National Aeronautics and Space Administration

Automatic Per-Pixel Classification of UAVSAR Imagery for Hurricane Flood Detection

Second AI and Data Science Workshop for Earth and Space Sciences, February 9-11, 2021
Poster ID: 4

Authors: Michael Denbina (michael.w.denbina@jpl.caltech.edu),
Zaid J. Towfic, Matthew Thill, Brian Bue, and Yunling Lou
Jet Propulsion Laboratory, California Institute of Technology

1. Introduction and Objectives

- The main objective of this project was to assess the viability of convolutional neural network-based image classifier architectures to automatically detect flooded areas in polarimetric radar imagery collected by the NASA/JPL Uninhabited Aerial Vehicle Synthetic Aperture Radar (UAVSAR) instrument. UAVSAR is a fully polarimetric L-band synthetic aperture radar, flown aboard a NASA Gulfstream III aircraft (see Figure 1).
- The objectives of our project were as follows:
  1. Build a sizable data set of matched UAVSAR imagery and flood masks, to use as training and testing data for the machine learning algorithms. We are using UAVSAR data collected after Hurricane Harvey for this purpose.
  2. Train two pixelwise predictors, U-Net and SegNet (both established convolutional neural network architectures) using the flood labels from objective #1, and evaluate the pixelwise prediction accuracy of each candidate model on blind validation data.

Fig. 1. UAVSAR mounted on a NASA Gulfstream III aircraft. The UAVSAR radar pod is mounted under the body of the aircraft.

2. Automated Flood Mapping Techniques

- There are pre-existing methods for detecting flooding in UAVSAR imagery based on manual interpretation (time-consuming), or simple decision tree or thresholding (low accuracy). However, we believed these methods were likely to be outperformed by more modern image classification approaches, including convolutional neural networks.
- Convolutional neural network architectures such as fully-convolutional-networks and encoder-decoder networks, as shown in Figure 2, immediately lend themselves to our per-pixel labeling effort. These networks utilize only convolutional and other spatial feature-preserving layers throughout the architecture in order to capture discriminative spatial patterns represented in the input images associated with the pixelwise labels.
- We tested two well-known architectures from the literature, SegNet and U-Net, and compared the results to a pre-existing decision tree classifier used as a baseline.

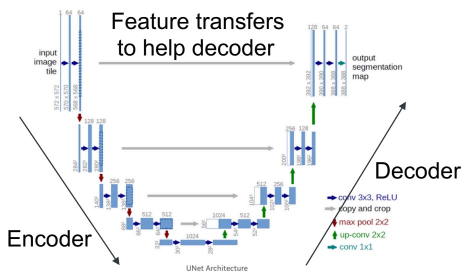

Fig. 2. Illustration of the U-Net architecture, from Ronneberget et al., "U-Net: Convolutional Networks for Biomedical Image Segmentation," MICCAI 2015, https://arxiv.org/abs/1505.04597

3. Manual Labelling of Flooded UAVSAR Data

- We developed a simple labelling tool which shows the user an image segment from the UAVSAR imagery, alongside various reference data to aid in determining the appropriate label. The user can click the button corresponding to the desired label for that segment, and the tool will then prompt the user with the next segment. An example is shown in Figure 3.
- We labelled 10873 total image segments (covering over 3.5 million UAVSAR image pixels). After manual labelling, we extracted 40000 randomly located, overlapping 64x64 pixel image patches from two flight lines to use for training. 300 randomly located, non-overlapping patches were extracted for testing/validation and kept separate from the training set.

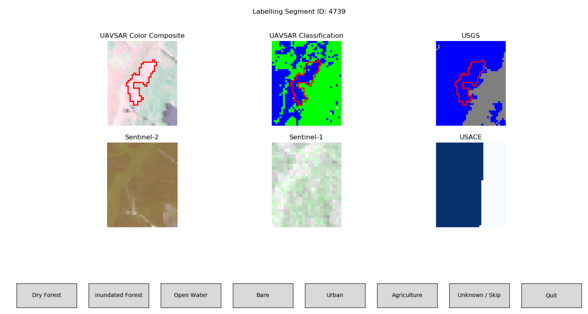

Fig. 3. Example of the UAVSAR manual flood labelling tool, for a flooded forest area.

