Early Identification of Oil Spills in Satellite Images Using Deep CNNs

Marios Krestenitis[1], Georgios Orfanidis[1], Konstantinos Ioannidis[1], Konstantinos Avgerinakis[1], Stefanos Vrochidis[1], and Ioannis Kompatsiaris[1]

Centre for Research & Technology Hellas, Information Techologies Institute, Thessaloniki, Greece
{mikrestenitis,g.orfanidis,kioannid,koafgeri,stefanos,ikom}@iti.gr

Abstract. Oil spill pollution comprises a significant threat of the oceanic and coastal ecosystems. A continuous monitoring framework with automatic detection capabilities could be valuable as an early warning system so as to minimize the response time of the authorities and prevent any environmental disaster. The usage of Synthetic Aperture Radar (SAR) data acquired from satellites have received a considerable attention in remote sensing and image analysis applications for disaster management, due to the wide area coverage and the all-weather capabilities. Over the past few years, multiple solutions have been proposed to identify oil spills over the sea surface by processing SAR images. In addition, deep convolutional neural networks (DCNN) have shown remarkable results in a wide variety of image analysis applications and could be deployed to overcome the performance of previously proposed methods. This paper describes the development of an image analysis approach utilizing the benefits of a deep CNN combined with SAR imagery to establish an early warning system for oil spill pollution identification. SAR images are semantically segmented into multiple areas of interest including oil spill, look-alikes, land areas, sea surface and ships.The model was trained and tested using multiple SAR images, acquired from the Copernicus Open Access Hub and manually annotated. The dataset is a result of Sentinel-1 missions and EMSA records for relative pollution events. The conducted experiments demonstrate that the deployed DCNN model can accurately discriminate oil spills from other instances providing the relevant authorities a valuable tool to manage the upcoming disaster effectively.

Keywords: Oil spill identification, SAR image analysis, Deep Convolutional Neural Networks, Disaster management.

1 Introduction

Oil slicks have a significant impact to the ocean and coastal environments as well as to maritime commerce and activities. Early measurement is crucial in such cases to manage the disaster and prevent further environmental damage. Toward this direction, various algorithms and approaches have been presented

to identify automatically oil polluted areas over the sea surface. Most of these methods process satellite data and apply various remote sensing principles.

Considering the main objective of an early warning system, the accurate identification of oil slicks could assist the relevant authorities to have a more complete overview of the event. A wider dispersion of oil slicks on sea surfaces will result to major environmental problems not only to the maritime environment but also in coastal territories if its detection time is significantly large. Reversing the posed problem, a framework that provides a better understanding of the oil polluted areas and how its dispersion involves will decrease the response time and thus, manage the disaster more efficient. An all-weather solution will enhance even more the reliability of the system for such situations. Thus, proper satellite image analysis can potentially provide such solutions towards the required early disaster management.

Aiming at identifying oil spills by analyzing visual representations, the proposed model processes SAR images due to their independence regarding the weather conditions and the acquisition time. The method deploys a DCNN and semantically segments the regions of the input image into instances of interests (oil-spills, ships, land etc.). Due to the nature of the architecture, the model essentially learns the physics behind the oil spills, like size and shape, and so, it can accurately classify the required image regions.

The rest of the paper is organized as follows. In section 2, relevant works dealing with the oil spills identification problem are analyzed while in section 3, the proposed model is outlined. Section 4 presents the corresponding experimental results and finally, conclusion are drawn in Section 5.

2 Related Work

Incentive algorithms were focused on the utilization of images in the visible spectrum. Numerous approaches were proposed such as exploiting polarized lenses [14] and hyper-spectral imaging [5]. Researches proved that in visible spectrum oil slicks and water cannot be sufficiently distinguished while further limitations are inserted due to weather and luminosity conditions. Nevertheless, the field is considered still active due to the advancements of sensing technologies. To surpass optical sensor constraints, microwave sensors including radars were utilized. For early pollution detection, the acquired data rely on specialized sensors, namely Synthetic Aperture Radar (SAR), where successive pulses of radio waves are transmitted from some altitude and their reflection is recorded to produce a representation of the scene. SAR imagery was primarily used in [12] due to its invariance in lighting conditions and the occlusion caused by the existence of clouds or fog [3].

