



AUV navigation using cues in the sand ripples

Hadar Shalev¹ · Liav Nagar¹ · Avi Abu¹ · Alberto Testolin² · Roei Diamant¹

Received: 12 August 2021 / Accepted: 29 September 2022 / Published online: 21 October 2022
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2022

Abstract

Subsea navigation by autonomous underwater vehicles (AUVs) is a demanding task that involves the integration of inertial sensors, gyrocompasses, Doppler velocity loggers, and reference from acoustic beacons. In this paper, we propose to augment this information by providing an external measurement of heading change. We rely on the direction of sand ripples, which are abundant on the seabed near the shore and whose direction is, locally, constant. Thus, any apparent change in their directivity, as detected by the AUV, would reflect as a change in the vehicle's heading. Considering this, we developed a mechanism that detects regions of interest (ROIs) containing sand ripples within a synthetic aperture sonar (SAS) image, segments the ROI into highlight and shadow, and evaluates the angle difference between ROIs within two consecutive SAS images. For detection of sand ripples and estimation of angle difference, we employ two deep neural networks, while for segmentation we formulate a fuzzy-logic clustering. Taking advantage of a transfer learning approach, we trained the deep networks on simulated SAS images and on a large database of 2088 real SAS images, which we share for reproducibility. Results from real SAS images from three different sites show a good trade-off between precision and recall for sand-ripple detection, and an error of a few degrees in the heading change estimation, which well exceeds a geometrical-based benchmark. We also show performance from a real-time experiment for which we implemented our method on an AUV and estimated its heading change on-the-fly.

Keywords Underwater navigation · Sand ripples · Synthetic aperture sonar · Heading change · Deep neural networks · Segmentation

1 Introduction

Navigating underwater is a demanding task that involves integrating information from inertial sensors, compasses and gyrocompasses, Doppler velocity loggers, and acoustic beacons (Stutters et al., 2008; Paull et al., 2014). In addition, template matching to navigate - based on cues on the ground such as rock or infrastructures (Fallon et al., 2013), or by depth maps (Rupeng et al., 2020; Fezzani et al., 2019) — are also used, either directly or by means of simultaneous localization and mapping (Norgren & Skjetne, 2018). One of the key challenges in underwater navigation is to keep track of heading changes. Indeed, as we showed in an earlier work (Klein & Diamant, 2018), in such scenarios observability is

low and errors in the heading estimation propagate rapidly. To assist in estimating the heading change, in this paper we present a novel method that allows to evaluate the angle of a heading change in shallow water using cues from the seabed.

For estimating directional angles of motion while being surveyed, principle component analysis has been adopted to find the most dominant axes within a 3-D accelerometers data framework (Kunze et al., 2009; Deng et al., 2015). Another option is to model the forward and lateral accelerations, such that the heading angle is found to be the one that maximizes the correlation of the estimated acceleration and the pre-determined model (Combettes & Renaudin, 2015). Alternatively, frequency analysis of inertial signals could find the direction that maximizes the spectral density of the accelerometer and gyroscope signals' energy (Han & Ortiz, 2014; Kouroggi & Kurata, 2014). For underwater navigation, heading direction can be estimated by path following (Fossen et al., 2015), modeling of motion (Putri et al., 2017), and the prediction of directional angles by neural networks (Song et al., 2020) or by tracking (Wang et al., 2019). Another option

✉ Roei Diamant
roee.d@univ.haifa.ac.il

¹ Department of Marine Technologies, University of Haifa, Haifa, Israel

² Department of Mathematics and Department of General Psychology, University of Padova, Padua, Italy

Fig. 1 Example of an apparent shift in the orientation of the sand ripple as a result of an AUV maneuver. Left panel: Image 1, heading 89.8deg. Middle panel: Image 2, heading 15.6deg. Right panel: Track of the AUV



is to use visual feedback to estimate the heading direction by calculating changes in the angle-of-observation (Khan et al., 2017), or by fusing velocity sensors and Earth-positioning sensors (Costanzi et al., 2017).

While the above approaches rely on instruments onboard the submerged vehicle — which are prone to drift and errors due to the platform’s movements in all directions, or visual aid approaches that require the submerged vehicle to remain close to the seabed — we propose a method based on external information from the environment. Our method is based on the nominal direction of sand ripples, which remains constant within a wide area. Abundant in shallow areas (<40m depth), sand ripples are caused by the force of wind-generated waves that displace patches of soft sand composites on the seabed. The result is small ($\sim 20\text{cm}$), sandy hills with a periodic pattern, whose frequency depends on the nominal swell (Heathershaw, 1982). Due to the visible distinction between the highlight (hill-top) and shadow (hill-bottom) regions, these hills are well observed by sonar systems like synthetic aperture sonar (SAS) or multibeam. Sand ripples have a unique direction, determined by the dominant direction of the waves. As a result, on a regional scale, all sand ripples will point in the same direction. We utilize this characteristic to learn about changes in the AUV’s direction.

While without a reference heading direction we cannot directly estimate the direction of the sand ripples, we can evaluate apparent changes in this direction, i.e., changes caused by a shift in the heading direction of the submerged vehicle. An example of two SAS images observed at different angles is shown in Fig. 1. As a result, analyzing two sequential sonar images to find a shift in the observed direction of sand ripples allows a direct measurement of a heading change. This process is illustrated in Fig. 2. Feeding the estimated heading change to the vehicle’s navigation system as external information would thus improve accuracy.

To estimate the apparent changes of sand ripples between two sonar images, we designed and implemented a system comprised of the following components:




Fig. 2 Illustration of the process of estimating the AUV’s heading change by cues from sand ripples

1. A deep neural network (DNN) set to find regions of interest (ROIs) consisting of sand ripples within a sonar image.
2. A clustering method that distinguishes between the highlight and shadow of sand ripples within the ROI, yielding a binary ROI.
3. A second DNN, whose objective/purpose is to estimate angle differences between two binarized ROIs.

To train and test our DNNs, we exploited an SAS simulator that produced images of different seabed backgrounds obtained at different angles of observation. Taking a transfer learning approach, we then refined the resulting networks with a magnitude of 80,000 tagged sand-ripple regions originating from 2088 SAS images that we collected in three sea expeditions using the SAS system of our own autonomous underwater vehicle (AUV). To train and test our second DNN, we analyzed results from over 80,000 pairs of ROIs obtained in three designated sea experiments, including deliberate heading changes of the AUV. We share this database with the research community. Finally, we implemented our solution on the AUV and tested our approach in real-time in a fourth designated sea experiment to evaluate the practicality of our

approach. The results show a trade-off between precision and recall for the sand-ripple classification that exceeds a geometrical benchmark, and an error of a few degrees in estimating the heading change of an AUV in real sea experiments. We therefore conclude that our approach can contribute valuable information to improve future undersea navigation systems.

