

An Active Learning Pipeline to Detect Hurricane Washover in Post-Storm Aerial Images

Evan B. Goldstein^{1*}, Somya D. Mohanty², Shah Nafis Rafique², Jamison Valentine²

¹Department of Geography, Environment, and Sustainability,
University of North Carolina at Greensboro

²Department of Computer Science, University of North Carolina at
Greensboro

*Corresponding Author:

twitter: [@ebgoldstein](https://twitter.com/ebgoldstein)
email: ebgoldst@uncg.edu

This paper was accepted & presented at the AI for Earth Sciences Workshop at NeurIPS 2020 (Neural Information Processing Systems Online Conference, 2020). The paper went through light peer review/assessment for the workshop (3 reviewers), but has not gone through peer review at a journal.

Please cite as:

Goldstein EB, Mohanty SD, Rafique SN, Valentine J, 2020, An Active Learning Pipeline to Detect Hurricane Washover in Post-Storm Aerial Images, AI for Earth Sciences Workshop at NeurIPS 2020, <EarthArXiv DOI>

An Active Learning Pipeline to Detect Hurricane Washover in Post-Storm Aerial Images

Evan B. Goldstein

Department of Geography, Environment,
and Sustainability
University of North Carolina at Greensboro
ebgoldst@uncg.edu

Somya D. Mohanty

Department of Computer Science
University of North Carolina at Greensboro
sdmohant@uncg.edu

Shah Nafis Rafique

Department of Computer Science
University of North Carolina at Greensboro
shah.nafis@alumni.uncg.edu

Jamison Valentine

Department of Computer Science
University of North Carolina at Greensboro
jtvalent@uncg.edu

Abstract

We present an active learning pipeline to identify hurricane impacts on coastal landscapes. Previously unlabeled post-storm images are used in a three component workflow — first an online interface is used to crowd-source labels for imagery; second, a convolutional neural network is trained using the labeled images; third, model predictions are displayed on an interactive map. Both the labeler and interactive map allow coastal scientists to provide additional labels that will be used to develop a large labeled dataset, a refined model, and improved hurricane impact assessments.

1 Introduction

Hurricanes can change coastal landscapes by redistributing large quantities of sediment. For storms that impact the US, coastal landscape change can be assessed using Emergency Response Imagery collected by the National Geodetic Survey Remote Sensing Division of the US National Oceanographic and Atmospheric Administration (NOAA; <https://storms.ngs.noaa.gov/>). This comprehensive aerial imagery is often obtained soon after the storm event, and is typically large both in terms of the number of individual images per storm and the size of each image. Additionally, this imagery is unlabeled.

Here we use post-storm imagery to identify hurricane impacts to coastal landscapes along the US Atlantic and Gulf coasts. Our work combines ideas from two lines of research — first, the use of crowd sourcing to label the impacts visible in post-storm imagery [Liu et al., 2014, Morgan et al., 2019], and second, the use of machine learning to classify and investigate coastal landscape dynamics [e.g. Buscombe and Ritchie, 2018, Ridge et al., 2019, Goldstein et al., 2019]. We develop an active learning [e.g., Settles, 2011] methodology to identify storm impacted landscapes from crowd-labeled imagery, specifically the binary classification task of identifying the presence of washover deposits in each image. Washover is a deposit of sediment left on land surfaces after elevated coastal water levels [e.g., Hudock et al., 2014, Lazarus, 2016]. We crowd-source labels for 528 post-storm images from a two recent hurricanes (Florence and Michael, both 2018) and develop a deep learning model to detect the presence of washover deposits in a corpus of unlabeled imagery. The model is used to determine which unlabeled images have the most uncertain label, and pass these images back to the labeler for human annotation. In addition, a public facing map is used to display geolocated ML predictions of washover presence in post-storm images and provide an additional route for researchers

to label imagery. Both the labeler and interactive map provide more training data for future model improvements.

2 Active Learning Pipeline

The active learning system is broken up into three components: an online labeler, a machine learning model, and an online interactive map (Figure 1). In the subsections below we discuss these components sequentially.




