KelpNet: Probabilistic Multi-Task Learning for Satellite-Based Kelp Forest Monitoring

Maximilian Pierzyna

m.pierzyna@tudelft.nl | maximilian-pierzyna.de

Delft University of Technology, Delft, The Netherlands

Drone photo: Tom Bell

Abstract

Kelp forests are critical for marine ecosystems. They harbor a diverse range of species and maintain ecological balance, which necessitates the accurate monitoring of their evolution. We propose a multi-task ensemble deep learning framework to **predict probabilistic maps of kelp forests from Landsat 7 satellite imagery**. We train parallel image classification and segmentation models to achieve **robust kelp predictions**. Both model types are created as ensembles of 25 members producing probabilistic outputs. A comparison of the classification and segmentation outputs **allows for human sanity checking** of the model predictions. Our approach yields a high accuracy with a mean dice score of 0.7047 on test data and **performed well in the DrivenData "KelpWanted" machine learning competition** (#38/671, 3.88% below winning solution).

Dataset

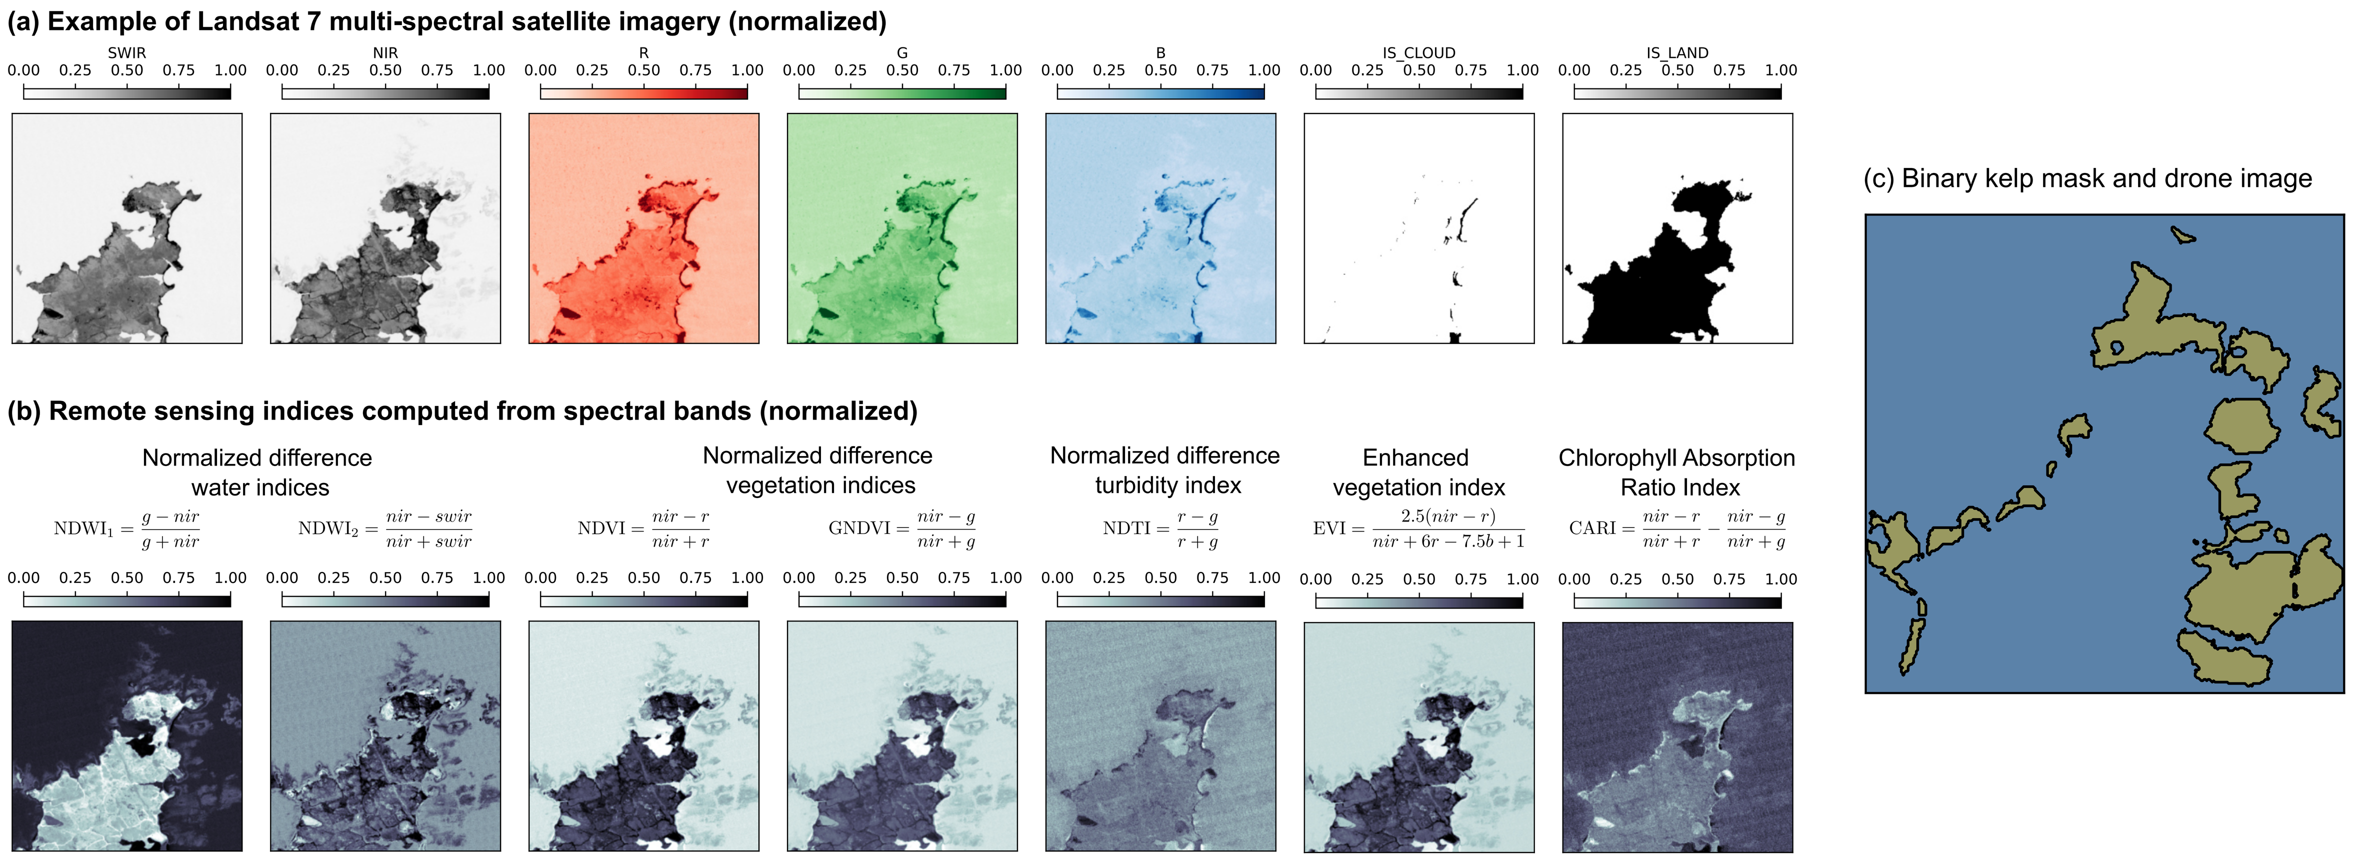

Figure 1: Sample of (a) Landsat 7 satellite imagery and (b) derived remote sensing indices used to predict (c) binary kelp mask.

Multi-task ensemble methodology

A high complexity semantic segmentation model (DeepLab V3, ~39.6M parameters) predicts binary maps of kelp occurrence from satellite imagery inputs. The kelp segmentation maps are post-processed by applying coarse kelp maps obtained from a lower complexity image classification model (CNN, ~2.75M parameters). Both model types are trained as ensembles where each member only utilizes 3 randomly selected feature channels out of the 12 available ones to produce probabilistic kelp maps.