The remainder of this paper is organized as follows. In Sect. 2, we present the state-of-the-art in machine learning approaches for the analysis of sonar images. In Sect. 3, we introduce our system's model and assumptions. Section 4 describes the SAS simulator and the way we formed our database. In Sect. 5, we explain the details of the two DNN models and the clustering method. Section 6 contains the evaluation and testing results for the tagged data set and for the real-time sea experiment. Conclusions are drawn in Sect. 7.

2 Approaches for analysis of SONAR Images

Detection of sand ripples and their orientation measure in SAS imagery has been explored in Williams and Coiras (2010), where the scene is filtered to explore repetitions or *ripplicities* for several highlight-shadow pairs. Normalized by the difference between the upper and lower bounds for ideal and random ripplicity indication, respectively, allows for the setting of an a-priori threshold. The orientation is determined relative to the surveying AUV as the direction repetitions are most dominant. Similarly, sand ripples can be identified by setting a threshold on the connection between the observed crest orientation and the backscatter variance in an ROI (Crawford & Skarke, 2014); the orientation of the ripple is found as the angle in which the greatest variation is observed in the backscatter. Another possible modeling approach for sand-ripple identification is to measure the local grazing angle and to relate it to the global grazing angle and the Lambert coefficient for evaluation of the ripple surface (Tang et al., 2009b). While tested over real SAS images, the results highly depend on the ripples' homogeneity, whereas - as the non-Gaussian model in Tang et al. (2009a) and as the examples in Fig. 1 show - the highlight-shadow structure may break throughout the ROI. Sand ripples can also be identified and their orientation estimated by feature-based detection (Williams, 2015), where centroids of ellipse-bounding shadow regions, mapped to angle and length features, are used to identify ripples by the density and orientation homogeneity of shadows. The positive results over a large number of real SAS images supports extending the feature-based approach towards machine learning classifiers. A first step in this direction is to formalize the task as an optimization problem, and the statistical search for detection thresholds (Klemm et al., 2015). However, the recent advances in image processing using deep learning models

highlight the potential for further improvement. For example, convolutional neural networks have been successfully used to estimate image orientation (Fischer et al., 2015), and have been exploited in underwater scenarios to classify targets in SAS images (Williams, 2016) or to detect moving objects from SONAR measurements (Testolin & Diamant, 2020; Testolin et al., 2020). Deep learning architectures based on Siamese networks have proven particularly useful for tasks requiring the comparison of two image patches (Zagoruyko & Komodakis, 2015; Zbontar & LeCun, 2016), which corresponds to the scenario discussed in the present paper.

3 System model

3.1 Setup and assumptions

Our setup includes an AUV scanning of the seabed with an SAS system, while performing a mission. While submerged, the AUV continues to update its navigation solution using internal sensors (e.g., accelerometers, gyrocompasses, barometric pressure, etc.) and external positioning-aid information (e.g., acoustic beacons, Doppler velocity loggers, or bathymetric maps). The motion is mostly in straight lines; however, occasionally, the AUV performs heading changes as part of its mission or to correct for water current-based drifts. The AUV periodically produces SAS images that are available for analysis onboard the AUV. Given two such sequentially collected SAS images, our goal is to evaluate whether a heading change has been made and, if so, at which angle. In turn, this information is then fed back to the navigation system as another source of information to improve navigation accuracy.

The evaluation for heading change requires the appearance of sand ripples in the two analyzed SAS images. We assume that, during the collection of the two explored images, the AUV did not make changes to its pitch, yaw, or depth. We also assume that the sand ripples in the two explored images are aligned to the same (unknown) direction, such that any apparent change in the sand ripples' direction is translated directly to a heading change of the surveying vehicle. Since the distance traveled by the AUV during the processing of each SAS image is on the order of 50 m, while changes in the direction of sand ripples is generally over a much larger scale (an area of hundreds of meters and more (Schnipper et al., 2008) — we argue that the latter assumption is weak.

We note that our approach is opportunistic by nature. That is, if the AUV is scanning an area that includes sand ripples, it can use these for its navigation. As such, we avoid directing the AUV towards a better scan of sand ripples. Still, since a SAS often includes many patches of sand ripples, local changes in the direction of the sand ripple are averaged out. This is performed by calculating the mean of the angle change

over all pairs of sand ripple patches between two adjacent SAS images.

3.2 Preliminaries: origin of sand ripples

Sand ripples are abundant in the sandy sea beds of any continental shelf characterized by directional flow such as an estuary, a coastline, or a tidal channel. Derived by the oscillatory force of the surface waves and water current, sand ripples show large bed roughness and are thus resistant to flow, are influenced by wave dissipation, and are enhanced by sediment transport (Soulsby et al., 2012). Sand ripples also lead to suspension, burial, and fatigue damage of submerged pipes, and migrating sandwaves have been shown to cause the collapse of drilling platforms (Cong et al., 2019). They are, therefore, the focus of a large body of literature exploring the mechanics of their formation and their characterization.

The periodicity, height, and direction of sand ripples are determined by the nominal wave and current characteristics, the water depth, and the sea-bed material. As a function of their magnitude, sand ripples are classed as mega-ripples, dunes, and sandwaves, characterized as either *rolling-grain* ripples with no flow separation behind the crest, or *vortex* ripples that include vortex shedding at the crest (Stegner & Wesfreid, 1999). As the response time of sand ripples decreases with the strengths of the waves and currents, the ripples become more homogeneous in deeper water (Soulsby et al., 2012), and our common practice is to find sand-ripples of up to 40 m depth. Reaching to a few meters in height and to hundreds of meters in wave length (Cong et al., 2019), the ripples are well observed by high resolution sonar like SAS, where the crest is observed as the highlight and the soft sand at the bottom of the hills as shadow.

4 Description of database

To properly train our deep learning models, we first created a large number of training images. We take a transfer learning approach, starting from synthetic SAS images produced by a designated simulator, and then fine-tuning the models on real SAS images collected in multiple experiments with our

own AUV. A description of the simulator and the collected images is given in this section.


4.1 Description of the SAS simulator

Our SAS simulator is based on statistical models of SAS image pixels (Cobb et al., 2010), which we implemented to yield synthetic SAS images of a given size and seabed type. The seabed texture models are based on the assumption that SAS pixels are identically independently-distributed (i.i.d). The pixel's statistical model is given by the product of a complex Gaussian component with a correlated Gamma function. The statistical differences between seabed textures are the number of components in the statistical models and its parameters. For sea-grass, rock, and sand ripples, a total of five components were used, while for sand, only three components were used. Figure 3 introduces three synthetic SAS images of sand, sand ripples, and sea-grass textures, and Fig. 4 shows a real SAS image with sand-ripple formations sideways to an output of sand-ripple texture from the simulator. While apparent differences are shown between the real SAS image and the simulated one, the basic shape of the ripple is still visible in the simulated image. For the task of sand-ripple classification, we generated a data set of 3000 synthetic images, divided into 1500 images of sand ripples and 1500 images of sand or seagrass texture.


