinations using equation 2. We observed that non-uniform weighting is necessary, with the highest F_1 score generated when $w_0 = 1$ and $w_1 = 40$ as shown in Figure 5. Note that using unweighted cross-entropy, as in previous work, is the equivalent of setting $w_0 = w_1 = 1$. In this condition we found that the model completely fails to learn ocean flow.

In some instances where the model did not converge, TPs were 0, which resulted in an undefined F_1 score. In such instances, we followed the default reporting convention of Scikit-learn python library [33] and reported those F_1 s as 0.00.

Pathline Length n and Sequence Length S : In RipViz, we use pathline sequences to learn flow behavior. In order to find the optimal pathline length, n , and sequence length, S , for our application domain, we trained multiple models while changing n and S , and observed the F_1 score as shown in Figure 6. We noticed two patterns from this experiment. First, increasing the pathline length resulted in a higher accuracy up to an optimal length; beyond this, the accuracy decreased. Second, using a sequence of pathlines instead of a single pathline per seed point produces more accurate results. We found that $n = 900$ and $S = 100$ were optimal for our data set.

Size of binary image $I = L \times L$: We represented each pathline in a binary image as discussed in 3.2. We trained several models with our training dataset where binary images were of different sizes as shown in Figure 7. We compared each model by calculating the F_1 score, similar to how F_1 score was calculated for tuning Threshold T . Our experiments indicated that small binary images ($L \times L = 5 \times 5$) cannot sufficiently capture the variations of pathline movements, resulting in a lower F_1 score. We also found that larger binary images ($L \times L = 128 \times 128$) also resulted in a slightly lower F_1 score. Furthermore, larger binary images resulted in a high memory usage and a longer training time. We found that $L \times L = 64 \times 64$ is sufficient at capturing the variations of pathlines resulting in a high F_1 score. In our experiments we used $L \times L = 64 \times 64$ as the size of the binary image I .

Seeding strategy: Various seeding strategies attempt to optimize different goals such as aesthetics, clutter reduction, capturing flow features, etc [39]. In this paper, the primary consideration for pathline seeding take into account where one can have the highest return in informational value regarding the water movement. Because most of

Pathline Length n	1	10	20	40	60	80	100	120	140
1	-0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
10	-0.31	0.32	0.38	0.30	0.34	0.39	0.35	0.33	0.32
300	-0.66	0.62	0.67	0.67	0.68	0.71	0.72	0.72	0.70
700	-0.76	0.77	0.80	0.80	0.81	0.81	0.81	0.83	0.83
900	-0.78	0.80	0.81	0.81	0.84	0.85	0.86	0.84	0.86
1300	-0.68	0.69	0.66	0.70	0.72	0.75	0.78	0.75	0.78
1500	-0.53	0.57	0.56	0.58	0.58	0.61	0.61	0.60	0.63
1900	-0.40	0.48	0.42	0.39	0.36	0.38	0.38	0.33	0.35
2300	-0.32	0.30	0.31	0.31	0.30	0.31	0.31	0.31	0.31
2700	-0.31	0.31	0.31	0.31	0.30	0.30	0.32	0.30	0.30
3000	-0.31	0.31	0.31	0.31	0.30	0.31	0.31	0.32	0.31

Fig. 6: **Pathline length (n) vs sequence length (S):** F_1 scores for different combinations of n and S . The first column represents traditional pathlines, while subsequent columns are for sequence of pathlines. Notice that increasing the pathline length resulted in higher accuracy up to an optimal length; beyond this, the accuracy decreased. Also notice that using a sequence of pathlines over a single pathline per seed point produced even more accurate results. We found that $n = 900$ and $S = 100$ were optimal for our data set.

our training data and random beach images that one can find on the web are taken in landscape mode, seeds are distributed uniformly along the horizontal axis. Conversely, the water body is usually in the middle, with the sky on top and the beach on the bottom halves of the framing. We therefore use a Gaussian distribution of seeds along the vertical axis. This is a form of importance sampling [4]. During the training process, we use importance seeding to maximize pathlines that track water movement, and reduce irrelevant information such as the beach or cloud movement in the sky. During the testing process, we use regular sampling since there is no guarantee that the test data will also be in landscape mode.

4.3 Comparison with FlowNet using Single Long and Sequences of Short Pathlines

We compared RipViz with FlowNet [17], a deep learning based streamline clustering method. In FlowNet, 3D streamline features were learned by an autoencoder. These 3D streamlines were traced across the full extent of the data. Additionally, they traced one streamline per seed point. The trained autoencoder was used to generate 1D representations of 3D streamlines. The dimensionality of these 1D representations is further reduced by t-distributed stochastic neighbor embedding (t-SNE) [43]. The resulting low dimensional vectors are then clustered by using DBSCAN [13]. An interactive user interface was used to find the desired clustering by changing the maximum distance between two feature descriptors (eps) and the minimum number of samples in each cluster (ms). The autoencoder in FlowNet is trained with a unweighted binary cross-entropy function (i.e., $w_0 = w_1 = 1$).

In order to compare with FlowNet, we transformed our long 2D pathlines to 3D by adding a time axis as the z axis [42]. However, since our pathlines are sparsely distributed in the 3D binary volume, we found that the FlowNet model with the unweighted binary cross entropy did not learn to differentiate the pathlines. Resulting in a single cluster for all the pathlines. We show the clustering result in Figure 9. Additionally, the training time was relatively long. A

Binary Image Size $L \times L$	1	10	20	40	60	80	100	120	140
5x5	0.16	0.17	0.23	0.15	0.19	0.24	0.20	0.18	0.17
32x32	0.31	0.32	0.38	0.30	0.34	0.39	0.35	0.33	0.32
64x64	0.78	0.80	0.81	0.81	0.84	0.85	0.86	0.84	0.86
128x128	0.65	0.66	0.64	0.67	0.69	0.72	0.75	0.72	0.75

Fig. 7: **Binary image size ($L \times L$):** Notice that a binary image size of 64×64 produced the highest F_1 score with least amount of computation.



Fig. 8: **Seeding strategies:** We used regular grid to seed pathlines when testing. For training we used importance seeding.

proper comparison requires modification of FlowNet to work with our application domain. The modifications are discussed in Appendix A.

We trained the modified FlowNet model with long pathlines, integrated across the entirety of the video, and with single short pathlines ($n = 900$). The output of the *modified* FlowNet is shown in Figure 9. The first column shows a representative frame of each video. We see that without modifications, FlowNet lumps all the pathlines into one cluster. Both the modified FlowNet and RipViz, when fed with single pathlines that run the entirety of the video, also fail to identify the rips. When fed with shorter pathlines ($n = 900$), modified FlowNet start to show signs of rip currents but includes significant FPs as well. Only, when sequence of pathlines as used in RipViz do we see the rip currents clearly.