Figure 1: Schematic of the components (boxes) and processes (arrows) involved in the active learning pipeline.

2.1 Labeler

The labeling interface [Rafique et al., 2020] serves a stream of images and a given set of questions to coastal researchers (Figure 2). The labeler is designed to accommodate multiple researchers labeling a single image (e.g., to ensure correct labeling via consensus, and assess inter-rater reliability), and is currently hosted on a virtual machine exposed via an URL <https://coastalimagelabeler.science/>. An administrator uploads images to be labeled (in this case — from NOAA), develops questions for the labeler to present to users, assigns sets of images to each user to label, and periodically exports the labeled data to an open data repository. We developed a bespoke software package to download and manage the collection of NOAA post-storm imagery [Moretz et al., 2020a,b].




Figure 2: The labeler interface. Users are presented with a series of images and questions.

To begin the project, 8 coastal researchers labeled 388 images from Hurricane Florence [Goldstein et al., 2020b]. Each image was labeled until 2 or more researchers agreed on the label. These images were used to train and test the deep learning model. After this initial task, the labeler now hosts a changing set of images for all users that are selected based on the results of the deep learning model (described in the next section).

2.2 Model

The 388 labeled images from Hurricane Florence (179 washover, 209 no washover) were used to develop a model for the binary classification task of detecting washover presence/absence [Goldstein and Mohanty, 2020]. We use VGG16 [Simonyan and Zisserman, 2014], as it has been used successfully for other remote sensing tasks [e.g., Sinha et al., 2019, Sumbul et al., 2020]. The base model was initialized with ImageNet weights, joined to 1 fully connected layer (with 50% dropout), and then fine tuned with a low learning rate (1e-5). We resize the post-storm images to 416 x 416, use image augmentation during training (rotation, width and height shift, shear, zoom, horizontal and vertical flip), and train the model for 200 epochs with an early stopping callback. We use 140 labeled images from Hurricane Michael (70 washover, 70 no washover), labeled by two coastal scientists, to test the model. The F1 score for the test set is 0.92, and the confusion matrix can be seen in Table I. Visual inspection of results using Grad-cam [Selvaraju et al., 2017] suggests that the model is correctly identifying washover deposits (Figure 3).

Table 1: Confusion Matrix for Testing Data (Hurricane Michael)

	Obs. Washover	Obs. No Washover
Pred. Washover	65	6
Pred. No Washover	5	64

The trained model is then used for inference with 9,700 unlabeled images from Hurricane Florence (2018), Hurricane Michael (2018), and Hurricane Isaias (2020). For each of the 3 storms we select the 100 images with the least label certainty — the 100 images with sigmoid output values closest to the decision value of 0.5. Therefore, the threshold certainty value for each storm is $T_s = \max_k(\text{abs}(0.5 - S(I_k)))$, where $S(I_k)$ is the sigmoid output of image k , and \max_k is the value of k^{th} image (where $k = 100$ in this case).




Figure 3: Three images with washover deposits from Hurricane Michael, and the associated Grad-cam heat map from the last model layer.

These 300 images are sent back to the labeler for expert annotation (i.e., an *uncertainty sampling* method of active learning). We select 100 images for each storm to balance the burden on labelers but have enough images to finetune the model further. Once these images have been labeled by a coastal experts we will retrain the model, perform this active learning step on unlabeled images, and send a new batch of images to the labeler. See [Goldstein et al., 2020a] for the most recent labels.

2.3 Map

Each image is geolocated. We visualize the output from the model inference on a map (<https://uncg-daisy.github.io/StormImpactMap/>) to observe the presence/absence of washover deposits contextualized with other data – e.g., hurricane tracks <https://www.nhc.noaa.gov/data/tcr/>, measured washover extents [Doran et al., 2019a], and coastal change forecasts [Doran et al., 2019b]. Post-storm images with their washover prediction are displayed on the online map as a marker (Figure 4). This online interactive map is another route for active learning – users can click the marker and see links to the post-storm image, the Grad-cam overlay on the image obtained from the last layer of the model, and a button to mark the ML prediction as correctly (or incorrect if no washover is actually seen). As washovers are a rare class in the dataset, the map aids in finding and labeling washover imagery and works against future class imbalance in the labeled data. This labeled imagery from the map interface will be incorporated in future model retraining.




