Intertidal Beach Classification in Infrared Images

Bas Hoonhout †‡, Fedor Baart †∞, Jaap van Thiel de Vries †

†Delft University of Technology
Department of Coastal Engineering
Delft, The Netherlands

‡ Deltares
Unit of Hydraulic Engineering
Delft, The Netherlands
bas.hoonhout@deltares.nl

∞ Deltares
Software Centre
Delft, The Netherlands

www.cerf-jcr.org

ABSTRACT

Hoonhout, B.M., Baart, F., Van Thiel de Vries, J.S.M. 2014. Intertidal Beach Classification in Infrared Images. *In:* Green, A.N. and Cooper, J.A.G. (eds.), *Proceedings 13th International Coastal Symposium* (Durban, South Africa), *Journal of Coastal Research*, Special Issue No. 66, pp. xxx-xxx, ISSN 0749-0208.

Digital imagery is a powerful datasource for coastal monitoring, maintainance and research. It provides high-resolution measurements in both time and space. The size and resolution of long-term imagery datasets provide great opportunities, but also poses problems of tractability in the data analysis. In order to fully use the possibilities of these datasets, reliable and automated classification of images is essential. This paper discusses an automated classification approach based on probabilistic graphical networks, in this case Conditional Random Fields (CRF). The algorithm is applied in pixel space only. Therefore it does not rely on in-situ measurements, nor is there a need for image rectification. The algorithm consists of three steps: segmentation, feature extraction and model training and prediction. We applied the method to a coastal thermal infrared image stream that monitors the wetting and drying of the upper intertidal beach in relation to tide an meteorological parameters. Classification of the upper intertidal beach provides information on the potential sources of Aeolian sediment. The use of 62 extracted features and structured learning proves to provide significantly better classification results compared to algorithms solely based on intrinsic features.

www.JCRonline.org

ADDITIONAL INDEX WORDS: *Coastal imagery, classification, infrared, intertidal beach, segmentation, features, conditional random fields, CRF.*

INTRODUCTION

Digital imagery is a powerful datasource for coastal monitoring, maintainance and research. It provides high-resolution measurements in both time and space. Remote sensing applications makes long-term monitoring feasible, providing unprecedented datasets. In the last couple of decades many coastal applications of in-situ imagery are developed. Holland et al. (1997) uses a generic coastal video system for nearshore monitoring. This system inspired many applications, like an intertidal beach mapper (Aarninkhof et al., 2003), subtidal bathymetry extraction (Holman et al., 2013), rip-current detection (Dongeren et al., 2013) and vegetation mapping (Schretlen and Wijnberg, 2012). Vousdoukas et al. (2012) uses a similar system for run-up monitoring. More advanced algorithms use multiple cameras for stereo rectifcation, for example of waves (de Vries et al., 2009).

The size and resolution of long-term imagery datasets provide great opportunities, but also poses problems of tractability in the data analysis. In order to fully use the possibilities of these datasets, reliable classification of images is essential. Traditionally classification of coastal images relies on algorithms tailored to a specific purpose (e.g. Aarninkhof et al., 2003). By definition such algorithms have limited applicability. More generic approaches exist, but are often not fully automated, limiting the feasibility of handling large datasets (e.g. Quartel et al., 2006). This paper discusses a fully automated classification approach based on probabilistic graphical networks. In contrast with most coastal classification algorithms, we not only use intrisic intensity features, but also include a large set of extracted features in our discriminative pixel descriptors. Although the approach is generic