In order to obtain a quantitative comparison between *modified* FlowNet with shorter pathlines and RipViz, we calculated their F_1 scores. If a seed point is flagged as a rip by either method is within the expert annotated rip current boundary then we mark it as TP, otherwise we mark it as FP. If seed points within the rip are not selected then we mark those points FN. We found that the F_1 score for *modified* FlowNet was 0.32 compared to the F_1 score of 0.85 for RipViz as shown in Table 1. We hypothesize the low F_1 score for the *modified* FlowNet was due to its high FP rate.

Furthermore, we found tuning the two hyperparameters eps and ms for the clustering step was crucial in finding the appropriate clusters. We noticed having some domain knowledge was helpful for the user when exploring these two hyperparameters to find the ideal clustering. In contrast, RipViz is automated and does not require any user input during run time.

Both *modified* FlowNet and RipViz are based on autoencoders. However, RipViz learns how each pathline of the same seed point was related in time by learning spatiotemporal features of pathline sequences. This allows RipViz to learn ocean flow behavior more effectively. On the other hand, *modified* FlowNet does not learn how pathlines are related in time. Therefore, we hypothesize that using pathline sequences is better for learning quasi-periodic behavior, such as in ocean flow.

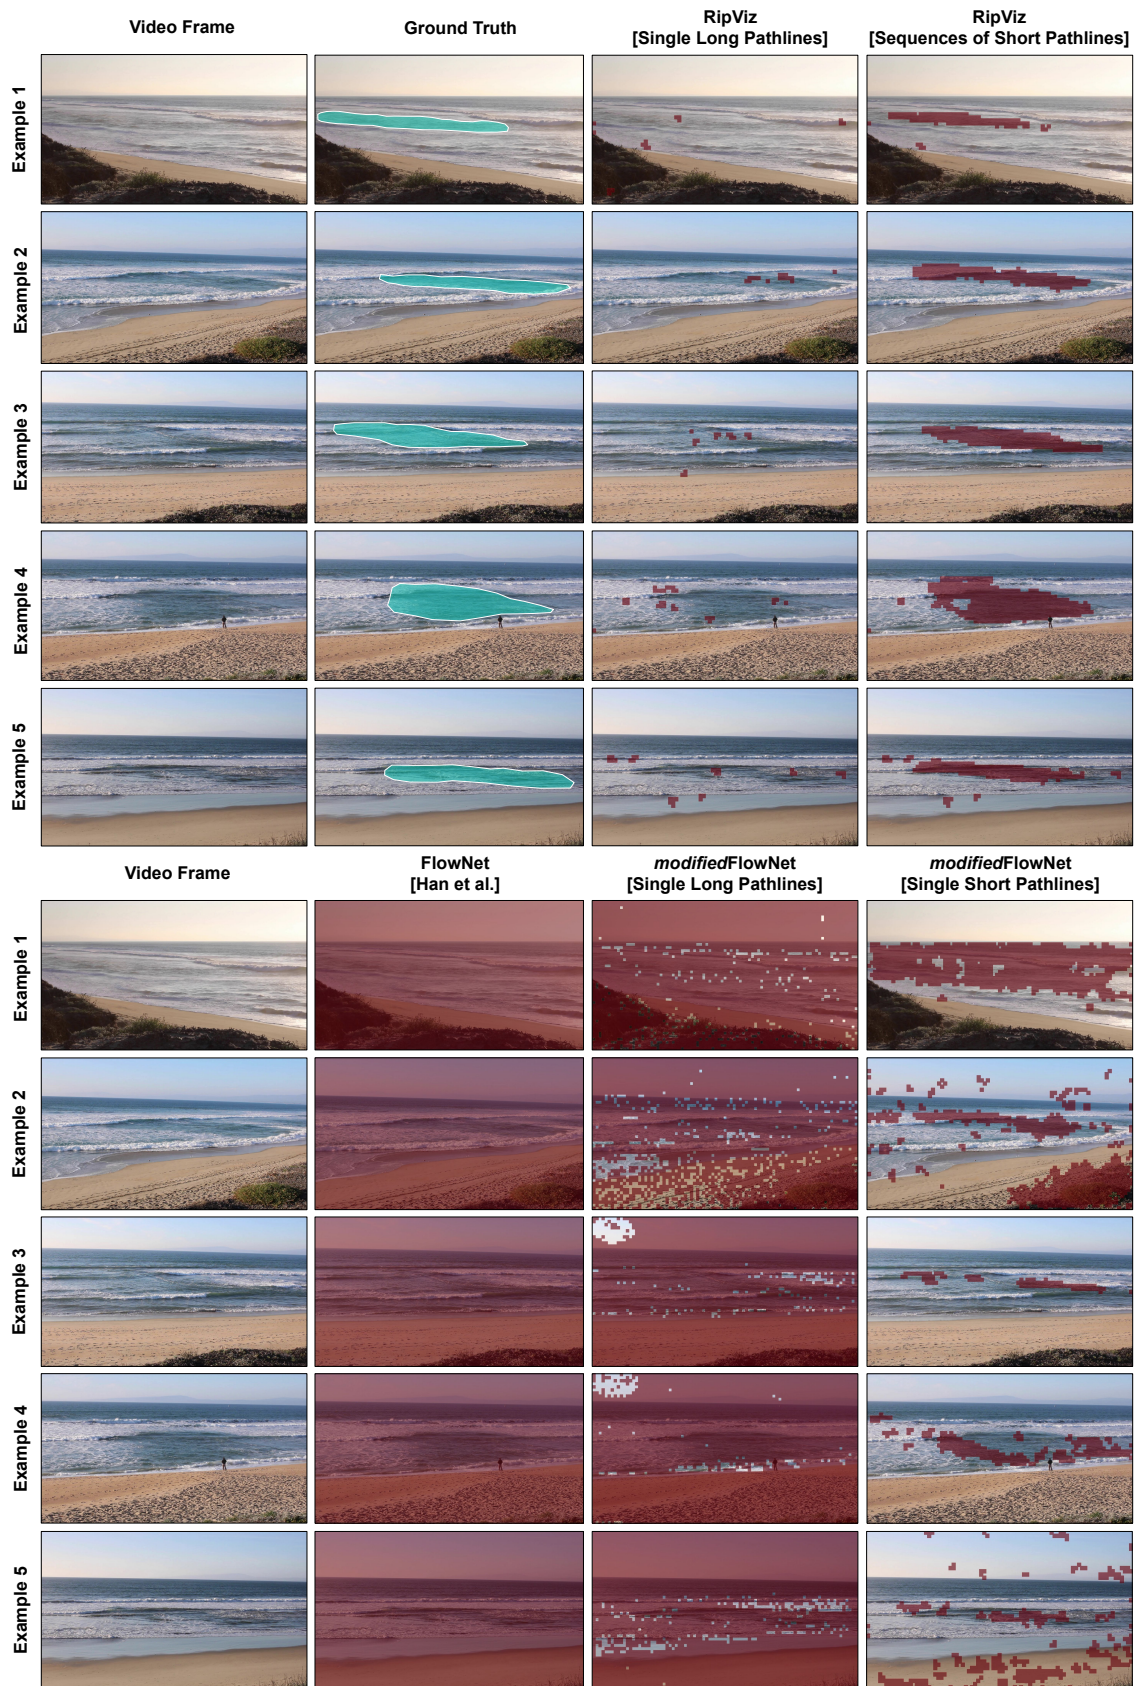


Fig. 9: **Comparison with FlowNet and different pathline parameters:** The *Video Frame* column shows a representative frame of the video clip while the *Ground Truth* column shows the expert's ground truth estimate. Notice that FlowNet without modifications (*FlowNet* Column) incorrectly clusters all seed points into a single group. Both *modifiedFlowNet* and RipViz when trained with longer pathlines were not able to detect the rip current but not as definitive. RipViz trained with shorter pathline sequences was able to detect the rip currents more precisely. Best viewed in color.

4.4 Comparison with existing rip detection methods

We compared RipViz with existing rip current detection methods as shown in Figure 10. The first and second columns of Figure 10 show a representative frame of the video and the expert drawn ground truth. The remaining columns show the output of the object detector [11], RipViz, filtered arrow glyphs [30], filtered color maps [30], and timelines [30], respectively.

4.4.1 Comparison with behavior based methods

Timelines: Mori et al. [30] proposed to use timelines to detect and visualize rip currents. They placed timelines parallel to the beach and traced the points on the timeline using optical flow to update its position. They observed the shape of the timeline as it gets dragged by the rip current. Timelines get deformed and extend to the rip channel above its initial position, as shown in Figure 10.

However, the initial placement of the timeline is a major contributing factor for timelines to be successful. All the timelines indicated in Figure 10 are placed optimally. To compare the timeline method with RipViz, we calculated the F_1 score. If the timeline was able to visualize the rip current in a video, then we counted it as TP, otherwise it is counted as FN. We found that properly placed timelines can visualize rip currents in most cases. However, the user has to specify the good initial placement for each video. In contrast, RipViz is automated, and no user input was needed during run time, as shown in Table 1.

Direction based clustering: Philip and Pang [34] hypothesized that the rip current is directly opposite the dominant flow in the vector field, which corresponds to the incoming wave direction. Places where the vectors were opposing the majority flow are potential rip zones. Areas with sufficiently large clusters of such vectors and large enough magnitude were highlighted as rip zones.

