Cover Statement

This is a non-peer reviewed preprint submitted to EarthArXiv. The manuscript represents original research that has not been previously published or accepted for publication elsewhere. Please note that this version of the manuscript has not undergone peer review and may differ from the final version. We welcome feedback and comments from the community to improve the quality of our work.

The application of CNN-based Image Segmentation for Beach Recovery Tracking

Byungho Kang[1,2], and Orencio Durán Vinent[1] *

March 30, 2023

[1]Department of Ocean Engineering, Texas A&M University, College Station, TX, USA
[2](Present Address) Department of Civil and Environmental Engineering, University of Houston, Houston, TX, USA

Abstract

Coastal erosion due to extreme events can cause significant damage to coastal communities and deplete beaches. Post-storm beach recovery is a crucial natural process that rebuilds coastal morphology and reintroduces eroded sediment to the subaerial beach. However, our understanding of this process is limited, which hinders effective management of coastal risks and vulnerability. In this study, we propose a new semantic segmentation method based on Convolutional Neural Networks (CNN) to detect changes in sand composition on the beach over time using high-temporal resolution video monitoring. We validate our model using a set of beach imagery recorded at Cedar Lakes, Texas, after Hurricane Harvey in 2017. We train and predict on image patches to minimize information loss from rescaling. Our CNN-based model achieves high accuracy (mean accuracy of 95.1% and mean IOU of 86.7%) and introduces a new metric to measure potential false positives and negatives. Our study also discusses how to identify blurry or rainy beach images in advance of semantic segmentation prediction, as our model is less effective in predicting these types of images. Our findings provide valuable insights into post-storm beach recovery and demonstrate the potential of high-frequency video monitoring and CNN-based semantic segmentation for improving our understanding of beach recovery processes.

*Corresponding author: oduranvinent@tamu.edu

1 Introduction

Beach recovery involves the restoration of eroded sediment to the shoreline and the reconstruction of subaerial features, such as berms and dunes, following high-energy events [1]. The beach serves as a natural buffer that protects communities and ecological habitats along the shore [2]. Thus, understanding the recovery process is crucial for coastal engineers and authorities to design and manage beaches [3] and estimate the risk associated with climate change and storm clustering [4]. However, gaining knowledge on the recovery of the beach to its pre-storm condition is challenging due to the complex interplay of various factors. These factors include the magnitude of storm-induced erosion, sediment availability, and local topography, as well as wave and wind processes that occur at different temporal and spatial scales. The interaction of these processes complicates the beach recovery process, making it difficult to understand the occurrence and dynamics of beach recovery [5].

Recent developments in video recording equipment and data storage devices have led to an increase in field observation data, shedding new light on beach recovery processes. For instance, Masselink et al. (2014) used Argus video images to investigate the morphological response of sand bars after a storm and found that wave conditions control bar dynamics [6]. Similarly, Senechal et al. (2015) utilized video observations to investigate the post-storm recovery of the sandbar and shoreline and found that the recovery processes exhibit seasonality [7]. In addition, Davidson and Turner (2009) and Splinter et al. (2014) developed empirical models for the shoreline position after a storm based on Argus images [8, 9].

As the volume of data increases, it becomes challenging to adequately represent beach recovery processes using only phenomenological observations and parameterization. Additionally, as empirical models become more complex with increasing data volumes, it becomes more difficult to identify which parameters are relevant, hindering their general applicability. To address these challenges, a complementary approach to the empirically-based coastal modeling (top-down approach) is pattern recognition (bottom-up approach) based on Machine Learning (ML).

The recent surge in processing capacity and methodological advancements in Machine Learning have led to the deep learning renaissance and nurtured the field of computer vision studies. The Deep Convolutional Neural Network (DCNN) was initially introduced for image classification [10], and it has proven to be useful in addressing several coastal science problems. Bus-

combe and Carini (2019) and Eadi Stringari et al. (2021) have used CNN to classify wave breaking in infrared [11] and video images [12]. Ellenson et al. (2020) proposed an automatic beach state recognition model based on CNN application to *timex* imagery [13], while Liu et al. (2021) developed a CNN model with eight different classic beach states [14]. However, the beach environment integrates different types of terrain, topography, and sand into a single zone, and to holistically analyze the beach recovery process, spatial information is necessary in addition to semantics to localize different objects and boundaries in images.

Image segmentation is the technique of partitioning a digital image into different subsets of image segments to allow additional image processing. The development of convolutional neural networks (CNNs) has greatly improved the accuracy of image segmentation tasks. The original study for pixel-level classification based on fully convolutional networks (FCN) [15] served as a basis for CNN-based image segmentation. Despite the usefulness of CNN-based image segmentation, however, the unique characteristics of the beach topography—different types of the sand layer with no apparent boundaries between interfaces, uncertainty related to spatial resolution due to photogrammetric distortion, and observational vulnerability to adverse weather and lighting conditions—make it challenging to apply the algorithm to analyze the short-range beach imagery for morphodynamic studies. Hence, a literature search revealed that relatively little research had been conducted on image segmentation for marine and coastal studies, and none of them had studied how the CNN-based image segmentation technique could monitor changes in beach morphology.

This paper investigates how CNN-based image segmentation can track sand changes during beach recovery after a storm. By tracking the sand composition in each image, we can determine the physical process acting at particular times during beach recovery. Sand texture can be used to infer water content (i.e. wet, dry or "rough" sand) and thus whether flooding, rain, water-driven sand transport or wind-driven (also called aeolian) sand transport took place. The physical processes in-turn affects beach recovery, which is essentially driven by aeolian sand transport behind dune formation and beach growth. We suggest ad hoc methods for training CNN for image segmentation and prediction for beach imagery analysis. We validate the prediction results by comparing them with the measurement, followed by qualitative validation of the prediction. Furthermore, we investigate instances where the model does not predict as intended and introduce possible

measures to determine those images prior to the image segmentation predictions.

2 Methods

We employed Convolutional Neural Networks (CNNs), a Machine Learning method, for image segmentation to analyze beach images. CNNs can effectively predict unlabeled images by learning the correlation between images and corresponding labels, outputting either a single discrete label (classification) or pixel-wise labels for an image (segmentation).

2.1 Dataset Acquisition and Annotation

To train and validate the CNN, we acquired images from three solar-powered stationary GoPro cameras installed in Cedar Lakes, Texas (N 28.819°,W 95.519°), which monitored beach recovery following Hurricane Harvey in 2017 (Kang et al., in review). The cameras captured over 51,000 images from November 2017 to May 2018. We selected 156 high-quality images from this dataset and hand-labeled them to train and validate the CNN.

To ensure labeling consistency, a single annotator labeled all the images in the dataset. The annotator labeled only the apparent part of each image (see Figure 1) and classified those regions into seven classes: water, sky, dry sand, wet sand, rough sand, objects, and vegetation. These classes correspond to different processes that control beach recovery.

We distinguished between "wet sand" and "rough sand" by their texture and moisture content. "Dry sand" refers to the sand region with a smooth texture and low moisture content (bright color), which can be considered a trace of aeolian transport. In contrast, "wet sand" is the sand region with a smooth texture but has a high moisture content (dark color) that indicates recent inundation. "Rough sand" is the sand region with a rough texture with an intermediate moisture content or organic crust, which can be regarded as a sign of precipitation in the near past. Hence, both "wet sand" and "rough sand" cannot be transported via aeolian transport.

We believe that this carefully annotated dataset will enable accurate and detailed analysis of beach recovery processes using image segmentation techniques.

2.2 Training CNN for Image Segmentation

The subsection presents the process of training a CNN model for semantic segmentation. The Deeplab v3+ model based on Resnet-18 was chosen as the backbone for the CNN. Out of the 156 images, 132 images were used to train Resnet-18, following the results of previous studies (Kang et al, in review) that showed that more than 100 images are sufficient for reliable semantic segmentation prediction through transfer learning.

2.2.1 Patchwise Training

To overcome the memory limitation of GPUs, the images were cropped into a number of patches, and these patches were used for training the CNN in this study. The cropping process allowed the training of the network without significantly reducing spatial resolution, which is important for investigating complex coastal regions that are often uncertain.

The high-resolution images used in this study are large, with a size of 2560 pixels (width) $\times$ 1920 (height). From these images, non-overlapping local patches with a size of 640 pixels $\times$ 640 pixels were extracted. This patch size was chosen as the greatest common divisor (GCD) of the width and height of the input image, which can minimize the zero padding from reshaping while keeping as much contextual information as possible in a single patch.

Each of the 132 images was split into 12 patches, resulting in a total of 1584 patches that were used to retrain the Imagenet pretrained Resnet-18 model with the Deeplab v3+ architecture. During training, these patches were augmented by flipping each patch horizontally with 50% probability, randomly translating them in both x and y-direction with the range of [-100, 100] pixels, and rotating them randomly within the [-20, 20] degree range.