"Bright" SAR image regions, known as sea clutters, are produced by capillary waves which, under the existence of oil spills, are depressed and depicted as dark formations. However wind slicks, wave shadows behind land, algae blooms, and so forth. [3] can result to similar formations minimizing the effectiveness of the oil spill detector. The most common procedure of similar detections includes

four discrete phases [18]. The first two phases include the detection of the dark formations in SAR images and the corresponding feature extraction, respectively. The features are compared with some predefined values in the third phase and finally, a decision making model classifies each formation. Several disadvantages accompany this method, originating from the restriction of extracting a set of features, the absence of a solid agreement over their nature and the lack of research over their effectiveness.

The majority of such detectors involve a two-class classification procedure, where one class corresponds to oil spills and a second, more abstract class that corresponds to dark formations of all similar phenomena [3]. The second class is usually considered as a group of subclasses like current shear, internal waves and so on. The characterization of the "dark spots" is highly affected by adjacent contextual information, like the presence of similar formations, ship routes etc. Considering the high resolution of the satellite SAR sensors, the acquired images may include not only maritime areas but coastal territories, also. Since SAR sensors operates under microwave frequencies, metallic objects are depicted as bright spots due to the beam reflectance. This explicit discrimination results to a fixed number of classes which comprises the main set in most relevant approaches. For example, decision trees were utilized in [18] to classify the extracted geometrical and textual features and so, oil spills could be discriminated from look-alikes. In addition, an object-oriented approach was used in [10] to radar image analysis and improve manual classification at the scale of entire water bodies. Conventional neural networks were also utilized to identify such environmental disasters [16] focusing on classifying the entire input image with one single label. Finally, a deep CNN model was used in [13] to discriminate oil spills with look-alikes nonetheless, the analysis was limited to a binary classification process.

Aiming at mitigating the limitations imposed by the relevant approaches, the proposed method utilizes a DCNN to segment semantically the processed SAR images instead of labeling local patches or marking the entire images. The classification result is applied at a pixel level and thus, the final image representation includes a map with all pixels annotated. In addition, most of the relevant methods prerequisite the extraction of some features that can describe the characteristics of the oils spills. Due to the sequential convolutional layers of the model, the requirement of computing initially a set of features to be classified as other relevant methods is not valid. Moreover, in relevance to similar DCNN-based methods [13], the presented scheme comprises a multi-scale architecture with four parallel DCNN branches resulting to more accurate classifications. Finally, the model was trained to identify more instances including land territories which eventually could increase the situational awareness of the operational personnel to manage more effectively the pollution disaster.

3 Methodology

The presented oil slick detector intends to segment semantically the input images and highlight the identified instances unlike the single labeling of the entire SAR image. Oil dispersion creates a wide range of irregular shapes that may also coincide with vessels or look alike objects. Thus, semantic segmentation could consist the most appropriate approach compared to other alternatives like defining multiple labels in input image [19] or bounding boxes over detected objects [9]. The presented method can analyze images containing multi-class objects without the need of breaking the image into multiple patches to label the identified instances.

The CNN model relies on the DeepLab model [1], which is reported to achieve high performances in various multi-class segmentation problems [7,15]. Following the DeepLab architecture, the presented oil slick detector consists of four DCNN branches to perform image semantic segmentation while convolution is applied with upsampled filters [8], originally introduced in [4]. Atrous convolution in combination with an Atrous Spatial Pyramid Pooling (ASPP) is deployed to provide parallel filters of different rate and meet the requirements for dense and wide *field-of-view* filters. Finally, bilinear interpolation is utilized to increase the resolution of the extracted feature maps and restore the initial resolution of the input image. In Fig. 1, a higher level representation of the overall procedure is presented.

The initially proposed model employs a fully convolutional architecture based on ResNet-101 model and pre-trained on MS-COCO dataset [11], which resulted to the highest performance in image semantic segmentation. Nonetheless, the repetition of max-pool blocks and strides through the network deteriorates the feature maps' resolution and increases the computational time requirements. To eliminate such constraints, atrous convolution was employed to control the feature maps' resolution over the network's layers.

