Deep learning applied to sea surface semantic segmentation: Filtering sunglint from aerial imagery

Teodor Vrećica [1] **Quentin Paletta** [2] **Luc Lenain** [1]

Abstract

Water waves are an ubiquitous feature of the oceans, which serve as a pathway for interactions with the atmosphere. Wave breaking in particular is crucial in developing better understanding of the exchange of momentum, heat, and gas fluxes between the ocean and the atmosphere. Characterizing the properties of wave breaking using orbital or suborbital imagery of the surface of the ocean can be challenging, due to contamination from sunglint, a persistent feature in certain lighting conditions. Here we propose a supervised learning approach to accurately detect whitecaps from airborne imagery obtained under a broad range of lighting conditions. Finally, we discuss potential applications for improving ocean and climate models.

1. Motivation and background

Exchanges between oceans and the atmosphere involve fifteen times as much carbon as human activities emit by burning fossil fuels, and are crucial in regulating temperature of the planet (Cavaleri et al., 2012). As water waves serve as a boundary between the ocean and the atmosphere, the understanding of wave processes is key to better model our changing climate. In particular, wave breaking governs many of the physical processes from air-sea momentum to gas and heat fluxes (Sullivan et al., 2007; Fairall et al., 2011). It was also shown to significantly increase albedo and aerosol concentrations (through generation of whitecaps and sea spray), the key variables in climate models (Lewis et al., 2004; Frouin et al., 2001).

As waves break, they entrain air in the water column leading to the formation of foam patches known as whitecaps. Analysis of wave breaking statistics is usually conducted by analyzing percentage of whitecap coverage of sea surface (e.g. Callaghan & White, 2009). While this approach was used to determine some of the relevant properties (e.g. energy dissipation rate), it does not provide information on the properties of breaking kinematics. In (Phillips, 1985), wave breaking statistics are analyzed in terms of distribution of length of actively breaking waves per unit of the sea surface area, $\Lambda(c_b)$, where c_b represents speed of the actively breaking waves. The main advantage of this approach is that moments of the $\Lambda(c_b)$ distribution are related to many of the important physical quantities. Whitecap coverage is in fact a function of the first moment of the distribution, and many other properties (breaking induced drift, momentum fluxes, and energy dissipation) can be derived from it.

Wave breaking is by nature temporally and spatially variable, and is especially difficult to observe and characterize. While the Phillips (1985) theory was established decades ago, it is only in relatively recent years that Λ distributions were computed from high resolution airborne imagery (Kleiss & Melville, 2010). Collected georeferenced images of the ocean surface are first gridded into a Cartesian coordinate system, pixels with brightness above certain threshold are classified as whitecaps (Kleiss & Melville, 2011), and whitecap kinematics are determined using Particle Image Velocimetry (PIV) and/or optical flow methods. An issue with this approach is that it can be difficult (depending on the lighting conditions) to differentiate between sunglint and whitecaps, which can contaminate the results.

Expected contributions In order to overcome this issue and to further the understanding of the role of wave breaking in air-sea interaction processes, here we propose a data driven approach that makes use of the recent development of deep learning techniques for computer vision. Sunglint and whitecaps have different spatial and temporal structures, which could be identified through neural network approaches. These techniques have been shown to outperform traditional approaches in other vision-based tasks (e.g. for autonomous vehicles Grigorescu et al., 2020). Due to its ability to classify objects at pixel resolution with great accuracy, the use of UNet model (Ronneberger et al., 2015) is proposed in this work.

[1]Scripps Institution of Oceanography, University of California, San Diego, USA [2]Department of Engineering, University of Cambridge, Cambridge, UK. Correspondence to: Teodor Vre ćica <tvrecica@ucsd.edu>.