To account for the imbalanced number of pixels for each class in the training image set, the weighted cross-entropy loss was used for the classification layer during training. The patches were shuffled randomly for every epoch during training, with a mini-batch size of 8. The optimization was performed using Stochastic Gradient Descent (SGD) with a momentum value of 0.9, an initial learning rate of 0.001, and a decrease ratio of 0.5 for every five epochs, which was repeated until 30 epochs.

Overall, the patchwise training approach allowed for the use of high-resolution images in training the CNN model, while also increasing the

dataset by generating a larger number of independent samples for training. The augmentation techniques further improved the model's performance and generalization ability.

2.3 Cut and Stitch Method for Prediction

The prediction accuracy of the trained network varies for different parts of coastal images. We observed that predictions made at the center of a patch were more accurate than those made near the edge of a patch, which is consistent with earlier findings [16]. In other words, semantic segmentation of a pixel at the edge of an image is more likely to be incorrect, while segmentation of the same pixel toward the center of the input layer produces a correct prediction.

This translational variance (pixel label prediction varying with respect to its position) is assumed to result from the interference of zero-padding on convolution [17]. Even though there is no zero padding for the input patches, the feature map size changes as information cascades through the pooling layers. Therefore, a few feature maps during the feed-forward process need to pad zeros to apply fixed-size filters for convolution.

To address the translational variance, we clipped 35 patch labels with different x and y translations and used only the central section of each predicted patch instead of simply concatenating the 12 patch labels as training. We discarded 160 pixels from the edge of each patch, except for those located at the perimeter of the original image, as recommended by [17]. This approach allowed more margin to deal with zero-padding during the feed-forward process than the 20 to 40 pixels suggested in earlier recommendations (Figure 2).

3 Validation

In the results section, we applied the trained model to predict different images in the validation set to evaluate the performance of the current model.

3.1 Validation with the labeled regions

We applied the trained model to predict different images in the validation set, and the predicted labels matched well with the hand annotation (Figure

3 (b) and (c)), despite the presence of some noise near the interface of each class.

Table 1: Metrics for semantic segmentation performance.

Global Precision[1] (%)	**Mean Precision (%)**	**Mean IoU (%)**	**Weighted IoU[2] (%)**
94.6	95.1	86.7	90.2

[1] Global Precision is the proportion between correctly classified pixels and total pixels, regardless of class.
[2] Weighted IoU is the mean IOU weighted by the number of pixels for each class.

The precision and Intersection of Union score of the predicted labels and the manual annotation in the validation image set indicated that the model correctly predicted the labeled regions (Table 1).

$$\text{Precision} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Negative}} \tag{1}$$

$$\text{Intersection of Union} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive} + \text{False Negative}} \tag{2}$$

The mean precision (95.1%) and Intersection of Union (86.7%) (Table 1) showed that the false-positive error, which arose from predicting the unlabeled region or incorrect predictions, was more than twice the false negatives. The false-positive error was larger in classes with smaller areas because the weighted Intersection of Union, which considered the number of pixels (size of each region), was greater than the mean Intersection of Union.

In addition to the metric that show the general performance of the model, Table 2 showed the Precision and Intersection of Union for each class (Figure 4). Interestingly, the "object" (Intersection of Union score: 74.8%) and "vegetation" (Intersection of Union score: 78.3%) classes showed low prediction performance compared to other classes that had Intersection of Union scores over 90%, although both classes had clear shapes and boundaries that could help the CNN to distinguish the regions more easily. One of the most plausible reasons for this low prediction performance was the limitation of hand annotation, which necessarily involved some degree of error due to a lack of skills to annotate the details. Since labeling objects and vegetation with different shapes was labor-intensive and more difficult than labeling the monotonous boundary, not only the actual objects and vegetation but

also the small areas surrounding them were labeled as the "object" or "vegetation" classes and were used for training and validation, resulting in less Intersection of Union for such classes. Additionally, we found that the primary source of false positives in the water region (Intersection of Union score: 79.0%) was related to the false prediction of high luminance areas as the water label (Figure 4).

Table 2: Precision and IOU for each class

	Precision (%)	**IoU (%)**
Water	95.9	79.0
Sky	98.8	98.3
Dry sand	94.6	91.2
Wet sand	95.8	92.2
Rough sand	93.7	92.8
Object	89.5	74.8
Vegetation	97.7	78.3

To observe how classes were correctly classified, we compared the predicted and measured numbers of pixels for each class. The confusion matrix (Table 3) showed that the small number of classes had contributed unequally to the false positives of a particular class. For instance, while most of the predicted "water" pixels corresponded to the measured water pixels (95.9% of measured water pixels), the primary source of the false positive error for the "water" pixel was the confusion of the "wet sand" (2.1%) and "sky" areas (1.7%) as the "water" region. The confusion between "dry sand," "rough sand," "vegetation" (0.1%), and the "water" class was very low. Similarly, the model accurately classified "dry sand" classes (94.6% of measured dry sand pixels), whereas most misclassification for "dry sand" was due to confusion of "wet sand" and "objects" as "dry sand."

3.2 Estimation of Uncertainty

The fact that certain classes contribute disproportionately to the confusion suggests that incorporating information on the "most confused" class for filtering could improve prediction performance.

Since it is impossible to determine the "most confused" class without labeling each pixel, we assumed that the class with the second-highest pre-

Table 3: Normalized confusion matrix.

	Water (%)	**Sky (%)**	**Dry sand (%)**	**Wet sand (%)**	**Rough sand (%)**	**Object (%)**	**Vegetation (%)**
Water (%)	95.9	1.7	0	2.1	0	0.2	0.1
Sky (%)	0.8	98.8	0	0	0.1	0.3	0
Dry sand (%)	0	0	94.6	1.4	0	3.1	0.9
Wet sand (%)	1.6	0	0.5	95.8	0.2	1.9	0
Rough sand (%)	0	0	0.5	1.3	93.7	3.5	1.0
Object (%)	0.1	0	1.5	3.3	2.2	89.5	3.3
Vegetation (%)	0	0	0.4	0	0.2	1.6	97.7

Classes in the column direction indicates predicted class and the row direction indicates measured class.

diction probability of any pixel predicted by the model corresponds to the "most confused" class. Similarly, we assumed that the difference between the highest and the second-highest prediction probability is an estimate of the uncertainty or confusion that we can use for filtering.

For the matrix notation, we converted the two-dimensional pixel coordinate (i, j) to the one-dimensional pixel coordinate as $d \in \Omega$ with the total number of pixels N. From the one-dimensional pixel domain d, we defined l_{dk} as a binary matrix such that $l_{dk} = 1$ if the class k is the class with the highest pixel score P_d, whereas we set $l_{dk} = 0$ for the other cases. Also, we defined m_{dk} as another binary matrix where $m_{dk} = 1$ if k is the class with the second-highest pixel score Q_d and $m_{dk} = 0$ otherwise.

Here, we can calculate the area fraction A_k of each given class k in $C : 1, 2, ..., 7$ as follows:

$$A_k = \frac{1}{N} \sum_{d=1}^{N} l_{dk} \tag{3}$$

In addition to the area fraction, we estimated the uncertainty or degree of confusion of prediction for each class.

As discussed previously, we only considered the domain of classes that are relevant to each other (e.g., "Water" ($k = 1$) and "Sky" ($k = 2$) can be confused, although "Sky" ($k = 2$) and "Vegetation" ($k = 7$) are difficult to confuse) when measuring the uncertainty, while classes that are unconnected to each other have been excluded.

We expanded the confusions between each region in Table **??**, and assumed that all adjacent classes would likely exhibit some level of confusion. To summarize the pair of relevant classes for each class k_u and k_v, we used a binary matrix R_{uv} as shown in Equation 4.

$$R_{uv} = \begin{pmatrix} 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 \end{pmatrix} \tag{4}$$

Based on the symmetric matrix of relevant classes R_{uv}, we estimated the uncertainty and potential error of each class. The potential false negative error (upper bound) of an arbitrary class a was defined using the Heaviside function Θ ($\Theta(x) = 0$ for$x < 0$ and $\Theta(x) = 1$ otherwise) as follows:

$$\delta_k^+ = \left[\sum_{d=1}^{N} m_{dk} \left(\sum_{n=1}^{C} l_{dn} R_{kn} \right) \Theta\{1 - (P_d - Q_d) \log_2 C\} \right] /N \tag{5}$$

Similarly, the potential false positive error (lower bound) was defined as:

$$\delta_k^- = \left[\sum_{d=1}^{N} l_{dk} \left(\sum_{n=1}^{C} m_{dn} R_{nk} \right) \Theta\{1 - (P_d - Q_d) \log_2 C\} \right] /N \tag{6}$$

where C was 7. We set the threshold for error uniformly as $\log_2 7$ for both Equation 5 and Equation 6, regardless of the number of relevant classes (i.e. sum for each row in Equation 4), in order to simplify the calculation.

