



ICUA2022
International Conference on
Underwater Acoustics

International Conference of Underwater Acoustics

Southampton, UK

20-23 June 2022



**Underwater Acoustics: Underwater acoustic detection
& classification & clearance of unexploded ordnance**

Reducing the false alarm rate of a simple sidescan sonar change detection system using deep learning

Yannik Steiniger

Department of Situational Awareness, German Aerospace Center DLR Institute for the Protection of Maritime Infrastructures, Bremerhaven, Bremen, 27572, GERMANY; yannik.steiniger@dlr.de

Sven Schröder and Jannis Stoppe

German Aerospace Center DLR Institute for the Protection of Maritime Infrastructures, Bremerhaven, Bremen, 27572, GERMANY; sven.schroeder@dlr.de, jannis.stoppe@dlr.de

Detecting changes on the sea floor between two sonar images is a challenging but important task. Typical image based change detection methods are prone to have a high false alarm rate. In this work we introduce a change detection processing chain which uses a convolutional neural network (CNN) to classify potential detections into the two classes object and non-object. Simulated data as well as real sidescan sonar images are used to analyse the proposed method. A new simulation pipeline based on Blender and image transformations is introduced to generate the synthetic dataset. Using the classification by the CNN as a filtering technique a reduction of the false alarm rate is achieved. In the most complex scenario the number of false alarms can be reduced by 57.5% compared to a filtering method based on the histogram of the pixel intensities inside the snippets.

1. INTRODUCTION

When searching for objects on the sea floor image sonar systems like sidescan sonar (SSS) or synthetic aperture sonars (SAS) are used. They are typically mounted on an autonomous underwater vehicle (AUV) to enable an automatic capturing of the data, i.e., scans of the sea floor. Using a special processing the data captured by the sonar system can be displayed as a georeferenced image. Since over the course of a scanning mission a large amount of data is automatically captured, an automatic processing to aid the operator interpreting the images is essential. Several research works have focused on implementing methods from the computer vision domain (mainly neural networks) for the detection [1–4], segmentation [5–7] and classification of sonar images [8–13].

If two images of the same location are captured at different time instances, information about changes in the environment can be extracted. This technique is known as change detection and can for example be used to detect sunken or manually placed objects in the later image. In general, change detection algorithms are divided into image-based and symbolic-based methods [14]. Image-based change detection directly compares two images while in symbolic-based change detection objects are detected in both images separately, georeferenced and then compared [15–17]. Image-based change detection using SAS can further be divided into non-coherent and coherent change detection, where the former considers only the magnitude information and the latter also the phase information [18]. Our work deals with SSS data and thus an image-based non-coherent change detection processing chain as shown in Figure 1 is applied. This change detection consists of the typical steps: image alignment, subtraction and detection [19, 20]. However, the result is prone to have a high false alarm rate caused by changes of the seafloor, noise or poor image alignment. Thus, the detections need to be filtered in order to reduce the number of false alarms. In our work, we propose to utilise a convolutional neural network (CNN) for this task. The CNN is trained to classify sonar snippets into the two classes object and non-object. To the best of our knowledge, deep learning methods like CNNs have not yet been considered for change detection in SSS or SAS images.

Furthermore, to account for the fact that sonar images suited to test change detection algorithms are hard to acquire, we set up a simulation pipeline based on Blender. Several transformations are applied to the image generated by Blender to generate a more realistic look. Objects can be added or removed and image transformations, like translation and rotation, be applied to simulate an image from a second mission with changes present. To this end, the baseline and deep learning based change detection processing chain are compared using simulated as well as real SSS images.

Our work comprises the following contributions:

- We propose a simple and efficient simulation pipeline to generate SSS images, which can be used to analyse the performance of a detector.
- We introduce the usage of deep learning for the task of non-coherent image-based change detection.
- We show that the deep learning based change detection significantly reduces the number of false alarms compared to a filtering without deep learning.

The remaining of this paper is organised as follows. Section 2 introduces the change detection processing chain in more detail. Next, in Section 3 the simulation of SSS images is explained. The real and simulated datasets are described in Section 4. In Section 5 the deep learning based change detection is compared to the baseline processing chain. Finally, the paper closes with a summary and outlook on future work in Section 6.

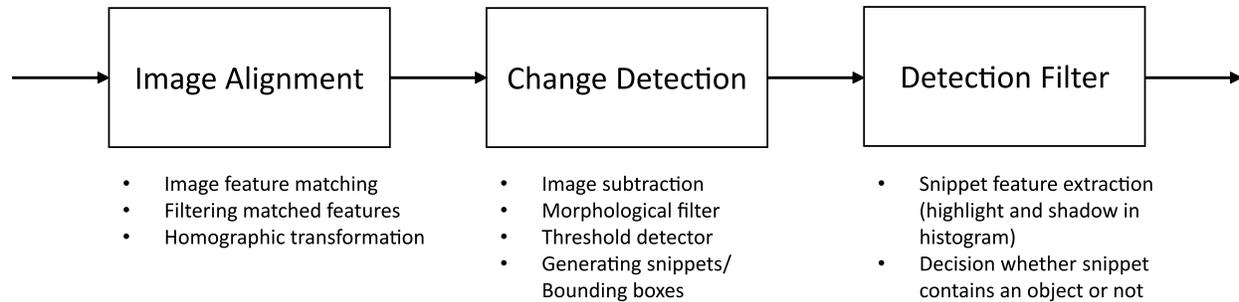


Figure 1: Structure of the baseline change detection.