As an example, in case of a 1-D input signal $x[i]$ atrous convolution with a filter $w[k]$ of length K, gives the output signal $y[i]$ as following:

$$y[i] = \sum_{k=1}^{K} x[i + r \cdot k]w[k], \tag{1}$$

where parameter r defines the stride with which signal $x[i]$ is sampled. Regular convolution can be considered a special case of (1) where $r = 1$.

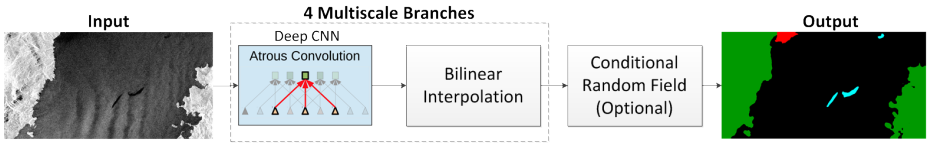

Fig. 1. High-level representation of the presented model.

DCNNs can identify similar objects of multiple scales and rotation due to the training procedure in similar representations. However, further robustness to scale variability is required for oil spill identification due to the orbit of the satellite. Scene and object representation in such images can display wide variety due to difference of the operational altitude. Moreover, oil spills present extreme diversity in shape and size due to the physics of dispersion. The model deploys atrous spatial pyramid pooling for managing the scale variability and was inspired from the corresponding R-CNN technique in [6]. As a result, the classification efficiency of the multi-scale regions is achieved by resampling the feature maps at a set of different rates and further processing them before fusing for the final output. At the final process stage, bilinear interpolation is utilized for the extracted feature maps to regain the initial resolution of the input image. To further enhance the scaling robustness, the model was extended with a multi-scale process to resolve the scale variability issue.

More specifically, four parallel DCNN branches that share the same parameters are used and extract separate score maps from the original image, two rescaled versions and a fused version of all of them. The four branches are combined in one by taking the maximum score across them at each position. It should be also noted that the impact on detector's performance of final CRF layer of the DeepLab model was also examined through the experiments. However, it was exclude it from the finalized model's architecture since it is mainly useful to refine the segmentation results, which in our case did not improve segmentation accuracy rates. For the oil spill detection, objects and regions in SAR imagery present ambiguous shape outlines, resulting in minor improvements of the segmentation accuracy when CRF module is employed, while a computational overhead is added.

4 Experimental Results

4.1 Dataset Description

One key challenge that the researchers has to confront in classification models is the absence of a public dataset which may be utilized for benchmarking. In previous works [2,10,18] the required datasets were developed manually making relevant works almost non comparable. This constraint motivated us to develop a new dataset by collecting satellite SAR images of oil polluted areas via the European Space Agency (ESA) database, the Copernicus Open Access Hub[1]. To ensure the validity of the data and the inclusion of oils spills in the images, the European Maritime Safety Agency (EMSA) provided the confirmed oil spill events through the CleanSeaNet service along with their geographic coordinates. By this approach, we guaranteed that the dark spots depicted in the SAR images correspond to oil spills.

After downloading the appropriate records, a set of preprocessing stages were conducted so that the products could be processed as common images:

[1] `https://scihub.copernicus.eu/`

- Localization of the confirmed oil spills.
- Cropped regions that contain both the oil spills and contextual information. Rescaling the images to a resolution of 1250x650.
- Radiometric calibration for projecting the images onto the same plane.
- Speckle noise suppression with a 7x7 median filter.
- Linear transformation from db to real luminosity values.

A sufficient number of SAR images were processed with the above procedure each of which may include instances of interest such as oil spills, look-alikes, ships and coastal territories. The representations were manually annotated based on the EMSA records accompanied with human identification. During the annotation process, every region was semantically marked with a specific colorization, producing a ground truth mask for every image. A training and a testing set consisting of 771 and 110 images,respectively, were created by randomly sampling the annotated images. Finally, it must be highlighted that the database is extended constantly and could be accessed by the community after receiving the proper confirmations by the relevant authorities.