For example, if "Dry sand"($k = 3$) had the highest pixel probability of 0.6 for an arbitrary pixel d, and "Rough sand" ($k = 5$) had the second highest pixel probability of 0.3, the pixel was deemed uncertain because the difference between the two probabilities was smaller than 0.35 ($= 1/\log_2 7$). However, if the difference between the pixel probabilities of "Dry sand" ($k = 3$) and "Water" ($k = 1$) for pixel i was less than 0.35, this pixel was not considered uncertain since these two classes were determined to be less likely to be confused based on the symmetric matrix R_{uv}.

Finally, we can estimate the uncertainty of a given class in each image by adding the potential false positive and false negative rates:

$$\delta_k = \delta_k^+ + \delta_k^-. \tag{7}$$

where δ_k^+ is the false positive rate and δ_k^- is the false negative rate of class k. By considering both types of errors, we can better capture the uncertainty of a given class in the image.

It is worth noting that the proposed approach of considering the second-highest pixel probability can be extended to loss functions for training. Specifically, it could be used to improve upon the commonly used cross-entropy loss function, which only considers the highest pixel probability when calculating the loss. However, this idea is beyond the scope of this study and requires further investigation.

3.3 Visual comparison: Sand detection

In this subsection, we evaluate how well the model performs on general-quality images with lower resolution and luminance than the ones it was trained on. The uncertain parts of general-quality images are larger, making it harder to annotate the ground truth and quantify the accuracy of the prediction. Therefore, we rely on visual conjecture of the object image and contextual information acquired from sequences of images captured before and after the object image to approximate the regions.

To evaluate the model's ability to detect changes in sand composition, which can be viewed as the trace of different beach processes, we applied it to images representing different sand area compositions. The model performed well in capturing the changes in sand composition (Figure 5). After inundation, the regions between debris and water were predicted as wet sand, while the regions between vegetation and debris were predicted as rough sand that resulted from rain during the inundation (Figure 5B). As the wet sand dries, the rough sand expands and replaces the wet sand (Figure 5E). After the rough sand expansion, the dry sand transported from the vicinity overlays on top of the rough sand and forms the dry sand regions (Figure 5F).

The proposed filtering method based on prediction uncertainty accurately identified false positives in wet sand, as shown by the filtered wet sand (opaque red regions) distributed at the center and edge of the dry sand area in Figure 5F. However, there was a dense distribution of the filtered dry sand region (opaque yellow area) around the interface, possibly because dry sand is less transparent than other sand regions, making model predictions less accurate.

Figure 6 illustrates how the model accurately predicts images captured from coastal cameras with different angles and tracks how dry sand replaces the rough sand region. The model was able to capture the details of objects and vegetation, subtle texture changes between rough and dry sand

regions, and the interface's uncertainty. However, it failed to identify the sand streamer, the narrow stream of saltating dry sand that overtops the wet sand region (Figure 6C and F). These results suggest that the model is more sensitive to texture changes between dry and rough sand (smooth-rough) than color differences between dry and wet sand (bright-dark). The model was also able to identify the difference between objects (debris) and vegetation with high precision (Figure 6F), reconfirming that most objects get buried in sand during the dry sand expansion, while the vegetation's shoot remains unburied.

3.4 Limitation for the current model

In this paper, we have discussed the usefulness of the proposed method in tracking how the sand areas change over time. However, there are cases where the model misidentifies the coastal image, and we provide a possible solution to overcome these challenges.

3.4.1 Blurry Images

One major problem with the model is that it fails to recognize blurry images. In such cases, the model incorrectly predicts the blurry area of the image as the dry sand or the sky (Figure 7A and D), while the inference on the clear part (center) was accurate. Therefore, to ensure the reliability of image segmentation results, it is necessary to determine whether or not the image is blurry.

There are various metrics for measuring the degree of image focus (sharpness), but it is challenging to select one that works accurately for all varying image conditions [18]. Among many operators for defining the blurriness of an image, we select the Laplacian-based operator, which is known to work accurately for most image conditions [19].

The Laplacian of the image ΔI_{ij} can be obtained by convolving the discrete Laplacian filter F with the gray scale of the original image space, I_{ij}, as

$$\Delta I_{ij} = I_{ij} \star F, \quad \text{where} \quad F = \begin{pmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{pmatrix}. \tag{8}$$

The variance of this Laplacian gives the measure of sharpness of an image

[20],

$$\sigma^2(\Delta I_{ij}) = \sum_{ij}(\Delta l_{ij} - \overline{\Delta I})^2. \tag{9}$$

The threshold sharpness level for determining blurry images would depend on the types of images we are analyzing, as the image's size or the type of object that the image is capturing could also affect the sharpness indices. In this research, we empirically discovered that the sharpness indices of blurry images typically fall below 100 (Figure 8). Figure 7A had a 58.66 sharpness index, consistent with the blurry image definition. In contrast, we discovered instances where the model provides a reasonably accurate semantic prediction for images with low sharpness indices. Despite having a low sharpness index of 63.33 (Figure 7B), semantic segmentation in Figure 7E was seemingly different from Figure 7D, as most of the features, such as objects, water, vegetation, and wet sand, had been correctly predicted by the model.

This difference in prediction quality suggests that images with low sharpness indices may be divided into two categories: "true" blurry images that are actually difficult to recognize, and images that can be examined using the model. We observed that haze-affected images often have a low sharpness index, despite the fact that they are not blurry and can be analyzed by the model (Figure 7E). This means that some images with low sharpness indices can still be analyzed through image segmentation and offer useful information about the beach processes.

We have developed a methodology to classify a batch of low-sharpness photos as blurry or hazy. Any image space I_{ij} can be decomposed into direct attenuation $R_{ij}t_{ij}$ and the air-light $A(1 - t_{ij})$ [21]. The transmission t_{ij} and the scene radiance R_{ij} denote the amount of light reaching the camera and the original scene without the influence of scatter, respectively. Using the dark channel prior, the original scene radiance can be recovered as follows (for details see [21]):

$$R_{ij} = \frac{I_{ij} - A}{\max(t_{ij}, t_0)} + A, \tag{10}$$

where t_0 is a threshold value.

To quantify the improvement in sharpness of the recovered scene radiance R_{ij} compared to the original image I_{ij}, we introduced the improvement ratio

υ, defined as:

$$\upsilon = \frac{\sigma^2(\Delta R_i)}{\sigma^2(\Delta I_i)}. \tag{11}$$

We found that the improvement ratio for hazy images υ_{haze} was much larger than that for blurry images υ_{blurry}. For example, the haze image in Figure 7B had an improvement ratio of 20.86, while the sharpness index for the blurry image in Figure 7A had only marginally improved after dehazing ($\upsilon = 3.21$).

Therefore, by examining the increase in the sharpness index after dehazing, it is possible to determine whether an image is "blurry" and unsuitable for semantic prediction or "hazy" and usable for semantic segmentation, which can be particularly useful for coastal studies vulnerable to recurring sea fogs or sea sprays.

3.4.2 Potential error estimation for different image condition

Using the proposed relationship, we classified a set of images into three categories: normal, blurry, and haze images, and predicted the area fraction and relative potential error for each image sampled every hour throughout the observation period. Figure 9 shows the distribution of relative false positive (δ_3^-/A_3) and false negative (δ_3^+/A_3) potential errors for the "normal condition" dry sand, which is less than 30% for dry sand area fractions larger than 20% ($A_3 > 20\%$). The error bounding with respect to larger areas can be found for all the other regions, including wet sand ($\delta_4^\pm/A_4 \lesssim 20\%$ for $A_4 > 10\%$), rough sand ($\delta_5^\pm/A_5 \lesssim 20\%$ for $A_5 > 20\%$), and objects ($\delta_6^\pm/A_6 \lesssim 20\%$ for $A_6 > 10\%$), suggesting that the model reliably analyzes larger regions that are important for capturing beach processes.

Furthermore, haze images generally had similar dry sand area fractions and relative potential errors ($A_3 \simeq 60\%, \delta_3^\pm/A_3 \simeq 10\%$) to the normal condition images. Wet sand area fraction and the potential error for haze images had a wide range of values ($10\% \leq A_4 \leq 30\%, 10\% \leq \delta_4^\pm/A_4 \leq 20\%$) but also showed much similar values to the normal images than the blurry images. These differences in wet area fraction values between haze and blurry images are consistent with empirical findings that a CNN can predict haze images more reliably than blurry images.

