

A Framework for Interactive Visualization and Classification of Dynamical Processes at the Water Surface

S. Wanner¹, C. Sommer¹, R. Rocholz², M. Jung², F. A. Hamprecht¹, B. Jähne^{1,2}

¹HCI - Heidelberg Collaboratory for Image Processing, University of Heidelberg, Germany

²Institute of Environmental Physics, University of Heidelberg, Germany

Abstract

A framework for the visualization and classification of multi-channel spatio-temporal data from water wave imaging is presented. Our interactive visualization tool, WaveVis, allows a detailed study of the water surface shape in reference to additional data streams, like thermographic images or classification results. This facilitates an intuitive and effective inspection of huge amounts of data. WaveVis was used to select representative training examples of events for a supervised learning approach and to evaluate the results of the classification. The interactive classification and segmentation software ilastik was used to train a Random Forest classifier. The benefit of the combination of both programs is demonstrated for two applications, the estimation of the rain rate from the segmentation of impact craters, and the detection of small scale breaking waves. The classification of the impact crater of raindrops on the water surface worked very well, whereas the detection of the breaking waves was satisfactory only under certain experimental conditions. Nevertheless, the combination of WaveVis and ilastik proved to be valuable in both cases.

1. Introduction

Modern imaging techniques have been proven to be useful to investigate dynamical processes that occur at the air-water interface and potentially enhance the exchange of heat and gas between ocean and atmosphere [ZAJ*04, SLR*01, SLAJ04]. In particular, imaging techniques are available to accurately measure the temperature at the water surface (thermography) and the shape of the waves (wave slope imaging) in a wind wave tank with high frame rates (e.g. 312 Hz) [RWSJ11]. High resolution, high frame rates and multi-channel data are desirable. It is, however, a challenging task to visually inspect these data sets, to detect and extract interesting features, and to further use these features for automatic classification of events. Due to this, the huge amount of data (Tbytes) is often times reduced to only a few statistical measures. This statistical approach might conceal interesting processes that could be observed within the multichannel data if a detailed inspection was feasible. To look closer into the data, one could use side-by-side displays of still images (e.g. [ZAJ*04]). But this is no longer practicable for dynamical sequences with many frames. Also simple 2D false color overlays are inconvenient for the observer. For the inspection of thermographic images of the

water surface with respect to the shape and temporal evolution of the waves, the visualization tool *WaveVis* was developed [Jun08, Wan10], which interactively render the waves in 3D with OpenGL and maps the temperature in false color onto that surface. The color overlay could in principle represent any kind of geometrically registered surface data. Another tool, *ilastik* [SSKH11, JSJ10], was developed to interactively train a Random Forest classifier with an intuitive graphical user interface for a broader range of applications. This tool can be used for segmentation within independent multi-channel data sets, based on the Random Forest classifier.

In this paper, we present results from a combination of *WaveVis* and *ilastik*. We attempt to train a Random Forest classifier to detect certain events that are clearly visible in the thermographic images, if we use the shape information of the waves as input alone. If this was possible, we could on the one hand learn more about the wave dynamics and on the other hand reduce the experimental costs and effort, since merely the wave slope imaging would be needed for future experiments. For instance, the state of the art techniques for the detection of surface renewal by microscale breaking waves are performed with either expen-

sive infrared cameras [ZAJ*04], or by flow field measurements (e.g. PIV) which need close range optical access from the side [SLJA01, LKS06] that is not always provided. A set of co-localized and synchronized measurements of the wave shape and the surface temperature field (section 2) is used to explore the possibilities of an automatic classification of events. For this study, *WaveVis* (section 3.1) was modified to allow for an interactive labeling and to export labeled images into *ilastik* for processing (section 3.2). The segmentation result can then again be inspected with *WaveVis*. Results from two applications are shown, first the detection of raindrop impacts in section 4.1, and second the detection of small scale breaking waves in section 4.2. Section 5 gives a conclusion. The source code of the software is available to the public, see section 6.