4.2 Results

For the conducted experiments, three foreground classes were defined i.e. oil spills, look-alikes and ships as well as two classes for the background pixels corresponding to land and sea areas. The overall performance was measured in terms of pixel intersection-over-union (IoU) for every class and averaged for all classes (mIoU). Since the dataset will be meant for future methods benchmarking, a predefined training and testing set should be established. Thus, we decided not to cross-validate the dataset in order to produce comparable results with future proposed methods, following the approach of models benchmarking as in [1], where the model is evaluated in 4 benchmarking datasets. Furthermore, considering the stochastic nature of oil spills shape and size,representative training and testing sets can be produced by single splitting the dataset. Thus, cross-validation would add an exhaustive computational overhead.

Our initial experiments were conducted using the aforementioned dataset and by deploying a simple DCNN network [1] without any multi-scale approach implemented.The selected batch size is equal to 16 image patches while every batch fed into the model is considered as one step of the training process. The corresponding results are provided in Table 1. Based on the numerical results, it can be concluded that that the background areas can be detected with high accuracy when the steps are increased, while oil spills and look-alikes drop bellow 50%. The latter is justifiable since the model cannot generalize without multi-scale analysis and thus, the dominant classes overfit the remaining classes. One interesting result is that "look-alike" class achieves its highest accuracy with 5K steps and drops gradually when they are increased, contrary to the oil spill accuracy. This behavior occurs because the pixels of both classes are usually misclassified. Comparing the results of the basic DCNN model with the results of the CRF expanded model (DCNN-CRF), no significant improvement was

Table 1. Segmentation results of simple model using mIoU/IoU.

Intersection-over-union (IoU)					
Steps	Sea Surface	Oil Spill	Look-alike	Land	mIoU
DCNN					
5k	93.4%	14.9%	**39.8%**	70.3%	54.6%
10k	93.3%	19.8%	36.5%	70.9%	55.1%
15k	93.1%	20.6%	36.5%	67.9%	54.5%
20k	**93.8%**	**21.4%**	35.8%	**75.1%**	**56.5%**
DCNN-CRF					
5k	93.7%	10.8%	**40.9%**	72.0%	54.4%
10k	93.5%	12.6%	36.7%	72.2%	53.8%
15k	93.2%	13.9%	36.2%	69.0%	53.1%
20k	**94.2%**	**15.4%**	36.9%	**77.5%**	**56.0%**

Table 2. Segmentation results of multiscale model using mIoU/IoU.

Intersection-over-union (IoU)					
Steps	Sea Surface	Oil Spill	Look-alike	Land	mIoU
Multi-scale DCNN					
15k	95.3%	43.4%	34.0%	85.1%	64.49%
20k	95.1%	47.6%	33.2%	85.8%	65.49%
25k	95.5%	40.3%	**53.5**%	82.5%	68.0%
30k	**96.0**%	**48.0**%	50.9%	**89.9**%	**71.2%**
35k	95.6%	42.3%	45.0%	87.3%	67.6%
Multi-scale DCNN-CRF					
15k	95.19%	37.0%	30.3%	89.8%	63.1%
20k	95.1%	38.1%	30.3%	89.8%	63.3%
25k	95.6%	30.6%	**53.3**%	84.7%	66.1%
30k	**96.0**%	**39.6**%	50.2%	**92.9**%	**69.7**%
35k	95.7%	34.1%	44.2%	90.8%	66.2%

achieved since the mask does not contain substantial background noise due to the speckle filtering preprocessing stage.

The second set of our experiments included the testing of a multi-scale DCNN scheme to deal with the semantic segmentation of the SAR images. Due to the computational overhead of the four parrallel DCNN branches the initial batch size (16) was reduced in this case to 2 image patches per training step. The results are presented in Table 2, where multi-scale DCNN and DCNN-CRF were evaluated. Regarding the CRF addition, the module did not improve significantly the performance of the model as observed also in the case of a single DCNN branch. In addition, the segmentation rates were increased according to the

Table 3. Segmentation results of simple model using mIoU/IoU.

