GIMA
Geographical Information Management and Applications

Detecting boats using satellite images and deep learning: The case of the Netherlands

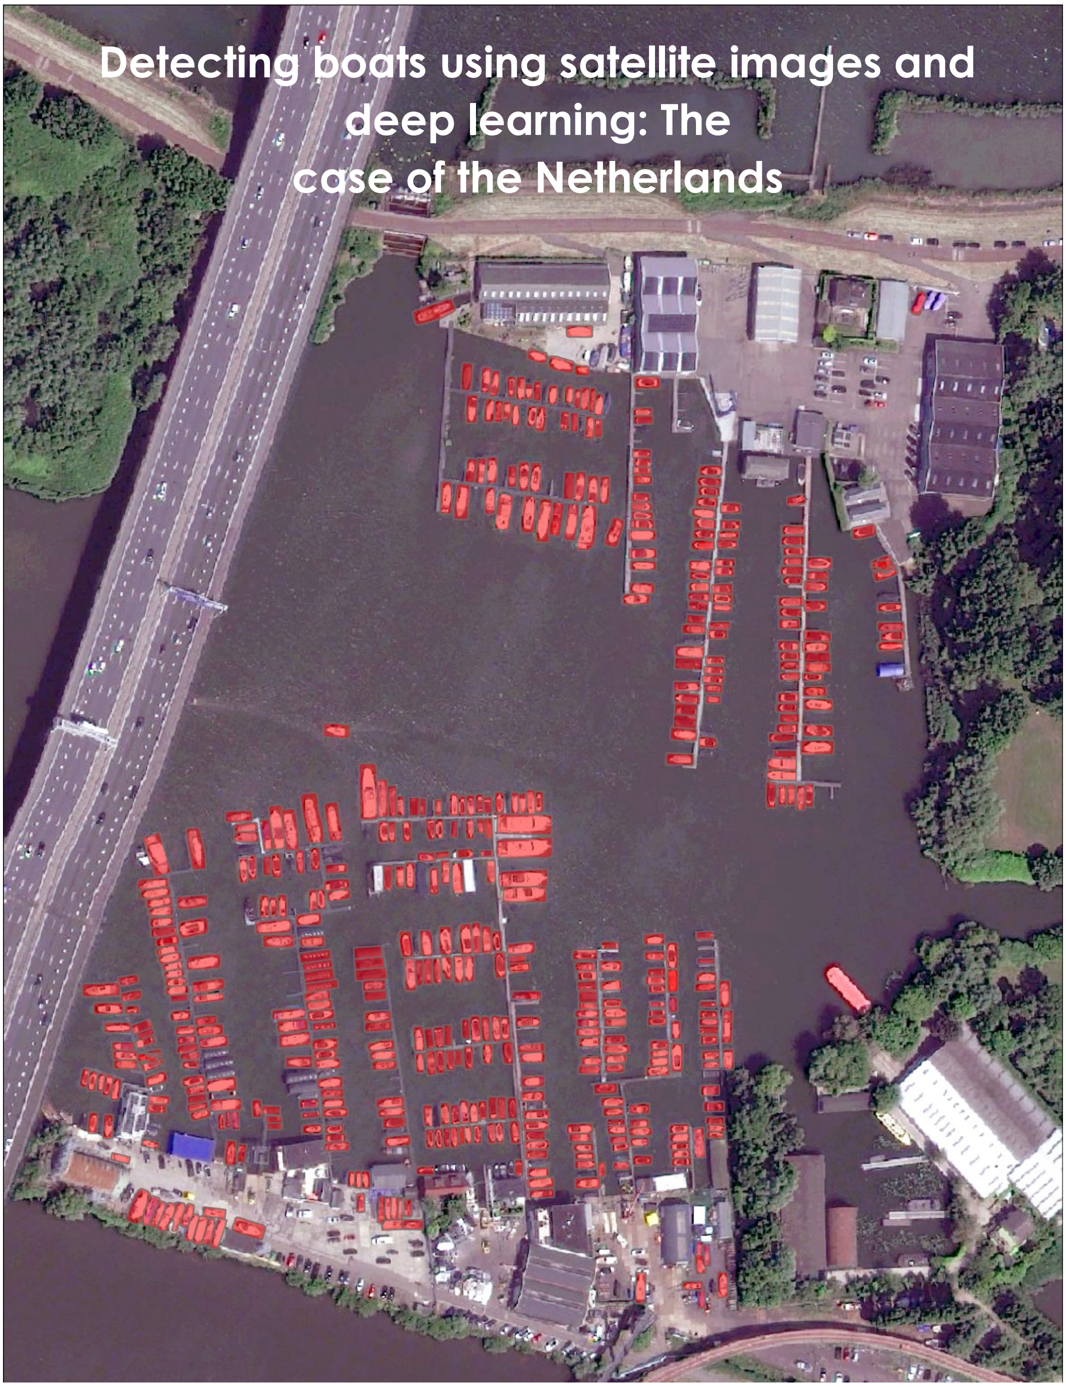

Internship summary

Name: Apostolos Koutsaris
Organisation: Waterrecreatie Nederland

Summary of Internship Report

This internship, conducted at Waterrecreatie Nederland as part of the GIMA (Geographical Information Management and Applications) MSc program, investigated the potential of combining deep learning, remote sensing, and GIS technologies for a critical task: detecting and counting recreational boats in satellite imagery. Waterrecreatie Nederland, as a knowledge organization, is committed to improve water recreation safety, water network management, and promote overall sustainability. To meet their objectives, they sought to explore the feasibility of an automated method for accurately counting recreational boats. To address this need, this internship aimed to develop a robust deep learning model capable of accurately identifying and counting boats in satellite imagery across the Netherlands.

This model would then be applied to satellite images covering water areas and key locations like ports and marinas throughout the country. To assess the model's accuracy and reliability in estimating boat populations, its performance was evaluated against ground truth data.

The methodology employed a combination of remote sensing techniques, GIS tools and deep learning algorithms. More than 330 high resolution (30cm) satellite images were gathered totalling 837 GB, prioritizing images that captured all ports and marinas in the Netherlands. The selection also targeted water networks and major urban areas to maximize image acquisition while minimizing coverage limitations. These collected images then underwent preprocessing steps to prepare them for model training.

A convolutional neural network (CNN), a type of deep learning model, was trained using annotated datasets tailored to identify and count boats within satellite imagery. The annotation process involved labeling the location and type of boats present in the images, in total more than 11.000 ships were annotated. More than eleven models were created differing on the amount of the training images, the training epochs, and other settings. Training the CNN model involved feeding it these annotated images, allowing the model to learn the characteristics of recreational boats and distinguish them from other objects in satellite imagery.

The final model was trained over 6500 training images for 35 epochs using the ArcGIS Learn framework with a PyTorch backend. It employed a Mask R-CNN architecture for boat detection, built upon a ResNet50 backbone model. The learning rate was set between 6.3096e-06 and 6.3096e-05. This configuration achieved an average precision score of 0.798 for the "Boat" class.