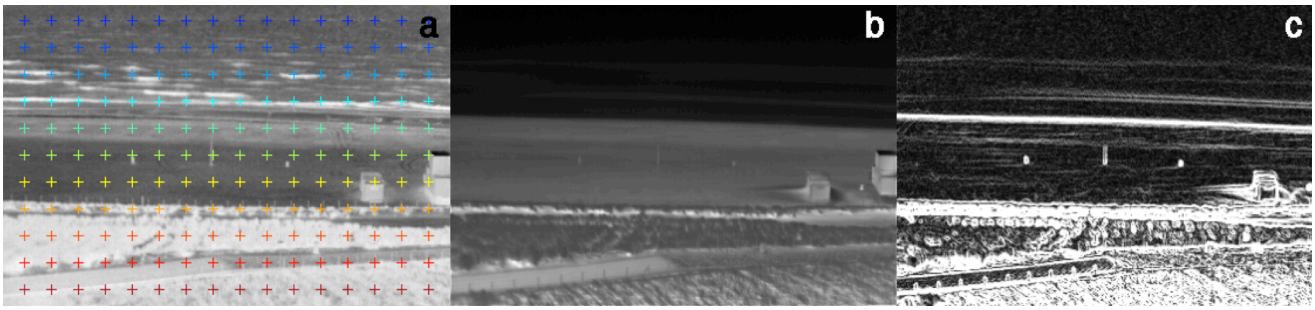

Figure 1. a) Raw infrared image. b) Arbitrary time-variance image. c) Spatial gradient in time-variance.

DOI: 10.2112/SI65-xxx.1 received Day Month 2013; accepted Day Month 2013.

for any kind of imagery, we focus on the use of coastal infrared imagery. Coastal infrared images are a valueable data source for monitoring beach moisture content as supply limiter for Aeolian sediment transport.

COASTAL INFRARED IMAGERY

Coastal infrared imagery provides high resolution and long-term information on surface temperatures. We use coastal infrared imagery to monitor beach temperatures, which are related to beach moisture content (Edwards et al., 2013). Moisture content is important for Aeolian sediment transport in supply-limited environments (de Vries et al., 2014). Large moisture content increases the threshold wind velocity for Aeolian transport and hence decreases the sediment transport volumes compared to transport-limited environments (Pye and Tsoar, 1990).

The intertidal area is of particular interest since its moisture content changes continuously under influence of tide, solar radiation, rain and wind. The intertidal area is also expected to be an important source of sediment for Aeolian transport (de Vries et al., 2010). The delicate balance between meteorology and tide is expected to explain when and what sediment is transported from the intertidal area to the dry beach and dunes.

A long-term monitoring campaign is initiated in the southern part of Holland, The Netherlands. The research area includes the 20Mm^3 mega-nourishment known as the Sand Motor (or Sand Engine; Stive et al., 2013) that will feed large amounts of sediment to the Holland coast in the coming years. In context of this campaign we monitor accretional Aeolian processes, amongst others, using a thermal infrared camera looking down at the dune, beach and sea (Figure 1a). In view of the camera we have about 120m dry sandy beach, an intertidal area with a tidal range of about 1.5m, a relatively low vegetated dune and a small road. In summertime a construction of the coastguard is present as well. Tide and weather conditions are measured at nearby meteorological stations. First results show indeed clear correlations between tide, meteorology and the wetting and drying of the beach. Tide and rain increase the moisture content in the intertidal area, and thereby decrease the temperature. In contrast, either high solar radiation or strong winds decrease the moisture content and thus the threshold velocity for Aeolian transport (Figure 2). These correlations are not obvious in regular CCD images that are available at the same site.

These first results are obtained by tracking individual pixels in time. For a full analysis of intertidal moisture content it is necessary to know the location of the intertidal area. This location varies in time due to variability in tide and morphology. We use a generic and automated classification algorithm for coastal imagery to recognize the area that is a potential source for Aeolian sediment in an infrared variance image. Infrared images are single channel, non-normalized and relatively low resolution images that provide special challenges for classification that are not widely covered in other disciplines. The algorithm is applied in pixel space only. Therefore it does not rely on in-situ measurements, nor is there a need for image rectification.

CLASSIFICATION ALGORITHM

Like in normal coastal CCD images, the intertidal is not always very distinct in coastal infrared images either. Unlike CCD images, however, thermal infrared images are non-normalized. Their output is always in degrees of temperature and the frame of reference is fixed as long as the view is fixed. Therefore subsequent thermal infrared images can easily be compared over a fairly long time span and sequences of images can be used for classification. Such sequences hold much more information than

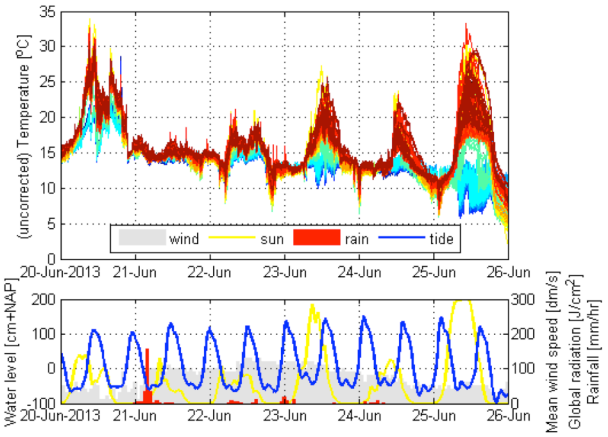

Figure 2. Uncorrected infrared surface temperatures with influence of 24 hours cycle, tide, solar radiation, wind and rain. Each line corresponds to a specific pixel indicated by the raster in Figure 1a.