Intersection-over-union (IoU)						
Steps	Sea Surface	Oil Spill	Look-alike	Ship	Land	mIoU
10k	93.0%	19.5%	**36.5**%	**12.7**%	67.5%	45.9%
20k	**93.6**%	**21.0**%	35.8%	11.5%	**72.0**%	**46.8**%

Table 4. Segmentation results of multiscale model using mIoU/IoU

Intersection-over-union (IoU)						
Steps	Sea Surface	Oil Spill	Look-alike	Ship	Land	mIoU
10k	95.6%	49.7%	**54.9**%	15.7%	86.9%	**60.6**%
20k	95.4%	50.1%	45.8%	20.1%	78.4%	58.0%
30k	95.1%	**50.3**%	35.4%	22.4%	87.5%	58.1%
40k	**95.8**%	38.1%	48.2%	**25.4**%	**88.8**%	59.3%

training steps, resulting to state-of-the-art outcomes when comparing to Table 1. Similar to the results of the first set of experiments, the background classes were identified more accurate than the foreground but, in comparison with the the simple model, the accuracy rates were improved for the foreground regions, as well. This result occurs due to the lesser number of the foreground ground truth pixels in comparison with the corresponding background pixel. Nonetheless, a more efficient sampling approach could improve the results.

Additional experiments were performed to examine the model's segmentation capability in an extended version of our dataset, where ships and vessels were separately identified and thus, a new class was inserted. The extracted results are presented in Table 3 and Table 4 for the simple and the multi-scale analysis, respectively. Results proved once again the advantage of the multi-scale technique over the simple DCNN. The corresponding results display a minor decrement in comparison with the four classes case nonetheless, as far it concerns the oil spill detection, the model can still identify accurately the polluted territories. On the contrary, ship localization rates are considered low, as expected, since the corresponding image regions are too narrow/small and therefore, difficult to be sufficiently identified. Moreover, the number of the ship samples was insufficient for the training process since the main objective was to identify the polluted areas and not directly their potential source. In order to have a more generic model that could deal pollution related tasks in general, the database will be augmented with further samples of objects of interest.

The results of both techniques can be visually compared in Figure 2 for 4 and 5 classes (*simple* refers to one branch [13] while *msc* corresponds to the four branches approach). Analyzing the two figures, we can conclude that the proposed multi-scale DCNN outperforms the simple DCNN when oil spills and look-alikes are concerned. Both techniques perform high segmentation rates for

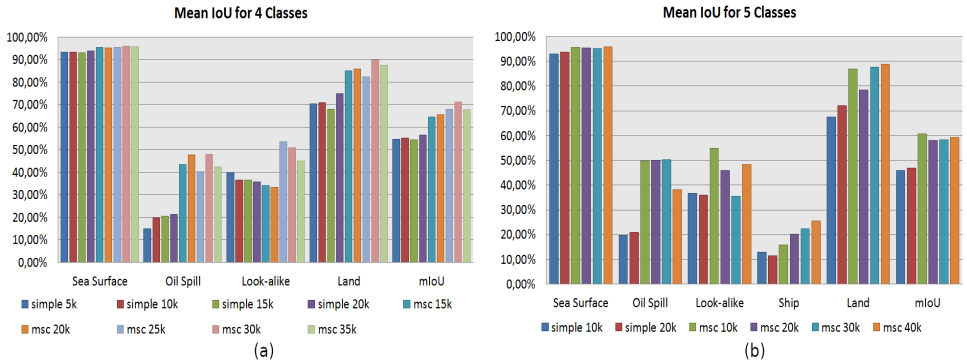

Fig. 2. Mean IoUs of **(a)** 4 classes and **(b)** 5 classes.

background pixels, i.e. sea surface and land area, leading over 90% and 80%, respectively. Similarly, the foreground regions are identified equally sufficient, regardless the complexity of the task (oil spills and ship localization).

For comparisons, we additionally utilized the segmentation accuracy of the highest performance model in the four classes problem, so as it could be compared with simple models of image classification. Therefore, a determined amount of overlapping patches were cropped from each pair of ground truth and predicted masks. In order to arise a credible dataset, the following restrictions were introduced:

1. Apart from sea surface class, each patch should contain at least a minimum amount of pixels classified in one of the rest three classes. A threshold of 2% was selected implying that the class containing the most pixels should contain at least 2% of the amount of pixels of the sea surface class.
2. An image patch is labeled only if one of the three classes dominates. So, a threshold was defined equal to 50%, meaning that the dominant class should contain at least 50% of the amount of the pixels of a non dominant class.
3. If a patch does not satisfy the aforementioned rules, it is excluded from the accuracy estimation.

