

Seeing Green

Computer Vision for Kelp Mass Detection

Greenwave is a non-profit specializing in regenerative farming techniques for aquaculture. They expand the kelp farming industry by providing farmers with foundational training, innovative tools, and financial support. Combining these last two efforts, Greenwave approached the DSI to develop a computer vision pipeline that would accelerate subsidy delivery by deriving total kelp mass from an image of the crop. A sufficiently accurate model could automate Greenwave's subsidy verification and delivery process, reducing operational overhead and improving outcomes for participating farmers.

The team began with a control image set. Kelp was hung against a white background demarcated with measuring lines. An example can be seen in figure one, below. The approach involved two steps. Firstly, models were developed to extract feature vectors from the images. These numerical representations of the image are computer readable definitions of important characteristics that vision models use to understand what an image contains. The team employed three different models during the extraction process: rule-based methods, transformers, and convolutional neural networks (CNNs). The rule-based method involved the selection of specific features, followed by the setting of thresholds on pixel values of the kelp image to predict the location of the kelp in the image. The transformer approach leveraged the models' context-based analysis to derive kelp mass based on the intersection of concomitant characteristics of Kelp - such as color and shape. Finally, the CNN approach made use of a pre-trained segmentation model to learn high-level features and parameters to identify the presence of kelp. These extracted features were then fed through a trained neural network attached to the segmentation model to obtain the final weight prediction.

The results are promising. Feeding the team's CNN feature-vector into the neural network provided the best results, showing a 23% error rate and an average loss of 0.33kg. This means that for a given photo of kelp, on average the model predicts the weight within 0.33kg of the actual weight. The team believes this to be a strong starting point for the project, especially given that most of the dataset consists of smaller bunches of kelp, which the model performed even more accurately on. From here, Greenwave or future clinic teams will be able to tune the model to a higher degree of accuracy, or potentially test and fine-tune the model with photos directly from the field.



Fig 1. Kelp segmentation derived by CNN




Fig 2. CNN Prediction Accuracy




Fig 3. Transformer Prediction Accuracy