Blurry images had a dry sand area fraction of over 60% with very high certainty ($\delta_3^\pm/A_3 < 10\%$), while the area fraction of wet and rough sand was mostly limited to less than 10% (Figure 9). This finding confirms that, with

a very low level of uncertainty, the model falsely recognizes the blurry part of an image as the dry sand region while failing to capture any other features of other regions.

4 Application to coastal flooding and after-storm recovery

In the previous section, we outlined our approach to image segmentation and potential false negative and positive estimation. We used this approach to measure the time series of area fraction $A_a(t)$ and error bounds $\delta_a^{\pm}(t)$ of each class to track the beach's recovery after the storm. We conducted the analysis for a period of 220 days from November 16th, 2017, to June 23rd, 2018. To ensure the quality of our semantic segmentation prediction, we excluded images captured during sunrise and sunset when the sun's altitude changes rapidly. Only images captured between 7:00 AM and 6:00 PM each day were used for the analysis.

As discussed earlier, we removed blurry images while retaining hazy images when measuring the time series $A_k(t)$ and $\delta_k^{\pm}(t)$. We defined a blurry image as any image that satisfied two conditions: (a) its sharpness index was less than 100 and (b) the sharpness index for its recovered (dehazed) image was less than 1,000. The condition (b) was based on the assumption that, due to the large improvement factor υ for haze images, all dehazed pictures, regardless of their original sharpness indices, would have a sharpness index of at least 1,000 if the image was a haze image.

To obtain daily time series for each region and their potential error, we averaged $A_k(t)$ and $\delta_k(t)$.

4.1 Daily time series

We tracked changes in the composition of sand and other regions by measuring the daily time series of area fractions. After the camera was replaced on day 117, the field of view shifted from cross-offshore to cross-shore direction, resulting in a significant shift in the time series (Figure 10). The replacement caused the camera to capture more of the sky, water, and wet sand regions, while vegetation, objects, and rough sand regions, which previously comprised the majority of recorded images, were narrowed or not presented in the image. During coastal inundation on the 19th observation day, the

object area fractions increased significantly from 5 percent to 20 percent as large debris was transported to the back beach (Figure 11), and this inundation also transformed most of the rough sand regions into wet sand, which roughly doubled after flooding (Figure 10).

What stands out in the time series in Figure 10 is the increase in dry sand area from day 120 to day 214 when there was another coastal inundation. During this observation period, the proportion of dry sand area fluctuates between zero and local maxima, although the peak value of each cycle slowly increases from around 40 percent on day 125 to around 55 percent on day 205. Furthermore, the summation of dry, rough, and wet sand increased from 70 to 80 percent from day 120 to day 214, while the vegetation area increased from 0 to 5 percent. Objects, however, had decreased from 15 to 0 percent, which indicates that the objects were replaced by sand and vegetation as the sand buried them and that the beach may have risen vertically (Figure 10).

We speculate that the transition from a non-recovering to a recovering state may have influenced the trend in dry sand area, leading to vertical accretion of the beach. The dry sand area pattern changed after a massive aeolian transport on day 125 (Figure 10 and middle Figure 12). It is plausible that the vertical accretion due to the massive transport event shifted the beach's state to where it could accumulate more dry sand (bottom Figure 12).

5 Conclusion

In conclusion, this study presents a novel methodology for accurately analyzing coastal images during beach recovery using CNN-based image segmentation, while addressing label uncertainty, information loss, and memory issues during CNN training. The proposed measure for potential false positives and false negatives in the prediction provides a useful tool for quantifying prediction uncertainty. The study also proposes techniques to address circumstances that negatively impact segmentation prediction, such as blurry images, based on sharpness indices and image dehazing.

The application of the CNN to monitor beach evolution over time represents the first attempt to monitor beach recovery with high-temporal scale. The study suggests that a small number of massive aeolian transport events are the primary driver of the beach recovery process, which vertically grows the beach and changes the beach condition afterward.

Future work may involve further validation of the proposed methodology, including cross-checking prediction results with additional cameras and exploring the potential of the new false positive/negative measure for other image segmentation tasks. The study's findings may have broader applications beyond beach sand imagery analysis, particularly for tasks requiring high-resolution images and addressing and quantifying the uncertainty of chaotic layers without clear boundaries, such as analyzing geological or glaciological mélange and defining chaotic structures in sedimentary basins.

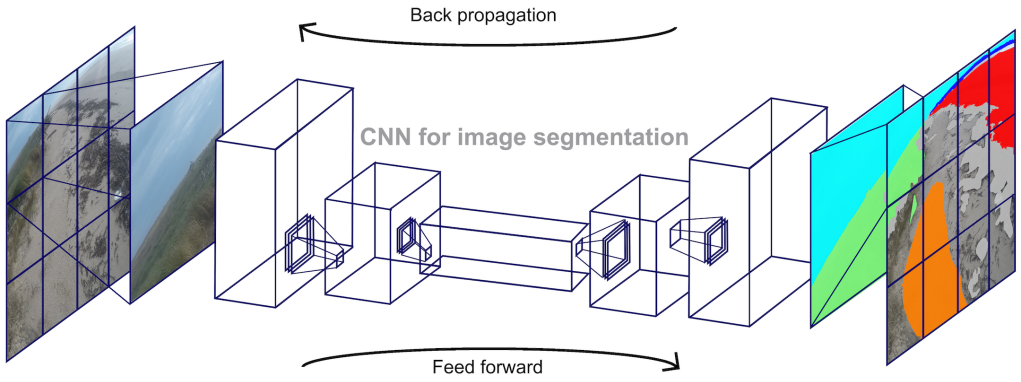

Figure 1: An illustration of the patch-wise training approach. High-resolution images (left) and their corresponding manual annotations (right) were divided into 12 patches at the pre-processing stage. Only the verifiable regions (colored in the manual annotations) were labeled for training the CNN, while ambiguous areas were not included to avoid influencing the training process.

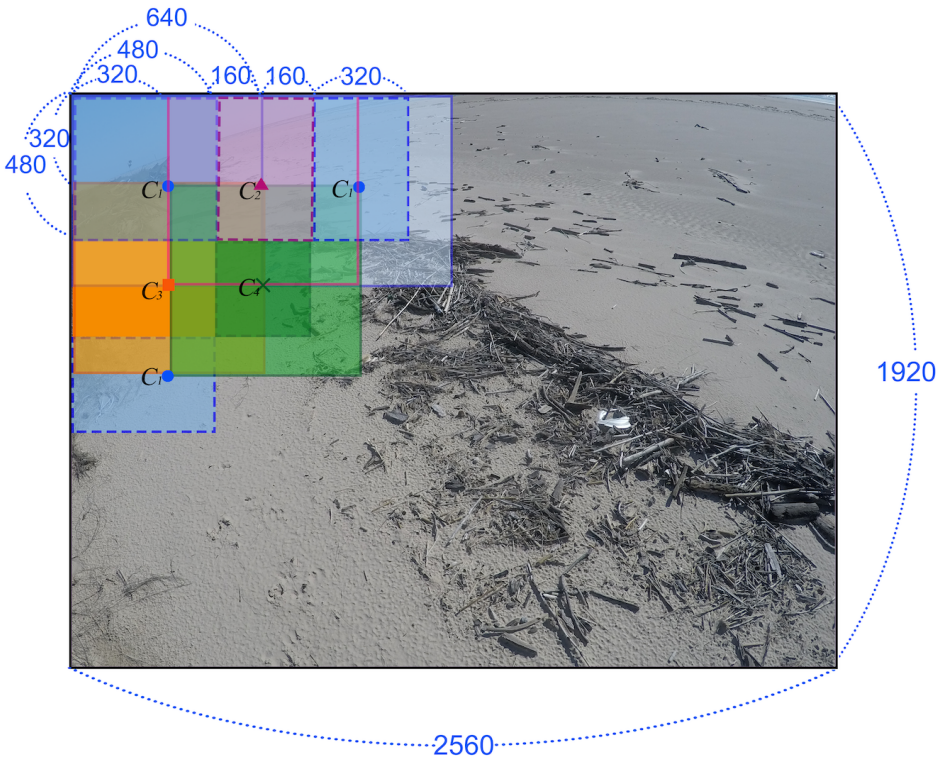

Figure 2: An illustration of the cut and stitch method used for refining the prediction. The method involves dividing the original image into 35 patches, including 12 original patches (C1), 9 patches with x-translation of 480 pixels (C2), 8 patches with y-translation of 480 pixels (C3), and 6 patches with both x and y translations of 480 pixels (C4). The patches are then added up to create a single image segmentation prediction.