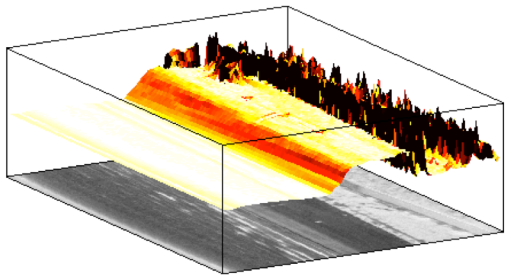

Figure 3. Raw infrared image (bottom), time-variance (surface height) and spatial gradient (surface color). The red sloped area corresponds to the upper part of the inter-tidal area.

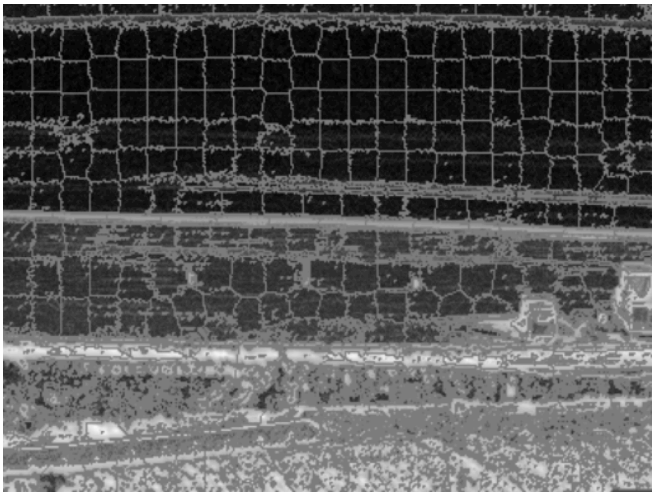

Figure 4. Segmented artificially constructed two-channel infrared variance image.