Tackling Climate Change with Machine Learning Workshop at ICML 2021

2. Observations

Aerial images of sea surface obtained from several previous, ongoing and planned airborne campaigns conducted at the Air-Sea Interaction Laboratory (SIO) will be used for training and validation of the model. Millions of sea surface images, obtained under a broad range of environmental conditions, were collected over the course of these expeditions. The database of aerial imagery will be georeferenced and gridded in Cartesian coordinate system. The discretization step (pixel size) depends on the altitude of the aircraft. Therefore, in order to ensure consistency, all of the images will be discretized with a resolution corresponding to the highest altitude flight under consideration (generally in the 0.1-0.5 meter range). In order to label images for training and validation, they will be segmented into whitecaps, water surface, and sunglint according to the methods established in (Kleiss & Melville, 2011), with manual correction of sunglint where necessary.

Wave breaking is a function of currents, wind, and wave statistics (with wind direction being the most important quantity), which constrain the shape and direction of breaking waves. Most of the considered airborne measurement campaigns were done in conjunction with in-situ measurements which are able to provide detailed descriptions of these quantities.

3. Processing of airborne images of sea surface

Method for determining wave breaking statistics In (Phillips, 1985), wave breaking statistics were defined in terms of $\Lambda(c_b)$ distribution, defined as the length of active breakers, per unit area, per unit of the breaker velocity (c_b). What is relevant for purposes of this proposal, is the fact that the distribution is determined from motion of whitecaps in two subsequent images.

Whitecap velocities are estimated using two subsequent images (with a time step of 0.2 seconds) using the optical flow method of (Liu et al., 2009). The method applies initial downsampling of the image and gradually upsamples it to the original resolution, which enables for accurate computation of velocities for both fast and slow moving breakers. A schematic showing a breaking wave is shown in Figure 1. Based on a set of criteria (Kleiss & Melville, 2011), some of the pixels on the boundary of the whitecap are classified as part of an actively breaking waves. They (and their velocities) are indicated with red quivers.

Contamination of images by sunglint While the data presented in figure 1 were collected in near-perfect lighting conditions, this is often not the case, as sunglint can contaminate large portions of the image. An issue that arises from this is that sunglint can also be classified as an actively breaking wave using the existing methods, especially in the case of smaller scale breaking waves (see Figure 2 for an example). The current approach for eliminating sunglint is to ensure that the detected breaking waves are present in subsequent images (e.g. over more than 0.4 sec). This assumes that sunglint is typically short lived within a given area, however this is not always the case. The exact amount of persistent sunglint depends on cloud coverage, angles of observation and of the Sun, and surface roughness. While there are methods for eliminating sunglint by using bandwidth information from hyperspectral imagery (Kay et al., 2009), these methods were found to not be applicable for differentiating between whitecaps and sunglint.

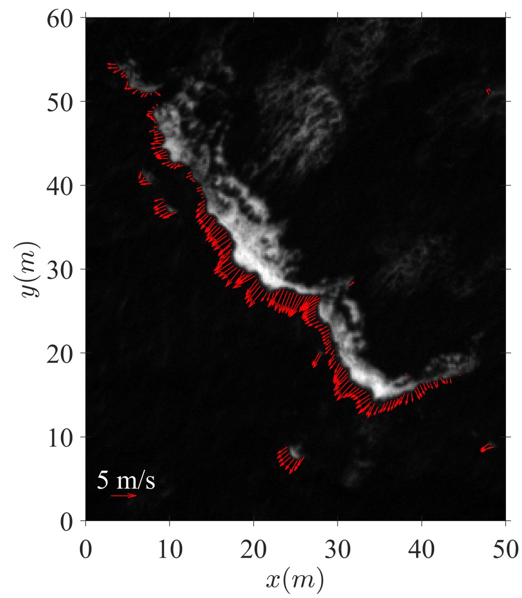

Figure 1. A georeferenced sea surface image showing an actively breaking wave. Breaking velocities are plotted in red quivers.

4. Applications of deep learning models

The proposed workflow for tackling the issue of detecting sunglint from airborne imagery is discussed in detail below, and is summarized as a flowchart in Figure 2.