For the task of angle difference estimation, we obtained a data set of roughly 20,000 pairs of sand-ripple images with angle difference randomly uniformly generated in a range of $[-90, 90]$ degrees. To this end, we modified the simulator to generate sand-ripple textures at different angles. Specifically, the covariance of the simulated sand-ripple textures is composed of a mixture model of gamma function components. Each component is shifted by a rotation angle, θ_i , along the x-axis. Formally, let parameters l_{x_i} and l_{y_i} define the major and minor axes of the ellipse that represents the intensity of the i th component in the covariance model of the gamma noise, and β_x, β_y defines the vertex and co-vertex of the ellipse. The covariance of the gamma function mixture with N components is given by

$$\rho_g(x, y) = \frac{1}{v} \sum_{i=1}^N \eta_i \cdot \exp \left[-\frac{1}{2} (x \ y) \Sigma_i^{-1} (x \ y)^T + R_h^2(x, y) \right], \quad (1)$$


Fig. 3 Simulated SAS textures. **a** sand; **b** sand ripples; **c** sea-grass



(a)



(b)



(c)




Fig. 4 **a** Example of SAS image with sand ripples. **b** Output from the SAS simulator for sand ripple background

where

$$R_h^2(x, y) = \exp\left[-\frac{x^2}{4\beta_x^2} - \frac{y^2}{4\beta_y^2}\right], \quad \Sigma_i = \Psi_i \Gamma_i \Psi_i^T,$$

and

$$\Gamma_i = \begin{bmatrix} l_{x_i}^2 & 0 \\ 0 & l_{y_i}^2 \end{bmatrix}, \quad \Psi_i = \begin{bmatrix} \cos\theta_i & \sin\theta_i \\ -\sin\theta_i & \cos\theta_i \end{bmatrix}.$$

4.2 Process of data collection

To complement the synthetic images, we analyzed real SAS images. The images were obtained by our ECA Robotics A18 5.5 m AUV and its two-sided Kraken MINSAS 120 SAS, whose center frequency is 337 kHz. The obtained resolution is 3 cm per pixel, and the range covered is a rectangle of 170×50 m². We analyzed a set of 2088 SAS images obtained at three locations near the shores of the Israeli coast. Since the database was collected near shore sites, we found sand ripples in nearly 60% of the tagged images. We share this database in Diamant (2022).

We manually tagged the database of SAS images to a binary mask, indicating places with and without sand ripples. This process included the following steps:

1. We start by separating the port and starboard sides of the SAS image, and converting the raw measured values in dB scale into a linear scale saved in a standard image format.
2. Using the VGG open-source image annotator (Dutta et al., 2016), we had an expert manually tag the images into polygon borders in Cartesian coordinates around areas with sand ripples.
3. Based on the tagged locations, we form a binary mask over the image, with white areas representing sand ripples




Fig. 5 Example of the process of sand-ripple tagging. Upper-middle image shows the original SAS image; Bottom-middle image shows the tagging mask; Left images are two examples of resulting ROIs with sand ripples; Right images are two examples of resulting ROIs without sand ripples

and black areas representing areas with no sand ripples. An example of such a mask is presented in Fig. 5.

4. Setting a threshold over the number of white pixels within a binary mask, the corresponding SAS image is split into windows of areas with sand ripples and without sand ripples.

For the labeling of angle changes between pairs of SAS images, we divided the ROIs that were labeled to include sand-ripple pairs, each member taken from a different SAS image. The result was a large data set of more than 80,000 pairs of ROIs. Since each SAS image can be related to the AUV’s heading direction, relative to the north, as was logged upon obtaining the image, the angle change for a pair of ROIs was calculated as the difference between the two heading directions.

5 Sand-ripple Identification and angle estimation

The full process leading to heading estimation is illustrated in the block diagram in Fig. 6. For a given pair of SAS images, the process begins with the division of each SAS image into ROIs using a sliding window approach, and an analysis is individually performed for each ROI. The potential ROIs are of 128×128 pixels, whereas the full SAS images is on the order of 4000×6000 pixels. Next, each ROI is identified as either containing or not containing sand ripples. If such ROI are identified, a clustering algorithm is used to segment each ROI to highlight and shadow yielding a binary ROI. Then, for all combinations of ROI pairs, where each member of the pair must be from a different SAS image, we estimate the degree of angle change in the direction of the sand ripples. Finally, the heading direction is calculated as the median of all angle change estimations obtained from all ROI pairs. In

Fig. 6 Block diagram of the process for head change estimation

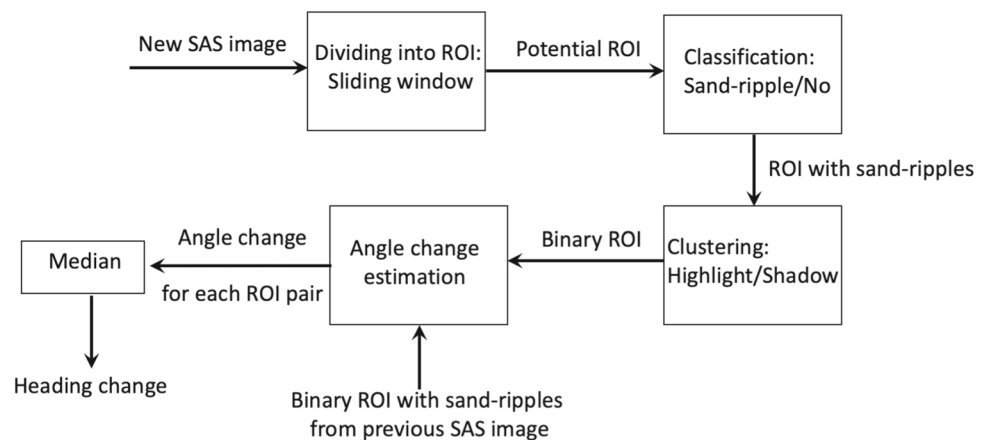


Table 1 Resulting hyperparameters for the sand-ripple classification network

Layer	Input size	Filters	Kernel size	Stride	Padding	Output size
Input	–	–	–	–	–	128×128
Conv2D, ReLU, BatchNorm, MaxPool (H=29)	128×128	13	(27,27)	(1,1)	(0,0)	$13 \times 29 \times 29$
Conv2D, ReLU, BatchNorm, MaxPool (H=27)	$13 \times 29 \times 29$	24	(21,21)	(1,1)	(0,0)	$24 \times 27 \times 27$
Conv2D, ReLU, BatchNorm, MaxPool (H=18)	$24 \times 27 \times 27$	29	(21,21)	(1,1)	(0,0)	$29 \times 18 \times 18$
Flatten	$29 \times 18 \times 18$	–	–	–	–	9396
Linear, ReLU, Dropout ($p=0.7$)	9396	–	–	–	–	9396
Linear	9396	–	–	–	–	1

the following, we describe the details of each of these three components.

