Google Research

Team: EO Science team

Cross Modal Distillation for Flood Extent Mapping


Shubhika Garg


Ben Feinstein


Shahar Timnat


Vishal Batchu


Gideon Dror


Adi Gerzi Rosenthal


Varun Gulshan

Problem Description

Segment the extent of flooding from Sentinel-1 SAR imagery.

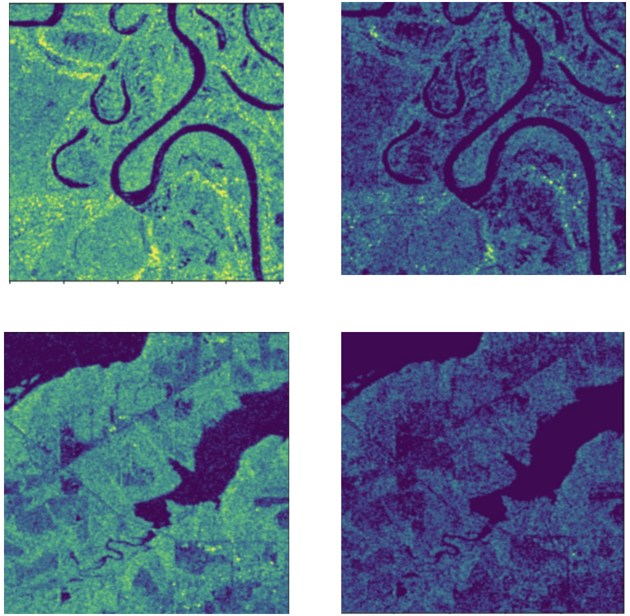

VV BAND VH BAND

SAR INPUT (2 CHANNEL)

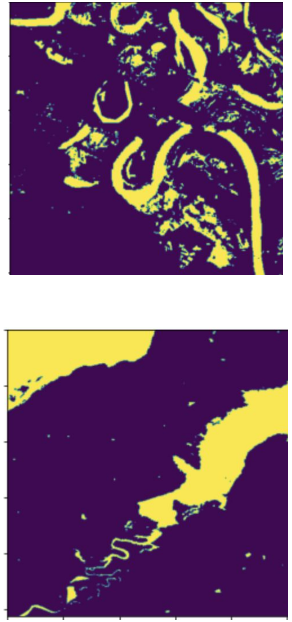

WATER LABEL OUTPUT

Why are we solving this?

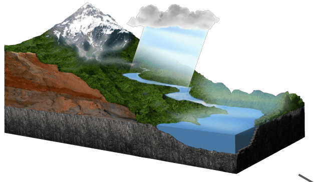

Hydrological Model

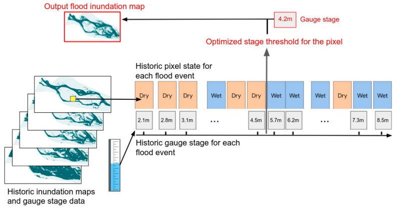

Inundation Model

Warning distribution

Flood warning systems platform

- Floods affect between 85 million to 250 million people annually
- **Aim***:* Alleviate harms by improving the accuracy of early flood warning systems

Why are we solving this?