Classification results for image patches are presented in Table 5. Since calculated accuracy is dependent of the amount of patches extracted from every image, two different values for the horizontal-vertical ratio of patch size were examined. It should be noted that the results included in the tables are somehow dissimilar since, for the second metric, a single label is evaluated for every patch.

A possible comparison with relevant approaches would be somehow iniquitous due to the lack of a common image dataset as a base. Moreover, most of other relevant approaches attempt to solve a binary classification problem (oil spills and look-alikes), in contrast with the proposed method, excluding other information that may be valuable in disaster management. In addition, our algorithm annotates each pixel with one valid state where other methods designate image

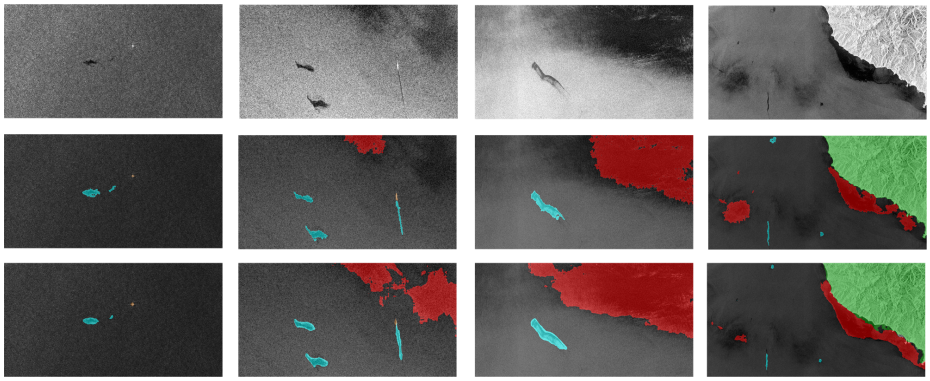

Fig. 3. Example of 4 testing images (from top to bottom): SAR images, ground truth masks and resulted detection masks overlaid over SAR images.

regions and so, accuracy is determined in a completely different basis.Thus, comparing classification approaches for different posed problems is somehow invalid. Nonetheless, some results of comparison are provided. The method in [16] which exploits a neural network resulted a 91.6% and 98.3% accuracy for oil spills and look-alikes (without considering the ship instances), respectively. Highest accuracy was achieved by the method in [18] which deployed a decision tree forest and achieved an accuracy equal to 85.0%. Finally, the probabilistic based method in [17] achieved accuracy equal to 78% and 99% for oil spill and look-alikes classes, respectively. Without the constraint of extracting features, initial results of the proposed approach are comparable to those of state-of-the-art methods and with the merit of semantically annotated regions.

Table 5. Segmentation accuracy results.

Image patch classification accuracy results			
Number of patches: 3,3			
Overall 85.2%	Oil Spill 89.1%	Look-alike 69.2%	Land 97.4%
Number of patches: 5,3			
Overall 84.1%	Oil Spill 91.0%	Look-alike 67.6%	Land 93.8%

For representation and qualitative purposes, Figure 3 includes some examples of semantically annotated images in order to demonstrate the accuracy of the model and the distinctiveness of the problem. The cyan colored pixels denote the identified oil spill regions while the red marked pixels corresponds to

look-alike areas. In addition, green marked territories resembles the coastal regions while black colored corresponds to the sea surface. Finally, the detected vessels are marked with brown color and cover the smallest image regions in the representations. Oil spills are very similar to look-alikes as both are represented by black masses and so, they can easily be misclassified. Nonetheless, the model was properly trained to discriminate these instances due to the differences they display as natural phenomena (size, shape etc.). Eventually, their accurate identification relies on the fact that the model itself learned their physical attributes providing a valuable discrimination for the disaster management authorities.