Figure 4: Examples of the interactive web map. Left: a view of North and South Carolina, US with the Hurricane Florence track (blue), NOAA post-storm image tiles (stitched and rectified images), and markers (orange and green circles) denoting the number of model predictions in a given region (markers are aggregated at this level of map zoom, color refers to the number of markers in the region). Right: A detailed view of Hurricane Florence storm impacts in North Carolina. Background imagery is NOAA storm tiles, blue markers are the centroids of individual aerial images. Markers are clickable and show image name, acquisition date, sigmoid output for the image, links to the image and Grad-cam overlay, and a button to allow labeling.

3 Future Directions

We have developed a three part active learning pipeline with unlabeled post-storm aerial images. The labeler, model, and map all work in conjunction toward two overarching goals: 1) to crowd source the development of a large labeled dataset of storm impacts, and 2) continually improve a model to detect storm impacts in sandy, low-sloped coastal regions. A successful model can potentially be used for rapid and automated assessment of post-storm imagery. The pipeline is currently focused on identifying hurricane washover deposits, but can be expanded to look at other impacts (e.g., damage to the built environment, flooding, dune erosion, ecosystem impacts, etc.).

This pipeline does not currently incorporate pre-storm imagery, so the model is susceptible to classifying past storm impacts (i.e., washover from previous storms) and other sandy bowl-shaped geomorphic features (e.g., dune blowouts) as current washover deposits. Though note that washover deposits can be long-lived features that are reactivated in subsequent storms [e.g., Hosier and Cleary, 1977]. Future work will focus on testing different model architectures (to allow for active learning using *query-by-committee*) and also providing additional data to the learner (i.e., pre-storm imagery, storm water levels, storm wave heights) that can be used to improve model performance.

Acknowledgments and Disclosure of Funding

We thank E. Lazarus, D. Buscombe, T. Beuzen, N. Cohn, H. Williams, P. Limber, L. Taylor, A. Ellenson, J. Warrick for helpful discussions about the project and for labeling images. The authors gratefully acknowledge support from the USGS (G20AC00403 to EBG and SDM), NSF (1939954 to EBG), Microsoft AI for Earth (to EBG and SDM), and an Early-Career Research Fellowship from the Gulf Research Program of the National Academies of Sciences, Engineering, and Medicine (to EBG). The content is solely the responsibility of the authors and does not necessarily represent the official views of the funders. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the opinions or policies of the U.S. Government or the funders. Mention of trade names or commercial products does not constitute their endorsement by the U.S. Government or funders.

References

Daniel Buscombe and Andrew C Ritchie. Landscape classification with deep neural networks. *Geosciences*, 8(7):244, 2018.

Kara S. Doran, Justin J. Birchler, and Kirsten J. Bendik. Storm-induced overwash extent. 2019a. doi: <https://doi.org/10.5066/P9BW6CG6>.

Kara S. Doran, Justin J. Birchler, Heather A. Schreppel, Hilary F. Stockdon, and David M. Thompson. Storm-induced coastal change forecasts—archive of individual storm events. 2019b. doi: <https://doi.org/10.5066/P9Z362BC>.

Evan B. Goldstein and Somya D. Mohanty. Washoverml: Initial model, November 2020. URL <https://doi.org/10.5281/zenodo.4277581>.

Evan B Goldstein, Giovanni Coco, and Nathaniel G Plant. A review of machine learning applications to coastal sediment transport and morphodynamics. *Earth-Science Reviews*, 194:97–108, 2019.