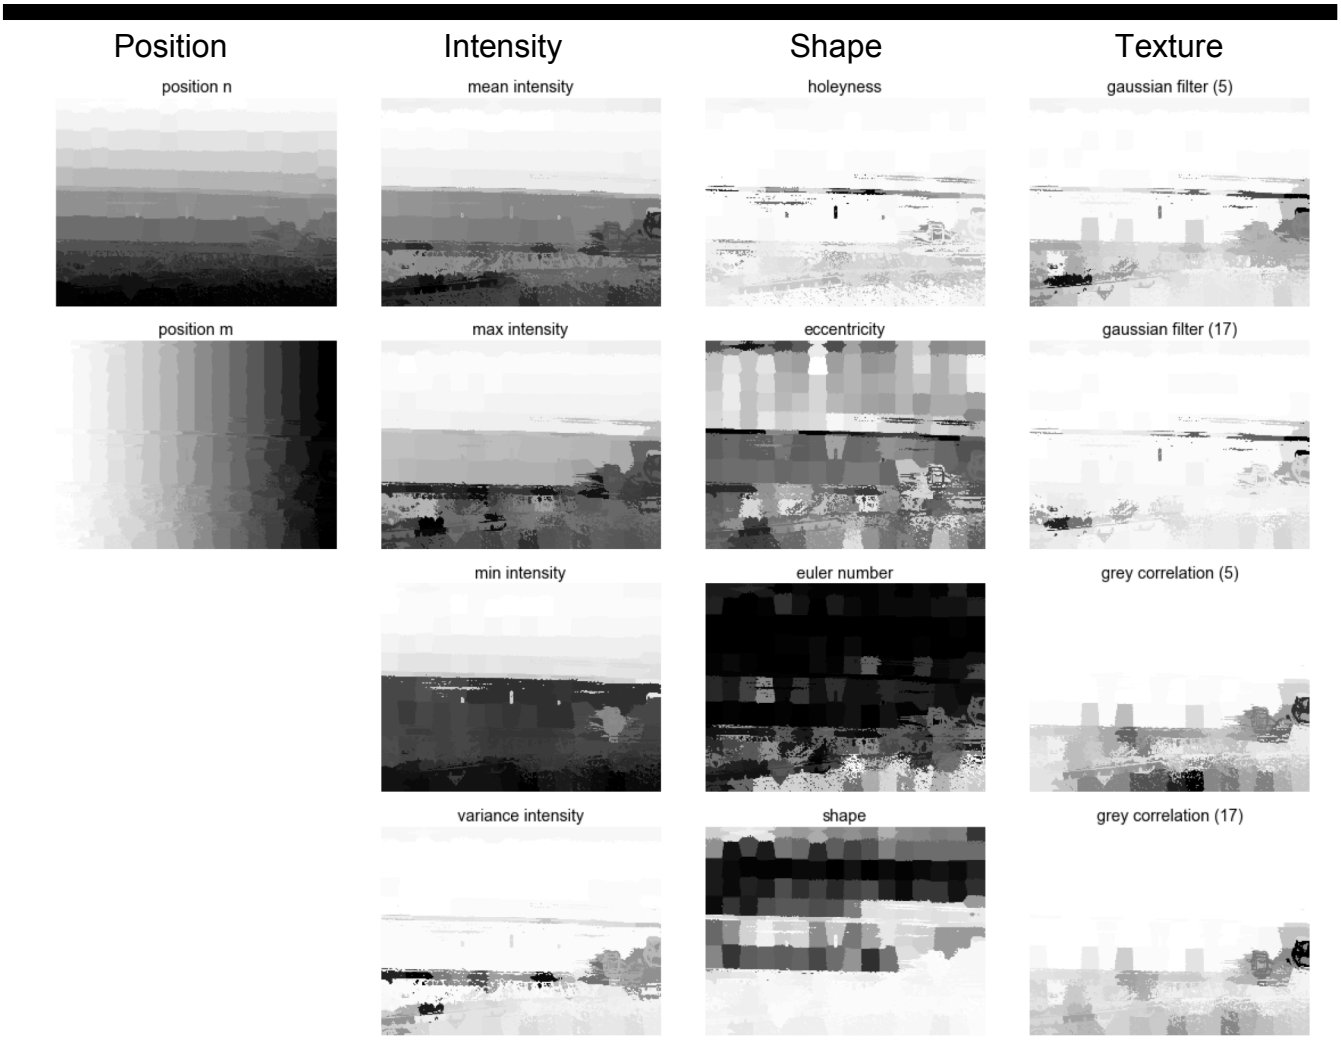

Figure 5. Several features extracted from the infrared variance image in Figure 1b. Low values are white, high values are black. The position and intensity features are referred to as intrinsic features, whereas the shape and texture features are referred to as extracted features.

single images, which makes classification of the intertidal area feasible.

The classification algorithm is based on Conditional Random Fields (CRF), which is a probabilistic graphical model with pairwise potentials (Koller and Friedman, 2009). The algorithm is based on 3 steps that are explained in the following:

1. Segmentation
2. Feature extraction
3. Model construction and training

Segmentation

Building an algorithm that classifies individual pixels will most likely result in an intractable algorithm that is very sensitive to local scatter. Instead, the image is first segmented into superpixels. A superpixel is a cluster of pixels with similar intrinsic features[1]. The classification algorithm subsequently classifies all superpixels. All pixels in a superpixel by definition share the same class. The algorithm discussed uses four intrinsic features in particular for segmentation:

1. Pixel coordinate M
2. Pixel coordinate N
3. Time-variance of infrared pixel
4. Spatial gradient in the time-variance

Figures 1b and 1c show an example of the latter two features for an arbitrary 24 hours sequence of infrared images. Figure 3 shows both features in a single 3D plot. A part that coincides with the upper part of the intertidal area is clearly visible, especially in the spatial gradient. Besides, the sea and beach part of the image are distinct through their difference in time variance. Vegetation is characterized in highly scattered values in both time-variance and spatial gradient. The differences in variance related to tide and meteorology are also visible from the plain time series depicted in Figure 2.