The developed deep learning model was seamlessly integrated within the ArcGIS Pro environment. This integration allowed for the model to be used alongside the powerful spatial analysis and visualization functionalities offered by GIS software.

By implementing a custom python script, we were able to automate the detection process and iterate the use of the 'Detect objects using deep learning' tool within ArcGIS Pro for each satellite image of the dataset. The script generated polygons around detected boats and saved them as separate shapefiles per image. We merged these 330 shapefiles to produce national-level results.

The implementation of the deep learning model showed promising results in the boat detection and counting task. The model demonstrated high accuracy in identifying boats within water bodies and coastal regions across the Netherlands. A total of 330 high-resolution satellite images were processed, detecting approximately 258,000 boat objects. From those 258.000 detections, 170.000 boats were found in the water and 88.000 recreational boats in the land. After a validation process using ground truth data, an estimated population of 179,000 boats on water and 29,000 boats on land was derived.

To calculate the accuracy metrics of the model we utilized an Intersection over Union (IoU) threshold of 0.5. For water detections, the model achieved high precision, recall, F1 score and AP values of 0.97, 0.93, 0.95 and 0.96 respectively. For land detections, the model achieved a precision of 0.26, recall of 0.86, F1 score of 0.40 and an average precision (AP) of 0.32.

The model's performance revealed a noticeable distinction between water and land detections. While water detections achieved high precision and recall, land detections showed lower accuracy due to the complexity of the land background in relation to the water areas. Additional improvements in the model's training dataset and settings may be necessary to improve performance in distinguishing boats within complex terrestrial settings

Overall, this internship successfully showcased the potential of integrating advanced technologies including deep learning, remote sensing, and GIS for accurate boat detection and population estimation. The developed model provides a valuable tool for Waterrecreatie Nederland to monitor and manage recreational boat activities effectively. The findings underscore the importance of leveraging modern geospatial technologies for environmental monitoring and resource management. Further refinements in model training and data preprocessing can enhance the accuracy and scalability of boat detection deep learning models, showing the way for broader applications in marine resource management and conservation efforts.