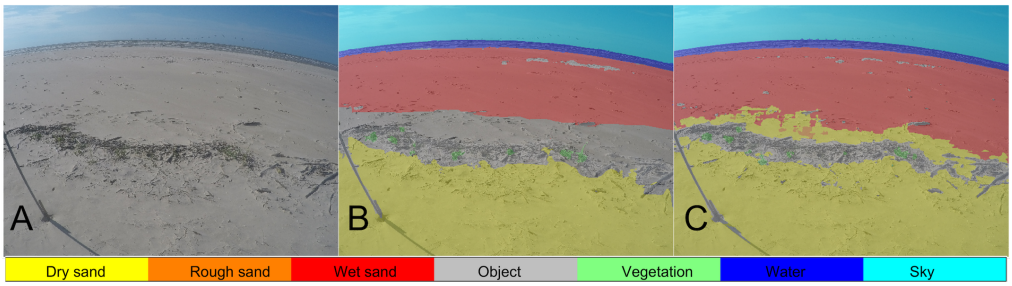

Figure 3: (A) Beach image (B) Beach image overlayed with the hand annotated ground truth (measured label) (C) Beach image overlayed with the prediction from the model. The central part of (B) was left unlabeled, while every pixel was predicted in (C).

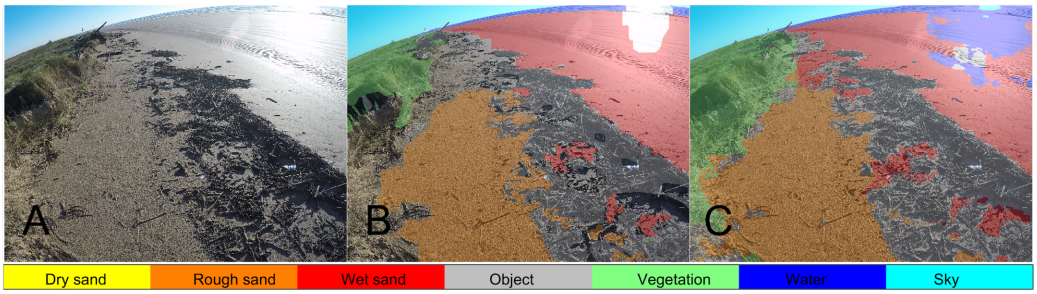

Figure 4: (A) Beach image (B) Beach image overlayed with the hand annotated ground truth (measured label) (C) Beach image overlayed with the prediction from the model. Due to the high luminance, the upper part of the region was left unlabeled in (B), where the model falsely identifies this as the water region.

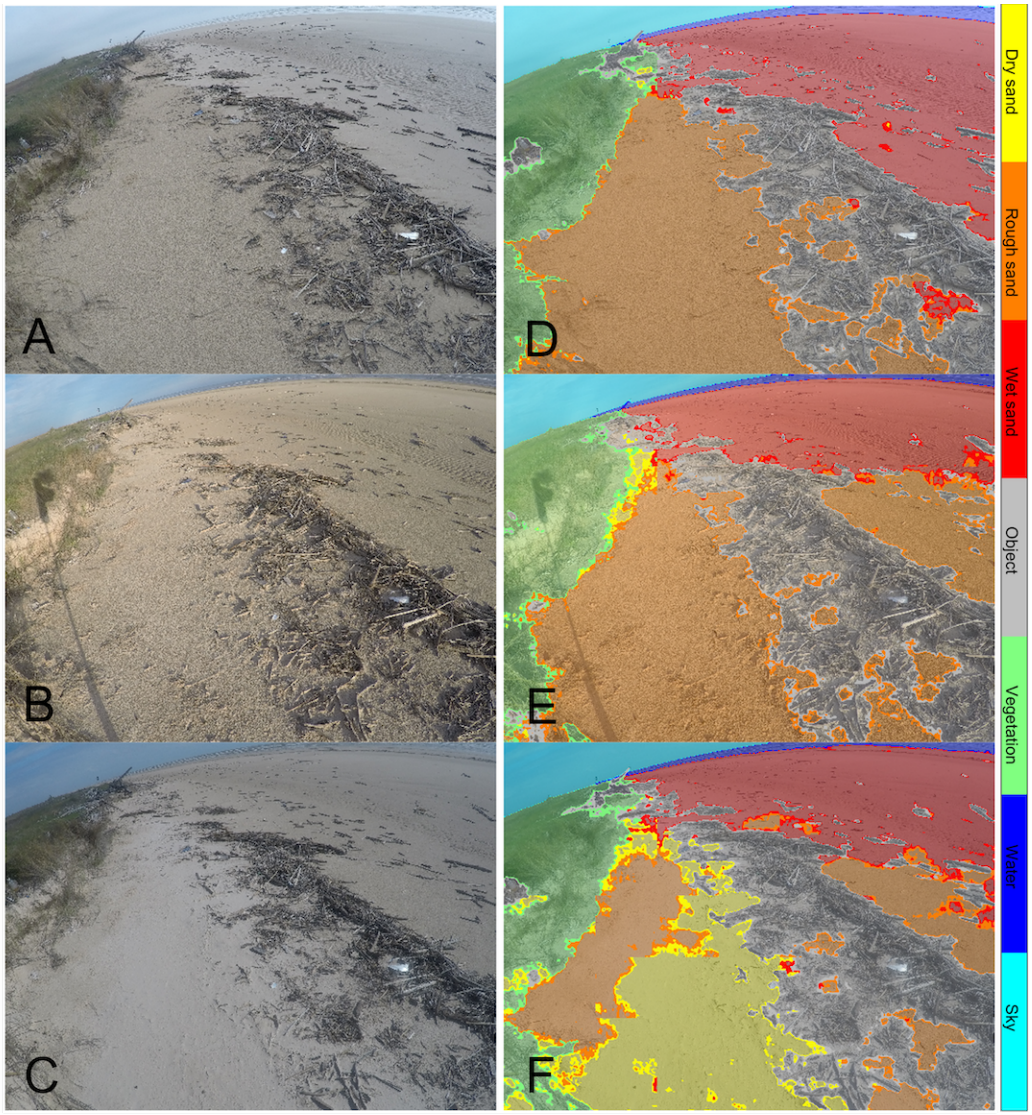

Figure 5: (A) Beach image with large wet sand, (B) beach image with large rough sand, (C) beach image where the dry sand is overlayed on the rough sand; (D),(E), and (F) are semantic segmentation predictions overlayed on the original images (A),(B) and (C). Transparent labels indicate regions predicted with certainty $(A_a - \delta_a^-)$, while opaque labels represent the potential false-positive part (δ_a^-) of the semantic segmentation.

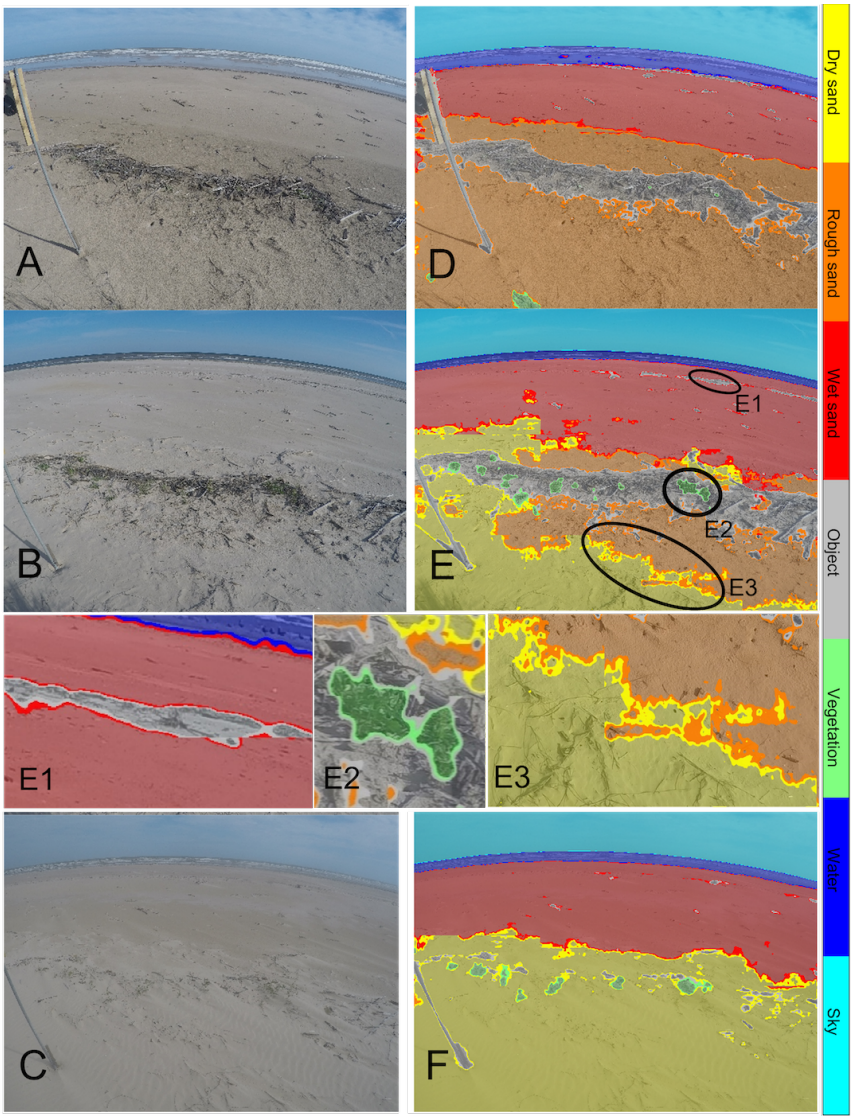

Figure 6: (A) Beach image with rough sand, (B) beach image where dry sand has replaced half of the rough sand, and (C) beach image where the dry sand has replaced all of the rough sand after aeolian transport; (D), (E), and (F) are semantic segmentation predictions overlaid on the original images (A), (B), and (C). Enlarged prediction results for the objects (E1), vegetation (E2), and the interface between the dry and rough sand regions (E3) demonstrate the accuracy of the model. Transparent labels indicate regions predicted with certainty ($A_a - \delta_a^-$), while opaque labels represent the potential false-positive part (δ_a^-) of the semantic segmentation.