From the time-variance and spatial gradient features we artificially construct a two-channel infrared variance image. We use the SLIC segmentation algorithm (Achanta et al., 2010) for

[1] We define intrinsic features as those features directly related to the input image and hence to either pixel location or pixel intensity (first 2 columns in Figure 5). In the following sections we will introduce extracted features that are related to second-order properties of the image, like cluster size and shape (last two columns in Figure 5).

segmentation of these images (Figure 4). Segmentation consists of clustering pixels in a space stretched by the four features listed above using a K-means algorithm. The number of superpixels (clusters) is given as input. We use approximately 400 superpixels per image. A second parameter to the clustering algorithm is the compactness, which determines the weight of the pixel coordinate features M and N compared to the time-variance features. A high compactness results in relatively square and heterogeneous superpixels whereas a low compactness results in scattered, but homogeneous superpixels. We use a relatively high compactness of 40.

Feature extraction

In order to classify a superpixel we need features to distinct a superpixel of one class from a superpixel of another. For segmentation we already used four features regarding location and time-variance of the pixels. These features were determined based on individual pixels. When dealing with superpixels more features can be extracted that help us discriminate between classes of pixels. An important property of CRFs is that features are independent of each other. Therefore dependent features may be used without risking the result to be biased (Koller and Friedman, 2009). We use 62 features from the main categories shown in the columns of Figure 5. This figure shows a subset of the features in use.

The successful classification depends on the distinctive quality of the different features. Several features show a strong correspondence with the upper intertidal zone. The example features of position and intensity are computed using trivial functions. The shape and texture definitions are computed as follows. Holeyness is defined as the superpixel convex hull area divided by the total pixel area. The eccentricity is the ratio of the minor and major axis of an ellipsis fitted to the shape of the superpixel. The Euler number is computed as one divided by the number of holes. Shape is defined as the area divided by the squared perimeter.

The Gaussian filter (5) and (7) correspond to the variance left in a superpixel after a Gaussian filter with sigma 5 and 17 has been applied. The grey correlation correspond to the correlation with a Grey Level Co-occurrence Matrix (GLCM; Haralick et al., 1973) with interval of 5 and 17 pixels. In the full feature set the angles of the grey patterns vary by angle. Figure 5 shows that the Gaussian filter textures feature correspond to both intertidal and vegetation, whereas the GCLM texture corresponds with vegetation only.

Model construction and training

The final step of the algorithm is the construction of a model that classifies all superpixels (and thereby all pixels) in a new unseen segmented infrared variance image. The output of the model is thus an assignment for each pixel to a single class taken from a set of predetermined classes. The set of predetermined classes we use is: sky, sea, intertidal, beach, vegetation, object.

The model we construct is a Conditional Random Field (CRF). A CRF is a probabilistic graphical model. A probabilistic graphical model is a type of Bayesian network and constructed as a graph: with nodes and edges (Figure 6). The nodes represent (discrete) conditional probability distributions. The graph as a whole represents the joint conditional probability function. We use a node for each superpixel. Each node represents the discrete probability of the assignment of that superpixel over all classes. We define the class with the largest probability to be the assignment of that superpixel.

The nodes are connected by edges. The edges represent dependencies of the nodes. We use a two-dimensional grid as graph. Each node is connected to its neighboring superpixels: two at the corners, three at the edges and four elsewhere. The dependencies of adjacent superpixels are governed by pairwise potentials. Potentials are located in between ordinary nodes. Like ordinary nodes, potentials are a factor in the joint probability function of the graph, but unlike ordinary nodes, potentials do not host a probability distribution itself. Potentials in a CRF are used to increase the probability of adjacent superpixels to be assigned an equal class. The reasoning behind this is that the probability of a superpixel next to a superpixel classified as “sea” is more likely to be sea itself. Also the probability of a superpixel to be classified as “dune” is less likely when their neighbors are all classified as “sea”. The class “beach” or “sky” is more likely to be appropriate. In order to stimulate adjacent superpixels to be of equal class, potentials penalize the difference between adjacent superpixels by