Mori et al. [30] used a similar approach and mapped direction to color and magnitude to hue for better visualization. As shown Filtered Color Map column of Figure 10, the method can detect the location of the rip current. However, this method also highlighted the swash zone, the shallow part of the beach, as evident in examples 1-3. We found that while the rip current was highlighted in most cases, the number of false positive pixels tend to be high. In order to compare with RipViz, we calculated the F_1 score. As shown in Table 1 filtered color map resulted in a F_1 score of 0.28, due to its high false positive rate.

Mori et al. [30] also proposed arrow glyphs to visualize rip currents. The output of this method is shown in Arrow Glyph column of Figure 10. We noticed that the arrow glyph method is more susceptible to noise in the flow field compared to other methods and RipViz. This is evident by the arrows projected on the sky and the beach although there is a little movement on those parts of the video. Although arrow glyph method highlighted the rip current in the majority of the videos, there were also a large number of false positive detections. Similarly we calculated the F_1 score. As indicated in Table 1, filtered arrow glyph had a F_1 score of 0.16 due to its relatively high false positive rate.

4.4.2 Comparison with appearance based methods

Object Detectors: de Silva et al. [11] used object detectors to detect rip currents. They trained a deep learning based object detector, Faster R-CNN [37], with images of rip currents. The images they used had a clear visual signature for bathymetry controlled rip currents with a darker region between breaking waves. Then they used the model to predict the location of rip currents, by overlaying a bounding box on the rip currents on videos. The output of the method is shown in the *Object Detector* column of Figure 10.

We noticed that the object detector can detect parts of rip currents as shown in examples 2-5 in Figure 10. However, rip currents have different appearances [6]. As evident in example 1, it is possible for an object detector to miss the actual rip current location. In contrast, RipViz is trained on pathlines that capture the *behavior* rather than appearance of rip currents. Hence, it was able to detect the rip current in all the examples.

Furthermore, we noticed that bounding boxes only detects the general area of the rip current. Sometimes the bounding boxes can cover

Method	Automated?	F_1 Score
Object Detector [11]	yes	0.43
Timelines [30]	no	0.72
Filtered color map [30]	yes	0.28
Filtered arrow glyph [30]	yes	0.16
FlowNet [17]	no	0.16
modified FlowNet [single long pathlines]	no	0.19
modified FlowNet [single short pathlines]	no	0.32
RipViz [single long pathlines]	yes	0.08
RipViz [sequence of short pathlines]	yes	0.85

Table 1: **Results summary:** Notice that RipViz has the highest F_1 score, and does not rely on user input at run time.

parts of the beach or miss parts of the rip especially with long elongated rips. We also noticed that the information such as the curvature of the rip current cannot be gleaned from bounding boxes alone. In order to compare with RipViz, we calculated the F_1 score. We counted how many pixels were covered by the bounding box. Out of those, rip pixels are counted as TP, the remaining pixels were counted as FP. If the object detector does not predict a bounding box, then we counted the pixels within the boundary as FN. As shown in Table 1, the bounding boxes resulted in a F_1 score of 0.43 compared to 0.85 for RipViz. We attribute this lower F_1 score to the large number of false positives and false negatives generated respectively by the non-rip areas of the bounding box and by part of the rip current not being covered by the bounding box.

Additionally, we tested RipViz on other types of rip currents where currently there are no published object detector available. In Figure 11, we show a few examples of such rip current types. Experts categorize the rip currents in the first and second row as sediment and structural rips. However, as shown in Figure 11 RipViz could detect these rip currents regardless of their appearance. More importantly, no additional training data were needed for sediment and structural rips. Examples 3 and 4 illustrate two rip currents where object detectors failed due to a lack of expected visual features. In these two examples, RipViz detected the rip currents because it uses *behavior* rather than appearance to detect rip currents.