5.1 Sand-ripple classification

The main challenge in classifying sand ripples is a non-homogeneous seabed, which may contain patches of sand ripples, rocks, grass, and sand in the same SAS image. An example of this is shown in Fig. 5, where - in the same SAS image - rocks, sand ripples, and sand exist. Therefore, to identify sand ripples, we employed a two-dimensional Convolutional Neural Network (LeCun et al., 1998). This architecture has a basic deep learning structure which is compatible with 2D data such as images, where nearby values (pixels) are usually related to one another. Convolutional neural networks use this assumption and learn weights and biases, which are shared with different areas of the image, assuming the same geometric features are relevant to several image areas.

Following common practice, in the proposed network architecture each convolutional layer is followed by a layer of rectified linear units (ReLU) (Agarap, 2018), a layer of batch normalization which helps in regularize training and preventing overfitting, and a max pooling layer which allows to reduce the spatial size of the representation, thus requiring fewer parameters and speeding up the computa-

tion. All network’s hyperparameters were optimized using the Optuna toolbox (Akiba et al., 2019), which is an open source hyperparameter optimization framework to automate hyperparameter search. In particular, Optuna enables users to adopt state-of-the-art Bayesian algorithms for sampling hyperparameters and pruning unpromising trials, greatly improving over more traditional hyperparameter optimization techniques based on grid search or random search. The resulting optimal configuration is reported in Table 1. Learning was carried out using the Adam optimizer (Zhang, 2018) for 100 epochs, using an initial learning rate of 0.001, and a batch size of 200.

We started from the simulated database. The data was divided into proportions of 80%, 10%, 10% for network’s training, validation, and testing, respectively, while balancing the same number of samples per class. The results of the binary classification into “with sand ripples” or “without sand-ripples” classes led to an accuracy of 98%. While this is a great result, we were suspicious about the extent to which the simulation reflects classification performance for real sonar images; and indeed, results show that, for real SAS images, classification accuracy reduces to roughly 75%. Hence, the results called for further training on real SAS images. Yet, since the majority of our SAS images had no sand-ripple areas, causing imbalanced data, we turned to transfer learning to augment the model learned from the

simulation data with real SAS images. We opted for a simple transfer learning approach, where the search for the best hyperparameters starts from those determined from the synthetic data. During the transfer learning, all weights and biases remain unfixed, so they can be further customized to real data samples. As for the synthetic data set, the data set of real SAS images was randomly uniformly divided into 80%, 10%, and 10% for training, validation and test sets, respectively.

5.2 Highlight/shadow clustering

The separation between the highlight and shadow of the sand ripple is performed using a binary fuzzy-based clustering that generates the degrees of membership, u_{ki} . That is, the probability that the label of pixel i belongs to the k -th cluster, either highlight, $\{H$ or shadow, $S\}$. The full details of this clustering are given in Abu and Diamant (2020), and are summarized here for completeness.

Our fuzzy clustering is obtained by the optimization problem

$$(U, V) = \underset{u_{ki}, V_k}{\operatorname{argmin}} J_m(U, V) \quad (2)$$

$$\text{s.t. } \sum_{k=1}^c u_{ki} = 1, \forall i,$$

where c is the number of clusters. The objective function is defined by

$$\begin{aligned} J_m(U, V) = & \sum_{i=1}^N \sum_{k=1}^c u_{ki}^m \|\Phi(\widehat{x}_i) - \Phi(V_k)\|^2 \\ & + \sum_{i=1}^N \sum_{k=1}^c u_{ki}^m \|\Phi(\bar{x}_i) - \Phi(V_k)\|^2 \\ & + \sum_{i=1}^N \sum_{k=1}^c u_{ki}^m \|\Phi(\tilde{x}_k) - \Phi(V_k)\|^2 + G_{ki}, \quad (3) \end{aligned}$$

where G_{ki} is referred to as the *fuzzy factor* (Krinidis & Chatzis, 2010), \widehat{x}_i is the sonar image, $\Phi(\bar{x}_i) - \Phi(V_k)$ is a local second moment term in the kernel space and \bar{x}_i is given by

$$\bar{x}_i = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} y_j^2, \quad (4)$$

where \mathcal{N}_i is the local window of size $(2\beta + 1)^2$ centered at pixel i , $\{V_k\}_{k=1}^c$ are the center of the clusters. The term $\Phi(\tilde{x}_k) - \Phi(V_k)$ is a *between-cluster* term in the kernel space that represents the error between the fuzzy prototypes and the

empirical cluster centers, and \tilde{x}_k is the average of all pixels assigned to the k th cluster, such that

$$\tilde{x}_k = \frac{1}{n_k} \sum_{\{j|l_j=k\}} y_j, \quad (5)$$

where n_k is the number of pixels with an assigned label that equals the k th cluster.

5.3 Evaluation of AUV heading change

Given two binarized ROIs, the estimation of the angle change is carried out by another convolutional neural network. As for the classification of sand ripples, we took a transfer learning approach, starting from the synthetic data set and fine-tuning the model using the real SAS images. As before, during the first stage, weights and biases are let free to change rather than fixed, such that customization for real data samples is possible. The input for the network consists of two binarized ROIs, and the output is the angle change. We chose a Siamese architecture that successfully worked for stereo data (Zagoruyko & Komodakis, 2015; Zbontar & LeCun, 2016), where the training phase alternates between:

1. Convolution layers applied on first ROI
2. Convolution layers applied on second ROI
3. Concatenation of features extracted from both ROIs.
4. Linear layers.
5. Backwards step on the result of phase 4.

As for the classification task, the convolutional layers included an ReLU activation function, batch-normalization, and max-pooling, and the hyperparameters were optimized using Optuna. The resulting network architecture, after transfer learning, is displayed in Table 2.

6 Results

The results of our deep learning method are compared with two geometrical approaches: the method in Williams and Coiras (2010), referred to as *Geometric1*, and the method in Williams (2012), referred to as *Geometric2*. Both benchmarks identify sand ripples by detecting periodicity in the highlight-shadow relationship. To show the need for transfer learning, we also consider another configuration of the DNN, one that is produced after training on the synthetic data only, referred to as *DNN no transfer learning*. Both versions of the DNN schemes and both geometric schemes receive as input the outcome of the clustering solution (see Sect. 5.2). All four methods are given the same manually-tagged ROIs.

