Flood Segmentation on Sentinel-1 SAR Imagery with Semi-Supervised Learning

Siddha Ganju*, Sayak Paul*

Code: bit.ly/etci-code | Paper: bit.ly/etci-paper

* = Equal Contribution

Data Analysis

- 66k tiled images from Nebraska, North Alabama, Bangladesh, Red River North, and Florence.

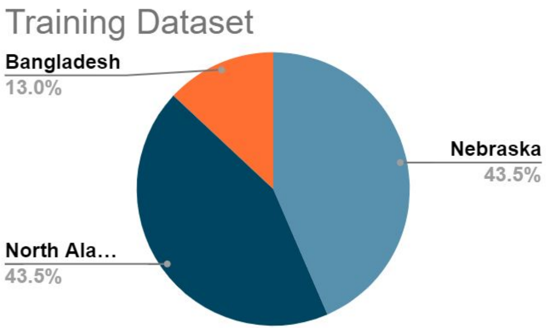

- Validation = Florence
- Test = Red River North

Data Processing - Removing images with swath gaps

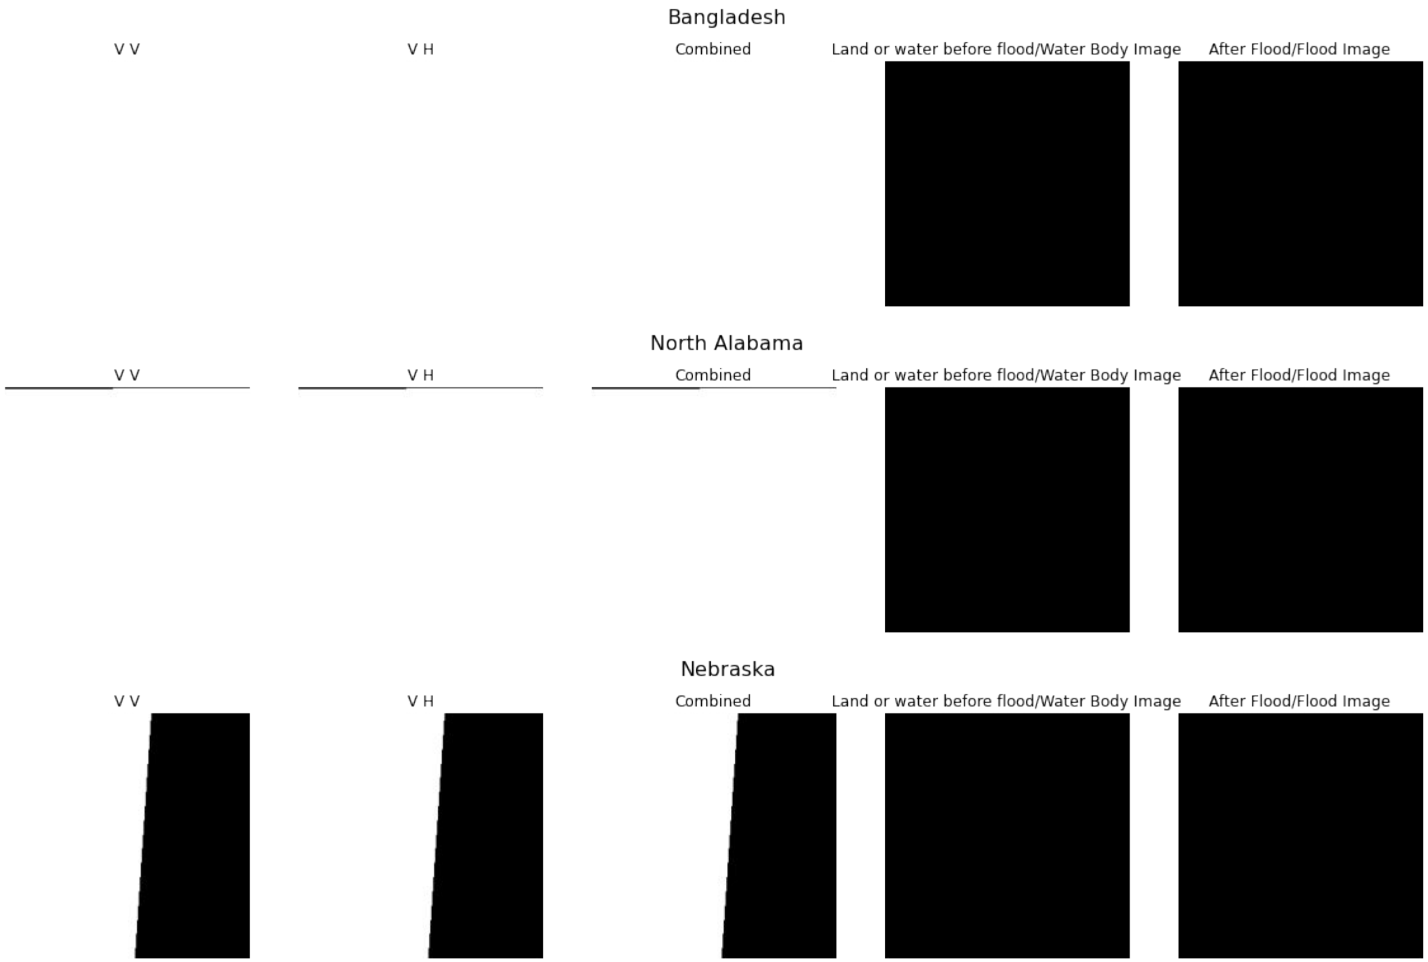

Data Processing - Generating RGB tiles with ESA Polarimetry

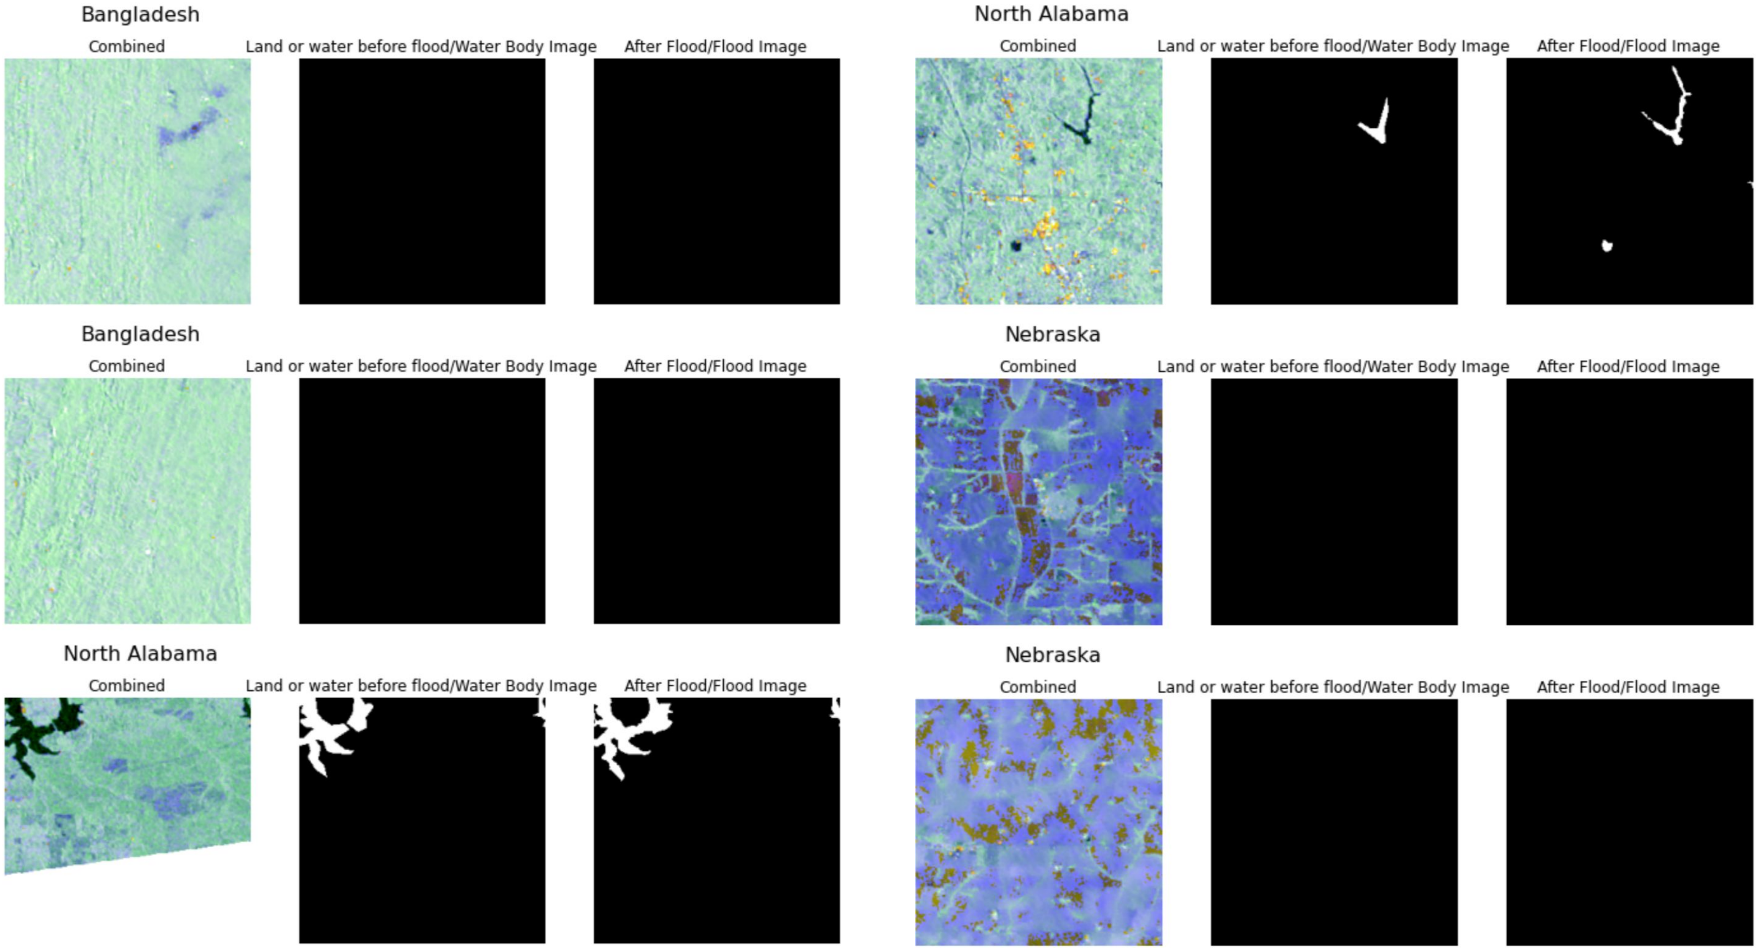

Sampling data

- Timestamps
  - Intuitively could tell the progress of flooding events.
  - Empirically no impact with random sampling or maintaining ordered pairs.

- Remove swath gaps or missing data
- Stratified sampling
  - Each training batch contained at least 50% of samples having flood levels.
  - This also helps mitigating the class imbalance problem in the dataset.

Augmentation

- **Training**: horizontal flips, rotations, and elastic transformations.
- **Test-time**: horizontal, vertical flips, transpositions and 90 rotations (*Dihedral Group D4*).

Method	IoU
U-Net	0.52
U-Net + TTA	**0.57**

Training details

- **Encoder backbone**: MobileNetV2 due to its pointwise convolutions and performance consistency.

Model Architecture + Encoder Backbone	**IoU**
U-Net + ResNet34 [13]	0.55
U-Net + RegNetY-002 [29]	0.56
DeepLabV3Plus + MobileNetV2	0.52
DeepLabV3Plus + RegNetY-002	0.46
U-Net + MobileNetV2	**0.57**

- **Segmentation architectures**: UNet and UNet++.
- **Loss function**: Performance improvement with *Dice loss* alone compared to Focal loss and the two combined.