4.5 Discoveries made by Experts using RipViz

Rip current researchers use their expert knowledge to identify the boundary of rip currents by observing the appearance of visual features such as gaps in breaking waves or sediment plumes. However, sometimes parts of the rip current lack these visual features, making it challenging for the expert to identify the entire extent of the rip current.

In contrast to appearance-based identification, RipViz uses the behavior of rip currents, not appearance, to detect rip currents. Experts can use RipViz as a visualization tool to determine the entire boundary of the rip current when visual features are lacking. In Figure 12, we demonstrate a few use cases where the rip current expert's original boundary estimate was updated after examining the visualization provided by RipViz. The updated boundaries now include feeder currents near shore and what appears to be a circulating pattern farther offshore in Discovery 1 and Discovery 3, as well as a weaker neighboring rip that merged with the dominant rip in Discovery 2. Incidentally, Discovery 1 and 3 both show circulating rips where the guidance from beach signage to swim parallel to shore may not always work [27]. We offer the complete table of discoveries made on our test data in the supplementary materials.

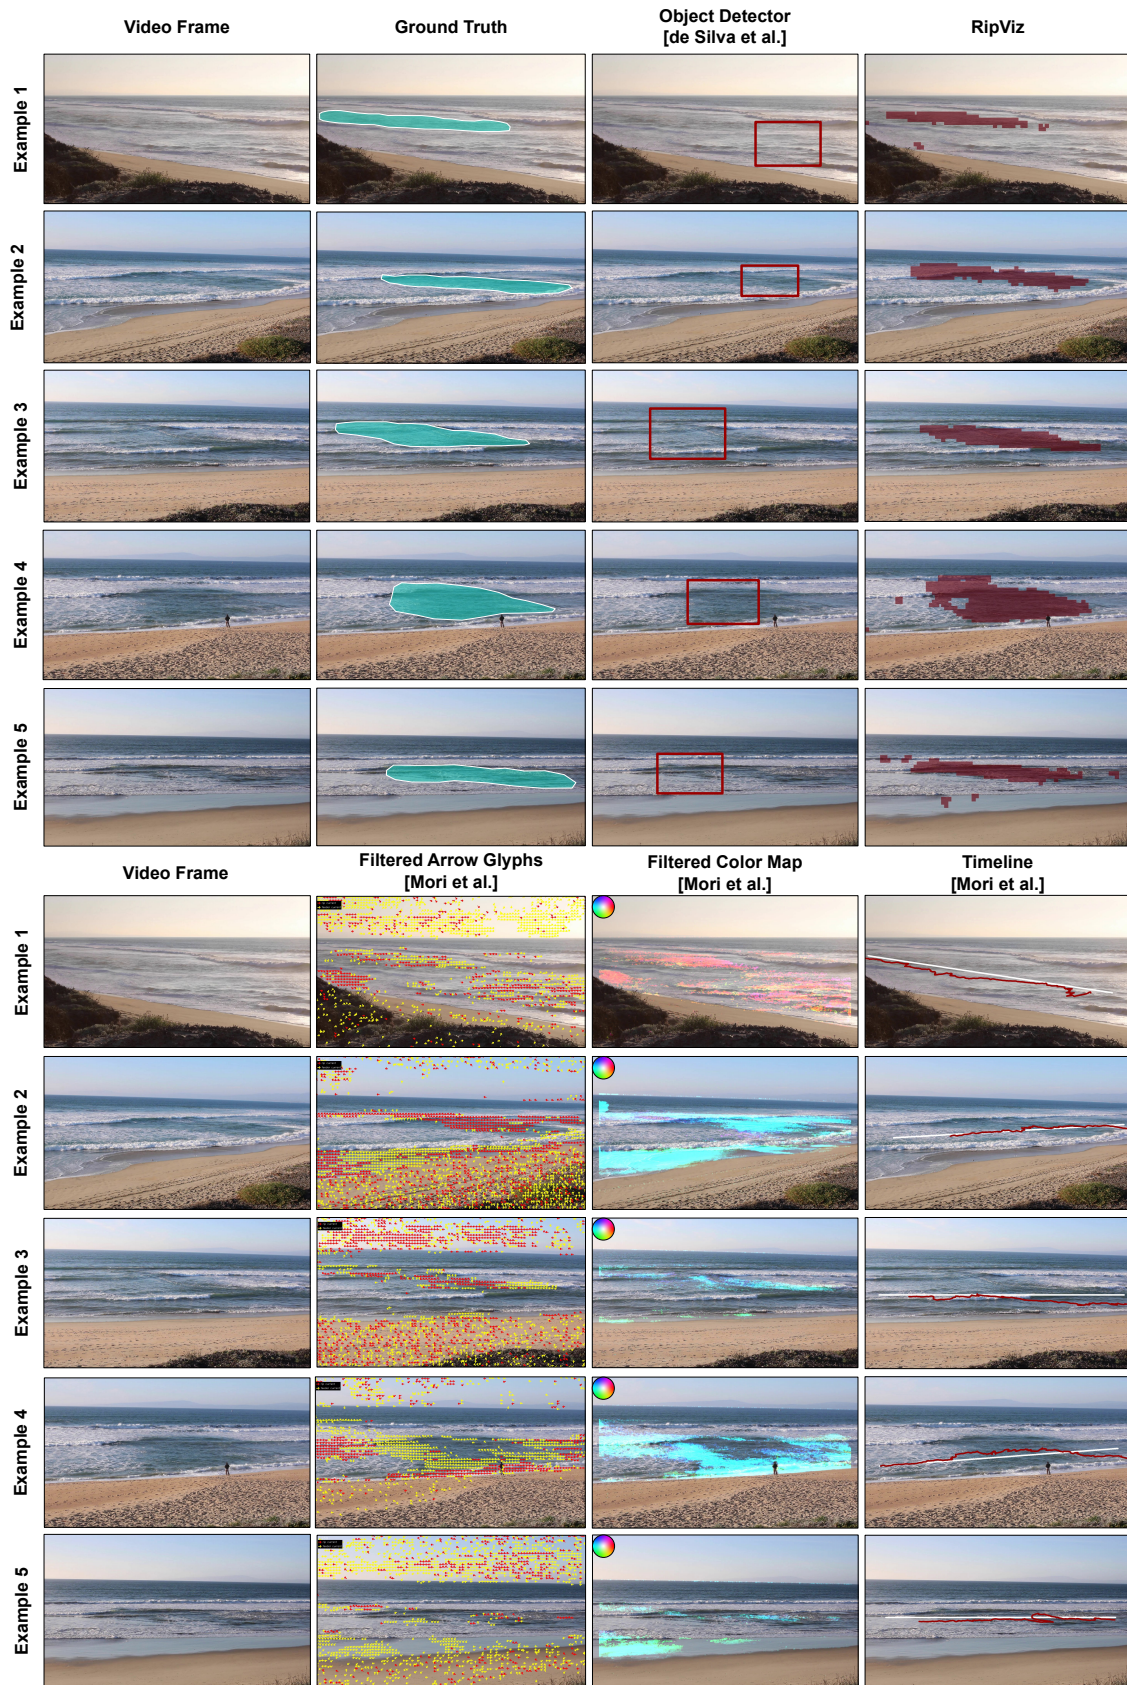


Fig. 10: **Comparison with prior rip current detection methods:** The *Video Frame* and *Ground Truth* columns show a representative frame of the video clip and the expert ground truth estimate respectively. The remaining columns show object detector, timeline, colormap, arrow glyph, and RipViz, respectively. Note that the timeline method requires user input to specify its initial placement, and that Filtered Arrow Glyphs did not filter out erroneous glyphs from the sky, beach, and non-rip area of the water. Best viewed in color.