2. CHANGE DETECTION PROCESSING CHAIN

A. BASELINE

The basis for performance comparisons of the methods proposed in this paper is a conventional change detection algorithm, which was developed at our institute for the automatic detection of changes between two SSS images. The structure of the algorithm can be divided into three stages, which are shown in Figure 1. Since the accuracy of navigation data of the platform on which the SSS is mounted is limited, differences (translation, rotation) between two sonar images usually occur. Therefore, the images must be aligned in the first stage. For this purpose, a feature matching based on openCV is performed. This provides a list with translation vectors of individual features, from which a homographic transformation matrix H can be generated. Before this, however, feature mismatches are filtered out. By a homographic transformation of the input image with the matrix H the output image of the first stage is generated.

The second stage contains a pixel-based change detection between the aligned input images. First, both images are smoothed using a median filter to reduce the present noise and normalised. A subtraction between the sonar images is determined to form the difference image. From this, regions in which changes have occurred are detected using a threshold detector and marked by bounding boxes. Misalignment or intensity differences between the input images can thereby lead to a high false positive rate. However, since in our use case the only interest is to detect changes in the position and number of objects, a filtering of the detections is applied in the third stage to reduce the false positive rate.

Here, a simple feature extraction is performed on the previously detected snippets to distinguish between object (true-positive) and non-object (false-positive). The presence of an object highlight and object shadow in the respective snippet is used as an object-specific feature. If these can be detected in the histogram of the snippet, it is classified as an object and passes the filter. More precisely, if there are no two distinct peaks in the histogram, indicating in the presence of a highlight and shadow area, the detection is removed.

B. DEEP LEARNING

The robustness of the filtering in the baseline change detection highly depends on the tuning of several parameters. Furthermore, past research has shown that CNNs are better suited for the classification of sonar snippets than hand-crafted features [21–23]. Thus, in our proposed deep learning based change detection processing chain we replace the filtering with a CNN trained for the classification of sonar snippets into object and non-object. The remaining processing chain stays as described before.

We use three different configurations of a CNN which differ in the number of layers to investigate the influence of the depth on the change detection performance. Table 1 specifies the three architectures. The configuration S uses only one convolutional layer while M and L use two and four, respectively. All three architectures use the same fully connected network with dropout to prevent overfitting.

Table 1: Architecture of the CNNs used for classification.

CNN	Layer	Kernel	Kernel size	Activation
S, M, L	Conv2D	16	3×3	relu
S, M, L	BatchNorm	-	-	-
L	Conv2D	16	3×3	relu
L	BatchNorm	-	-	-
S, M, L	MaxPooling	-	2×2	-
M, L	Conv2D	32	3×3	relu
M, L	BatchNorm	-	-	-
M, L	MaxPooling	-	2×2	-
L	Conv2D	32	3×3	relu
L	BatchNorm	-	-	-
L	MaxPooling	-	2×2	-
S, M, L	Dropout	-	-	-
S, M, L	Dense	100	-	relu
S, M, L	Dropout	-	-	-
S, M, L	Dense	2	-	softmax

3. SIMULATION

A. SIDE SCAN SONAR IMAGE GENERATION

The collection of SSS data is costly and time consuming. In addition, scenarios in which the same region is covered several times by SSS and changes in the location and number of objects on the seafloor between the acquired sonar images are rare. This results in a small dataset for testing change detection algorithms: Thus, we introduce a simulation pipeline for SSS images through which the proposed methods can be tested efficiently.

The simulation presented within this paper is based on optical ray tracing carried out in Blender. For this purpose, an arbitrarily complex scenario is created, consisting of a seabed and objects, e.g., tires and stones, distributed on it. The scenario is illuminated from the side by a light source, which emits plane light waves at an arbitrary angle. An orthographic projection camera is placed above the scenario, which generates the output of the optical ray tracing. In this way, typical features of a side scan sonar image such as acoustic highlights and shadows from arbitrary surface structures can be simulated.

Figure 2 shows the setup in Blender as well as an exemplary output image. In addition to the intensity of the light scattering, the normal vector to the light source and the (x, y) coordinates are stored for each pixel in the output. The result of the optical ray tracing is a high resolution image without noise, which is not yet a realistic simulation of a sonar image. In SSS images the intensity of a pixel is affected by neighbouring scatterers, because the angular resolution is limited by the aperture of the array and the range resolution is limited by the frequency bandwidth of the system. An ideal point scatterer would therefore be smeared onto the neighbouring pixels in the sonar image, depending on the characteristics of the system used, and would represent the so-called point spread function (PSF) of the system. This property is therefore to be added to the sonar image in a second step. For this purpose the PSF of a SSS system is generated in Python (cf. Figure 3.a) which is then convolved with the output from Blender. Furthermore, the contrast of highlight