2. Materials

The data set was provided by Rocholz and Schimpf [RWSJ11, Roc08] and consists of 150 sequences, 5000 images each, with 4 channels, and a resolution of 149×320 pixels, corresponding to an area of $223 \text{ mm} \times 104 \text{ mm}$ on the water surface. The sequences were recorded at 312 Hz inside the rain tower section of the wind wave tank of the University of Hamburg, Germany. The four channels are: the gradients of the water surface in x- and y-direction, the water surface elevation, and the temperature at the surface. The first three channels (see examples in figure 1 a)-c)) are obtained with a Color Imaging Slope Gauge (CISG), which measures both slope components of the waves simultaneously [ZC94, Roc08, BKJ97]. The surface elevation was then reconstructed from the 2D-gradient, except for an unknown offset [Zha96, Roc08]. The fourth channel, i.e. the water surface temperature, was measured by an infrared camera that was synchronized with the CISG camera. The thermographic images were geometrically registered to give pixel wise correspondence of the four individual channels. A high power CO₂-laser was used to superimpose artificial heat patterns onto the water surface, for instance the straight line to the right in figure 1d). These heat patterns help to make the horizontal or vertical transport of heat visible [SPJ06, SGJ04, RWSJ11].

3. Methods

3.1. Interactive visualization with *WaveVis*

The interactive visualization software *WaveVis* was written in C++ and OpenGL[®] by M. Jung [Jun08] and further developed by S. Wanner [Wan10]. *WaveVis* is designed for the display of wave data as described in section 2. In addition, data from a 3D flow field simulation [TH07] can be visualized. The description of the 3D functionality is beyond the scope of this paper and only the functionality for $2\frac{1}{2}$ D data is described, here. The program is controlled via plain text configuration files. As primary input, image sequences for




Figure 1: Example frame of the four input channels.

the surface elevation and the two surface slope components are needed. The surface elevation defines a mesh to model the water surface. The slope channels are used as input for the surface shading. The shading highlights small scale wave features with too low amplitudes to be visible in the height map otherwise, see figure 2a). Both, the virtual camera position and the virtual light source position can be changed interactively. A ring buffer is used to handle the different data streams. The frame rate of the visualization is adjustable and the animation can be stopped and stepped through in both time directions frame by frame. Because of to the ring buffer, there is no limitation to the total sequence size.

As secondary input, several additional image sequences (also of type float and with pixel-correspondence to the primary input) can be handled to be displayed as color overlays. I.e. the additional data is mapped in false color onto the virtual model of the waves, see figure 2 b). For scalar data, the color map can be customized (off-line) and the range can be changed interactively (at runtime). For 2D vectorial data, a HSI-color scheme is implemented in which the polar angle is represented by hue and the magnitude is represented by