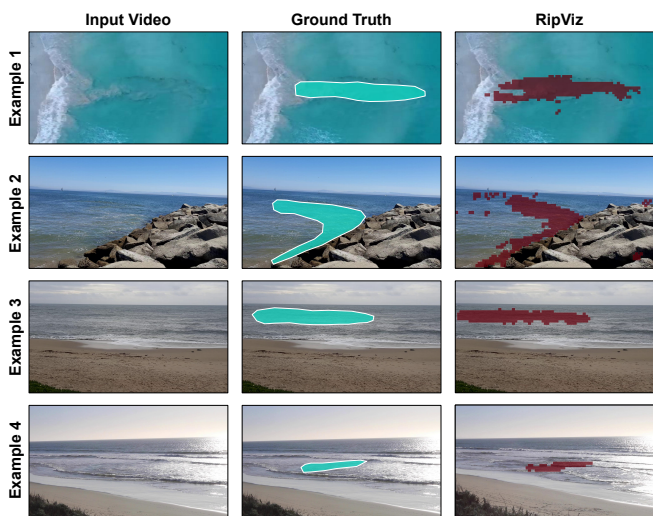


Fig. 11: **RipViz on novel rip current types:** Examples 1 and 2 are sediment rip and structural rip, respectively. Note that they have very different visual characteristics compared to bathymetry rips. No published object detectors exist for these two types of rips. However, notice that RipViz could detect rip currents regardless of their varying appearance. Examples 3 and 4 illustrate two rip currents where object detectors failed due to a lack of expected visual features. In these two examples, RipViz can still detect the rip current because it uses *behavior* rather than appearance to detect rip currents. Best viewed in color.

4.6 User Study

We conducted a user study to better understand how non-experts can use RipViz as a tool to become more aware of rip currents. We grouped 400 non-experts into two groups; one group was trained only with beach warning signs of rip currents, the other group was trained with RipViz. We then showed each participant a randomly picked video of a rip current and asked them if there was a rip current present. In the group trained with beach signs, 24% failed to recognize the existence of the rip current in the video. In the group trained with RipViz, only 14% were unable to recognize the presence of the rip current in the video.

We performed a second experiment to better understand if the participants could locate the rip current within the video. Both groups were given multiple choices of rip current boundary estimates and were asked to pick the correct one. In the group trained with beach signs, only 34% could choose the correct boundary. In the group trained with RipViz, 78% could select the correct boundary. The non-experts were acquired using Mechanical Turk [8, 32], with basic screening for reliable workers, and paid \$0.10-\$0.20 per task.

4.7 Sensitivity to video framing

Some of the parameters are dependent on the camera framing which would include distance of camera to the water and focal length or zoom factor of the lens. Alternatively, we can think about the width of the beach that is visible in the frame. We base our framing by examining a number of existing surf webcams (e.g. webcoos.com) and planned webcam installations. We experimented on how sensitive our parameters are to changes in camera framing. In the Figure 13, we show effects of varying the framing while keeping our current set of parameters. Recall that during testing, we trace the same number of seed points that are distributed in a regular grid. Based on the results shown in Figure 13, we believe the parameters are still valid within 20-30% change in camera framing. We found the F_1 score to be 0.82 ± 0.3 for all the examples, without much deviation.

5 SUMMARY AND REMARKS

As Ben Schneiderman succinctly captured in his quote: “The purpose of visualization is insight, not pictures”, our goal is to make apparent what may not be visible to the untrained eyes. The focus of this work is to first find the feature of interest (rip) and then present them in a simple, easy to understand manner by highlighting their location directly on the video. This capability is encapsulated in RipViz, a hybrid feature detection method that combined ML and flow analysis to extract rip currents from stationary videos. We used shorter pathline sequences to capture the flow behavior in a noisy quasi-periodic flow field. Then we used an LSTM autoencoder to learn the behavior of normal ocean pathline sequences. The trained model allowed us to label pathline sequences belonging to rip currents as anomalous by comparing the reconstruction error. By framing the rip detection problem as an anomalous flow detection problem, the onerous task of finding and labeling training datasets for each type of rip current is also greatly reduced.

The Visualization community has used deep learning methods to learn flow behavior from pathlines. In particular, existing literature uses autoencoders to learn from flow data. The authors use single long pathlines/streamlines per seed point as input to their deep learning models. However, straight-forward learning based on traditional pathlines did not work for our application due to the quasi-periodic flow fields such as those found near the surf zone. Therefore, we adapted and innovated on the existing deep learning methods to use a sequence of pathlines per seed point instead

The main contribution of this paper is:

- A hybrid feature detection method that combines machine learning and flow analysis techniques to automatically find and visualize dangerous rip currents.

In order to realize this, the following innovations are necessary:

- For flow fields with a quasi-periodic behavior, such as ocean flow, working with a *sequence* of shorter pathlines is better than a single long pathline.
- A weighted binary cross entropy, unlike the unweighted binary cross entropy used in previous works, is more effective when learning from sparsely distributed pathlines generated from ocean scenes.

A key assumption about the stationary videos is that the camera is sufficiently close to the water in order for the optical flow algorithm to pick up measurable velocities. For similar reasons, we assume that the camera is pointed mostly seaward and not parallel to the beach. Pointing the camera down a long stretch of beach will create an optical flow that is not representative especially for points farther away from the camera. Rectifying the frames prior to optical flow calculations may extend the usable range a bit farther but does not justify the extra computational cost. For these reasons, RipViz works best when the camera is close to the water and pointed mostly seaward.

Our training data were videos from mostly uncrowded beaches. We did study seed points along the path of a jogger. Such seed points were not marked as anomalous. This is likely due to the fact that the subsequent pathlines from the sequence marked the initial seed point as mostly normal. We believe that seed points that are placed on surfers or birds or other objects in the scene will likely be marked as normal as well. We plan to study this further with more testing data.