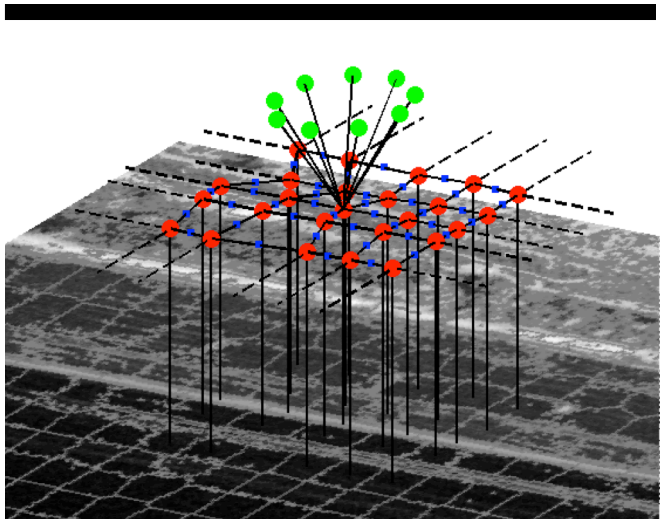

Figure 6. Impression of partial Conditional Random Field for segmented infrared variance image. Each node (red dot) is connected to a specific superpixel. All nodes are connected to adjacent nodes by edges. Each of those edges hosts a potential function (blue squares) that penalizes differences between the nodes connected by the edge. For the middle node a small number of feature nodes are shown (green dots).

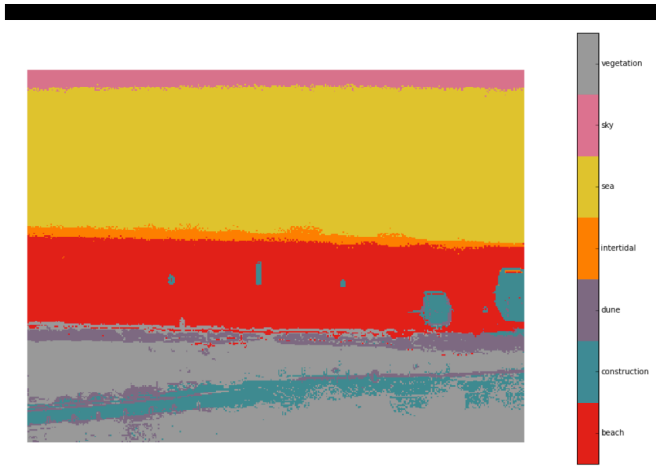

Figure 7. Infrared variance image classified using trained Conditional Random Field.

adding a simple energy function to the joint probability distribution, for example:

$$\psi_{potential} = e^{-[x_i - x_j]^2}$$

Where x_i and x_j are classes assigned to two adjacent superpixels.

The final step in constructing a CRF is to add the feature nodes. These nodes are used to condition the graph based on input (e.g. an unseen segmented infrared variance image). The probability distribution of each node is, besides its neighbors, dependent on the extracted features of that particular superpixel. Therefore each superpixel node is also connected to multiple feature nodes. (Figure 6). The feature nodes are observed (i.e. extracted from the input image) and thus these nodes are deterministic.

So far we constructed the structure of the joint probability function over all classes in the form of a graph. In order to classify images we need to fill the conditional probability distributions by training the graph with data. There are several methods available for structured learning, like message passing or structured Support Vector Machines. We use the latter because of its efficiency and availability (Andersen et al., v1.1.5). The learning algorithm iteratively approximates the maximum a-posteriori (MAP) class value for each node in the grid. This is the class assignment for all nodes with the largest joint probability.

The data we use for training are 24 hours infrared variance images from Kijkduin, The Netherlands. These images are segmented and manually classified. We used about one month of data for training (27 images).

RESULTS

Figure 7 shows an example of a classification result of the trained Conditional Random Field (CRF). The major parts of the image are well recognized. This is mainly due to the use of pairwise potentials and a large set of extracted superpixel features like shape and holeyness. Figure 8 shows a classification result where only intrinsic intensity features and no pairwise potentials are used. Vegetation and intertidal area, beach, surfzone and sky and even constructions and vegetation are easily interchanged if we only rely on intensity related features. This shows the great potential of structured learning of segmented images with extracted features.