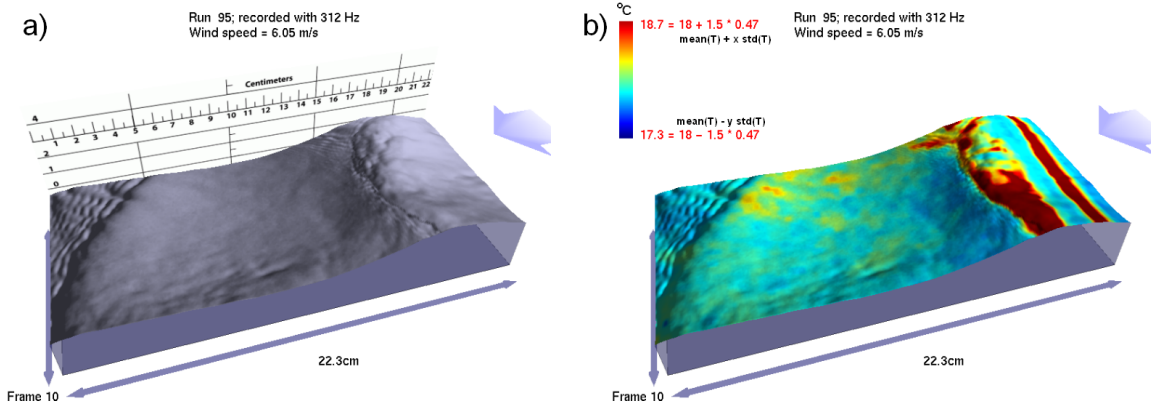


Figure 2: a) WaveVis display of the data from figure 1a)-c). Shading is exploited to make the small scale features visible. b) The temperature channel (figure 1d)) is mapped in false color onto the surface.




Figure 3: The interactive labeling function of WaveVis allows to assign different class values (represented by colors) to elliptic regions on the surface. In this example two ellipses were drawn to distinguish between a smooth wave trough (yellow) and a breaking wave crest (green).

saturation.

Here, the color overlay feature was used for the scalar thermographic images. For comparison Figure 2 b) shows the same data as Figure 1a)-d) as visualized with WaveVis.

WaveVis provides some other functionalities like a recording function to automatically take screenshots of each frame of the interactive visualization, a function to display customized reading scales (see figure 2a)), and a labeling function. The interactive labeling is implemented in a rather simple manner, i.e. the user can assign class values to regions on the surface that are represented by ellipses, see example in figure 3. The resulting class values, center points, semi-axes, and orientation of the ellipses, and the corresponding frame

numbers are stored in plain text. This is useful to effectively reduce the complexity of the data set and to facilitate further processing with different programs.

Using the visual inspection and the labeling function, a set of training images for the supervised learning algorithm inside *ilastik* was generated. The labels were afterwards interactively refined within *ilastik*.

3.2. Interactive Image classification with *ilastik*

ilastik (Interactive Learning and Segmentation Toolkit) [SSKH11,JSJ10] is a program that combines interactive machine learning and the ability to cope with complex textures and colors in a single convenient framework. We use *ilastik* to learn and classify events such as raindrops or microscale breaking waves based on user examples (labels) provided via WaveVis and the standard image feature set (color and texture) of *ilastik*. These image features are computed in the full 2D/3D or temporal pixel neighborhoods. The provided set of standard features includes color, edge and texture descriptors such as Gaussian filters, structure tensor, Hessian matrix, gradient magnitude, difference of Gaussians and Laplacian of Gaussian. All filter responses are computed on the elevation and the slope channels and aggregated on a pixel level to form the input for the classification step. The classification of *ilastik* is based on Random Forests [Bre01] which are fast to compute, inherently parallel and support multiple classes while showing classification performance competitive to boosting and SVMs [SUP*09, BZM07, SWA08]. After training, the classifier is used to predict the object class probability for every pixel. Once a classifier has been trained on a set of representative images, it can be used in batch mode to automatically process a very large number of images.




Figure 4: a) Characteristics of raindrop impacts. b) Example for the surface characteristics of a regular wave crest (highlighted on the left) and for a microscale breaking wave crest (highlighted on the right).

4. Applications

In the following, we introduce two applications: a) the detection of raindrops to estimate the rain rate, and b) the detection of small scale wave breaking, see figure 4. Both processes are important for air-sea gas exchange. The aim was to automatically detect the events by only using the surface shape (see section 1). While the detection of raindrop impacts is rather easy and could be implemented without using elaborated classification techniques, the discrimination of the dimpled structure of small scale breaking waves from other surface features is not as straight forward.

4.1. Raindrop detection and estimation of the rain rate

The craters that originate from raindrop impacts on the water surface are clearly visible within the surrounding wave field, even at high wind speeds (figure 4a)). The *ilastik* workflow for the training of the classifier is illustrated in figure 5 a). Craters were marked with a red label, and the background was marked with a green label (top panel). The labeling was performed in this way for a set of 10 images (i.e. 0.2% of one single sequence). This was sufficient to train a Random Forest classifier using the *ilastik* standard image feature sets *color*, *texture* and *orientation*. The segmentation result is illustrated in lower panel of figure 5 a).