Extensions and Future Works: Some of the contributions listed above are not specific to rip current detection. We plan to investigate how pathline sequences of unsteady flow fields coupled with the weighted cross-entropy loss function can be used to distinguish between laminar and non-laminar flows, and possibly train a model to detect vortices as anomalous behavior.

We surmise that analysis using pathline sequences may also benefit certain classes of flow data aside from those in this paper. In fact, we are exploring that avenue and plan to report on the results in a separate paper. In short, the particular needs of rip current detection led us to

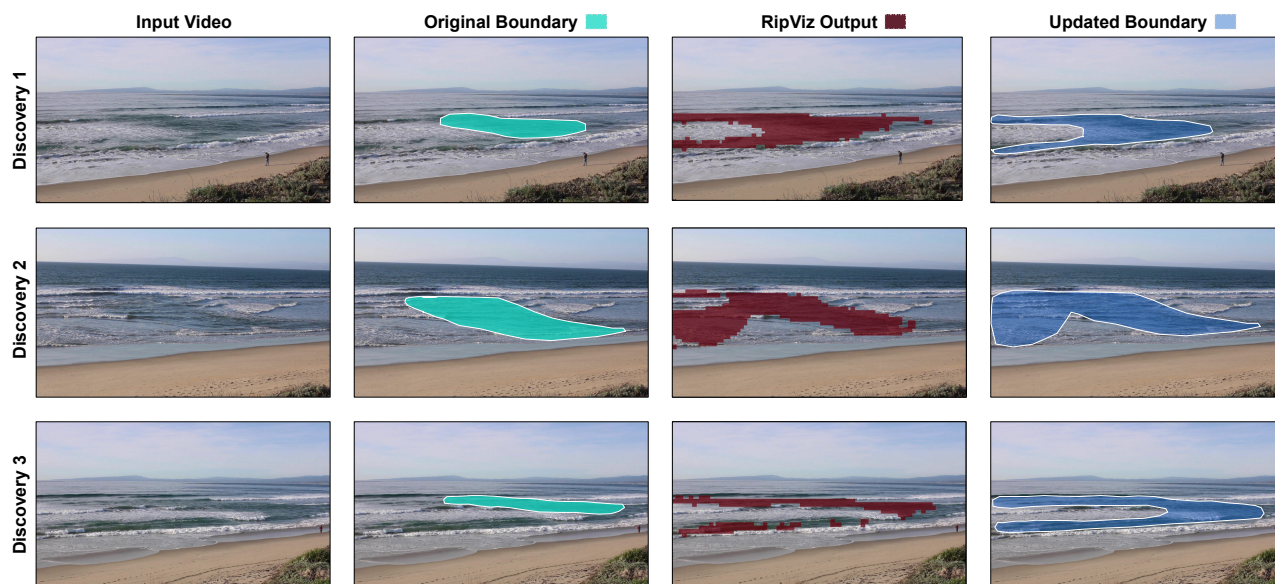


Fig. 12: **Discoveries made by Experts using RipViz:** Experts estimated the rip current boundary (*Original Boundary* Column) by observing only the input video (*Input Video* Column). After examining the output of RipViz (*RipViz Output* column), the experts updated their original rip current boundary estimate to include parts of the rip current that are not readily observable in the input video due to a lack of distinct visual features (*Updated Boundary* column). In the supplementary materials, we included a complete table of discoveries made on our test data. Best viewed in color.

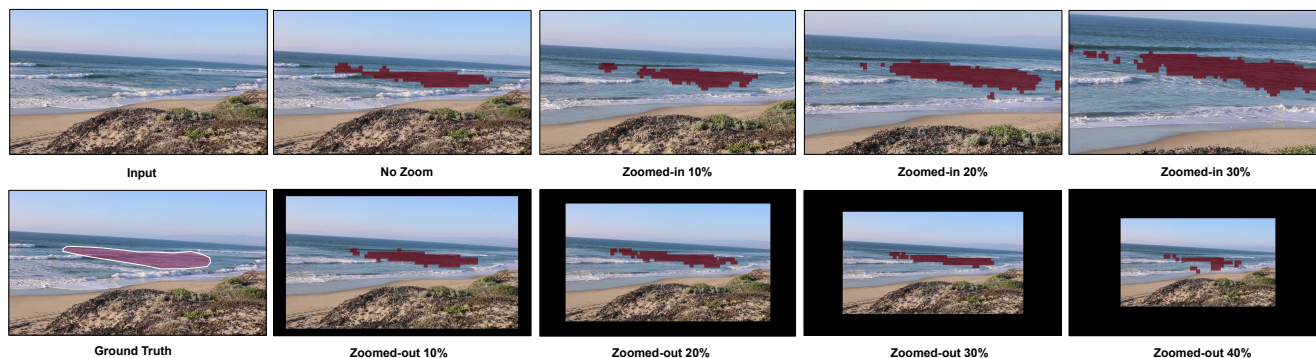


Fig. 13: **RipViz is less sensitive to different video framing:** We changed the scale of the rip current feature by zooming in and out from the original video, while keeping the aspect ratio of the video unchanged. Notice that for instances where the rip current feature is larger or smaller than the training data (within reasonable bounds), RipViz was still able to detect the rip current feature.

develop the approach presented in this paper, and which also provides another potential tool for the visualization community to study certain classes of flow fields.

APPENDIX A: MODIFIED FLOWNET

Han et al. [17] proposed, FlowNet, a deep learning based method to cluster and select streamlines. They first transformed each streamline into a 1D feature vector of length 1024. Then they used the t-distributed stochastic neighbor embedding (t-SNE) [43] to reduce the dimensionality of the 1D vector from 1024 to 2. The resulting vectors are then clustered using DBSCAN [13], a density-based spatial clustering method for applications with noise. The authors provided a user interface to vary the clustering by changing the maximum distance between two feature descriptors (*eps*) and the minimum number of samples in each cluster (*ms*) in the DBSCAN algorithm.

We started with the FlowNet architecture for this work, but modified it according to the requirements of our application. Their primary focus was to cluster streamlines from 3D steady flow fields. Therefore, the authors designed the deep neural network architecture to take in a 3D streamline as input. However, in our application domain, the

pathlines are 2D. In order to adapt FlowNet to our application domain, we changed the architecture of FlowNet to take in a 2D pathlines as input. Likewise, we changed the original neural network architecture by replacing 3D convolutional layers and 3D batch normalization layers, respectively, with 2D convolutional layers and 2D batch normalization layers to accommodate 2D pathlines. We also updated the input size and output size of the fully connected layers accordingly. We kept the same number of layers, same activation functions, and the same number of convolutional filters as defined in the FlowNet paper. We refer to the new neural network architecture as *modified* FlowNet.