Unfortunately we didn't obtain distinctive results as shown in Figure 7 for every image. Solving a 7-class classification problem with only 3 or 4 intrinsic features (and many extracted from the segmentation result) appeared to be an extremely complex task. Even manual classification was sometimes difficult. Compared to regular CCD image classification, infrared images are very scarce datasets for several reasons:

- Only two channels (time and spatial variance) compared to many channels in CCD images (RGB, HSV, LAB, etc.).
- Single view at single station for a couple of months compared to many views over many stations for years in a row.
- Low resolution (380x290) compared to high resolution (2000x1500 or more).

Since the purpose of developing an automated classification algorithm was to spot potential sources for Aeolian sediment, we trained another, simpler model that better fits the amount of data available and limits the risks of overfitting. This model distinguishes between two classes only: the upper intertidal area and all other areas. This model appears to be very robust (Figure 9). Still, on rainy days no clear distinction can be made between wet sand due to rain and wet sand due to tide, but then no potential source of Aeolian sediment is found, which is adequate behavior.

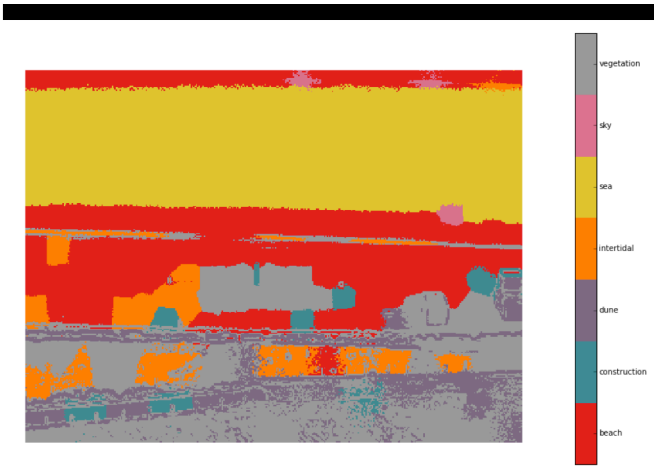

Figure 8. Infrared variance image classified using trained Conditional Random Field with intrinsic features only and no pairwise potentials.

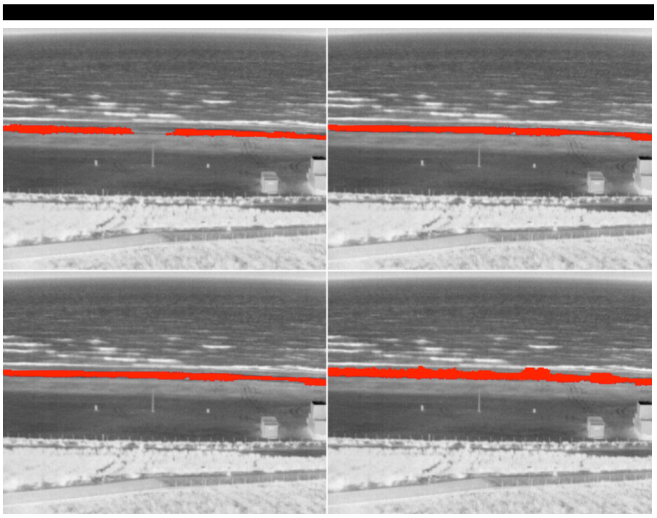

Figure 9. Four examples of classifying the upper part of the intertidal area as a potential source for Aeolian sediment. All classification is based on 24 hours infrared variance images, except the lower right image. The latter is based on a 14 days variance image and consequently shows a larger area as potential sediment source.

CONCLUSIONS

Image segmentation using Conditional Random Fields (CRF) appears to be a very generic and versatile approach also for coastal imagery. The segmentation and feature extraction routines used are generic and can be used in any classification algorithm based on CRFs. Both the segmentation and feature extraction are likely to perform even better with full color images and larger heterogeneous datasets, which in turn will increase the applicability and robustness of the approach.

A very specialized application on coastal infrared variance images for monitoring of intertidal beach moisture content and morphology appears to benefit from this general approach.

However, the limitations of the data available for this specialized application are reflected in the complexity of tasks the model can perform.

For more extensive validation and comparison of classification algorithms for coastal images, a large and heterogeneous benchmark dataset of coastal images is needed. This benchmark dataset should contain a variety of coastal images from different locations, seasons and view angles. Currently we are working on such benchmark dataset by manually classifying a large number of coastal images.