Arguing that the images produced by the SAS simulator cannot effectively capture the complexity of the real under-

Table 2 Resulting hyperparameters for the angle difference regression network

Layer	Input size	Filters	Kernel size	Stride	Padding	Output size
Input	–	–	–	–	–	128 × 128
Conv2D, ReLU, BatchNorm, MaxPool (H=29)	128 × 128	11	(17,17)	(1,1)	(0,0)	11 × 29 × 29
Conv2D, ReLU, BatchNorm, MaxPool (H=21)	11 × 29 × 29	20	(17,17)	(1,1)	(0,0)	20 × 21 × 21
Flatten	2 × 20 × 21 × 21	–	–	–	–	17640
Linear, ReLU, Dropout (p=0.3)	17640	–	–	–	–	17500
Linear	17500	–	–	–	–	1

water scene, we show evaluation and testing results for the 2088 SAS images we collected with our own Eca Robotics A18-D AUV. The AUV is equipped with a Kraken-based Miniature INterferometric SAS, which, during the making of the SAS image, corrects for orientation angles (pitch, yaw and roll) as long as these are below 3 degrees. We report that this limitation was kept during our data collection. The map in Fig. 7 shows the locations of the four experiments and details the number of SAS images collected at each site. During the four experiments, the sea level was 1. Further, since the carrier frequency of our SAS system is around 300 kHz, the quality of the SAS image is not effected by sediments that exists in the water column. For ground truth, in all four experiments we used the AUV's navigation solution. This solution fuses DVL and a strategic-level inertial unit. The short dive to depth of 10–20 m until the DVL data is available ensures that the dead-reckoning drift from the GPS position is limited. In addition, in our trials we used USBL fixes from a surface vessel. Due to these fixes, we report that in all our experiments the navigation error, as measured when the vehicle surfaces, was smaller than 20 m albeit surveys of sometimes 10 hours.

6.1 Data set results of SAS images

We start by exploring the performance of the sand-ripple classification. The results in Fig. 8 show the precision-recall trade-off of the sand-ripple classification task. Here, DNN evaluation refers to the results during the evaluation step. We observe that the DNN obtains better results than the two Geometric benchmarks, especially in terms of the recall. This means that while the two Geometric schemes are able to rule out non-relevant ROIs, relative to the DNN, they miss patches of sand ripples. A similar conclusion is drawn when comparing the classification accuracy of DNN and the two Geometric schemes, i.e., the percentage of correct classification of any type to the overall number of ROIs. Here, for a choice of threshold set at the 'knee' point of the two precision-recall curves, the DNN obtained 95% accuracy while the Geometric schemes obtained 83% accuracy. We also observe that the difference between the evaluation and testing results of DNN are small. This reflects the ability of the DNN, which after




Fig. 7 A map showing the locations of the experiments made, and the number of SAS images collected at each site. A picture of the AUV during one of these experiments is also included

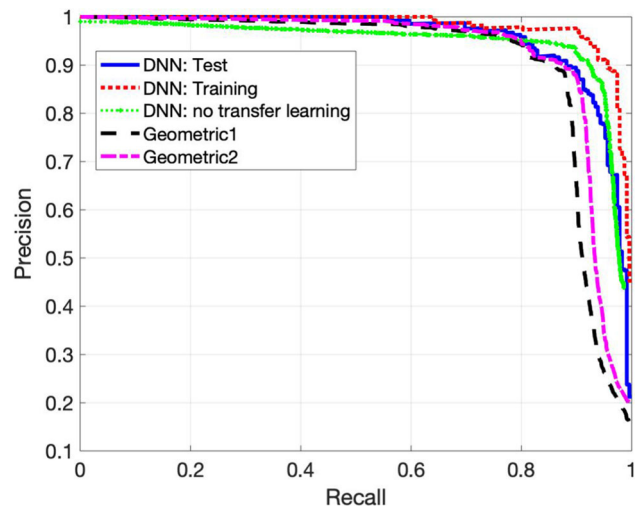


Fig. 8 Precision versus Recall for sand-ripple classification. Results are shown for 2281 ROIs including sand ripples and 9177 ROIs of type 'other'




Fig. 9 Empirical CDF for absolute error of angle prediction. Evaluation performed over 18,000 pairs of ROIs. Testing performed for 14,400 pairs of ROIs

training, was also able to capture robust parameters for different seabed types. Comparing the performance of the two DNN configurations, we observe a slight advantage for the transfer learning approach. This advantage increases for the angle estimation problem, as we show next.

Next, we explore performance in terms of angle estimation. Here, since the SAS simulator may effectively capture the angle shift observed for the sand ripple, and since the data from the real SAS images is ground-truthed by the reported heading of the AUV itself and may thus include errors, we also show prediction results of the numerical investigation. The input for the two schemes are pairs of ROIs, including only sand ripples that were randomly uniformly shifted in angle between 0 degrees and 45 degrees, to yield a Monte-Carlo set of 18,000 and 14,400 pairs of synthetic ROIs for evaluation and testing, respectively. Results in Fig. 9 are shown in terms of the empirical cumulative density function (CDF) of the absolute angle estimation error. In over more than 50% of the explored data set, the results of the two Geometric benchmarks are more than an absolute error of 5 degrees. In contrast, performance of the DNN shows an error on the order of a few degrees. For example, in 90% of the explored cases, the error was less than 2 degrees. A significant advantage is shown when using transfer learning, which reflects on the difference between the synthetic data and the real SAS images. Also notable is the almost exact match between the evaluation and testing results, which reflects the robustness of the DNN.

Figure 10 shows the results of the angle prediction for real SAS images. These results are obtained after the transfer learning stage for a training set of 88,000 ROI pairs and a testing set of 70,400 ROIs pairs. We obtained such a large data set by analyzing, for each pair of SAS images, all possible com-

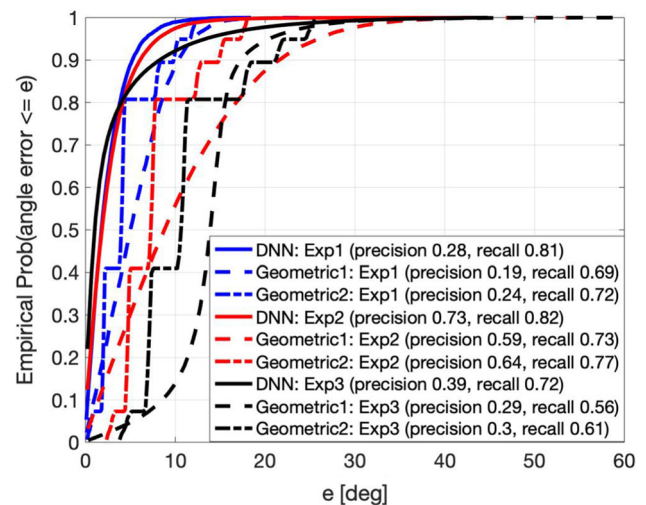


Fig. 10 Empirical CDF for absolute error of angle prediction for three sea experiments. Testing performed for 350 SAS images. Classification threshold set on 0.5