From the segmentation of the raindrops we are able to calculate the rain rate. The rain rate can be loosely defined as raindrop impacts per unit time, because for our experiments the drop size distribution is known a priori. To get an estimate of the rain rate it is not sufficient to just count craters in a random set of images, because then the time information is lacking. The sampling rate must be sufficiently fast, to ensure that no drops were missed in between the images. On the other hand, we need to avoid double counting of craters. In our data set, the craters are visible for a number (≈ 40) of consecutive frames, due to the high acquisition rate. For the human eye it is then relatively challenging to correctly count the individual impacts, because younger and older

craters co-exist.

With the segmentation output from *ilastik*, this task can be solved easily. After segmenting the craters in every single frame we used connected component analysis for a stack of 200 images, see figure 6 a). To suppress spurious detections a threshold on the raindrop size was sufficient. If the crater segmentation showed overlap in the time dimension, the crater was attributed to the same raindrop [Wan10]. In this way, the number of individual raindrops in intervals of 200 frames (≈ 0.6 s) was determined to give an estimate of the rain rate with a relative error of less than 0.1%. The result for one data record is illustrated in figure 6 b), showing the variability of the rain rate within a time span of 16 s.

4.2. Detection of Microscale Breaking Waves

Microscale breaking waves are short gravity waves (with wavelengths in the order of decimeters) that break without air-entrainment. The breaking process induces near surface turbulence, which becomes apparent from a dimpled surface structure in the wake of the breaking crest, see figure 4 b). There is no simple criterion (such as a threshold for the wave slope) to distinguish microscale breaking waves from regular waves [LKS06]. The turbulent wake can be detected with an Infrared camera if the temperature difference between the surface and the water bulk is sufficiently high so that the associated surface renewal can be seen in the temperature image. In this respect, the quality of our thermographic images was limited due to strong spatial inhomogeneities. However, using the combined visualization of the wave field and the temperature field in *WaveVis* we were able to learn how to detect microscale breaking by eye from the evolution of the surface shape. Our attempts to resemble this with standard image processing were not successful. Thus, we tried to use the Random Forest classifier from *ilastik* in order to classify breaking waves on basis of the surface shape data. In figure 5 b) the *ilastik* workflow is illustrated. In the upper image, the dimpled surface structure is labeled in red, and