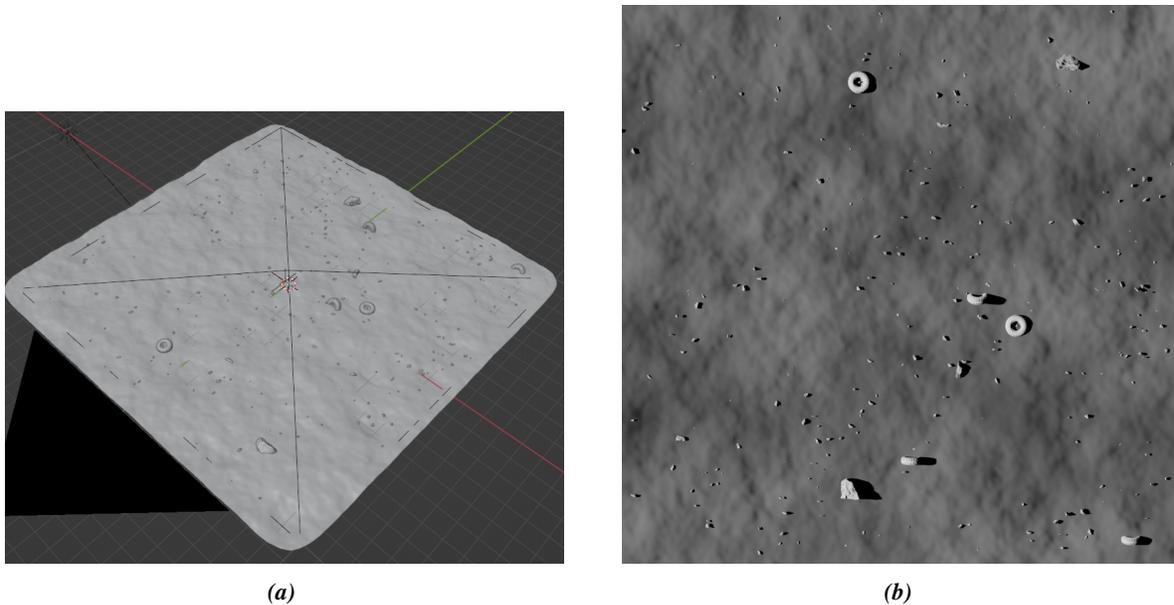


Figure 2: Simulation in Blender. (a) Scenario setup. (b) Generated output image.

regions in the image is increased by using the normal vector information of each pixel in order to simulate the strong backscattering from surfaces, which are orthogonal to the sonar system. To get further realism, noise is added to the simulated sonar image and a downscaling to an adjustable spatial resolution (e.g., 10×10 cm pixel size) is performed to ensure comparability to the experimental data. The whole simulation chain is shown in Figure 3.b, whereas the final output of the simulation is shown in Figure 4.

4. DATASETS

A. SIDESCAN SONAR DATA

Over the course of several sea and harbour expeditions, we have collected SSS data with the SeaCat AUV. The Edgetech 2205 sidescan sonar mounted on the SeaCat AUV operates at a centre frequency of 850 kHz with a bandwidth of 45 kHz. An experimental signal processing chain is used to generate the sonar images with a pixel resolution of 10 cm. In Figure 5 an example of a processed mosaic image is shown. We consider sonar images from three expeditions in this work. The images SSS1 and SSS2 were captured in a lake in Bremen, Germany, on two identically planned missions. The data for SSS3 was collected in a harbour in Bremerhaven, Germany. Finally, SSS4 and SSS5 are two images from two missions in a second harbour in Bremerhaven. Here again the missions were planned identically. Only the missions for SSS4 and SSS5 were designed as a change detection experiment, where an object was manually places on the sea floor. In order to increase the number of change detection experiments we synthetically insert objects from a reference sonar image into SSS1 and SSS3. These two images are denoted as SSS1* and SSS3* in the following. Figure 5 displays the images SSS1 and SSS1*.

From the captured and processed SSS images, five experiments are set up which are summarised in Table 4. By comparing SSS1 and SSS1* the number of detections and false alarms for a perfect aligned image can be studied. In Experiment 2 the change detection chain should not detect any objects. Compared to Experiment 1 the third experiment studies the case where the alignment is not perfect. This better reflects a real mission where, e.g., due to current, the tracks of the AUV are not identical. In Experiment 4 the image SSS3* was processed with slightly different settings to investigate the effect of changes in the image

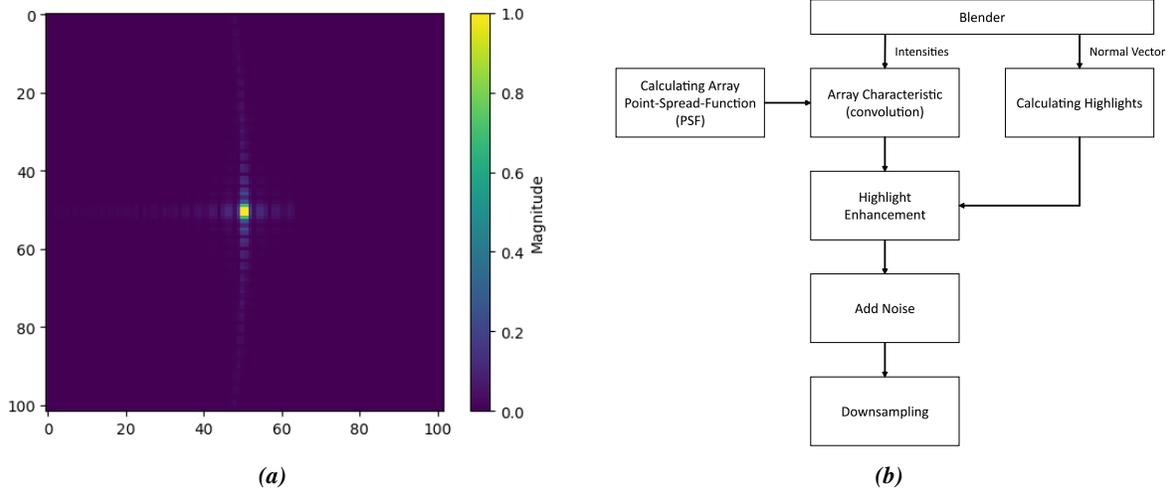


Figure 3: (a) Generated PSF to simulate the SSS antenna. (b) Overall processing chain of the simulation.