binations of the identified sand ripple ROI pairs. Then, results are obtained by taking the median out of the outcomes of all ROI pairs. To comment on the sensitivity to seabed type, we separate the testing results for the three sea experiments, each performed with a different seabed type. Specifically, Experiment 1 (Exp1) was performed on a sandy seabed at a depth of 30 m across Ashdod, Israel. Experiment 2 (Exp2) was performed along a rocky seabed at a depth of 40 m across Haifa, Israel. Experiment 3 (Exp3) was performed on a reef environment at a depth of 25 m across Acre, Israel. Results were obtained for the full analysis, including both the sand ripple ROI identification and the prediction of the angle difference. Classification thresholds for the DNN and the two Geometric schemes were obtained at the 'knee' point of the precision-recall trade-off curves in Fig. 8. The precision and recall for each method and for each experiment appears in the legend in Fig. 10. We observe a notable advantage of the DNN over the two Geometric benchmarks, which is even larger than that observed for the synthetic database. Moreover, in terms of robustness to seabed environment as reflected by comparing the results from the three experiments, we also observe an advantage of the DNN approach. We explain this by the non-homogeneity of the seabed, which is reflected as spatial differences of the sand-ripple pattern, even within the ROI. Comparing the results of the three different experiments, we note that the results deteriorate as the seabed becomes more complex. This is expected, since a complex seabed including more than sand ripples may confuse the network predictions. As a result, in absolute numbers, the error in angle difference estimation is roughly twice that shown for the synthetic data set.

In terms of the angle estimation, the performance difference between the DNN and the two benchmark schemes

may be explained by the fact that the two benchmarks are aimed to find the angle of the sand ripple while the DNN is trained to find differences between angles, much like a change detector which is arguably an easier task. Yet, also in terms of detection and false alarm rates we observe a significant improvement of our scheme over the benchmark. We thus claim that, at least for our datasets, the complexity of the sand ripples is hard to manage using a geometrical approach, while its features are well captured by the neural network.

6.2 Results of the real-time demonstration

In this section, we explore the performance of our heading change estimation scheme in a real AUV mission. This analysis tests the applicability of our approach in terms of real-time capability and as a whole system: starting from sand-ripple detection, following with highlight-shadow clustering, and ending with angle difference estimation. To this end, we implemented the two DNN schemes and the clustering algorithm over our AUV. The DNNs were implemented in Python, and the clustering algorithm in Julia. The software were operated by a bash script, running on a designated Jetson board Tx2 controller that was connected to the AUV main frame via ethernet, and was able to pull SAS images once they were formed. In particular, an SAS image is formed every 50 m the vehicle passed, which is roughly every $\Delta t = 25$ s. For a given pair of SAS images, obtained at time instances $t - \Delta t$ and t , the software was able to produce an angle difference estimation before the next image was received, i.e., in 25 s. That is, the system operated in real-time (Fig. 11).

Our trail was conducted southwest of the shores of Acre, Israel, at a water depth of approximately 25 meters. The bottom type was sandy with some rocky patches. Two examples of obtained SAS images are shown in Fig. 12. We let the AUV travel in a spiral-like motion, where every 100 m the heading angle changed by 10° . The result was a sequence of 32 SAS images, of which each pair corresponds to no heading change or to a heading change of 10° . This was verified using the AUV's Doppler velocity logger (DVL) system. We analyzed the results for forward and backward run. The former relates to analysis of images in a sequence from outside the spiral to its inner side, while the backward run performs the analysis in a reverse order. Since, from the perspective of the DNN, the angles in both directions are different, the results are not symmetric. The results of the heading change estimation for this run are shown in Fig. 11, with numbers marked in black and red indicating the angle error relative to the vehicle log in the forward and backward directions, respectively. We observe that results were not obtained for each image pair. This is because a necessary condition is the evaluation of ROI containing sand ripples in each image pair - a condition that was not always met in the explored sand-rock seabed environment. Still, out of the explored 31 image pairs,

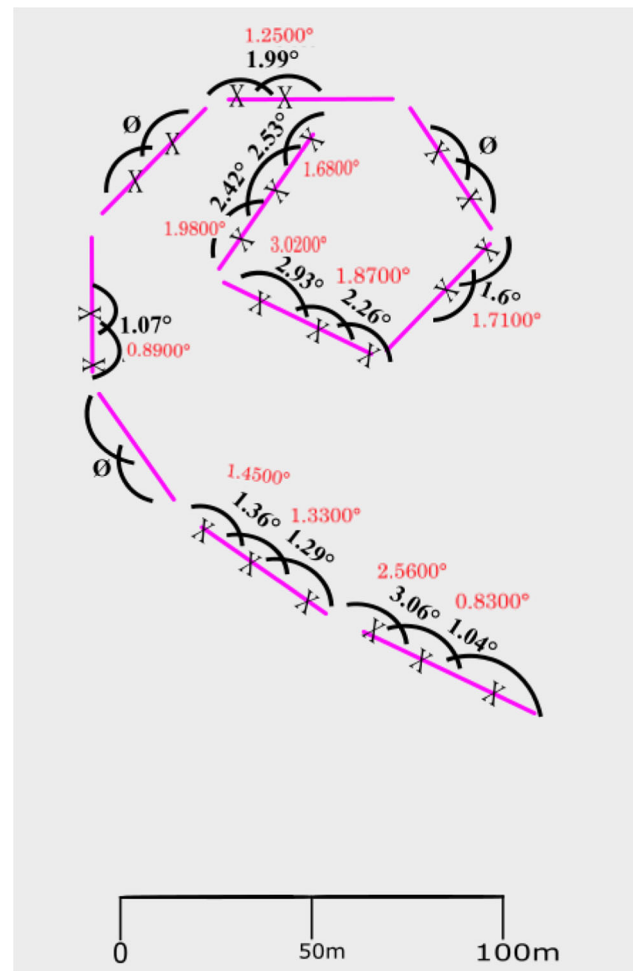



Fig. 11 Track of the AUV in the real-time experiment. Arches show the borders between pairs of SAS images with numbers indicating the absolute error of the heading change estimation, relative to the AUV log. Empty groups indicate pairs where no sand ripples were detected. Numbers marked in black and red indicate results for forward and backward run, respectively (Color figure online)

22 yielded a heading change estimation. The results show a minimum absolute error of 1.04° and 0.83° ; a maximum absolute error of 3.07° and 3.02° ; and a mean absolute error of 1.88° and 1.68° for the forward and backward directions, respectively. While these errors may seem high, we argue that they are comparable to the errors of compass or gyro-compass that are less than a strategic level. Further, the heading change prediction obtained from the structure of the sand ripples is not effected by calibration errors nor to proximity to ferromagnetic materials: interference sources that significantly impact the performance of a compass. The external information about the AUV's heading change can be used to either bound the navigation solution or to improve it by adding the information about the estimated heading change within the tracking filter. Furthermore, the results are scalable. That is, in contrast to the estimation of

Fig. 12 Examples of SAS images from the real-time sea experiment **a** Organized sand ripples. **b** Patches of sand ripples



an inertial system, ours is not affected by the magnitude of the performed heading change, and the results are expected to remain similar for both a small or large heading change.