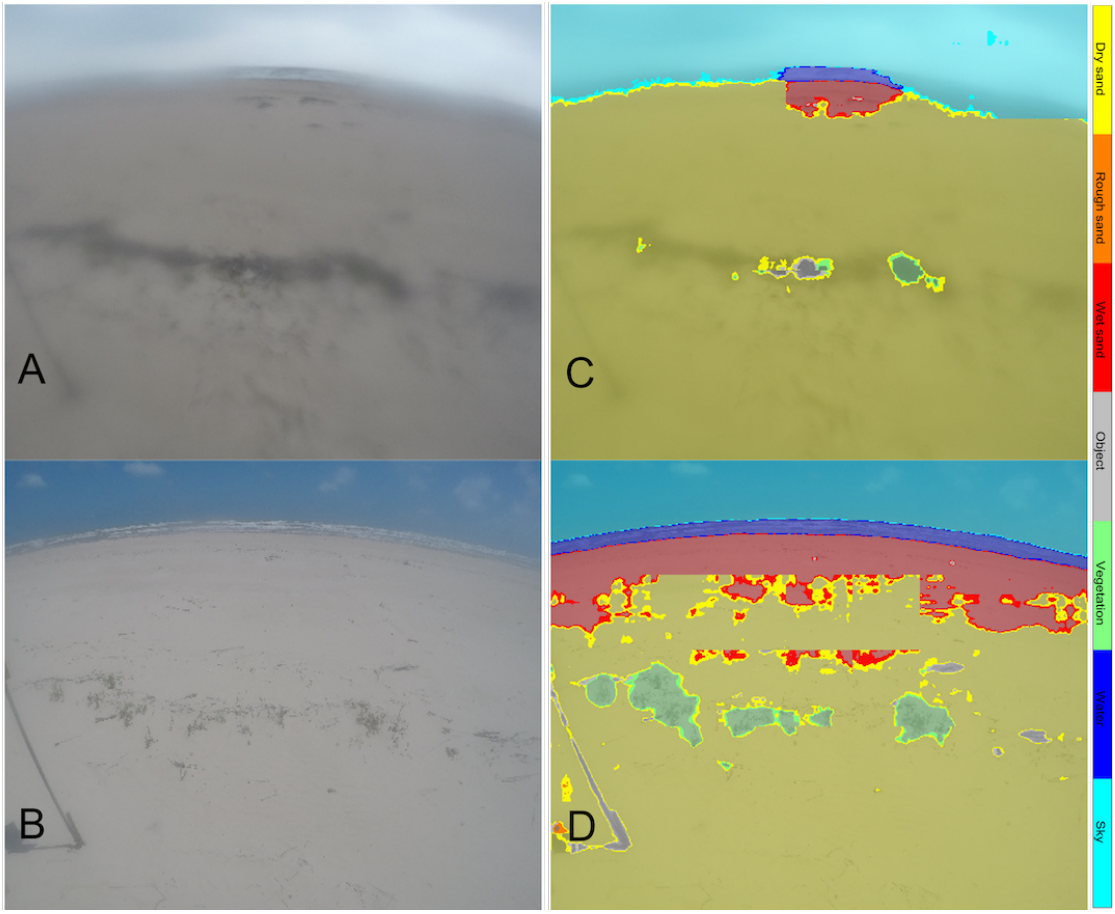

Figure 7: (A) Blurry beach image, (B) beach image degraded by haze; (C),and (D)are semantic segmentation predictions overlaid on the original images (A), and (B). Semantic segmentation (C) misses most of the beach features, whereas (D) can predict most of the features, such as water, sky, objects, and vegetation. Transparent labels indicate regions predicted with certainty ($A_a - \delta_a^-$), while opaque labels represent the potential false-positive part (δ_a^-) of the semantic segmentation.

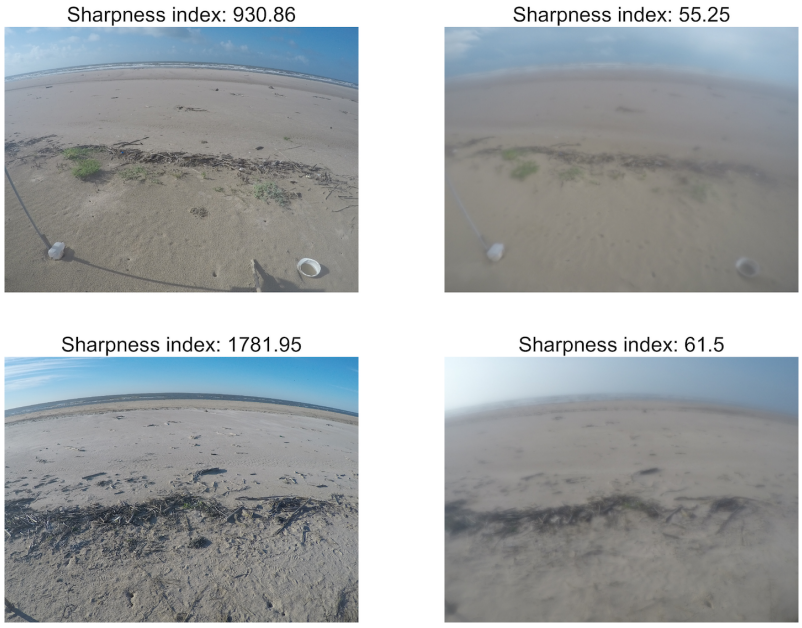

Figure 8: Images with sharpness indices. Images with a sharpness index higher than 100 can be regarded as "normal" images (left column), while images with a sharpness index of less than 100 are "blurry" images (right column).

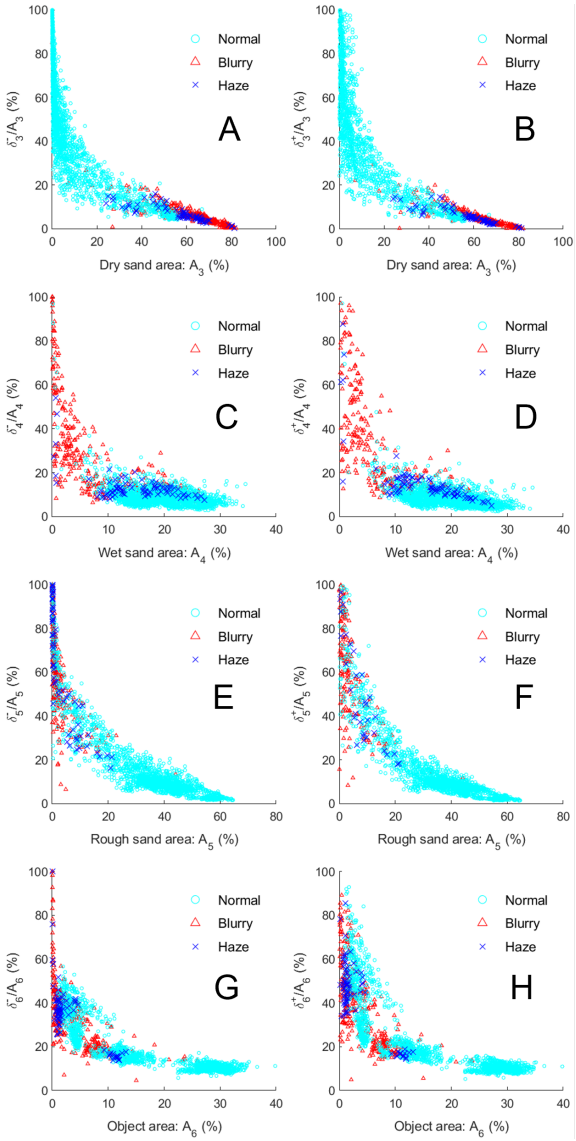

Figure 9: The left column shows the predicted area fraction and relative potential false positive (δ^-) errors for each region, while the right column displays the false negative (δ^+) errors for each region. Images with a sharpness index greater than 100 were categorized as "normal," while those with a sharpness index less than 100 were classified as either "haze" or "blurry" images, depending on whether the dehazed (recovered) image had a sharpness index greater than 1000.

