CREWES www.crewes.org

Oil Spills Identification on Satellite Radar Data using Deep Learning

Marcelo Guarido*, David Emery, and Kristopher Innanen

mguarido@ucalgary.ca

ABSTRACT

Oil spills in oceans are a major pollutant endangering oceanic and coastal marine life, and their detection is of high environmental importance. Manually detection presents a challenging and lengthy task. We presented a deep learning model based on the U-Net structure to identify oil bodies in satellite radar images with promising results. Our model successfully classified larger oil bodies with moderate success on smaller ones. Image feature engineering, such as a four-directional cumulative sum, brought important information to the model and performed more accurate predictions. Limited by computer resources, our model was relatively simple. We used pre-trained weights from the MobileNetV2 model. Although initial results are unsatisfactory, we will continue to explore the transfer learning methodology to generate more accurate oil detection algorithms.

U-NET MODEL

A U-Net model is a modified structure of the *Convolutional Neural Networks.* It consists of a sequence of 2D convolutional (feature extraction) and pooling (size reduction) layers in the encoder (first half) part extracting small parts of different figures and keeping them in the central part of the model. The decoder (second half) of 2D transpose convolution (increase image size, concatenate with previous layers) and 2D convolutional layers to localize the small extracted features in the image. The output is a matrix of probabilities for each class in the classification. As we have only one class (oil, not oil), the output is a single matrix with the probabilities for each pixel being oil.

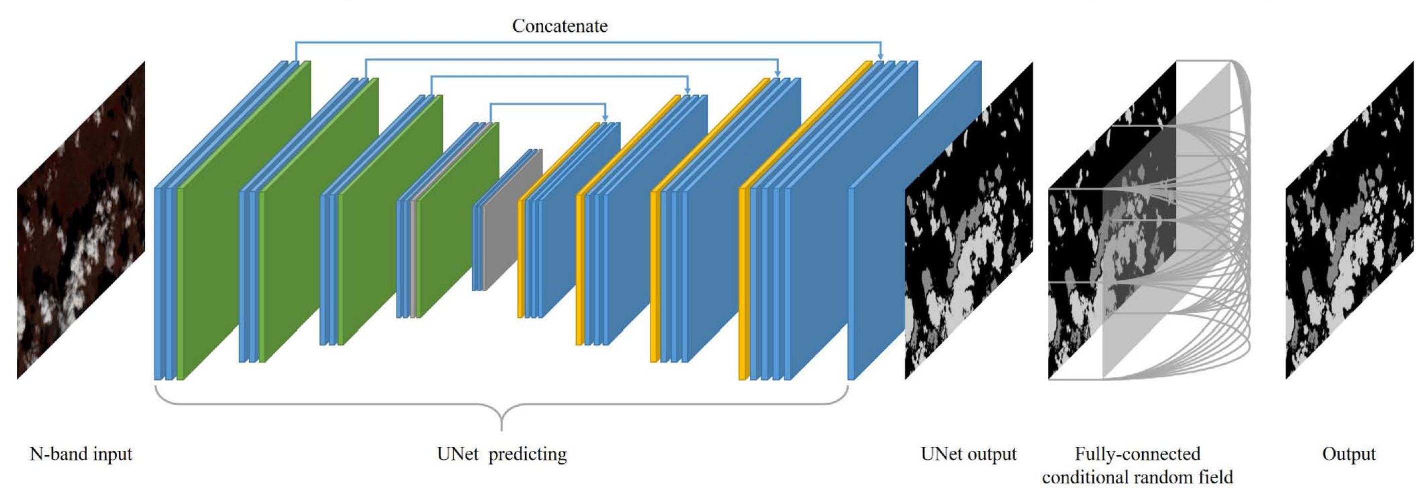

PREDICTIONS

The results using a U-Net model are presented in the figure below. The masks are in green, the predictions are in red, and the overlap presents a purplish colour. Although the masks are rectangular marks, predictions follow worm-shaped oil patterns. Another observation is that the predictions are more frequent on larger oil bodies. A few misclassifications are present in darker areas that are not related to spills but to some other event. Results so far are promising.