OPEN SOURCE

The three main steps of the classification algorithm (segmentation, feature extraction and model construction and training) are part of an open-source toolbox for coastal image analysis hosted by the OpenEarth repository (van Koningsveld et al., 2010). The toolbox relies on the Scikit Image toolbox (Scikit-image, v0.7.2) for segmentation and feature extraction and on the PyStruct toolbox (Mueller, v0.1) for model construction and training. The toolbox also provides a tool for efficient manual classification.

ACKNOWLEDGEMENTS

For their work discussed in this paper the authors are supported by the ERC-Advanced Grant 291206 – Nearshore Monitoring and Modeling (NEMO).

LITERATURE CITED

Aarninkhof, S.G.J., Turner, I.L., T, D.T.D., Caljouw, M. and Nipius, L., 2003. A video-based technique for mapping intertidal beach bathymetry. Coastal Engineering, 49(4): 275-289.

Achanta, R. et al., 2010. Slic superpixels. Ecole Polytechnique Federal de Lausssanne (EPFL), Tech. Rep, 149300.

Andersen, M., Dahl, J. and Vandenberghe, L., Python Software for Convex Optimization (CVXOPT), v1.1.5. http://cvxopt.org/.

de Vries, S., de Schipper, M.A., Stive, M.J.F. and Ranasinghe, R., 2010. Sediment exchange between sub-aqueous and sub-aerial coastal zones. Coastal Engineering Proceedings, p.^pp.

de Vries, S., Hill, D.J., de Schipper, M.A. and Stive, M.J.F., 2009. Using stereo photogrammetry to measure coastal waves. Journal of Coastal Research, SI 56: 1484-1488.

de Vries, S., van Thiel de Vries, J.S.M. and van Rijn, L.C., 2014. Aeolian sediment transport in supply limited situations. Aeolian Research.

Dongeren, A.P. et al., 2013. Rip current prediction through model-data assimilation on two distinct beaches. Coastal Dynamics 2013, Arcachon, France, p.^pp. 1775 - 1786.

Edwards, B.L., Namikas, S.L. and D'Sa, E.J., 2013. Simple infrared techniques for measuring beach surface moisture. Earth Surface Processes and Landforms, 38(2): 192-197.

Haralick, R.M., Shanmugam, K. and Dinstein, I.H., 1973. Textural features for image classification. Systems, Man and Cybernetics, IEEE Transactions on, SMC-3(6): 610-621.

Holland, K.T., Holman, R.A., Lipmann, T., Stanley, J. and Plant, N.G., 1997. Practical use of video imagery in nearshore oceanographic field studies. Journal of Oceanographic Engineering, 22: 81-92.

Holman, R., Plant, N. and Holland, T., 2013. Cbathy: A robust algorithm for estimating nearshore bathymetry. Journal of Geophysical Research: Oceans, 118(5): 2595-2609.

Koller, D. and Friedman, N., 2009. Probabilistic graphical models: Principles and techniques. The MIT Press.

Mueller, A., Pystruct - structured learning in python, v0.1. http://pystruct.github.io/.

Pye, K. and Tsoar, H., 1990. Aeolian sand and sand dunes. Unwin Hyman, London.

Quartel, S., Addink, E.A. and Ruessink, B.G., 2006. Object-oriented extraction of beach morphology from video images. International Journal of Applied Earth Observation and Geoinformation, 8(4): 256 - 269.

Schretlen, J.L.M. and Wijnberg, K.M., 2012. Argus video imagery in support of aeolian transport studies: An exploratory study at vlugtenburg beach, Twente University.

Scikit-image - image processing in python, v0.7.2. The scikit-image development team. http://scikit-image.org/.

Stive, M.J.F. et al., 2013. A new alternative to saving our beaches from sea-level rise: The sand engine. Journal of Coastal Research, 29(5): 1001 - 1008.

van Koningsveld, M. et al., 2010. Openearth - inter-company management of: Data, models, tools & knowledge. Proceedings WODCON XIX Conference, Beijing, China, p.^pp.

Vousdoukas, M., Wziatek, D. and Almeida, L., 2012. Coastal vulnerability assessment based on video wave run-up observations at a mesotidal, steep-sloped beach. Ocean Dynamics, 62(1): 123-137.