In order to distinguish complex structures such as sunglint and wave breaks, the use of a convolutional architecture such as UNet is proposed. Originally developed for biomedical image segmentation (Ronneberger et al., 2015), this popular segmentation model can be trained to classify pixels of the image individually, which is crucial for the proposed application. To train and evaluate the model, images would be labeled into three classes: water, whitecaps, and sunglint. To assess the generalisation ability of the model, airborne data would be split into two sets, containing data from distinct days. The training and validation data would be selected from one, and the test data would be selected from the other set. This is hoped to showcase that the model is capable of filtering sunglint under various conditions.

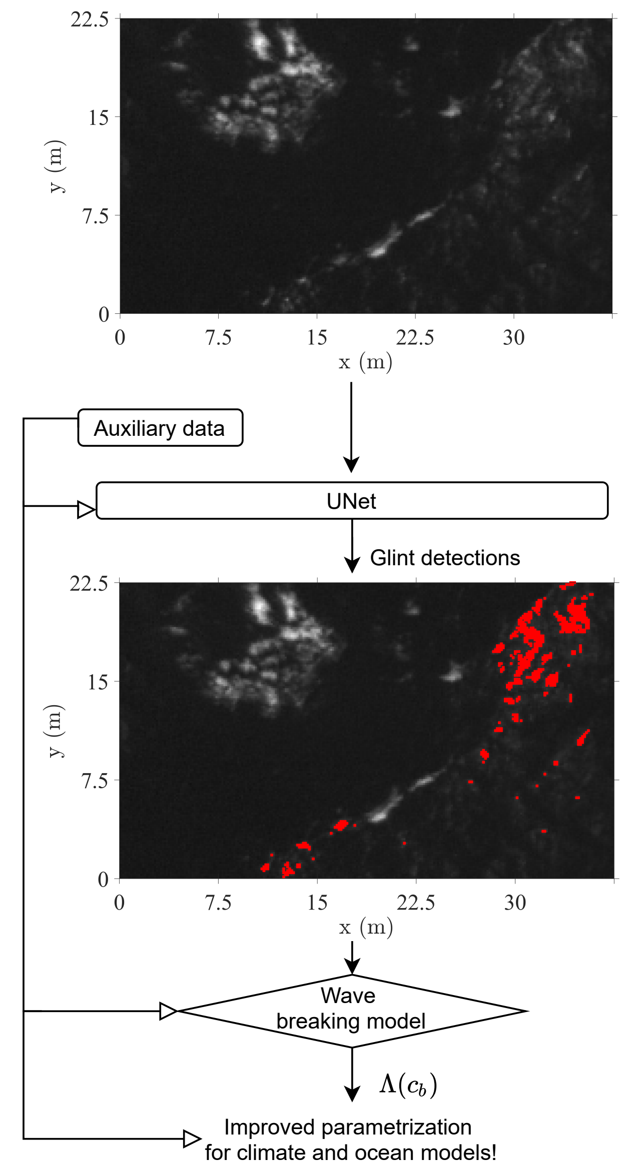

Figure 2. Proposed workflow. Auxiliary data (information on wind, currents, and wave statistics) and airborne sea surface imagery will serve as inputs. Note that the image contained high concentration of persistent glint, some of which was classified as actively breaking waves using the existing approach.

In order to improve the models' performance, various loss functions (cross entropy, focal loss function Lin et al. (2017)) will be employed, as well as hyper-parameter tuning and dataset balancing. Image gradients, and image differences (with single or multiple time steps) could be added as separate image channels to guide the model. As breaking waves with higher velocities are far more dynamically significant, the models performance will be evaluated as a function of exponent (probably fourth or fifth) of velocities corresponding to miss-classified pixels.