4. Results and Discussion

- **We generated predicted maps of flooded and non-flooded areas for UAVSAR data collected over the Texas Gulf Coast during flooding caused by Hurricane Harvey in August 31 - September 1, 2017, as shown in Figure 4.**
- We assessed the classifier accuracy using manually labelled testing data held out from training, as well as NOAA aerial imagery (used for validation manually at specific points shown in Figure 4 as black circles, limited by the cloud-free coverage of the imagery). While the map shows five different classes, for the purposes of assessing the classifier accuracy, we broke the results down into only two classes, flooded (containing open water and flooded vegetation) and non-flooded (containing all other classes).
- The overall accuracy of our U-Net classifier on the manually labelled testing data was 87%, with 85% accuracy for our SegNet classifier. The Kappa coefficient of U-Net was 0.73, with F1 score for the non-flooded class of 0.89 and F1 score for the flooded class of 0.84. SegNet had similar but slightly lower accuracies than U-Net. The baseline classifier (decision tree method) had overall accuracy of 75%, with Kappa coefficient of 0.48, non-flooded F1 score of 0.81, and flooded F1 score of 0.66. **Both U-Net and SegNet outperformed the baseline classifier.**
- **Using the NOAA aerial imagery, U-Net had overall accuracy of 82%, with Kappa coefficient of 0.65.** SegNet had similar but slightly lower accuracy (by ~2%). None of the classifiers are able to accurately identify flooding in urban areas, however, which is an area of future work (to collect more training data in these spatially complex, heterogeneous environments).

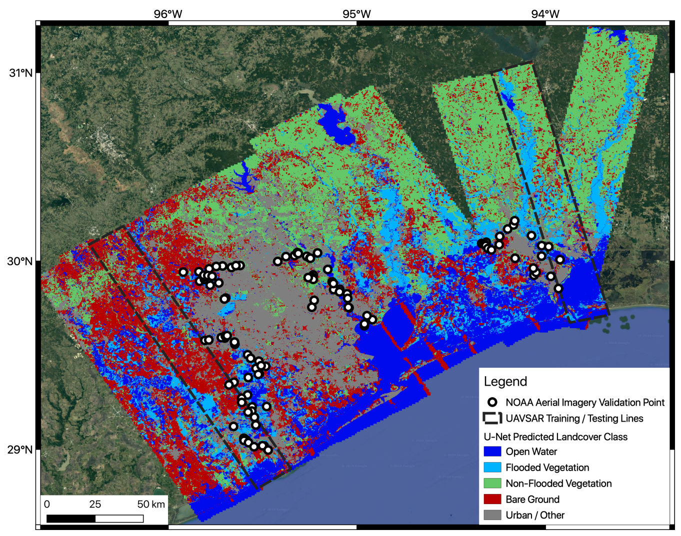

Fig. 4. Map of the study area affected by Hurricane Harvey, showing the U-Net predicted classes generated from the UAVSAR data as input. The dashed black boxes show the UAVSAR flight lines used to collect training data. The black circles show points used to validate the map using NOAA aerial imagery.

5. Conclusions and Future Work

- **We demonstrated potential of CNN-based image classifiers for classification of flooded areas in UAVSAR data.** With U-Net, 87% overall accuracy using manually labelled testing data, and 82% overall accuracy (outside of urban areas) when validated using NOAA aerial imagery (0.65 Kappa coefficient).
- **Future Work**: Apply trained classifiers to Hurricane Florence data (see Figure 5) to test transferability of trained classifiers to other study areas and assess classifier accuracy. **Collect more training and testing data, particularly in urban environments**, which is a current limitation of this approach. Test other input to classifiers, and increased patch size.

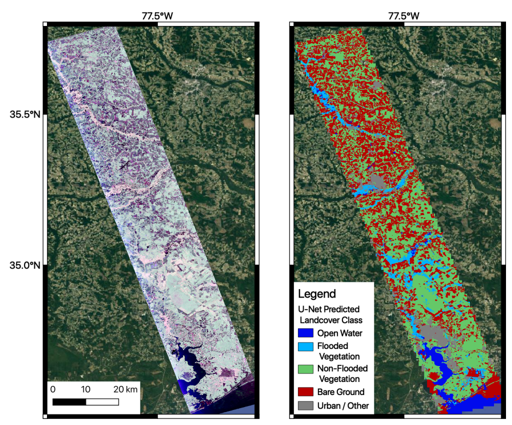

Fig. 5. UAVSAR data (left) and map of predicted classes (right) when applying the Hurricane Harvey trained U-Net classifier to UAVSAR data from Hurricane Florence. The end goal is to have a trained classifier ready to go that can be applied to data from future hurricanes.

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

www.nasa.gov

ACKNOWLEDGEMENTS: We thank the Data Science Working Group for funding this 2019 Data Science Pilot. We would also like to thank Gordon Wells of the Center for Space Research at The University of Texas at Austin for providing the flood inundation forecast data. We also acknowledge the UAVSAR team for collecting and processing the UAVSAR data used in this project.