Table 2: Conducted change detection experiments with real SSS images.

Experiment	Pre-image	Post-image	Number of changes	Goal
Exp 1	SSS1	SSS1*	5	detect inserted objects in the same image
Exp 2	SSS1	SSS2	0	no false alarms
Exp 3	SSS2	SSS1*	5	detect inserted objects in a different post-image
Exp 4	SSS3	SSS3*	5	detect inserted objects in a different post-image
Exp 5	SSS4	SSS5	2	detect objects in a realistic scenario

generation chain between the two captured images on the change detection performance. Experiment 5 is the planned change detection experiment with one object being present in the first image which was then moved when capturing the second one. This is the most realistic and complicated test scenario.

B. SIMULATED DATA

Because a large dataset for change detection application is hard to obtain, simulated data is used additionally. The purpose of this simulated dataset is to investigate the performance of the change detection chains on a large number of changes. The simulated data consists of three main components: seabed, unchanged objects and changed objects. The structure of the seabed is created by a wavy surface on which smaller stones are placed randomly. As in the real data, tires and larger stones are used as objects. For the following experiments, five simulated scenarios are created, each with a pre and post image. Thereby, three to four objects are added to the post image. In addition, the image is rotated and translated, and the angle of the light source is changed, which results in a change of the object shadows. These modifications are necessary to reproduce the navigation inaccuracy of an AUV and to create a misalignment of the sonar images. This ensures comparability to experimental data and reflects the influence of the image alignment.

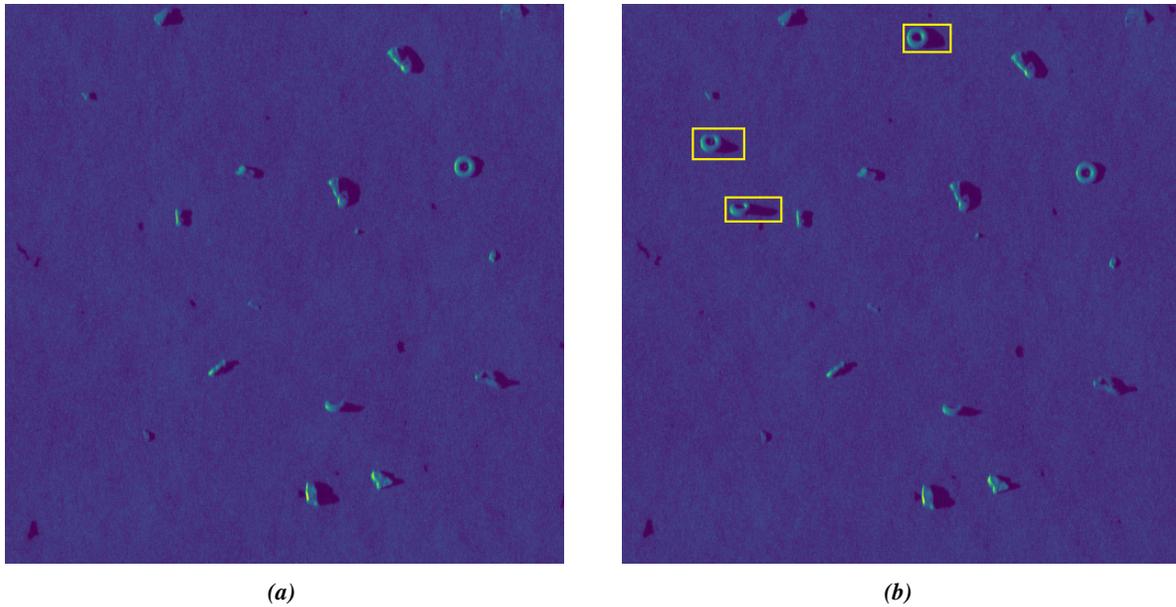
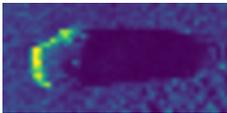
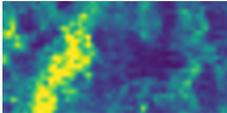


Figure 4: Example outputs of the simulation pipeline. (a) Serving as pre-image in experiment Sim 1. (b) Serving as post-image in experiment Sim 1.

Table 3: Examples and number of sidescan sonar snippets in the training and test set.