7 Conclusions

In this paper, we explored the estimation of sand-ripple direction to evaluate a heading change of a surveying AUV in near-shore environments. In particular, assuming that on a small scale of several kilometers, the nominal direction of sand ripples is homogeneous, an observed change in the direction of sand ripples between two consecutive SAS images reflects a change in the AUV heading. Using a combination of deep learning and fuzzy-logic clustering, we detected ROIs including sand ripples, segmented them into highlight and shadow, and estimated a change in sand-ripple direction between two consecutive SAS images. Results from a large set of 2088 SAS images obtained in three different sea environments, and results from a real-time analysis, demonstrate that our method obtains high accuracy in detecting sand ripples, and accuracy of a few degrees in the estimation of heading change.

An interesting venue for future research would be to investigate the performance of other recent deep learning architectures specifically tailored for change detection tasks, such as those based on dual attention mechanisms (Chen et al., 2020). Future work should also integrate data from other sea environments as well as demonstrate how to use the extra information about the heading change within the navigation solution, and eventually extend the proposed approach to incorporate a simultaneous localization and mapping framework.

References

- Abu, A., & Diamant, R. (2020). Enhanced fuzzy-based local information algorithm for sonar image segmentation. *IEEE Transactions on Image Processing*, 29, 445–460.
- Agarap, A. F. (2018). Deep learning using rectified linear units (ReLU). arXiv preprint [arXiv:1803.08375](https://arxiv.org/abs/1803.08375)
- Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. In: Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, 2623–2631.
- Chen, J., Yuan, Z., Peng, J., Chen, L., Huang, H., Zhu, J., Liu, Y., & Li, H. (2020). Dasnet: Dual attentive fully convolutional siamese networks for change detection in high-resolution satellite images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14, 1194–1206.
- Cobb, J. T., Slatton, K. C., & Dobeck, G. J. (2010). A parametric model for characterizing seabed textures in synthetic aperture sonar images. *IEEE Journal of Oceanic Engineering*, 35(2), 250–266.
- Combettes, C., & Renaudin, V. (2015). Comparison of misalignment estimation techniques between handheld device and walking directions. In: International Conference on Indoor Positioning and Indoor Navigation, 13–16.
- Cong, W., Yongfu, S., Zhenwen, L., Qikun, Z., & Zhipeng, Z. (2019). Experimental study on the formation and migration of submarine sandwaves under wave-only and combined wave-current conditions. *Marine Georesources & Geotechnology*, 39(3), 312–321.
- Costanzi, R., Fenucci, D., Caiti, A., Fanelli, F., Monni, N., Ridolfi, A., Allotta, B. (2017). Magnetometers independent heading estimation strategy for UUV based on position and speed observations. In: OCEANS 2017 - Aberdeen, 1–5, <https://doi.org/10.1109/OCEANS.2017.8084994>.
- Crawford, A., & Skarke, A. (2014). Automatic detection of sand ripple features in sidescan sonar imagery. In: 2014 Oceans - St. John's, 1–5, <https://doi.org/10.1109/OCEANS.2014.7003117>.
- Deng, Z., Guofeng, W., Ying, H., & Di, W. (2015). Heading estimation for indoor pedestrian navigation using a smartphone in the pocket. *Sensors*, 15, 21518–21536.
- Diamant, R. (2022). Link to the dataset of sand-ripples. <https://drive.google.com/drive/folders/1TBPvrziLJqW8FZw4KzKMhyPPySIVI?usp=sharing>. Accessed: 2021 Jun 27.
- Dutta, A., Gupta, A., & Zissermann, A. (2016). VGG image annotator (VIA). <http://www.robots.ox.ac.uk/vgg/software/via/>, version: 2.0.11, Accessed: 9,6,2022.
- Fallon, M. F., Folkesson, J., McClelland, H., & Leonard, J. J. (2013). Relocating underwater features autonomously using sonar-based SLAM. *IEEE Journal of Oceanic Engineering*, 38(3), 500–513. <https://doi.org/10.1109/JOE.2012.2235664>
- Fezzani, R., Zerr, B., Mansour, A., Legris, M., & Vrignaud, C. (2019). Fusion of swath bathymetric data: Application to AUV rapid environment assessment. *IEEE Journal of Oceanic Engineering*, 44(1), 111–120. <https://doi.org/10.1109/JOE.2017.2773139>
- Fischer, P., Dosovitskiy, A., & Brox, T. (2015). Image orientation estimation with convolutional networks. In: German Conference on Pattern Recognition, Springer, 368–378.
- Fossen, T., Pettersen, K., & Galeazzi, R. (2015). Line-of-Sight path following for Dubins paths with adaptive sideslip compensation of drift forces. *IEEE Transactions on Control Systems Technology*, 23, 820–827.
- Han, R., Ortiz, M. (2014). Smartphone based gait analysis using STFT and wavelet transform for indoor navigation. In: International Conference on indoor positioning and indoor navigation.
- Heathershaw, A. (1982). Seabed-wave resonance and sand bar growth. *Nature*, 296(5855), 343–345.
- Khan, A., Ali, S. S. A., Meriaudeau, F., Malik, A. S., Soon, L. S., & Seng, T. N. (2017). Visual feedback-based heading control