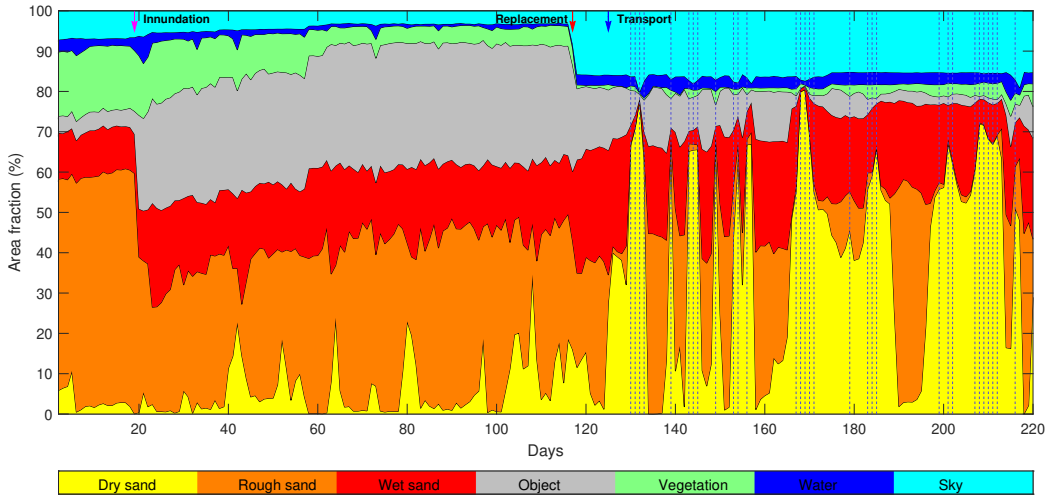

Figure 10: Time series of daily average area fractions for different classes. Blurry image days with a daily average sharpness index less than 100 and a daily average sharpness index of dehazed images less than 1000 were marked with dotted lines.

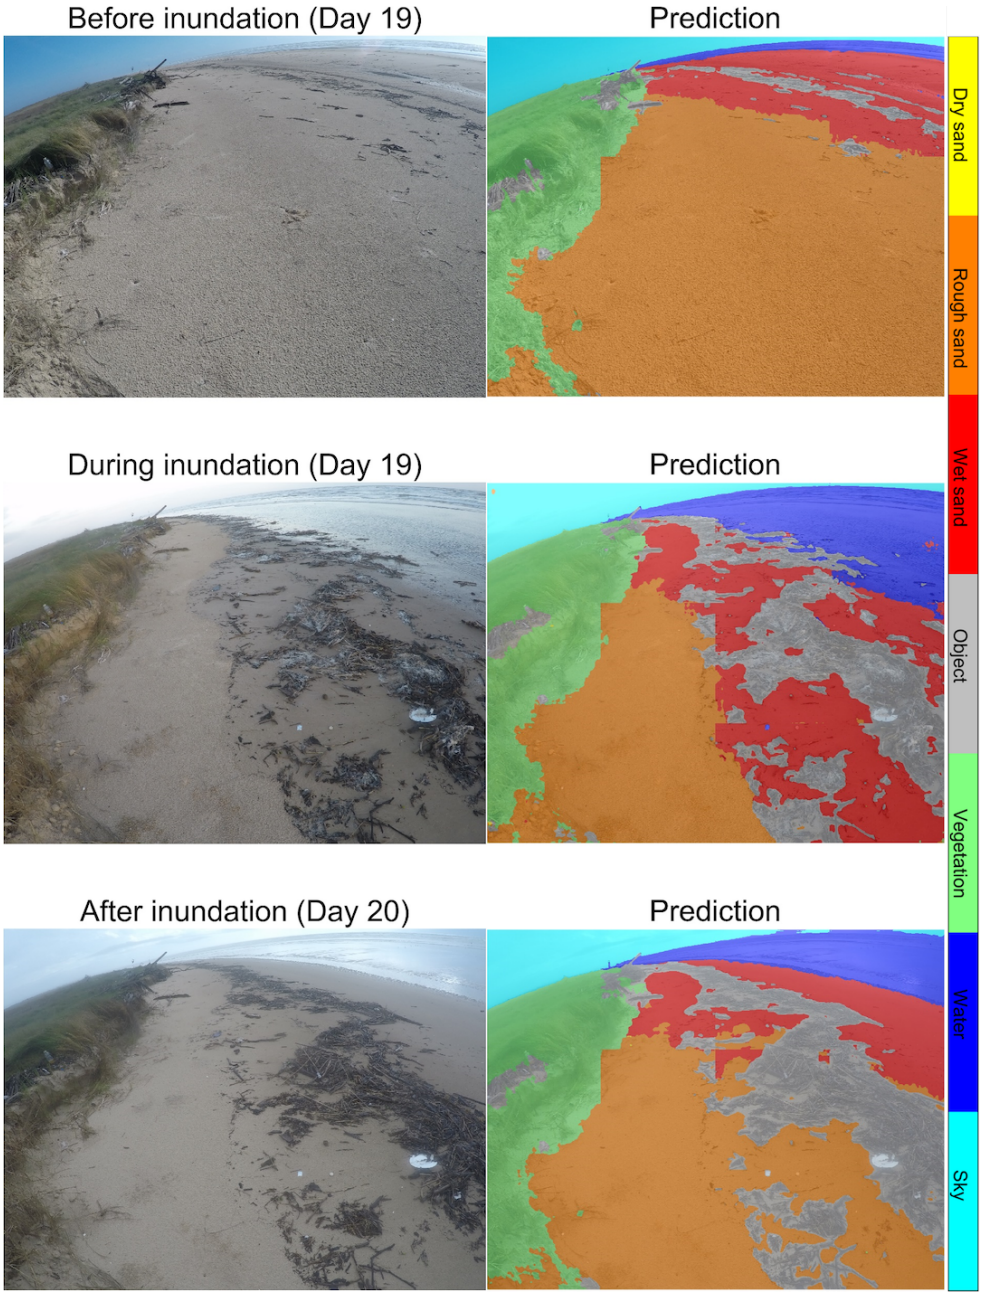

Figure 11: The inundation at the day 19. Debris that later trapped the sand was brought to the beach surface during the inundation. Following a significant increase in wet sand during the flooding, the rain on day 20 transformed the some of the wet sand into rough sand.