Class	Example	Number of training snippets	Number of test snippets
object		2,197	9
non-object		2,145	12

C. CNN TRAINING DATASET

In order to train the CNN for classification between object and non-object we use a dataset of snippets from previous experiments [3]. Those snippets are extracted from so called waterfall images where the acoustic shadow of an objects is always pointed to the starboard side. Because the change detection is carried out in mosaic images, the shadow can point in an arbitrary direction. Thus, the available dataset is heavily augmented using rotation of the snippets. When a snippet is samples for training it is rotated by a random angle in the range of $[0^\circ, 360^\circ)$. The number of training and test samples are given in Table 3. Note that the objects inserted to the mosaic images to form the post-image SSS1* and SSS3* are contained in the test set so they are not used for training.

The CNNs described in Table 1 are trained for ten epochs using the Adam optimiser. We set the learn rate to 0.001 and the batch size to 32. All networks have an image input size of 64×64 pixel. Resizing to this shape and normalization of the pixel intensities to the range $[0, 1]$ using min-max-scaling is applied to all sonar snippets prior to training. The training progress in terms of training loss and training accuracy is shown in Figure 6. Around ten epochs the loss and accuracy start to saturate, indicating that this amount of

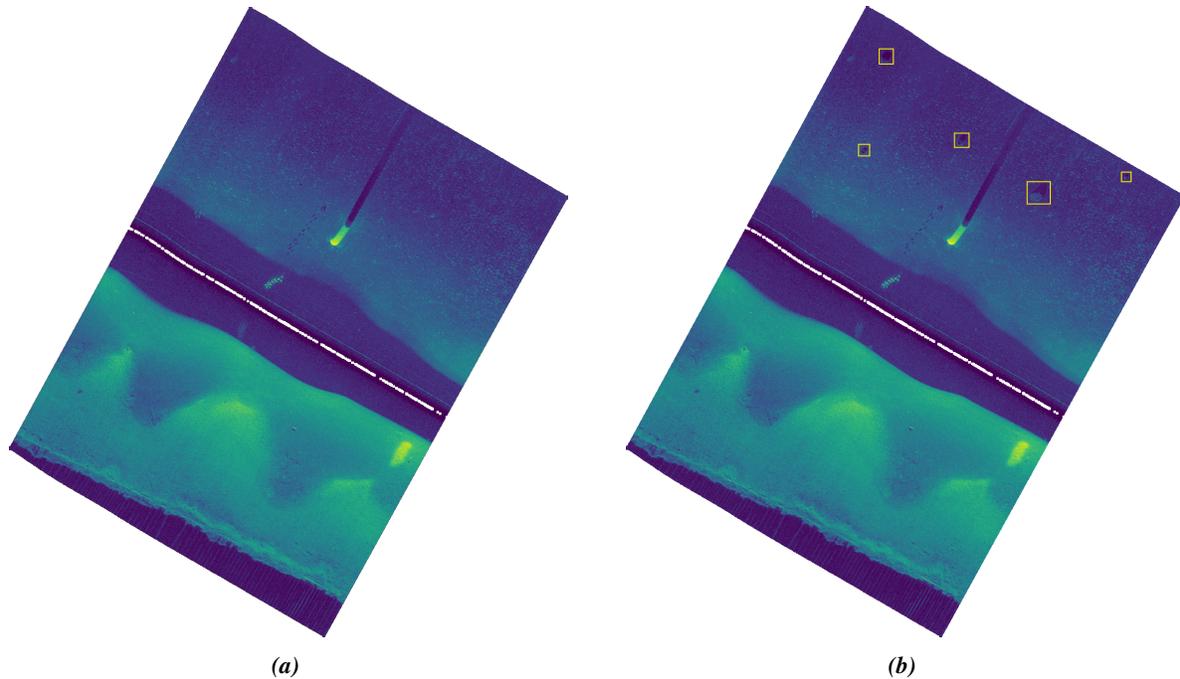


Figure 5: Example of a sidescan sonar image. (a) Mosaic image SSSI serving as a pre-image. (b) Mosaic image SSSI* with inserted objects serving as a post-image. Inserted objects are marked with a green bounding box.

training is sufficient. On the test data all networks reach a accuracy of at least 95%. Note however, that we are more interested in the reduction of the false alarm rate of the whole change detection processing chain which is analysed in the following.

5. EXPERIMENTAL RESULTS

The performance of the baseline change detection without filtering (CD), with histogram based filtering (Hist) and the deep learning based processing chain is measured using the number of true-positives (TP) and false-positives (FP). We calculate the distance from the centre of the ground truth (GT) and predicted bounding box and count a TP if this distance d is smaller than 1 m. Table 4 summarises these results for all experiments.

In all experiments with simulated data a similar or reduced number of false alarms is achieved by the deep learning based change detection. Only in Simulation 2 it produces up to 3 more false alarms compared to the baseline. The missed detections in Simulation 4 and 5 for the CNNs are due to missing detections in the previous detection step (see Figure 1) which cannot be recovered. In the baseline two ground truth objects are erroneously filtered.

Experiment 1 used the exact same sonar image, whereby objects are placed manually in the post image, to reflect the case of a perfectly identical navigation between both missions. Thus, the alignment is perfect and no false alarm should be generated. However, the change detection generated two bounding boxes close to each other for one object resulting in one false alarm. This is filtered by all three CNNs. With the histogram based filtering two ground truth contacts are filtered. CNN_S does not remove any ground truth object.

In Experiment 2 the two images come from two different runs and no object is present. The detection step produces 122 contacts which are reduced to 41 in the baseline processing chain with histogram based filtering. All CNNs reduce this number even further to 6 for CNN_S and 4 for CNN_M and CNN_L .