We also found that the unweighted binary cross-entropy function used in the FlowNet paper was insufficient to learn the variations of pathlines of the ocean flow for reasons discussed in Section 4.2. Therefore, we used the weighted binary cross-entropy loss function defined in Section 3.4 when training the *modified* FlowNet. We trained the *modified* FlowNet using the same approach as discussed in the FlowNet paper [17].

We used the trained *modified* FlowNet to find 1D representations of the pathlines by extracting the latent feature descriptor from the autoencoder. We reduced the dimensionality of the 1D vector by using t-SNE

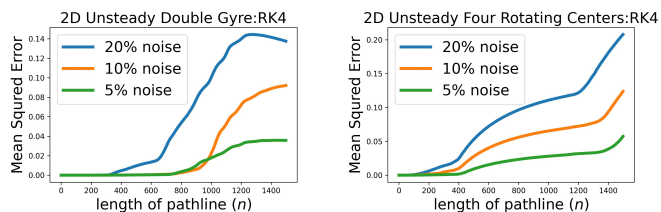


Fig. 14: **Comparison of pathlines with and without noise:** Mean Squared Error (MSE) of pathlines calculated at varying lengths, averaged over 200 randomly seeded pathlines. Notice that for both datasets longer pathlines have higher MSE compared to shorter pathlines.

to a vector with a length of 3. Then we clustered the resulting vectors using DBSCAN, which requires the user to set the maximum distance between two feature descriptors (*eps*) and the minimum number of samples in each cluster (*ms*) before clustering. The cluster with the largest overlap with the ground truth was selected.

APPENDIX B: MAKING THE CASE FOR SHORTER PATHLINES IN NOISY DATASETS.

We estimated the flow field from videos using optical flow. We assumed a small error/noise associated with the optical flow estimate. Numerical integration, even well designed higher order methods, are subject to accumulation of error. This problem is aggravated when one is working with noisy datasets. For the case of pathline integration, we hypothesized that shorter sequence of pathlines will be more accurate in capturing the flow behavior in noisy flow fields than a single longer pathline.

We used two synthetic datasets to indirectly verify our hypothesis: the *2D unsteady Double Gyre* dataset [40] and *2D Unsteady Four Rotating Centers* dataset [15].

We added noise at 5%, 10%, and 20% levels to each vector component of each dataset. We compared the similarity of pathlines originating from the same seed point. We observed that shorter pathlines were more similar to those without any noise. To quantify our observations, we randomly seeded 200 pathlines in each dataset and traced those pathlines. We calculated mean squared error (MSE) for pathlines at varying lengths as shown in Figure 14. We observed that longer pathlines have higher MSE compared to shorter pathlines. Therefore, we anticipate that shorter pathlines will capture a more accurate representation of the flow behavior in reconstructed noisy flow fields such as ours.

Additionally, we visualized how pathline sequences accumulate errors due to noise over longer integration times as shown in Figure 15. Here, we show two pathline sequences, one seeded in the sky and the other on the beach. We expected these two pathline sequences to be very short and remain closer to the seed point due to the lack of motion around those seed points. However, as we integrate beyond 900 time steps, we notice that the pathlines seem to travel far from the seed point even though we don't see any visible motion on the video. We attribute this to the accumulation of noise/error when integrating over a more extended time. The noise can come from optical flow estimation or video compression loss.

ACKNOWLEDGMENTS

This report was prepared in part as a result of work sponsored by the Southeast Coastal Ocean Observing Regional Association (SECOORA) with NOAA financial assistance award number NA20NOS0120220. The statements, findings, conclusions, and recommendations are those of the author(s) and do not necessarily reflect the views of SECOORA or NOAA.

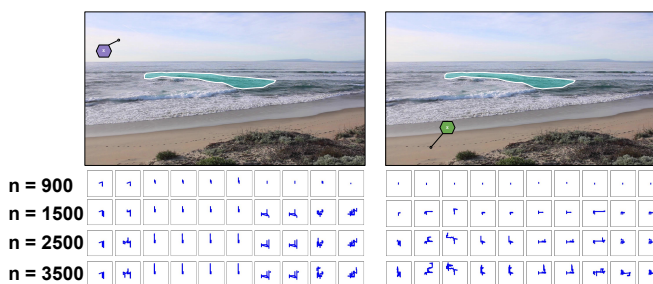


Fig. 15: **Visualization of noise in our dataset:** Here we visualize two pathline sequences, one seeded in the sky (left panel) and the other on the beach (right panel). Each sequence consists of 10 pathlines. We show the pathlines sequence at four integration time steps (n). We expected these pathlines to be very short and remain closer to the seed point. However, as the pathlines were integrated over a longer period (n), they traveled far from the seed point, contradicting our expectation. We attribute this to the accumulation of noise/error when integrating over a longer period in a noisy real world dataset.