5 Conclusions

In this paper, a novel approach was proposed for oil spill detection based on SAR image analysis aiming at a disaster management framework at early stages. Robust DCNN models can automate the detection of the polluted areas along with relevant objects like look-alikes, vessels or coastal regions. In addition, based on the performed analysis, initial results indicate that such models can provide an accurate estimation about the upcoming disaster targeting the best situation awareness of the relevant authorities. Thus, such models can be integrated to wider frameworks for disaster and crisis management. The extracted results are comparable to the state-of-the-art results, nonetheless, for general classification problems. More specific for the oil spill detection problem, the adaptation of relevant and more accurate deep learning methods may lead to further improvement of the identification accuracy. Larger training sets with sufficient samples and images acquired with improved SAR sensors could substantially improve the accuracy values, also. Current work can be extended to manage similar environmental disasters like floods. Thus, relevant image samples will be required to enhance the current database in order to refine and retrain the model.

6 Acknowledgments

This work was supported by ROBORDER and EOPEN projects funded by the European Commission under grant agreements No 740593 and No 776019, respectively.

References

1. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. arXiv preprint arXiv:1606.00915 (2016)
2. Cococcioni, M., Corucci, L., Masini, A., Nardelli, F.: Svme: an ensemble of support vector machines for detecting oil spills from full resolution modis images. Ocean Dynamics **62**(3), 449–467 (2012)
3. Fingas, M., Brown, C.: Review of oil spill remote sensing. Marine pollution bulletin **83**(1), 9–23 (2014)

4. Giusti, A., Ciresan, D.C., Masci, J., Gambardella, L.M., Schmidhuber, J.: Fast image scanning with deep max-pooling convolutional neural networks. In: Image Processing (ICIP), 2013 20th IEEE International Conference on. pp. 4034–4038. IEEE (2013)
5. Gonzalez, C., Sánchez, S., Paz, A., Resano, J., Mozos, D., Plaza, A.: Use of fpga or gpu-based architectures for remotely sensed hyperspectral image processing. INTEGRATION, the VLSI journal **46**(2), 89–103 (2013)
6. He, K., Zhang, X., Ren, S., Sun, J.: Spatial pyramid pooling in deep convolutional networks for visual recognition. In: european conference on computer vision. pp. 346–361. Springer (2014)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
8. Holschneider, M., Kronland-Martinet, R., Morlet, J., Tchamitchian, P.: A real-time algorithm for signal analysis with the help of the wavelet transform. In: Wavelets, pp. 286–297. Springer (1990)
9. Karpathy, A., Fei-Fei, L.: Deep visual-semantic alignments for generating image descriptions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3128–3137 (2015)
10. Konik, M., Bradtke, K.: Object-oriented approach to oil spill detection using envisat asar images. ISPRS Journal of Photogrammetry and Remote Sensing **118**, 37–52 (2016)
11. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: European conference on computer vision. pp. 740–755. Springer (2014)
12. Mastin, G.A., Manson, J., Bradley, J., Axline, R., Hover, G.: A comparative evaluation of sar and slar. Tech. rep., Sandia National Labs., Albuquerque, NM (United States) (1993)
13. Orfanidis, G., Ioannidis, K., Avgerinakis, K., Vrochidis, S., Kompatsiaris, I.: A deep neural network for oil spill semantic segmentation in sar images. In: Accepted for presentation in IEEE International Conference on Image Processing. IEEE (2018)
14. Shen, H.y., Zhou, P.c., Feng, S.r.: Research on multi-angle near infrared spectral-polarimetric characteristic for polluted water by spilled oil. In: International Symposium on Photoelectronic Detection and Imaging 2011: Advances in Infrared Imaging and Applications. vol. 8193, p. 81930M. International Society for Optics and Photonics (2011)
15. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
16. Singha, S., Bellerby, T.J., Trieschmann, O.: Satellite oil spill detection using artificial neural networks. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing **6**(6), 2355–2363 (2013)
17. Solberg, A.H., Brekke, C., Husoy, P.O.: Oil spill detection in radarsat and envisat sar images. IEEE Transactions on Geoscience and Remote Sensing **45**(3), 746–755 (2007)
18. Topouzelis, K., Psyllos, A.: Oil spill feature selection and classification using decision tree forest on sar image data. ISPRS journal of photogrammetry and remote sensing **68**, 135–143 (2012)
19. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. In: Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on. pp. 3156–3164. IEEE (2015)