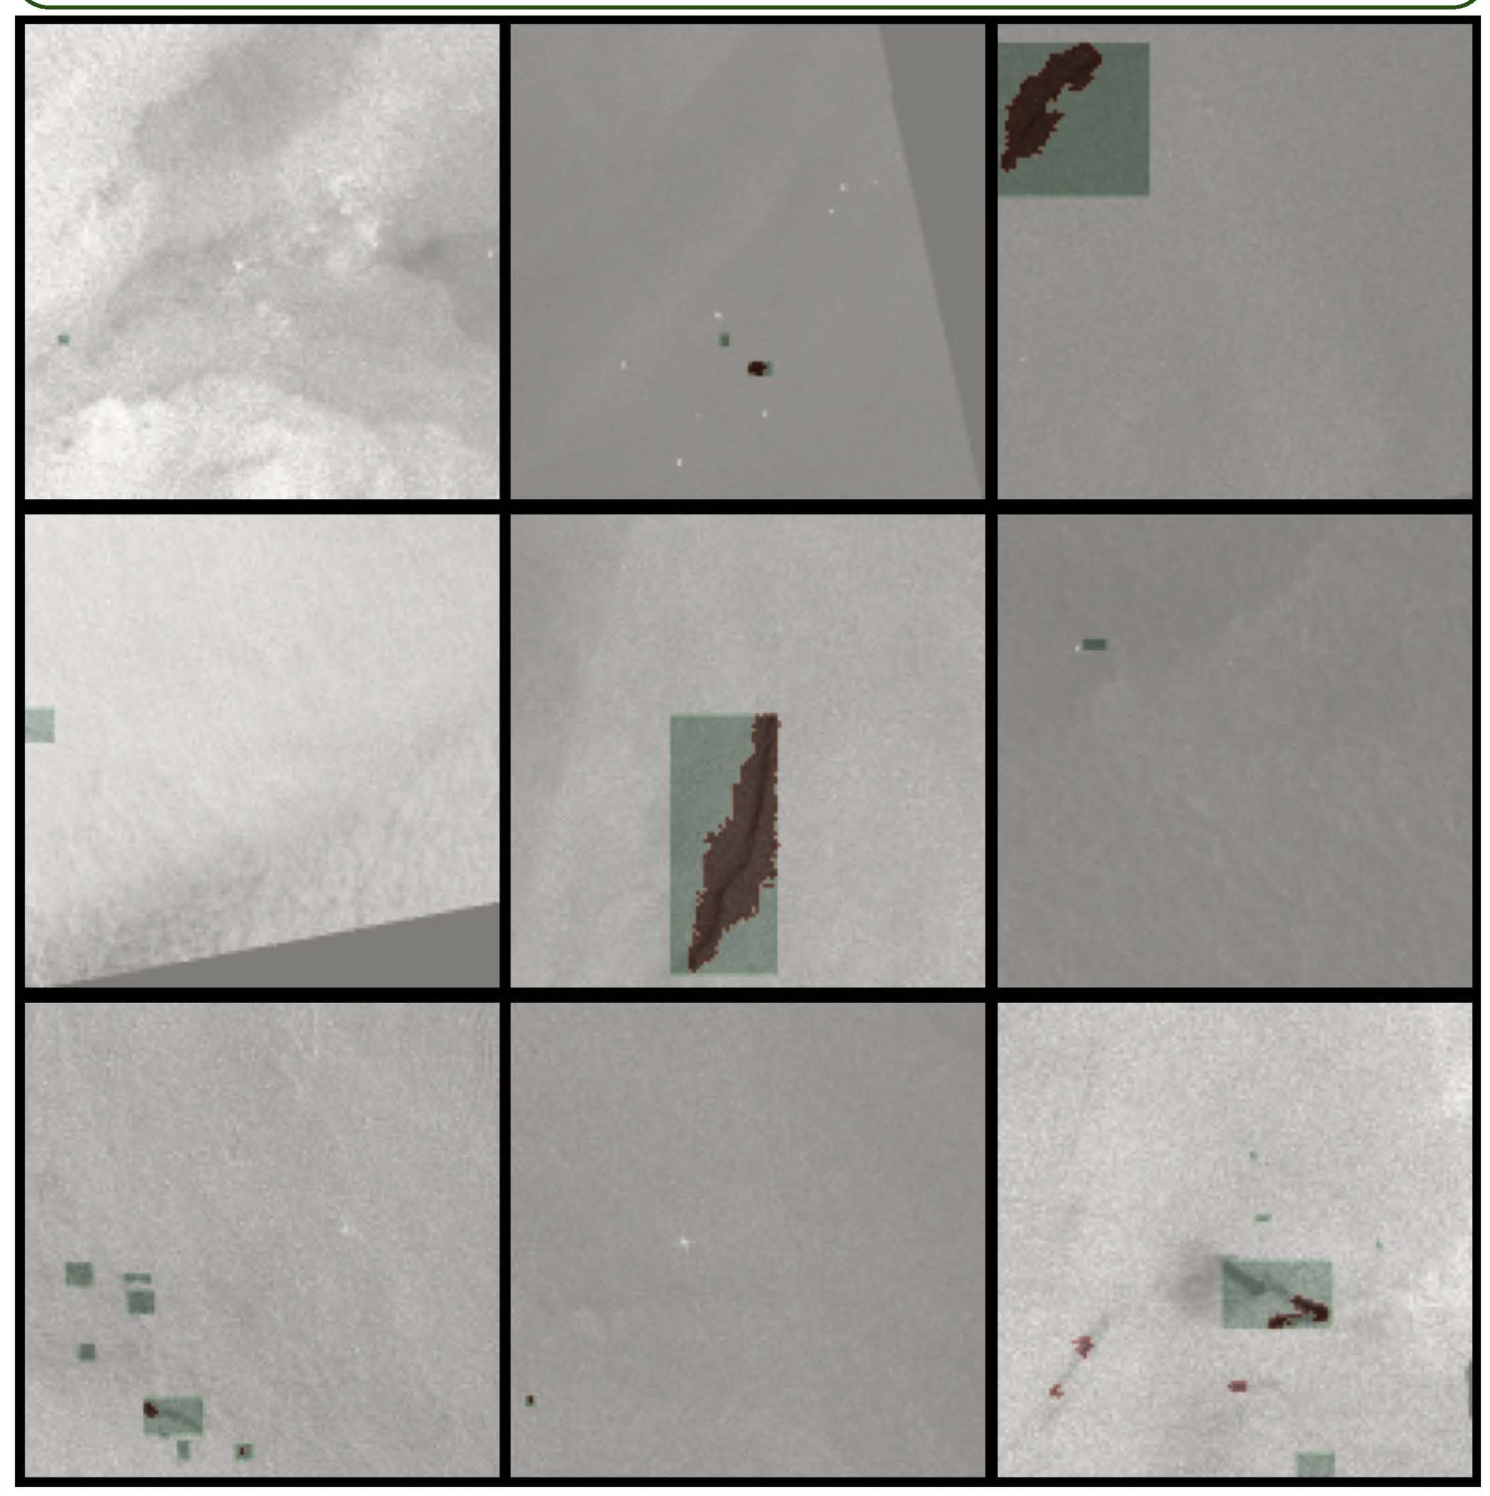

XEEK: SLICK IN A HAYSTACK – A LOCALIZATION CHALLENGE

In this project, we are using satellite radar data provided by the Xeek: Slick in a Haystack – a Localization Challenge competition to automatically identify oil spills in the ocean. They provided vast images covering a large area of the ocean. Manually finding an oil spill candidate (red marks in the figure in the right) can be an arduous job, as there are thousands of images like that. Those images are too large to feed to a deep learning model.

We are training a U-Net type model for image segmentation (pixel classification) by considering the model's size, the size of the input images, the number of images, and the available memory in the computer. Deep learning models work better when several images are fed at once, limiting their size. Our solution was to divide the large image into several patches of images and masks with 128x128 pixels and keep only the ones that contain our target.