REFERENCES

- [1] J. L. Ba, J. R. Kiros, and G. E. Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 3
- [2] M. Berenkjoub, G. Chen, and T. Günther. Vortex boundary identification using convolutional neural network. In *2020 IEEE Visualization Conference (VIS)*, pp. 261–265. IEEE, 2020. 2
- [3] A. J. Bowen. Rip currents: Theoretical investigations. *Journal of Geophysical Research*, 74(23):5467–5478, 1969. 1
- [4] K. Burger, P. Kondratieva, J. Kruger, and R. Westermann. Importance-driven particle techniques for flow visualization. In *2008 IEEE Pacific Visualization Symposium*, pp. 71–78. IEEE, 2008. 6
- [5] B. S. Carmo, Y. H. P. Ng, A. Prügel-Bennett, and G.-Z. Yang. A data clustering and streamline reduction method for 3d mr flow vector field simplification. In C. Barillot, D. R. Haynor, and P. Hellier, eds., *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004*, pp. 451–458. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. 2
- [6] B. Castelle, T. Scott, R. Brander, and R. McCarroll. Rip current types, circulation and hazard. *Earth-Science Reviews*, 163:1–21, 2016. 1, 8
- [7] J. Chen, B. Xie, H. Zhang, and J. Zhai. Deep autoencoders in pattern recognition: a survey. In *Bio-inspired computing models and algorithms*, pp. 229–255. World Scientific, 2019. 3
- [8] J. J. Chen, N. J. Menezes, A. D. Bradley, and T. North. Opportunities for crowdsourcing research on amazon mechanical turk. *Interfaces*, 5(3):1, 2011. 10
- [9] F. Chollet et al. Keras, 2015. 4
- [10] A. H. da F. Klein, G. G. Santana, F. L. Diehl, and J. T. 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*, pp. 107–116, 2003. 1
- [11] A. de Silva, I. Mori, G. Dusek, J. Davis, and A. Pang. Automated rip current detection with region based convolutional neural networks. *Coastal Engineering*, 166:103859, 2021. 1, 2, 3, 8
- [12] S. Ellis and S. McGill. Rip current and channel detection using surfcams and optical flow. *Shore & Beach*, 90(1):50–58, 08 2022. 3
- [13] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, vol. 96, pp. 226–231, 1996. 6, 11
- [14] B. Gary and A. Kaehle. Opencv. *Dr. Dobbs's journal of software tools*, 3, 2000. 3
- [15] T. Günther, M. Gross, and H. Theisel. Generic objective vortices for flow visualization. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 36(4):141:1–141:11, 2017. 12
- [16] M. Haller, D. Honegger, and P. Catalán. Rip current observations via marine radar. *Journal of Waterway Port Coastal and Ocean Engineering*, 140:115–124, 03 2014. doi: 10.1061/(ASCE)WW.1943-5460.0000229 2
- [17] J. Han, J. Tao, and C. Wang. FlowNet: A deep learning framework for clustering and selection of streamlines and stream surfaces. *IEEE transactions on visualization and computer graphics*, 26(4):1732–1744, 2018. 2, 3, 5, 6, 8, 11
- [18] R. Holman and M. Haller. Remote sensing of the nearshore. *Annual review of marine science*, 5, 12 2011. doi: 10.1146/annurev-marine-121211-172408 2
- [19] B. Kim and T. Günther. Robust reference frame extraction from unsteady 2d vector fields with convolutional neural networks. In *Computer Graphics Forum*, vol. 38, pp. 285–295. Wiley Online Library, 2019. 2
- [20] S. Leatherman and J. Fletemeyer, eds. *Rip Currents: Beach Safety, Physical Oceanography, and Wave Modeling (1st ed.)*. CRC Press, 2011. doi: 10.1201/b10916 1
- [21] S. Leatherman and S. Leatherman. Techniques for detecting and measuring rip currents. *International Journal of Earth Science and Geophysics (ISSN: 2631-5033)*, 3(1), 2017. doi: DOI: 10.35840/2631-5033/1814 2
- [22] L. Liu, W. Ouyang, X. Wang, P. Fieguth, J. Chen, X. Liu, and M. Pietikäinen. Deep learning for generic object detection: A survey. *International Journal of Computer Vision*, 128(2):261–318, 2020. doi: 10.1007/s11263-019-01247-4 1
- [23] S. Lodha, J. Renteria, and K. Roskin. Topology preserving compression of 2d vector fields. In *Proceedings Visualization 2000. VIS 2000 (Cat. No.00CH37145)*, pp. 343–350, 2000. doi: 10.1109/VISUAL.2000.885714 2
- [24] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. *International Joint Conference on Artificial Intelligence*, pp. 674–679, 1981. 3
- [25] J. B. Lushine. A study of rip current drownings and related weather factors. *National Weather Digest*, pp. 13–19, 1991. 1
- [26] J. Ma, C. Wang, and C.-K. Shene. Coherent view-dependent streamline selection for importance-driven flow visualization. In *Visualization and Data Analysis 2013*, vol. 8654, p. 865407. International Society for Optics and Photonics, 2013. 2
- [27] J. MacMahan, A. Reniers, J. Brown, R. Brander, E. Thornton, T. Stanton, J. Brown, and W. Carey. An Introduction to Rip Currents Based on Field Observations. *Journal of Coastal Research*, 27(4), 2011. doi: 10.2112/JCOASTRES-D-11-00024.1 1, 2, 8
- [28] S. Marchesin, C.-K. Chen, C. Ho, and K.-L. Ma. View-dependent streamlines for 3D vector fields. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1578–1586, 2010. 2
- [29] C. Maryan, M. T. Hoque, C. Michael, E. Ioup, and M. Abdelguerfi. Machine learning applications in detecting rip channels from images. *Applied Soft Computing*, 78:84–93, 2019. 1, 2
- [30] I. Mori, A. De Silva, G. Dusek, J. Davis, and A. Pang. Flow-based rip current detection and visualization. *IEEE Access*, 2022. 2, 3, 8
- [31] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Icml*, 2010. 3
- [32] G. Paolacci, J. Chandler, and P. G. Ipeirotis. Running experiments on amazon mechanical turk. *Judgment and Decision making*, 5(5):411–419, 2010. 10
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011. 5
- [34] S. Philip and A. Pang. Detecting and visualizing rip current using optical flow. In *EuroVis (Short Papers)*, pp. 19–23, 2016. 2, 3, 8
- [35] A. H. Rashid, I. Razzak, M. Tanveer, and A. Robles-Kelly. RipNet: A lightweight one-class deep neural network for the identification of RIP currents. In *Communications in Computer and Information Science*, pp. 172–179. Springer International Publishing, 2020. doi: 10.1007/978-3-030-63823-8_21 2
- [36] A. H. Rashid, I. Razzak, M. Tanveer, and A. Robles-Kelly. RipDet: A fast and lightweight deep neural network for rip currents detection. In *2021 International Joint Conference on Neural Networks (IJCNN)*. IEEE, jul 2021. doi: 10.1109/ijcnn52387.2021.9533849 2
- [37] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. 8
- [38] M. Ribeiro, A. E. Lazzaretti, and H. S. Lopes. A study of deep convolutional auto-encoders for anomaly detection in videos. *Pattern Recognition Letters*, 105:13–22, 2018. 3
- [39] S. Sane, R. Bujack, C. Garth, and H. Childs. A survey of seed placement and streamline selection techniques. *Computer Graphics Forum*, 39(3):785–809, 2020. doi: 10.1111/cgf.14036 2, 5
- [40] S. C. Shadden, F. Lekien, and J. E. Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D: Nonlinear Phenomena*, 212(3-4):271–304, 2005. doi: 10.1016/j.physd.2005.10.007 12
- [41] J. Tao, J. Ma, C. Wang, and C.-K. Shene. A unified approach to streamline selection and viewpoint selection for 3d flow visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(3):393–406, 2012. 2
- [42] H. Theisel and H.-P. Seidel. Feature flow fields. In *Proceedings of the symposium on Data visualisation 2003*, pp. 141–148, 2003. 6
- [43] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11), 2008. 6, 11
- [44] H. Yu, C. Wang, C.-K. Shene, and J. H. Chen. Hierarchical streamline bundles. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1353–1367, 2011. 2
- [45] J. Zhai, S. Zhang, J. Chen, and Q. He. Autoencoder and its various variants. In *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 415–419, 2018. doi: 10.1109/SMC.2018.00080 3