Evan B. Goldstein, Tomas Beuzen, Daniel Buscombe, Nicholas Cohn, Ashley Ellenson, Eli D. Lazarus, and Hannah E. Williams. Labels for Emergency Response Imagery from Hurricane Florence, Hurricane Michael, and Hurricane Isaias, November 2020a. URL <https://doi.org/10.5281/zenodo.4272064>.

Evan B Goldstein, Eli D Lazarus, Tomas Beuzen, Hannah Williams, Patrick Limber, Nicholas Cohn, Luke Taylor, Daniel Buscombe, Shah Nafis Rafique, Matthew Moretz, and et al. Labels for hurricane florence (2018) emergency response imagery from noaa. 2020b. doi: 10.6084/m9.figshare.11604192.v1. URL https://figshare.com/articles/dataset/Labels_for_Hurricane_Florence_2018_Emergency_Response_Imagery_from_NOAA/11604192/1.

Paul E Hosier and William J Cleary. Cyclic geomorphic patterns of washover on a barrier island in southeastern north carolina. *Environmental Geology*, 2(1):23, 1977.

Jessica W Hudock, Peter P Flaig, and Lesli J Wood. Washover fans: A modern geomorphologic analysis and proposed classification scheme to improve reservoir models washover fans: A modern geomorphologic analysis and classification. *Journal of Sedimentary Research*, 84(10):854–865, 2014.

Eli D Lazarus. Scaling laws for coastal overwash morphology. *Geophysical Research Letters*, 43(23): 12–113, 2016.

Sophia B Liu, Barbara S Poore, Richard J Snell, Aubrey Goodman, Nathaniel G Plant, Hilary F Stockdon, Karen L Morgan, and M Dennis Krohn. Usgs icoast—did the coast change? designing a crisis crowdsourcing app to validate coastal change models. In *CSCW Companion '14: Proceedings of the Companion Publication of the 17th ACM Conference on Computer Supported cooperative work & social computing*, pages 17–20, 2014. doi: 10.1145/2556420.2556790.

Matthew C. Moretz, Daniel Foster, John Weber, Rinty Chowdhury, Shah Nafis Rafique, Evan B. Goldstein, and Somya D. Mohanty. psi-collect: A Python module for post-storm image collection and cataloging, February 2020a. URL <https://doi.org/10.5281/zenodo.3687347>.

Matthew C Moretz, Daniel Foster, John Weber, Rinty Chowdhury, Shah Nafis Rafique, Evan B Goldstein, and Somya D Mohanty. psi-collect: A python module for post-storm image collection and cataloging. *Journal of Open Source Software*, 5(47):2075, 2020b.

Karen L Morgan, Nathaniel G Plant, Hilary Stockdon, and Richard J Snell. ICoast—did the coast change?: Storm-impact model verification using citizen scientists. *Coastal sediments 2019*, pages 1424–1438, 2019.

Shah Nafis Rafique, Evan B. Goldstein, and Somya D. Mohanty. Coastal-image-labeler v2, August 2020. URL <https://doi.org/10.5281/zenodo.4000221>.

Justin T Ridge, Patrick C Gray, Anna E Windle, and David W Johnston. Deep learning for coastal resource conservation: automating detection of shellfish reefs. *Remote Sensing in Ecology and Conservation*, 2019.

Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017.

Burr Settles. From theories to queries: Active learning in practice. In *Active Learning and Experimental Design workshop In conjunction with AISTATS 2010*, pages 1–18, 2011.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Saumya Sinha, Sophie Giffard-Roisin, Fatima Karbou, Michaël Deschartres, Anna Karas, Nicolas Eckert, Cécile Coléou, and Claire Monteleoni. Can avalanche deposits be effectively detected by deep learning on sentinel-1 satellite sar images? 2019.

Gencer Sumbul, Jian Kang, Tristan Kreuziger, Filipe Marcelino, Hugo Costa, Pedro Benevides, Mario Caetano, and Begüm Demir. Bigearthnet deep learning models with a new class-nomenclature for remote sensing image understanding. *arXiv preprint arXiv:2001.06372*, 2020.