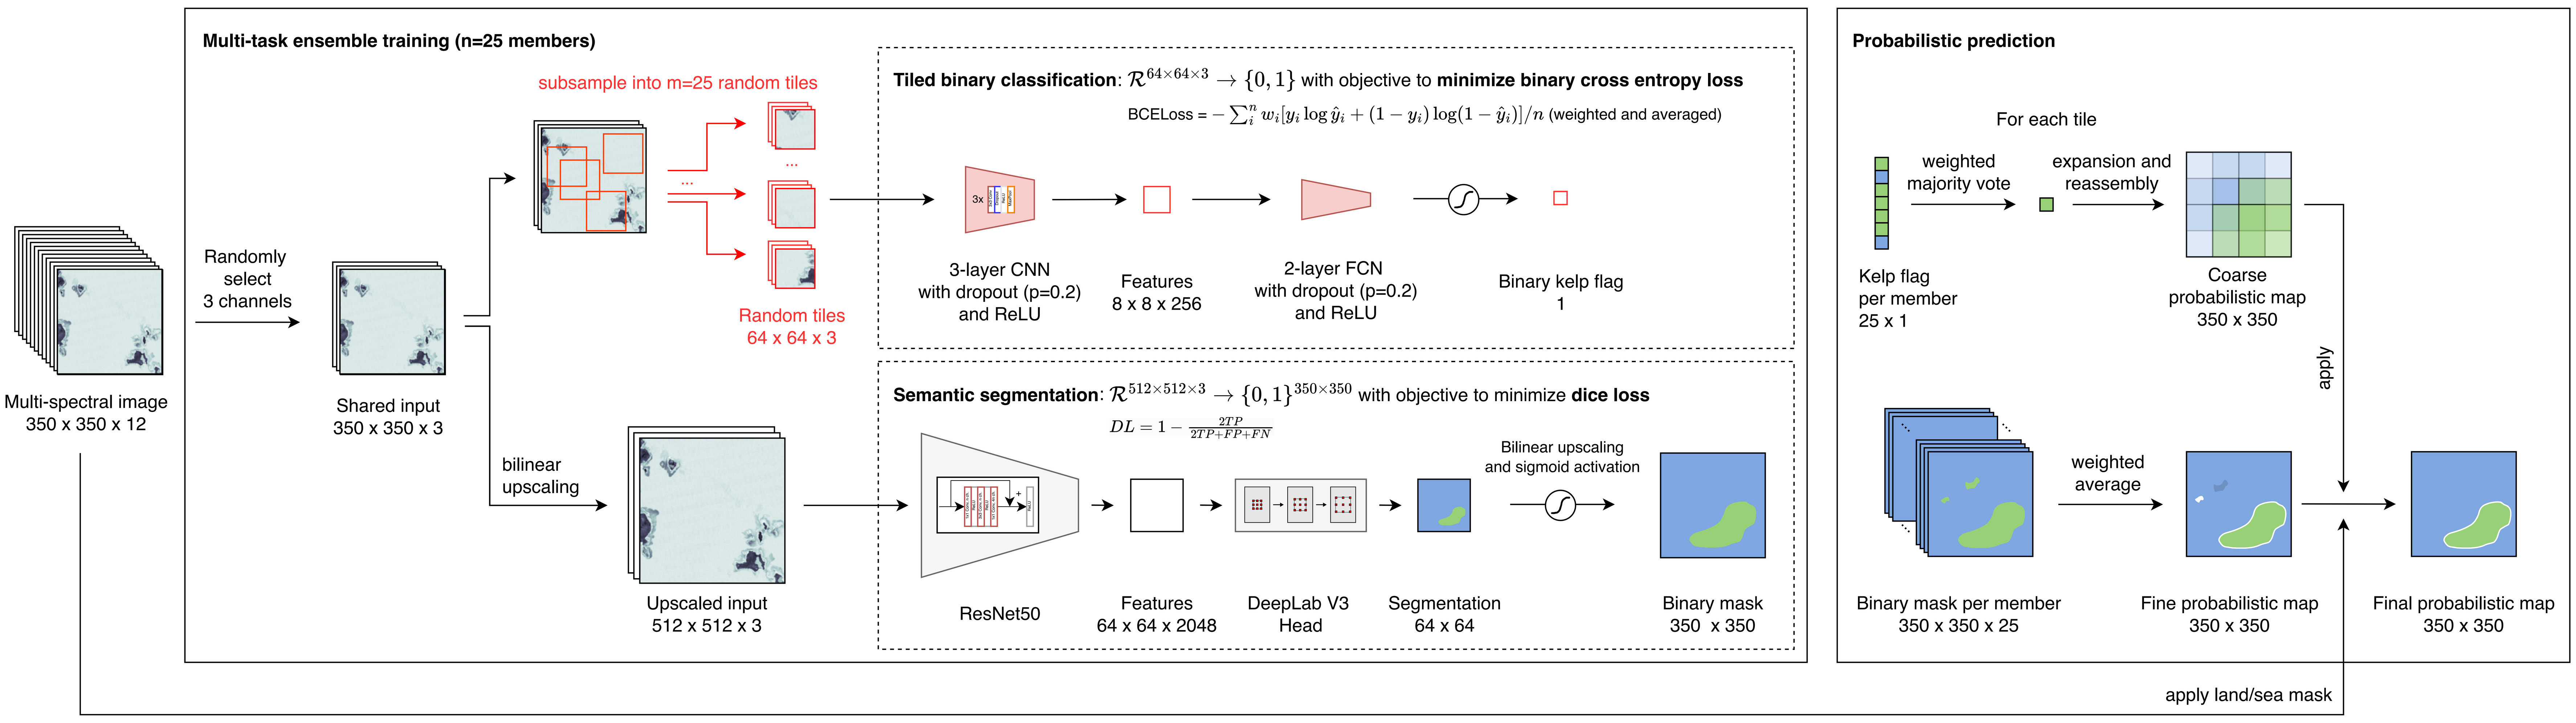

Performance on test data

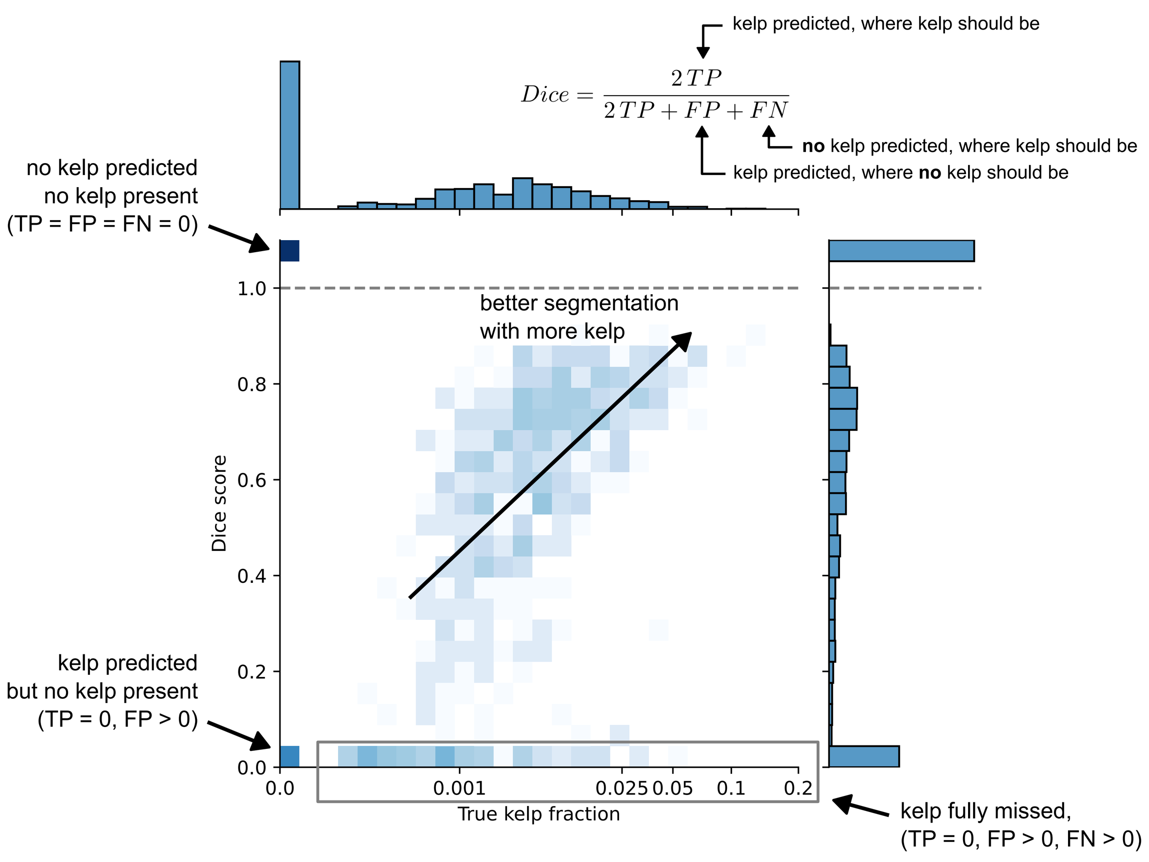

Figure 2: Joint probability distribution of dice score between estimated kelp mask (classification and segmentation combined) with kelp mask from test data.