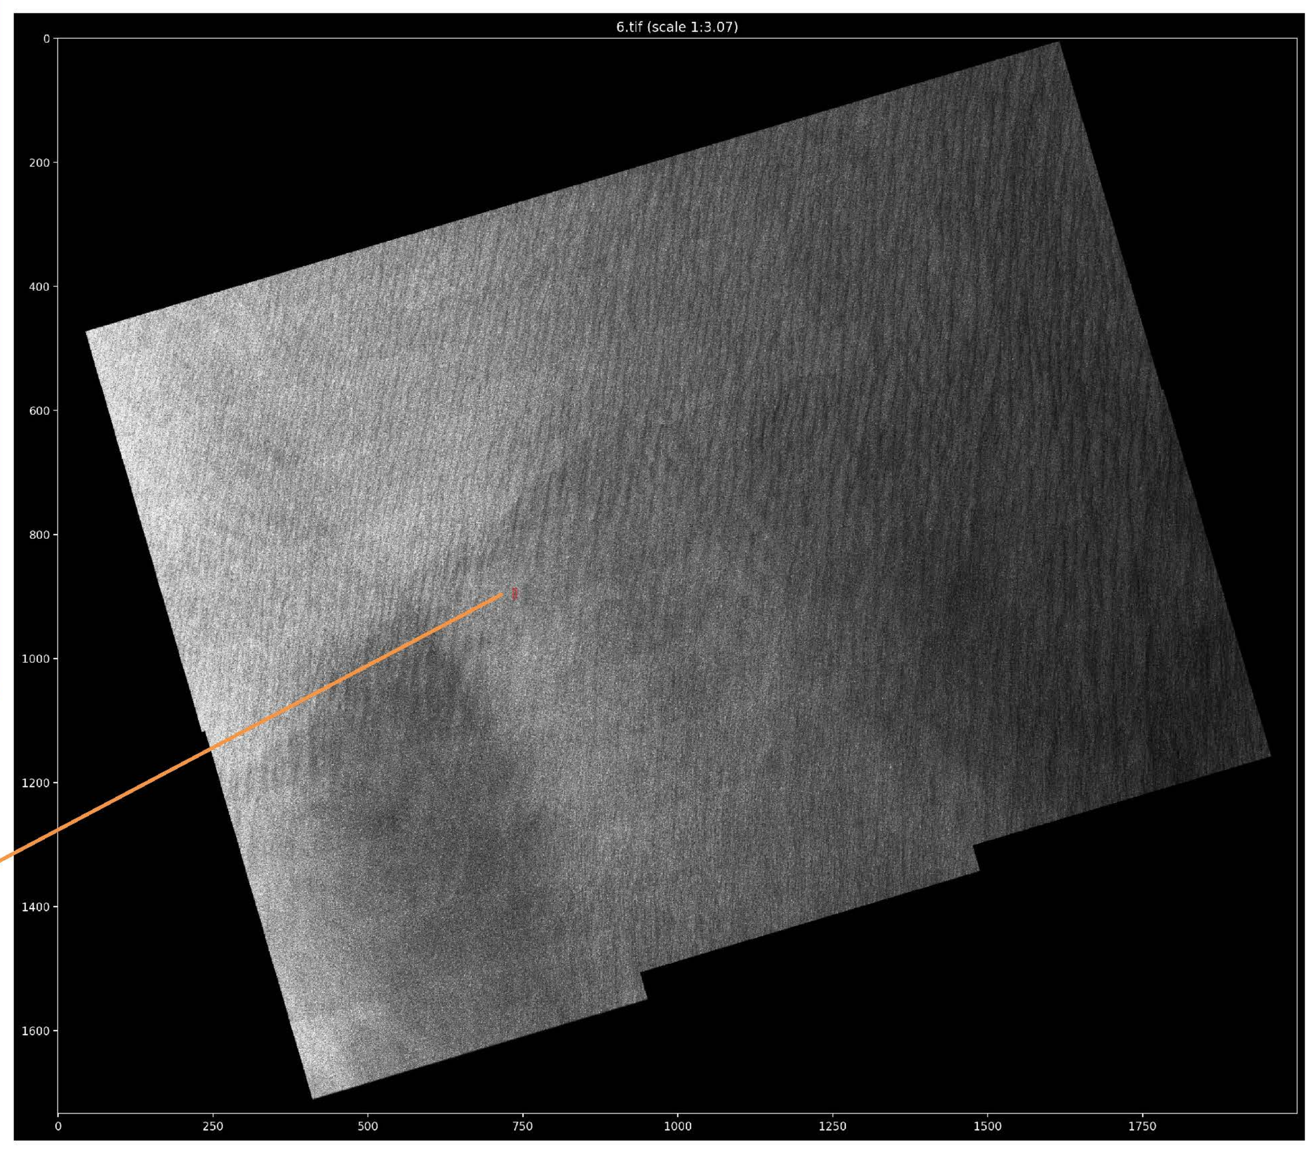

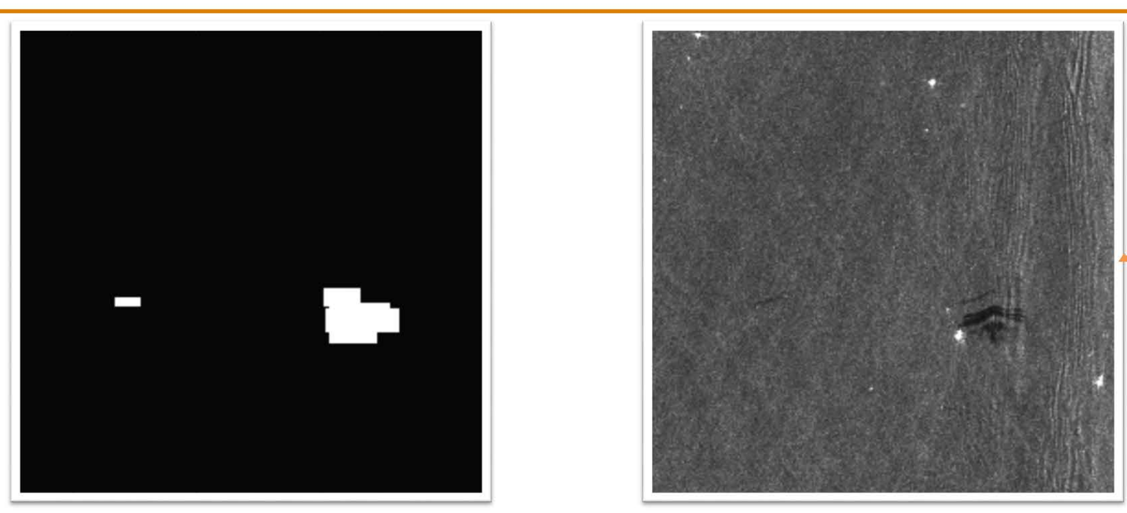

FEATURE ENGINEERING

Feature engineering means creating new features based on the ones we have. For image classification or segmentation, feature engineering is used to create more images for training (by rotating, zooming, flipping, and filtering) or making the target clearer (by filtering or manipulating the images). In this project, feature engineering was applied to highlight the target.