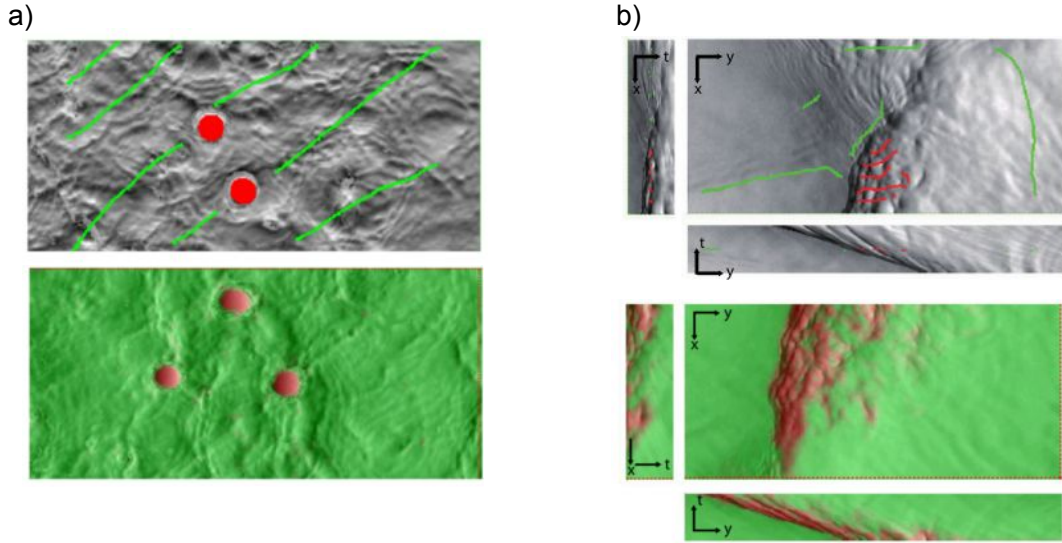


Figure 5: Interactive training in *ilastik*. **a)** top: Training image for raindrop detection with user labels, where red marks a raindrop crater. bottom: Test image with overlaid raindrop predictions. **b)** top: Training image for microscale breaking wave classification with user labels. Red indicates a breaking wave event and green indicates background. bottom: Resulting classification of an independent test frame.

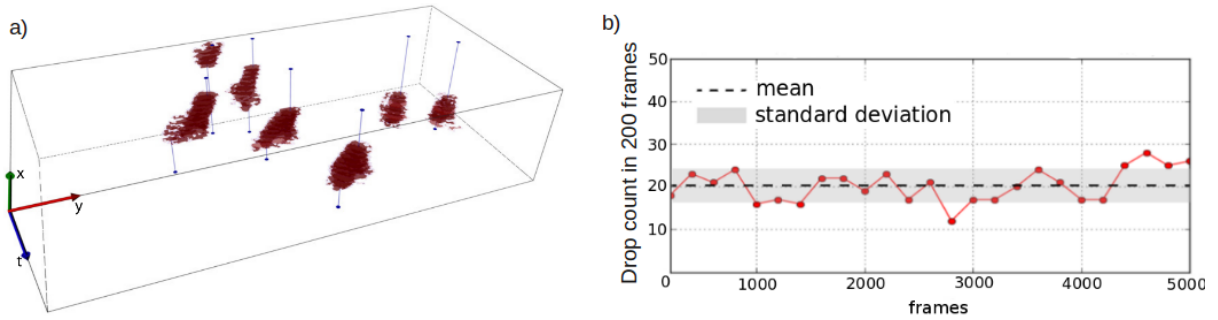


Figure 6: **a)** Clustering of raindrop detections in the space-time volume for input to the connected component analysis. **b)** Variability of the number of raindrops per 200 frames for one data record (5000 frames correspond to 16 s).

the background is labeled in green. The Random Forest was computed on basis of 10 representative training examples for breaking and non-breaking, each. The lower image in figure 5 b) shows the resulting segmentation of an independent frame; the dimpled surface structures are detected with reasonable performance. In contrast to the drop crater segmentation, it was necessary to introduce additional features for the Random Forest computation into *ilastik*. Simply using the same approach as in the raindrop detection (2D) requires a feature, which accounts for the long-range context of a wave. A canny edge filter with a broad support window was added using the *ilastik* plug-in mechanism. This filter was tuned to find the wave front with high surface slopes. It was then combined with a distance transform, to effectively

measure each pixels' distance to the wave front. This global feature helped to reduce false positive detections in regions far away from waves [Wan10]. Another option is to use the 3D capabilities of *ilastik* to incorporate temporal context directly. In our case, the third dimension is time. Figure 5 b) therefore shows additional slices through the space-time volume taken at the centerline of the x-y-plane. The temporal context is crucial to distinguish the dimpled surface structure from the wind driven waves. While the surface gradient and curvature is comparable for both, the propagation is not. Wind driven waves tend to propagate along with the wind at relatively high phase speeds, but the dimpled structures propagate in all directions, while being advected with the turbulent wake as it moves with the relatively slow surface