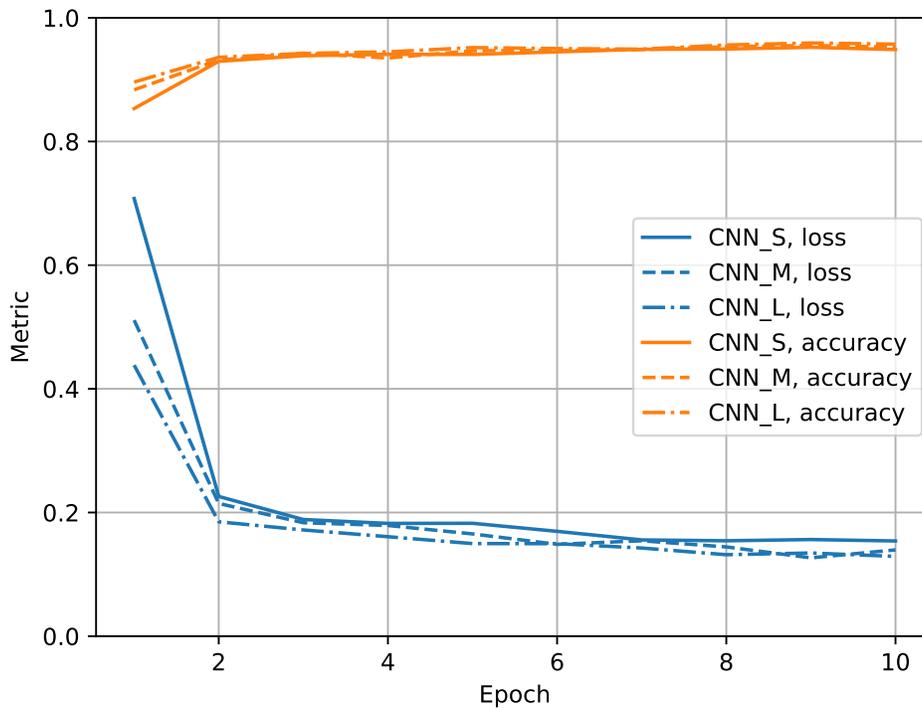


Figure 6: CNN training loss and accuracy during training.

In Experiment 3, 4 and 5 some ground truth objects are not detected prior to the filtering. Due to the preprocessing and inaccuracies in the alignment of the images our change detection chain struggles to detect small objects. Furthermore, in some cases only the highlight or shadow area of an object is detected. However, the ground truth bounding box encloses both areas of an object leading to $d > 1$ m which, at the end, results in a missed detection and a false alarm.

Histogram based filtering reduces the number of false alarms from 115 to 34 in Experiment 3. Again, all CNNs show a better reduction leading to 16 false alarms for CNN_S and CNN_M and 14 for CNN_L . The baseline change detection has less false alarms in Experiment 4 but at the same time also filters all ground truth objects. Several parameters, like the minimal peak distance in the histogram, influence the outcome baseline change detection. This result shows that a careful setting of these parameter is necessary to achieve a sufficient performance, since no ground truth object should be filtered. Deep learning based change detection is less sensitive to such parameter settings once a good network is trained.

A clear improvement in terms of false alarm reduction is observed for the most realistic test case in Experiment 5. After histogram based filtering 414 of the 664 false alarms remain. With CNN_L this can be reduced to 176. Investigating the remaining false alarms shows that often a highlight-shadow pattern is present. Furthermore, the predicted probability of belonging to class object is close to 0.5 indicating a high uncertainty. Some examples of false alarms generated by CNN_L are presented in Figure 7.

By sweeping through the threshold for object classification a receiver operating characteristic (ROC) curve can be generated. These curves can further be compared using the area under curve (AUC). Figure 8 displays the ROC curves for the three CNNs for Experiment 3 and 4. Other experiments either have a true-positive-rate or a false-positive-rate of zero and thus a non-meaningful ROC curve. All three CNNs show a similar performance with CNN_L being slightly better in Experiment 3 and CNN_M in Experiment 4. Thus, the size of the network does not have a large influence on the performance our analysis. However, it is expected that deeper networks would benefit more from additional training data, e.g., by including typical false alarm structures as shown in Figure 7.

Table 4: Performance of the change detection without filtering (CD), with histogram based (Hist) and with CNN based filtering in terms of true-positives (TP) and false-positives (FP). The second column states the number of ground truth (GT) objects in each experiment. Best results in terms of false alarm reduction are marked in bold.