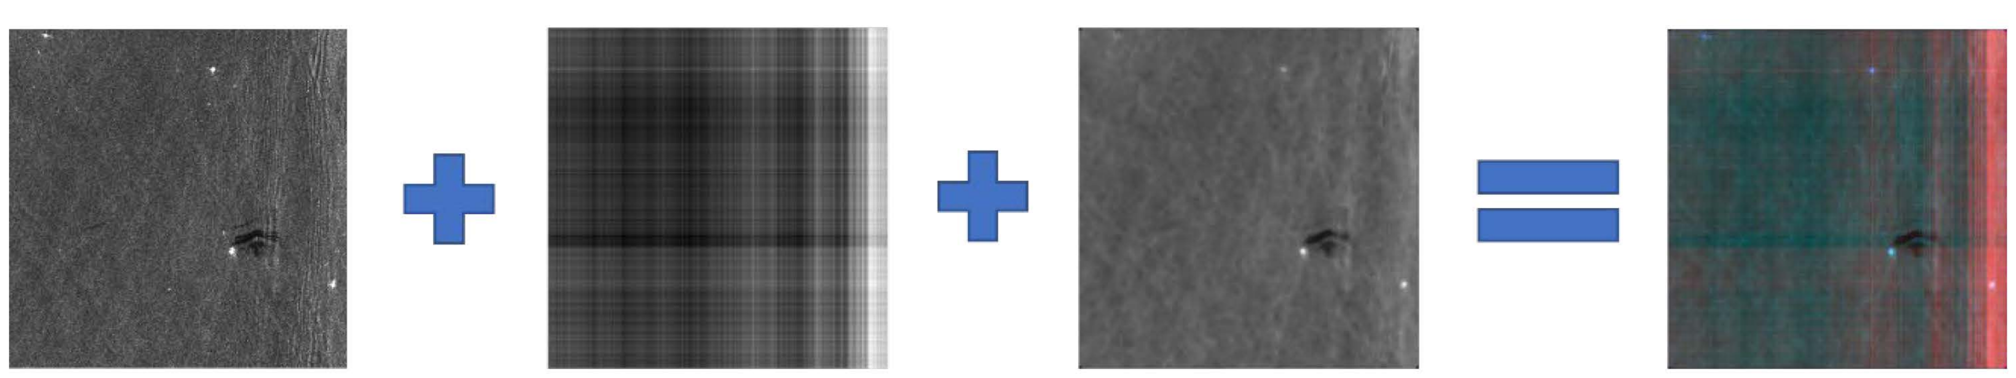

To generate the final RGB (red, green, and blue colour channels) image on the right, we started with the original grayscale image (one colour channel) as the first colour channel, a four-directional cumulative sum as the second channel, and a smooth version of the original one (from a 2D median filter) as the last channel. The goal is to keep the information of the original image where the black spots correspond to colour numbers close to zero and add scope aim with the cumulative sum. The median filter helps to regularize the colour values around the image. This process had a positive impact on the modelling part.

FUTURE WORK

Our model is presenting promising but limited predictions and improving it is challenging, as we would require a deeper and more complex model. With computer limitations, training a more complex model from scratch is not an option. A solution is to use a pre-trained model on another set of data, such as the MobileNetV2 for image segmentation. Our first tests were poor but gave us a starting point.

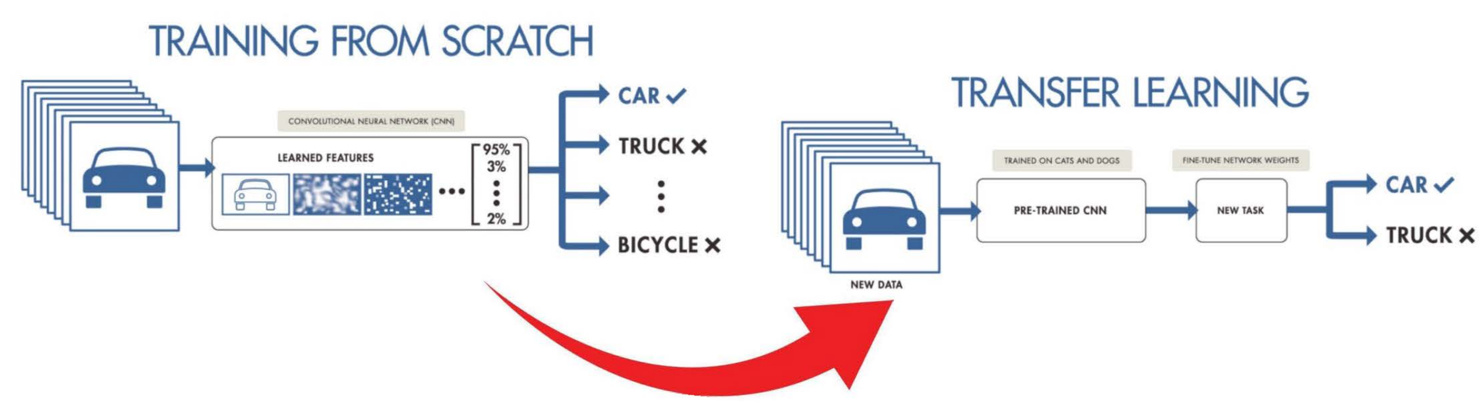

Acknowledgments

We thank the sponsors of CREWES for their continued support. This work was funded by CREWES industrial sponsors and NSERC (Natural Science and Engineering Research Council of Canada) through the grant CRDPJ 543578-19. The Canada First Research Excellence Fund supported the first author through the Global Research Initiative at the University of Calgary. We thank Soane Mota dos Santos for the suggestions, tips, and productive discussions.

NSERC CRSNG
UNIVERSITY OF CALGARY
FACULTY OF SCIENCE
Department of Geoscience

CREWES DSI