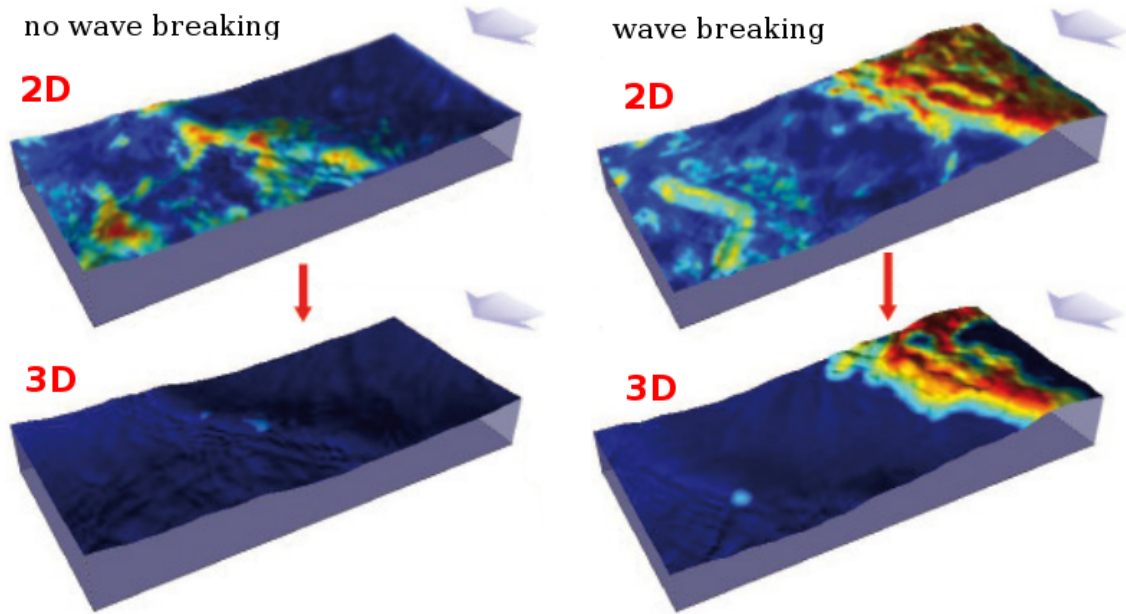


Figure 7: Wave breaking segmentation, based on single frames (2D) and using spatio-temporal information (3D). The left panel shows a wave that is actually not breaking. The right panel shows a microscale breaking wave that enters the field of view from the right. The 2D segmentation yields rather high rates of false positive detections. This is considerably improved using the temporal context in the 3D segmentation. The jet colormap is given the probability for wave breaking between $[0, 1] \rightarrow [blue, red]$.

drift current. This can hardly be observed in single snapshots but is easily visible in the animated visualization provided by *WaveVis*. For the same reason, the classification result on single frames (2D) gave high percentages of false positive detections. The false positive rate was effectively reduced by using 10 consecutive frames (3D) for the classifier, compare figure 7.

The resulting Random Forest classifier from *ilastik* was used to predict the probability of microscale breaking for different data records (different wind forcing) and the result was evaluated by means of a visual comparison to the thermographic images and wave evolution in *WaveVis*. This is illustrated in figure 8. It was found, that the quality of the prediction varied for the given data set. For moderate wind speeds the result was reasonable, but for higher and lower wind speeds, the prediction was less accurate [Wan10].

5. Conclusion

In this paper we introduced a framework for visualization and classification with application to multi-channel data sets from water wave imaging. The data – consisting of surface elevation, surface slopes and additional overlay data streams, like thermographic images or classification results)– can be interactively visualized using *WaveVis*. This allows an intuitive and effective inspection of huge amounts of data and fa-