The detailed information of the wave statistics, and wind speed and direction would also be given to a model. This would potentially be done by including additional image channels, or by conjoining it as auxiliary data through densely connected layers. If it is noticed that the model is struggling with certain classes or configurations of auxiliary parameters, additional imagery with such features will be labeled and included into the training sets in an iterative way. While the use of UNet model is initially planned, the use of other segmentation models (such as DeepLab, YOLO, or Mask R-CNN, He et al. (2017); Redmon et al. (2016); Chen et al. (2017)) would likely be considered to benchmark the approach. Furthermore, smaller scale glint and whitecaps can have similar structures, and it is their time evolution that could help to set them apart. To this end, use of attention or recurrent models could be pursued as well. Following successful training of the model, it is planned to compare the suggested method to traditional approaches, the codes and all other supporting information will be released to the general public.

5. Summary and outlook

The applications of deep learning techniques for the detection and classification of breaking waves collected from suborbital and orbital platforms are discussed in this paper. In order to prevent further climate change and to mitigate its effects, it is vital to model it as accurately as possible, and improving coupling of the ocean and the atmosphere would be a significant step forward. We hope that the proposed approach will enable rapid and accurate removal of sunglint from aerial and orbital remote sensing imagery. This would enable for determination of $\Lambda(c_b)$ distributions and other wave breaking statistics under a broad range of environmental conditions. This in turn could lead to improved parametrizations of wave breaking statistics and better estimates of momentum fluxes, surface kinematics, marine aerosols, heat and gas fluxes, albedo etc, all crucial to global ocean and climate models, and to many biological and chemical processes (transport of near surface pollutants for example).

References

Callaghan, A. H. and White, M. Automated processing of sea surface images for the determination of whitecap coverage. *Journal of Atmospheric and Oceanic Technology*, 26(2):383–394, 2009.

Cavaleri, L., Fox-Kemper, B., and Hemer, M. Wind waves in the coupled climate system. *Bulletin of the American Meteorological Society*, 93(11):1651–1661, 2012.

Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.

Fairall, C., Yang, M., Bariteau, L., Edson, J., Helmig, D., McGillis, W., Pezoa, S., Hare, J., Huebert, B., and Blomquist, B. Implementation of the coupled ocean-atmosphere response experiment flux algorithm with co2, dimethyl sulfide, and o3. *Journal of Geophysical Research: Oceans*, 116(C4), 2011.

Frouin, R., Iacobellis, S., and Deschamps, P.-Y. Influence of oceanic whitecaps on the global radiation budget. *Geophysical research letters*, 28(8):1523–1526, 2001.

Grigorescu, S., Trasnea, B., Cocias, T., and Macesanu, G. A survey of deep learning techniques for autonomous driving. *Journal of Field Robotics*, 37(3):362–386, 2020.

He, K., Gkioxari, G., Dollár, P., and Girshick, R. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.

Kay, S., Hedley, J. D., and Lavender, S. Sun glint correction of high and low spatial resolution images of aquatic scenes: a review of methods for visible and near-infrared wavelengths. *Remote sensing*, 1(4):697–730, 2009.

Kleiss, J. M. and Melville, W. K. Observations of wave breaking kinematics in fetch-limited seas. *Journal of Physical Oceanography*, 40(12):2575–2604, 2010.

Kleiss, J. M. and Melville, W. K. The analysis of sea surface imagery for whitecap kinematics. *Journal of Atmospheric and Oceanic Technology*, 28(2):219–243, 2011.

Lewis, E. R., Lewis, R., Karlstrom, K. E., Lewis, E. R., and Schwartz, S. E. *Sea salt aerosol production: mechanisms, methods, measurements, and models*, volume 152. American Geophysical Union, 2004.

Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988, 2017.

Liu, C. et al. *Beyond pixels: exploring new representations and applications for motion analysis*. PhD thesis, Massachusetts Institute of Technology, 2009.

Phillips, O. Spectral and statistical properties of the equilibrium range in wind-generated gravity waves. *Journal of Fluid Mechanics*, 156:505–531, 1985.

Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788, 2016.

Ronneberger, O., Fischer, P., and Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

Sullivan, P. P., McWilliams, J. C., and Melville, W. K. Surface gravity wave effects in the oceanic boundary layer: Large-eddy simulation with vortex force and stochastic breakers. *Journal of Fluid Mechanics*, 593:405, 2007.