Experiment	GT	TP					FP				
		CD	Hist	CNN _S	CNN _M	CNN _L	CD	Hist	CNN _S	CNN _M	CNN _L
Sim 1	3	3	3	3	3	3	23	18	15	16	15
Sim 2	3	3	3	3	3	3	25	14	16	16	17
Sim 3	4	4	3	4	4	4	43	23	17	18	21
Sim 4	3	2	1	2	2	2	33	17	13	14	14
Sim 5	3	1	1	1	1	1	13	12	5	5	6
Exp 1	5	5	3	5	4	4	1	1	0	0	0
Exp 2	0	0	0	0	0	0	122	41	6	4	4
Exp 3	5	4	3	3	2	2	115	34	16	16	14
Exp 4	5	3	0	2	2	1	237	24	52	52	35
Exp 5	2	0	0	0	0	0	664	414	195	201	176

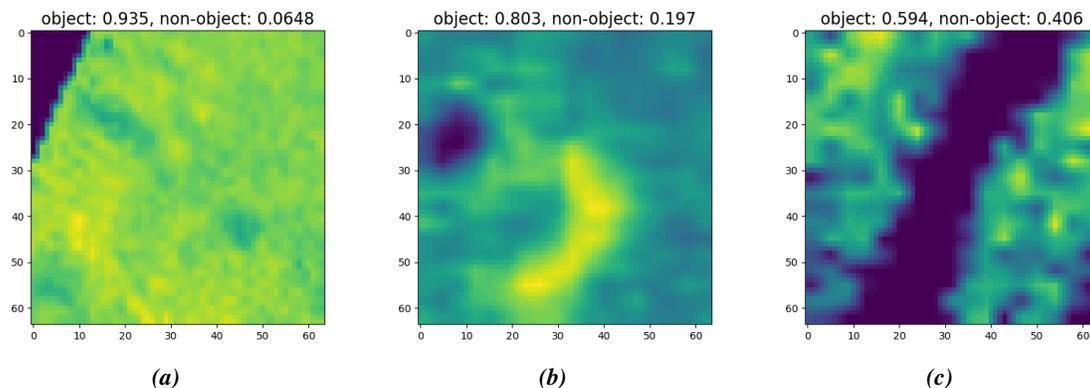


Figure 7: False alarms and classification probability of CNN_L in (a) Experiment 3, (b) Experiment 4 and (c) Experiment 5.

6. SUMMARY

This work has introduced CNNs for image based non-coherent change detection in SSS images. To test the deep learning based change detection processing chain, simulated data was generated using Blender and further transformations. Comparisons with a baseline which filters detections based on the histogram of pixel intensities, on simulated and real data shows a better performance for the CNN. In the most complex scenario our deep learning based change detection reduces the number of false alarms from 664 to 176 while for the baseline 414 false alarms are present which corresponds to a relative improvement of 95.2%.

More training data is needed to improve the performance of the CNN even further. In addition, the alignment of the two images is a critical step in the processing. Misalignment does not only leads to false alarms but also missed detection which cannot be recovered by our approach. In future work we will investigate symbol based change detection by applying deep learning based detectors on both images individually. By georeferencing and comparing these detections the need for a proper alignment can be avoided.

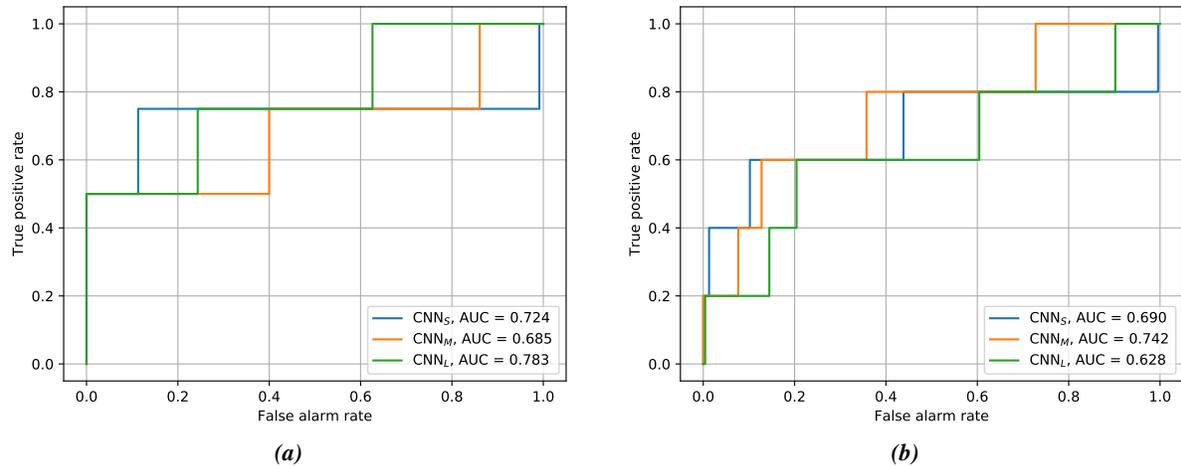


Figure 8: Comparison of the ROC curves for the three CNN. (a) Performance on Experiment 3. (b) Performance on Experiment 4.