Training with pseudo labeling

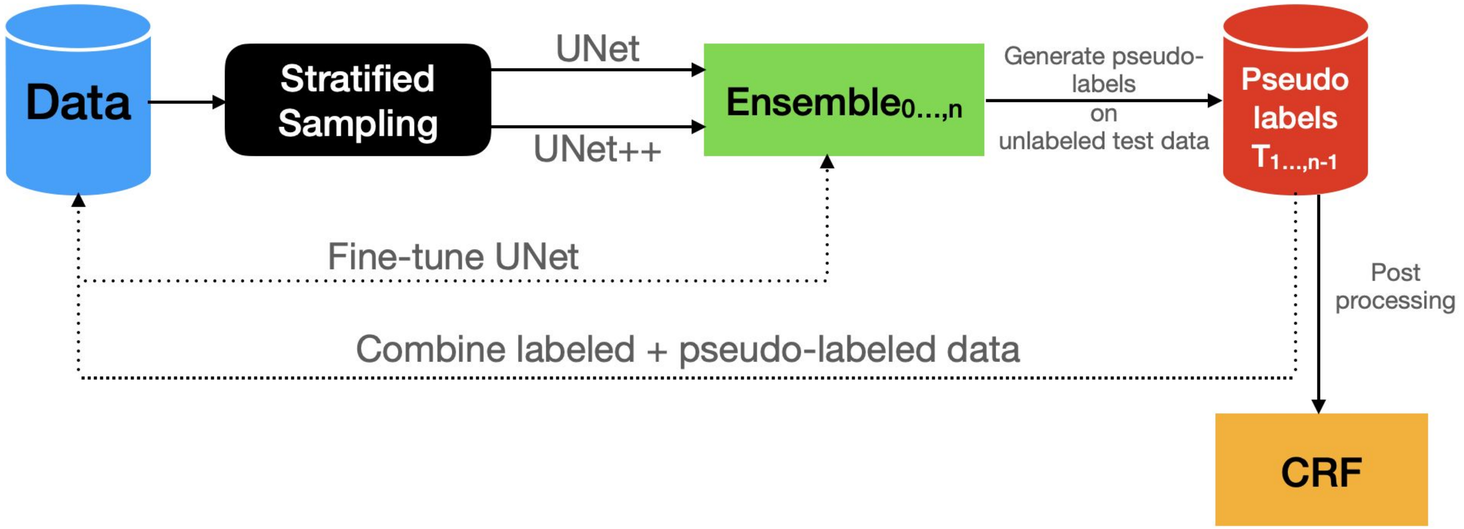

- **Step 1:** Training on available data, performing inference on entire test data, and generating Pseudo Labels
- **Step 2:** Filtering quality pseudo labels
- **Step 3:** Combining Pseudo Labels + Original Training data
- Repeat Steps 1,2,3
- Post processing with CRFs

Post processing with CRFs

- We used Conditional Random Fields (CRF) to post-process the predictions.
- CRFs helped refining the segmentation boundaries resulting in better final performance.
- Establishes a new SOTA on Sentinel-1 SAR Imagery.

Method Description	**IoU** $\uparrow$
Random Baseline (all zeroes)	0.00
Competition Provided Baseline	0.60
Standard U-Net	0.57
Ensemble with CRF post processing	0.68
Pseudo labeling + Ensembles with CRF post processing	**0.7654**

Benchmarking

- Segmentation masks are generated in approx 3 seconds per Sentinel-1 tile.
- Covers an area of approximately 63,152 squared kilometers!
- Larger than the area covered by Lake Huron, the second largest fresh water Great Lake of North America.

Future Work

- Eliminate CRFs because they are computationally expensive.
- Develop a single end-to-end training workflow to make the process more streamlined.
- More suitable architectures for segmenting satellite imagery.
- **Collaborating with the competition organizers and UNOSAT team to benchmark real time runtimes and to evaluate the scalability of our solution.**

Acknowledgements

- NASA Earth Science Data Systems Program, NASA Digital Transformation AI/ML thrust, and IEEE GRSS for organizing the ETCI competition.
- Google Developers Experts* program for providing Google Cloud Platform credits to support our experiments.
- Charmi Chokshi and domain experts Shubhankar Gahlot, May Casterline, Ron Hagensieker, Lucas Kruitwagen, Aranildo Rodrigues, Bertrand Le Saux, Sam Budd, Nick Leach, and, Veda Sunkara.

* https://developers.google.com/programs/experts/