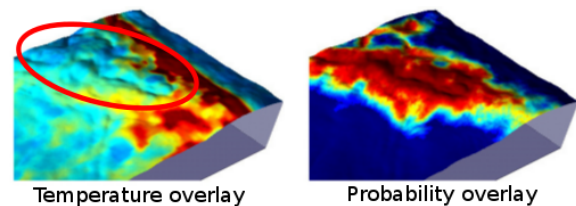


Figure 8: Left: Shows a region of surface renewal (marked with an ellipse) indicating a microscale breaking wave. Note that, the IR response is low in this area due to heat recirculation. Right: The result of the wave classification is depicted showing good correlation with the heat recirculation on the left. The jet colormap is coding the probability for wave breaking between $[0, 1] \rightarrow [blue, red]$.

cilitates the selection of representative training examples for a supervised learning approach. We used the classification and segmentation software *ilastik* to train a Random Forest classifier. This interactive tool is user-friendly and basically requires no experience in image processing. The work flow is demonstrated by means of two applications: a) computing the rain rate, and b) detection of microscale breaking waves. The classification of the impact crater of raindrops on the water surface worked very well with the built-in features of

ilastik. The results for the rain rate showed a good accordance to the marginal conditions of the experiments and enabled a more detailed study of the influence of rain on other variables that are important to understand the gas exchange under the influence of rain [RWSJ11, Wan10]. The estimation of the rain rate was carried out for more than 20 data records (100000 frames) [Wan10], which were obtained under very different conditions. The Random Forest classifier (*ilastik*), which was computed with only 10 labeled images, gave robust results for all of the records. This was easily assessed by visual inspection of the segmentation result as a color overlay in the animated wave display using *WaveVis*. The classification of microscale breaking waves is a much more demanding task. We showed that it is possible to perform a segmentation after introducing new features and taking the temporal context into account. However, we were unable to train a unique Random Forest classifier to performed evenly well for a wider range of experimental conditions (e.g. wind forcing). We suppose that this was partially due to limited main memory (32Gbyte), which constrained the number of the multi-channel and spatio-temporal training samples, as well as the number of consecutive frames that could be used for the segmentation. For this specific application we may conclude that the automatic classification does not yet perform as good as our visible perception with the aid of a good visualization tool. However, the proposed procedure incorporating spatio temporal context shows promising results and is subject to future research.

6. Software repositories

The software *ilastik* is open-source and freely available from <http://ilastik.org>. *WaveVis* is currently not open-source but can be obtained from the authors. Please send your request to bernd.jaehne@iwr.uni-heidelberg.de.

7. Acknowledgments

We thank the DFG (*Ja395/13 – 1&2*) and the Heidelberg Collaboratory for Image Processing (HCI) for financial support of this work. Thanks also to Christoph Garbe and Christoph Straehle from the HCI for their support.