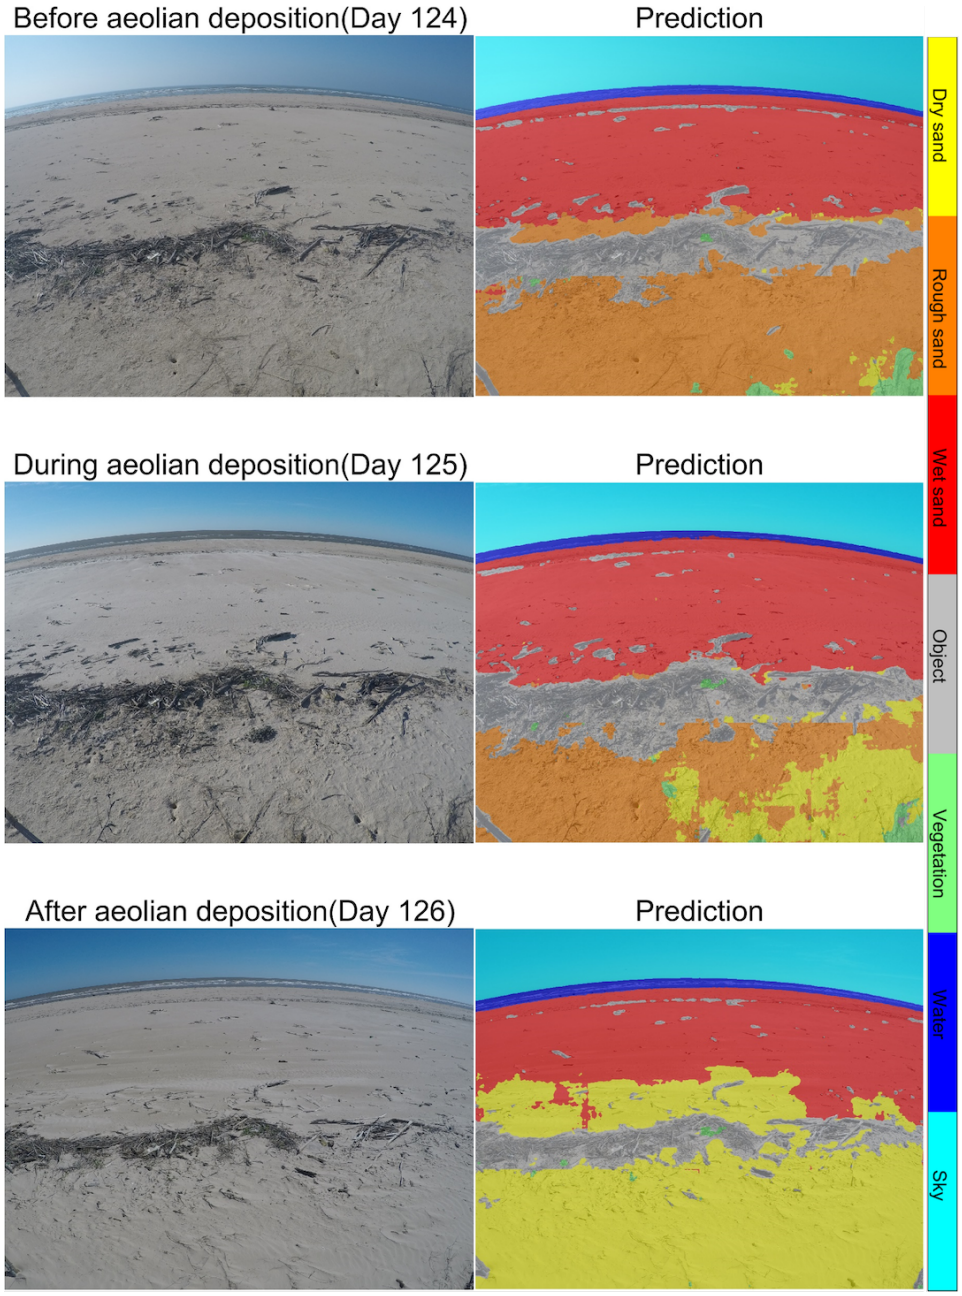

Figure 12: The aeolian transport at the day 124 to 126. Massive amounts of dry sand were moved from the berm to the beach on day 125, raising the beach elevation and changing the pattern of the dry sand area fraction after that.

References

[1] Robert A. Morton, G. Paine Jeffrey, and C. Gibeaut James. Stages and durations of post-storm beach recovery, southeastern texas coast, u.s.a. *Journal of Coastal Research*, 10(4):884–908, 1994. ISSN 07490208, 15515036. URL http://www.jstor.org/stable/4298283.

[2] CERC United States. *Shore protection manual*. Dept. of the Army, Waterways Experiment Station, Corps of Engineers, Coastal Engineering Research Center ; For sale by the Supt. of Docs., U.S. G.P.O., Vicksburg, Miss.; Washington, DC, 1984. (U.S.).

[3] N. Bascom Willard. Characteristics of natural beaches. *Coastal Engineering Proceedings*, 1(4), 1953. doi: 10.9753/icce.v4.10. URL https://icce-ojs-tamu.tdl.org/icce/index.php/icce/article/view/1588.

[4] Giovanni Coco, N. Senechal, A. Rejas, K. R. Bryan, S. Capo, J. P. Parisot, J. A. Brown, and J. H. M. MacMahan. Beach response to a sequence of extreme storms. *Geomorphology*, 204:493–501, 2014. ISSN 0169-555X. doi: https://doi.org/10.1016/j.geomorph.2013.08.028. URL https://www.sciencedirect.com/science/article/pii/S0169555X13004339.

[5] Matthew S. Phillips, Ian L. Turner, Ron J. Cox, Kristen D. Splinter, and Mitchell D. Harley. Will the sand come back?: Observations and characteristics of beach recovery. pages 676–682, 2015.

[6] Gerd Masselink, Martin Austin, Tim Scott, Tim Poate, and Paul Russell. Role of wave forcing, storms and nao in outer bar dynamics on a high-energy, macro-tidal beach. *Geomorphology*, 226:76–93, 2014. ISSN 0169-555X. doi: https://doi.org/10.1016/j.geomorph.2014.07.025. URL https://www.sciencedirect.com/science/article/pii/S0169555X14003900.

[7] Nadia Senechal, Giovanni Coco, Bruno Castelle, and Vincent Marieu. Storm impact on the seasonal shoreline dynamics of a meso- to macrotidal open sandy beach (biscarrosse, france). *Geomorphology*, 228:448–461, 2015. ISSN 0169-555X. doi: https://doi.org/10.1016/j.geomorph.2014.09.025. URL https://www.sciencedirect.com/science/article/pii/S0169555X14005005.

[8] M. A. Davidson and I. L. Turner. A behavioral template beach profile model for predicting seasonal to interannual shoreline evolution. *Journal of Geophysical Research: Earth Surface*, 114(F1), 2009. ISSN 0148-0227. doi: https://doi.org/10.1029/2007JF000888. URL https://doi.org/10.1029/2007JF000888. https://doi.org/10.1029/2007JF000888.

[9] Kristen D. Splinter, Ian L. Turner, Mark A. Davidson, Patrick Barnard, Bruno Castelle, and Joan Oltman-Shay. A generalized equilibrium model for predicting daily to interannual shoreline response. *Journal of Geophysical Research: Earth Surface*, 119(9):1936–1958, 2014. ISSN 2169-9003. doi: https://doi.org/10.1002/2014JF003106. URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2014JF003106.

[10] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4):541–551, 1989. doi: 10.1162/neco.1989.1.4.541. URL https://www.mitpressjournals.org/doi/abs/10.1162/neco.1989.1.4.541.

[11] Daniel Buscombe and Roxanne J. Carini. A data-driven approach to classifying wave breaking in infrared imagery. *Remote Sensing*, 11(7):859, 2019. ISSN 2072-4292. URL https://www.mdpi.com/2072-4292/11/7/859.

[12] Caio Eadi Stringari, Pedro Veras Guimarães, Jean-François Filipot, Fabien Leckler, and Rui Duarte. Deep neural networks for active wave breaking classification. *Scientific Reports*, 11(1):3604, 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-83188-y. URL https://doi.org/10.1038/s41598-021-83188-y.

[13] Ashley N. Ellenson, Joshua A. Simmons, Greg W. Wilson, Tyler J. Hesser, and Kristen D. Splinter. Beach state recognition using argus imagery and convolutional neural networks. *Remote Sensing*, 12(23):3953, 2020. ISSN 2072-4292. URL https://www.mdpi.com/2072-4292/12/23/3953.

[14] Bo Liu, Bin Yang, Sina Masoud-Ansari, Huina Wang, and Mark Gahegan. Coastal image classification and pattern recognition: Tairua

beach, new zealand. *Sensors*, 21(21):7352, 2021. ISSN 1424-8220. URL https://www.mdpi.com/1424-8220/21/21/7352.

[15] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. 2014. doi: 10.48550/ARXIV.1411.4038. URL https://arxiv.org/abs/1411.4038.

[16] Vladimir Iglovikov, Sergey Mushinskiy, and Vladimir Osin. Satellite imagery feature detection using deep convolutional neural network: A kaggle competition. *arXiv pre-print server*, 2017. doi: None arxiv:1706.06169. URL https://arxiv.org/abs/1706.06169.

[17] Bohao Huang, Daniel Reichman, Leslie, and Kyle. Tiling and stitching segmentation output for remote sensing: Basic challenges and recommendations. *arXiv pre-print server*, 2019. doi: None arxiv:1805.12219. URL https://arxiv.org/abs/1805.12219.

[18] Said Pertuz, Domenec Puig, and Miguel Angel Garcia. Analysis of focus measure operators for shape-from-focus. *Pattern Recognition*, 46(5): 1415–1432, 2013. ISSN 0031-3203. doi: 10.1016/j.patcog.2012.11.011. URL https://dx.doi.org/10.1016/j.patcog.2012.11.011.

[19] Megan J. Russell and Tania S. Douglas. Evaluation of autofocus algorithms for tuberculosis microscopy. IEEE. ISBN 1557-170X. doi: 10.1109/iembs.2007.4353082. URL https://dx.doi.org/10.1109/iembs.2007.4353082.

[20] J. L. Pech-Pacheco, G. Cristobal, J. Chamorro-Martinez, and J. Fernandez-Valdivia. Diatom autofocusing in brightfield microscopy: a comparative study. In *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, volume 3, pages 314–317 vol.3. ISBN 1051-4651. doi: 10.1109/ICPR.2000.903548.

[21] He Kaiming, Sun Jian, and Tang Xiaoou. Single image haze removal using dark channel prior. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1956–1963. ISBN 1063-6919. doi: 10.1109/CVPR.2009.5206515.