REFERENCES

- [1] D. Einsidler, M. Dhanak, and P.-P. Beaujean, “A deep learning approach to target recognition in side-scan sonar imagery,” in *OCEANS 2018 MTS/IEEE Charleston*. IEEE, 2018, pp. 1–4.
- [2] H. T. Le, S. L. Phung, P. B. Chapple, A. Bouzerdoum, C. H. Ritz, and L. C. Tran, “Deep gabor neural network for automatic detection of mine-like objects in sonar imagery,” *IEEE Access*, vol. 8, pp. 94 126–94 139, 2020.
- [3] Y. Steiniger, J. Groen, J. Stoppe, D. Kraus, and T. Meisen, “A study on modern deep learning detection algorithms for automatic target recognition in sidescan sonar images,” in *Proceedings of the 6th Underwater Acoustics Conference and Exhibition (UACE)*, ser. Proceedings of Meetings on Acoustics. Acoustical Society of America, 2021, p. 070010.
- [4] Y. Yu, J. Zhao, Q. Gong, C. Huang, G. Zheng, and J. Ma, “Real-time underwater maritime object detection in side-scan sonar images based on Transformer-YOLOv5,” *Remote Sensing*, vol. 13, no. 18, p. 3555, 2021.
- [5] J. L. Chen and J. E. Summers, “Deep convolutional neural networks for semi-supervised learning from synthetic aperture sonar (SAS) images,” in *173rd Meeting of Acoustical Society of America and 8th Forum Acusticum*, ser. Proceedings of Meetings on Acoustics. Acoustical Society of America, 2017, p. 055018.
- [6] K. Li, F. Yu, Q. Wang, M. Wu, G. Li, T. Yan, and B. He, “Real-time segmentation of side scan sonar imagery for AUVs,” in *2019 IEEE Underwater Technology*. IEEE, 2019, pp. 1–5.
- [7] Y. Song, B. He, and P. Liu, “Real-time object detection for AUVs using self-cascaded convolutional neural networks,” *IEEE Journal of Oceanic Engineering*, vol. 46, no. 1, pp. 56–67, 2021.
- [8] N. D. Warakagoda and Ø. Midtgaard, “Transfer-learning with deep neural networks for mine recognition in sonar images,” in *Proceedings of the 4th International Conference Synthetic Aperture Sonar Synthetic Aperture Radar*, vol. 40, 2018, pp. 115–122.

-
- [9] A. Bouzerdoum, P. B. Chapple, M. Dras, Y. Guo, L. Hamey, T. Hassanzadeh, H. T. Le, O. Nezami, M. Orgun, S. L. Phung, C. H. Ritz, and M. Shahpasand, "Improved deep-learning-based classification of mine-like contacts in sonar images from autonomous underwater vehicle," in *Proceedings fo the 5th Underwater Acoustics Conference and Exhibition (UACE)*, 2019, pp. 179–186.
- [10] X. Qin, X. Luo, Z. Wu, and J. Shang, "Optimizing the sediment classification of small side-scan sonar images based on deep learning," *IEEE Access*, vol. 9, pp. 29 416–29 428, 2021.
- [11] D. P. Williams, "On the use of tiny convolutional neural networks for human-expert-level classification performance in sonar imagery," *IEEE Journal of Oceanic Engineering*, vol. 46, no. 1, pp. 236–260, 2021.
- [12] I. D. Gerg and V. Monga, "Structural prior driven regularized deep learning for sonar image classification," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, pp. 1–16, 2022.
- [13] Z. Cheng, G. Huo, and H. Li, "A multi-domain collaborative transfer learning method with multi-scale repeated attention mechanism for underwater side-scan sonar image classification," *Remote Sensing*, vol. 14, no. 2, p. 355, 2022.
- [14] F. Nicolas, A. Arnold-Bos, I. Quidu, and B. Zerr, "Symbolic simultaneous registration and change detection between two detection sets in the mine warfare context," in *OCEANS 2019 MTS/IEEE Marseille*. IEEE, 2019.
- [15] E. Coiras, J. Groen, D. Williams, B. Evans, and M. Pinto, "Automatic change detection for the monitoring of cluttered underwater areas," in *Proceedings of the 1st International Conference & Exhibition on Waterside Security*, 2008, pp. 99–105.
- [16] M. Gendron, M. Lohrenz, and J. Dubberley, "Automated change detection using synthetic aperture sonar imagery," in *OCEANS 2009*. IEEE, 2009, pp. 1–4.
- [17] J. Ferrand and N. Mandelert, "Change detection for mcm survey mission," in *Proceedings of the 2012 International Conference on Detection and Classification of Underwater Targets*, V. Myers, Ed. Cambridge Scholars Publishing, 2014, pp. 193–206.
- [18] V. Myers, I. Quidu, B. Zerr, T. O. Sabo, and R. E. Hansen, "Synthetic aperture sonar track registration with motion compensation for coherent change detection," *IEEE Journal of Oceanic Engineering*, vol. 45, no. 3, pp. 1045–1062, 2020.
- [19] V. Myers, A. Fortin, and P. Simard, "An automated method for change detection in areas of high clutter density using sonar imagery," in *Proceedings of the 3rd International Conference and Exhibition on Underwater Acoustic Measurements*, 2009, pp. 287–294.
- [20] Ø. Midtgaard, R. E. Hansen, T. O. Saebo, V. Myers, J. R. Dubberley, and I. Quidu, "Change detection using synthetic aperture sonar: Preliminary results from the larvik trial," in *OCEANS 2011*. IEEE, 2011, pp. 1–8.
- [21] J. McKay, I. D. Gerg, V. Monga, and R. G. Raj, "What's mine is yours: Pretrained cnns for limited training sonar atr," in *OCEANS 2017 MTS/IEEE Anchorage*. IEEE, 2017, pp. 1–7.
- [22] D. P. Williams, "Demystifying deep convolutional neural networks for sonar image classification," in *Proceedings of the 4th Underwater Acoustics Conference and Exhibition (UACE)*, 2017, pp. 513–520.
- [23] G. Huo, Z. Wu, and J. Li, "Underwater object classification in sidescan sonar images using deep transfer learning and semisynthetic training data," *IEEE Access*, vol. 8, pp. 47 407–47 418, 2020.