References

- [BKJ97] BALSCHBACH G., KLINKE J., JÄHNE B.: Multichannel shape from shading techniques for reconstruction of specular surfaces. In *Tagungsband Herbsttagung des Graduiertenkollegs "3D Bildanalyse und -synthese"* (Universität Erlangen, Nürnberg, November 1997), H.-P. Seidel, B. Girod, H. Niemann (Hrsg.). 2
- [Bre01] BREIMAN L.: Random forests. *Machine Learning* 45 (2001), 5–32. 3
- [BZM07] BOSCH A., ZISSERMAN A., MUNOZ X.: Image classification using random forests and ferns. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on* (2007), pp. 1–8. 3
- [JSJ10] JEHLE M., SOMMER C., JÄHNE B.: Learning of optimal illumination for material classification. In *Pattern Recognition* (2010), Springer, pp. 563–572. 1, 3
- [Jun08] JUNG M.: *Entwicklung einer Visualisierung von Messdaten mittels OpenGL - WaveVis*. Tech. rep., Institut für Wissenschaftliches Rechnen, Universität Heidelberg, 2008. 1, 2
- [LKS06] LOEWEN M. R., KAMRAN SIDDIQUI M. H.: Detecting microscale breaking waves. *Measurement Science and Technology* 17 (Apr. 2006), 771–780. 1, 4
- [Roc08] ROCHOLZ R.: *Spatiotemporal Measurement of Short Wind-Driven Water Waves*. Dissertation, Institut für Umweltphysik, Fakultät für Physik und Astronomie, Univ. Heidelberg, 2008. 2
- [RWSJ11] ROCHOLZ R., WANNER S., SCHIMPF U., JÄHNE B.: Combined visualization of wind waves and water surface temperature. In *6th Int. Symp. Gas Transfer at Water Surfaces, Kyoto, May 17–21, 2010* (2011). in press. 1, 2, 6
- [SGJ04] SCHIMPF U., GARBE C., JÄHNE B.: Investigation of transport processes across the sea surface microlayer by infrared imagery. *Journal of Geophysical Research-Oceans* 109, C8 (2004), C08S13. 2
- [SLAJ04] SIDDIQUI M., LOEWEN M. R., ASHER W. E., JESSUP A. T.: Coherent structures beneath wind waves and their influence on air-water gas transfer. *J. Geophys. Res.* 109 (2004), C03024. 1
- [SLJA01] SIDDIQUI M., LOEWEN M., JESSUP A., ASHER W.: Infrared remote sensing of microscale breaking waves and near-surface flow fields. In *Geoscience and Remote Sensing Symposium, 2001. IGARSS 01. IEEE 2001 International* (2001), vol. 2, pp. 969 – 971. 1
- [SLR*01] SIDDIQUI M. H. K., LOEWEN M. R., RICHARDSON C., ASHER W. E., JESSUP A. T.: Simultaneous particle image velocimetry and infrared imagery of microscale breaking waves. *Phys. Fluids* 13 (2001), 1891–1903. 1
- [SPJ06] SCHIMPF U., POPP C., JÄHNE B.: Active thermography: a local and fast method to investigate heat and gas exchange between ocean and atmosphere. In *Verhandlungen der Deutschen Physikalischen Gesellschaft, Spring Conference, Heidelberg, 15.-17.03.2006* (2006), Deutsche Physikalische Gesellschaft. 2
- [SSKH11] SOMMER C., STRÄHLE C., KÖTHE U., HAMPRECHT F. A.: "ilastik: Interactive learning and segmentation toolkit". In *8th IEEE International Symposium on Biomedical Imaging (ISBI 2011)*, in press (2011). 1, 3
- [SUP*09] SANTNER J., UNGER M., POCK T., LEISTNER C., SAFFARI A., BISCHOF H.: Interactive texture segmentation using random forests and total variation. In *Proceedings of the British Machine Vision Conference (BMVC)* (London, UK, September 2009). to appear. 3
- [SWA08] STATNIKOV A., WANG L., ALIFERIS C. F.: A comprehensive comparison of random forests and support vector machines for microarray-based cancer classification. *BMC Bioinformatics* 9 (2008), 319. 3
- [TH07] TSAI W., HUNG L.: Three-dimensional modeling of small-scale processes in the upper boundary layer bounded by a dynamic ocean surface. *Journal of Geophysical Research* 112 (2007), C02019. 2
- [Wan10] WANNER S.: *Interaktives Rendering von Wellendaten windgetriebener Wasseroberflächen und Ereignisklassifizierung*. Diplomarbeit, Institut für Umweltphysik, Fakultät für Physik und Astronomie, Univ. Heidelberg, 2010. 1, 2, 4, 5, 6

- [ZAJ*04] ZAPPA C. J., ASHER W. E., JESSUP A. T., KLINKE J., LONG S. R.: Microbreaking and the enhancement of air-water transfer velocity. *J. Geophys. Res.* 109 (2004), C08S16. [1](#)
- [ZC94] ZHANG X., COX C. S.: Measuring the two-dimensional structure of a wavy water surface optically: A surface gradient detector. *Experiments in Fluids* 17 (Aug. 1994), 225–237. [2](#)
- [Zha96] ZHANG X.: An algorithm for calculating water surface elevations from surface gradient. *Experiments in Fluids* 21 (1996), 43–48. [2](#)