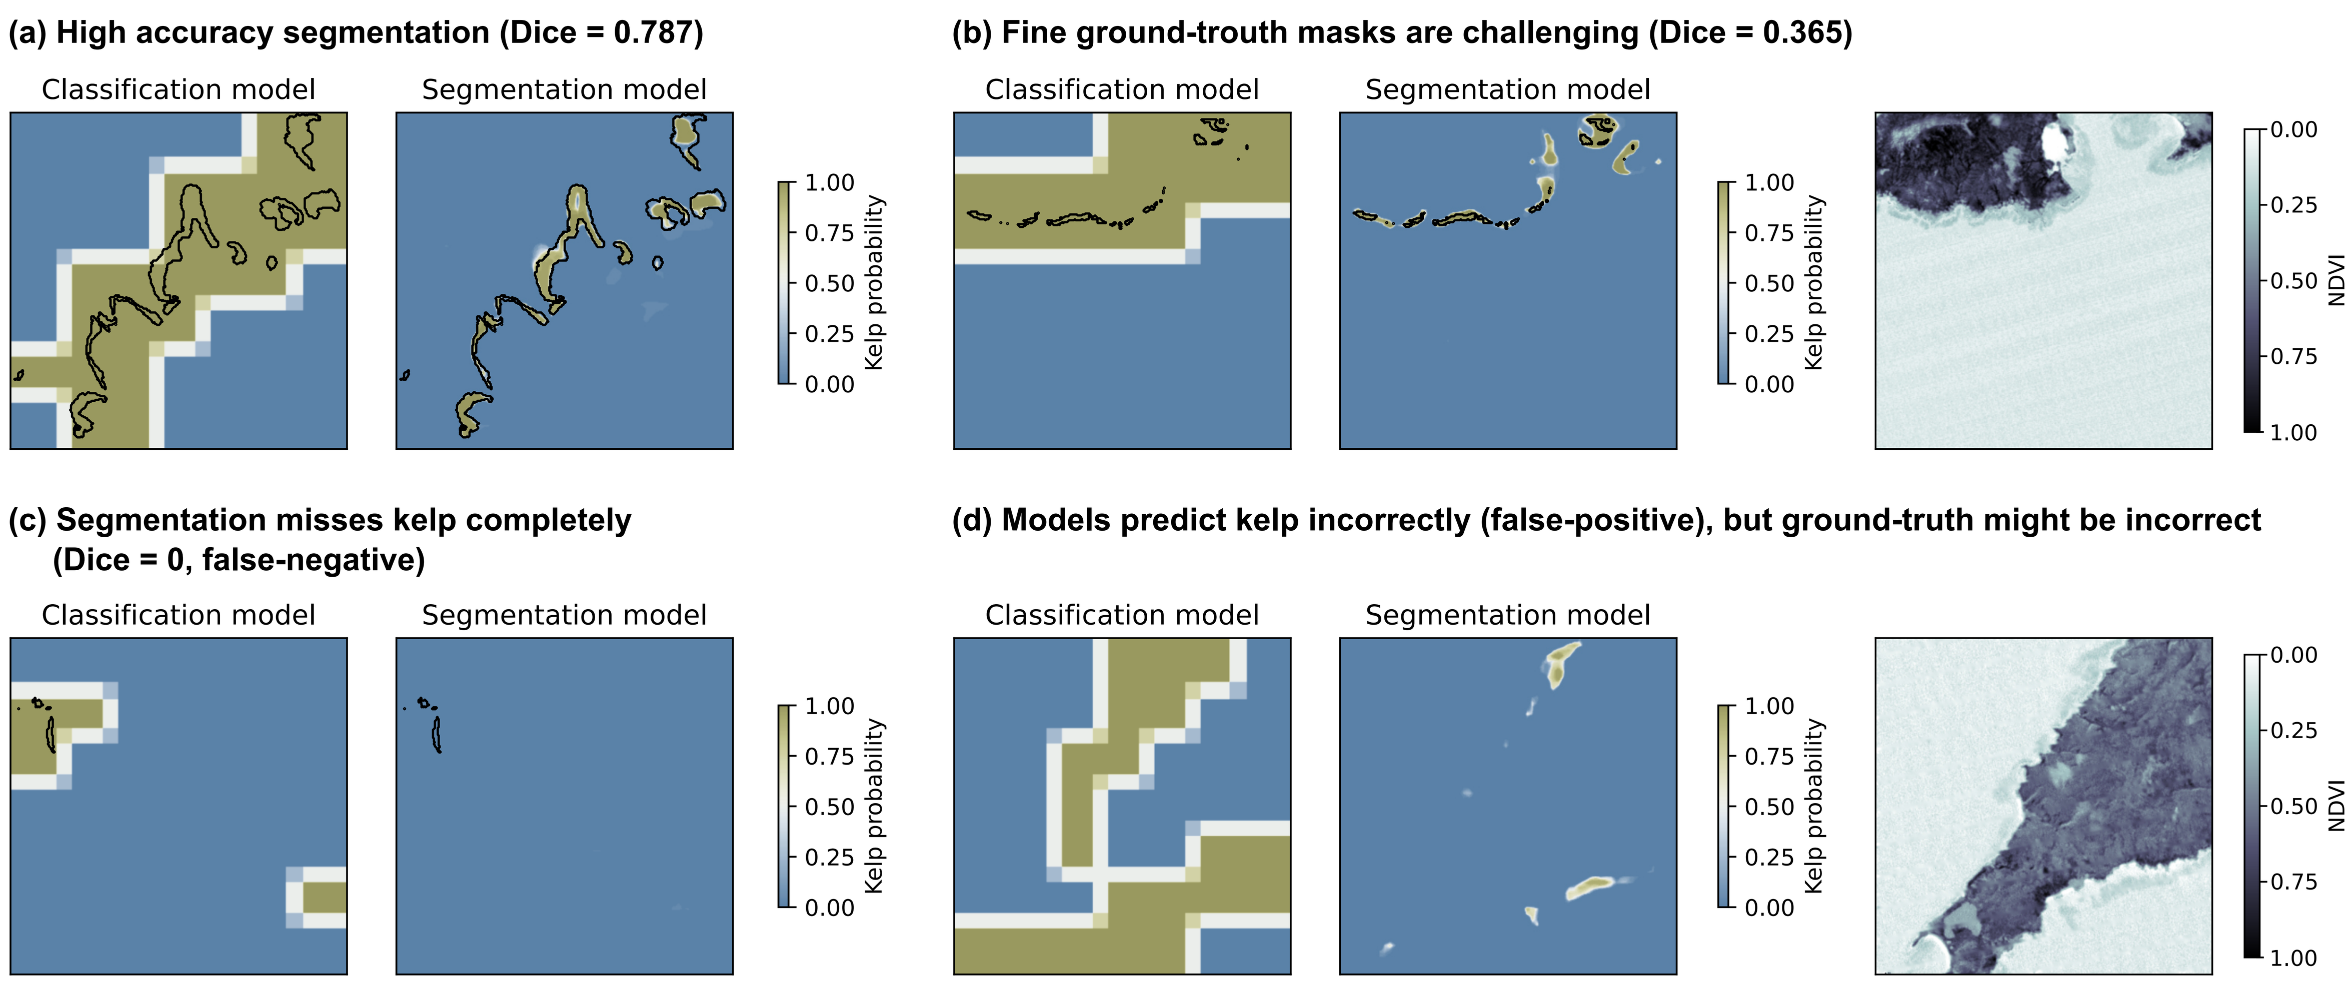

Figure 3: Characteristic example predictions from classification and segmentation model. Scores in panel titles are the scores of the combined prediction against test data.

KelpNet performs well on a large part of the test data (~850 samples). In most cases, KelpNet correctly predicts the presence of kelp (cf. Fig 3a+b) with increasing accuracy when more kelp is contained in the image, typically due to larger patches. Also, 86% of images not containing kelp are correctly labeled as such (true-negative) with only 14% false-positive kelp predictions. Manual investigation of some false-positive cases indicates that kelp might indeed be present but is not correctly labeled (cf. Fig 3d). For ~10% of test cases, KelpNet misses the existing kelp fully. Interestingly, the classification model often identifies the correct area of interest while the segmentation model misses it (cf. Fig 3c).

Future directions

- Normalize channels based on sea data only
- Increase interpretability, e.g., by assessing channel importance
- Use coarse mask already during training of fine segmentation model
- Expand multi-task learning, e.g., by adding regression head to predict auxiliary information, such as kelp fraction

Acknowledgement

The author is grateful to

- DrivenData, The Zooniverse, Woods Hole Oceanographic Institution, Tom Bell, Henry Houskeeper, and Jarrett Byrnes for setting up the "KelpWanted" competition and providing the dataset
- Sukanta Basu (University at Albany) for providing guidance and the computational resources
- Felix Dahle and Harish Baki for fruitful discussions