- of autonomous underwater vehicle for pipeline corrosion inspection. *International Journal of Advanced Robotic Systems*, 14(3), 1729881416658171.
- Klein, I., & Diamant, R. (2018). Observability analysis of heading aided INS for a maneuvering AUV. *Navigation, Journal of the Institute of Navigation*, 65(1), 73–82.
- Klemm, R., Lehmann, B., Kraus, D. (2015). Dual-tree complex wavelet transformation based sand ripple suppression. In: OCEANS 2015 - Genova, 1–6, <https://doi.org/10.1109/OCEANS-Genova.2015.7271570>.
- Kourogi, M., & Kurata, T. (2014). A method of pedestrian dead reckoning for smart phones using frequency domain analysis on patterns of acceleration and angular velocity. in *Position* (pp. 164–168). IEEE/ION: Location and Navigation Symposium (PLANS).
- Krinidis, S., & Chatzis, V. (2010). A robust fuzzy local information c-means clustering algorithm. *IEEE Transactions on Image Processing*, 19(5), 1328–1337.
- Kunze, K., Lukowicz, P., Partridge, K., Begole, B. (2009). Which way am i facing: Inferring horizontal device orientation from an accelerometer signal. In: Proceedings of the International Symposium on Wearable Computers, Linz, Austria, 149–150.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278–2324.
- Norgren, P., & Skjetne, R. (2018). A multibeam-based SLAM algorithm for iceberg mapping using AUVs. *IEEE Access*, 6, 26318–26337. <https://doi.org/10.1109/ACCESS.2018.2830819>
- Paull, L., Saeedi, S., Seto, M., & Li, H. (2014). Auv navigation and localization: A review. *IEEE Journal of Oceanic Engineering*, 39(1), 131–149. <https://doi.org/10.1109/JOE.2013.2278891>
- Putri, T. W. O., Latifa, U., Trilaksono, B. R., & Hidayat, E. M. (2017). Modelling and identification of underwater glider for heading angle correction. *International Journal of Modeling and Optimization*, 7(2), 65–69.
- Rupeng, W., Ye, L., Teng, M., Zheng, C., Yusen, G., & Pengfei, X. (2020). Improvements to terrain aided navigation accuracy in deep-sea space by high precision particle filter initialization. *IEEE Access*, 8, 13029–13042. <https://doi.org/10.1109/ACCESS.2019.2960563>
- Schnipper, T., Mertens, K., Ellegaard, C., & Bohr, T. (2008). Amplitude equation and long-range interactions in underwater sand ripples in one dimension. *Physical Review E*, 78(4), 047301.
- Song, S., Liu, J., Guo, J., Wang, J., Xie, Y., & Cui, J. H. (2020). Neural-network-based AUV navigation for fast-changing environments. *IEEE Internet of Things Journal*, 7(10), 9773–9783. <https://doi.org/10.1109/JIOT.2020.2988313>
- Soulsby, R., Whitehouse, R., & Marten, K. (2012). Prediction of time-evolving sand ripples in shelf seas. *Continental Shelf Research*, 38, 47–62.
- Stegner, A., & Wesfreid, J. E. (1999). Dynamical evolution of sand ripples under water. *Physical review E*, 60(4), R3487.
- Stutters, L., Liu, H., Tiltman, C., & Brown, D. J. (2008). Navigation technologies for autonomous underwater vehicles. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 38(4), 581–589. <https://doi.org/10.1109/TSMCC.2008.919147>
- Tang, D., Henyey, F. S., Hefner, B. T., & Traykovski, P. A. (2009). Simulating realistic-looking sediment ripple fields. *IEEE Journal of Oceanic Engineering*, 34(4), 444–450. <https://doi.org/10.1109/JOE.2009.2025905>
- Tang, D., Williams, K. L., & Thorsos, E. I. (2009). Utilizing high-frequency acoustic backscatter to estimate bottom sand ripple parameters. *IEEE Journal of Oceanic Engineering*, 34(4), 431–443. <https://doi.org/10.1109/JOE.2009.2015402>
- Testolin, A., & Diamant, R. (2020). Combining denoising autoencoders and dynamic programming for acoustic detection and tracking of underwater moving targets. *Sensors*, 20(10), 2945.
- Testolin, A., Kipnis, D., & Diamant, R. (2020). Detecting submerged objects using active acoustics and deep neural networks: a test case for pelagic fish. *IEEE Transactions on Mobile Computing*. <https://doi.org/10.1109/TMC.2020.3044397>
- Wang, X., Zhang, G., Sun, Y., Cao, J., Wan, L., Sheng, M., & Liu, Y. (2019). AUV near-wall-following control based on adaptive disturbance observer. *Ocean Engineering*, 190, 106429.
- Williams, D. P. (2012). AUV-enabled adaptive underwater surveying for optimal data collection. *Intelligent Service Robotics*, 5(1), 33–54.
- Williams, D. P. (2015). Fast target detection in synthetic aperture sonar imagery: A new algorithm and large-scale performance analysis. *IEEE Journal of Oceanic Engineering*, 40(1), 71–92. <https://doi.org/10.1109/JOE.2013.2294532>
- Williams, D.P. (2016). Underwater target classification in synthetic aperture sonar imagery using deep convolutional neural networks. In: 2016 23rd international conference on pattern recognition (ICPR), IEEE, 2497–2502.
- Williams, D.P., & Coiras, E. (2010). On sand ripple detection in synthetic aperture sonar imagery. In: 2010 IEEE International Conference on Acoustics, Speech and Signal Processing.
- Zagoruyko, S., & Komodakis, N. (2015). Learning to compare image patches via convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition, 4353–4361.
- Zbontar, J., LeCun, Y., et al. (2016). Stereo matching by training a convolutional neural network to compare image patches. *Journal of Machine Learning Research*, 17(1), 2287–2318.
- Zhang, Z. (2018). Improved adam optimizer for deep neural networks. In: 2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS), IEEE, 1–2.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.



Hadar Shalev is a programmer in the Underwater acoustic and navigation laboratory, University of Haifa. She finished her Bsc. In 2018 and her Msc. in 2020, both in the Dep. of Computer Science, on the topic of neural networks. Currently she is working on implementing convolutional deep networks for object detection in sonar images



Liav Nagar is an electronics and software P.E. As part of his work in the Underwater acoustic and navigation laboratory, he is in-charge on the electronic design of the THEMO moorings. Currently, he is developing autonomous vehicles



Avi Abu received his Ph.D. in the underwater acoustic and navigation laboratory (ANL), Haifa University, Haifa, Israel. He received his B.Sc and M.Sc degrees from Tel-Aviv University, Tel-Aviv, Israel, in 1998 and 2003, respectively. And He is currently a fellow at the underwater acoustic and navigation laboratory (ANL), Haifa University, Haifa, Israel. His research interests include sonar image segmentation, underwater object detection and sonar image de-noisin



Alberto Testolin received the M.Sc. degree in Computer Science and the Ph.D. degree in Psychological Sciences from the University of Padova, Italy, in 2011 and 2015, respectively. He is currently an assistant professor with a joint appointment at the Department of Information Engineering and Department of General Psychology at the University of Padova. He is broadly interested in artificial intelligence, machine learning and cognitive neuroscience. His main research interests are

statistical learning theory, predictive coding, sensory perception, cognitive modeling and applications of deep learning to signal processing, networking and optimization. He is an active member of the IEEE Task Force on Deep Learning



Roe Diamant received his PhD from the Department of Electrical and Computer Engineering, University of British Columbia, in 2013, and his B.Sc. and the M.Sc. degrees from the Technion, Israel Institute of Technology, in 2002 and 2007, respectively. From 2001 to 2009, he worked in Rafael Advanced Defense Systems, Israel, as a project manager and systems engineer, where he developed a commercial underwater modem with network capabilities. In 2015 and 2016, he was a visiting Prof. at the

University of Padova, Italy. In 2009, he received the Israel Excellent Worker First Place Award from the Israeli Presidential Institute. In 2010, he received the NSERC Vanier Canada Graduate Scholarship. Dr. Diamant has received three Best Paper awards, and serves as an associate editor for the IEEE Journal of Ocean Engineering. He is the coordinator of the EU H2020 project SYMBIOSIS (BG-14 track), and leads the underwater Acoustic and Navigation Laboratory (ANL) as an Assist. Prof. at the Dept. of Marine Technologies, University of Haifa. His research interests include underwater acoustic communication, underwater navigation, object detection, and classification