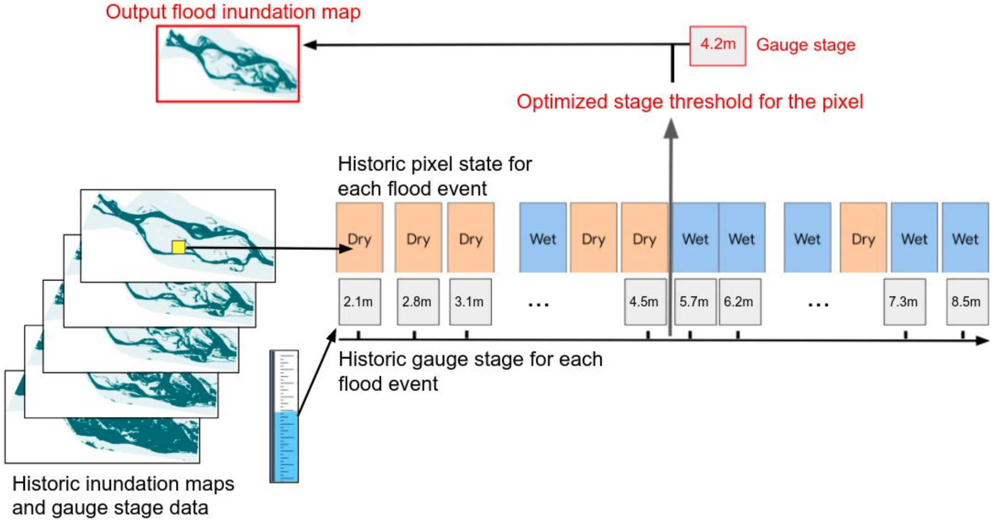

Inundation Model: module of flood warning system
Image taken from Flood forecasting model paper

Inundation Model

- Learns a mapping between historical river water gauge levels and the corresponding flooded area to predict future flooding extent
- Increase the accuracy of the algorithm used to generate historical flood maps.

What are SAR Images?

- Sentinel 1 satellite comprises a constellation of two polar-orbiting satellites and provides us with **free C-band synthetic aperture radar (SAR) image data.**

- Is an example of **active remote sensing** that allows to get image in both day and night.

- Sentinel-1 has a spatial **resolution of 10m** and a **return period of 6 days**.

Active Sensing

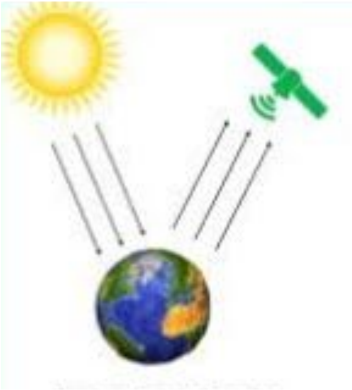

Passive Sensing

Why SAR Images?

- Clouds are heavily correlated to flooding events. A cloud-free optical image (e.g. Sentinel 2) is rarely available in such events.

- SAR images see through cloud due to the large wavelength (~cms) used for imaging.

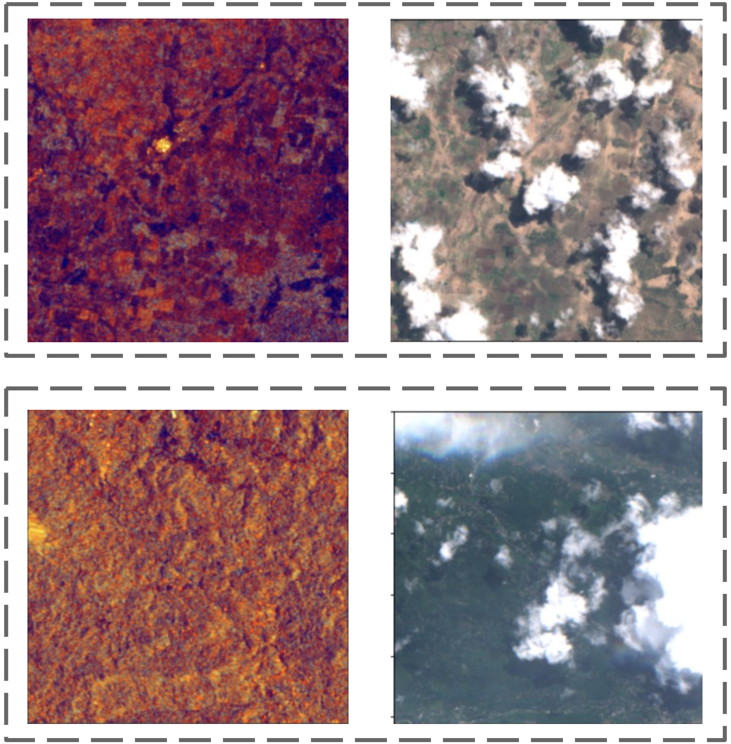

SAR Image   Optical Image

Existing Methods

Thresholding the VV/VH band

Used extensively as it is easy to segment water in SAR images because of its **low backscatter intensity**.

Failure modes are:

- Generates lot of false positives as background noise
- Sensitive to mountain shadows

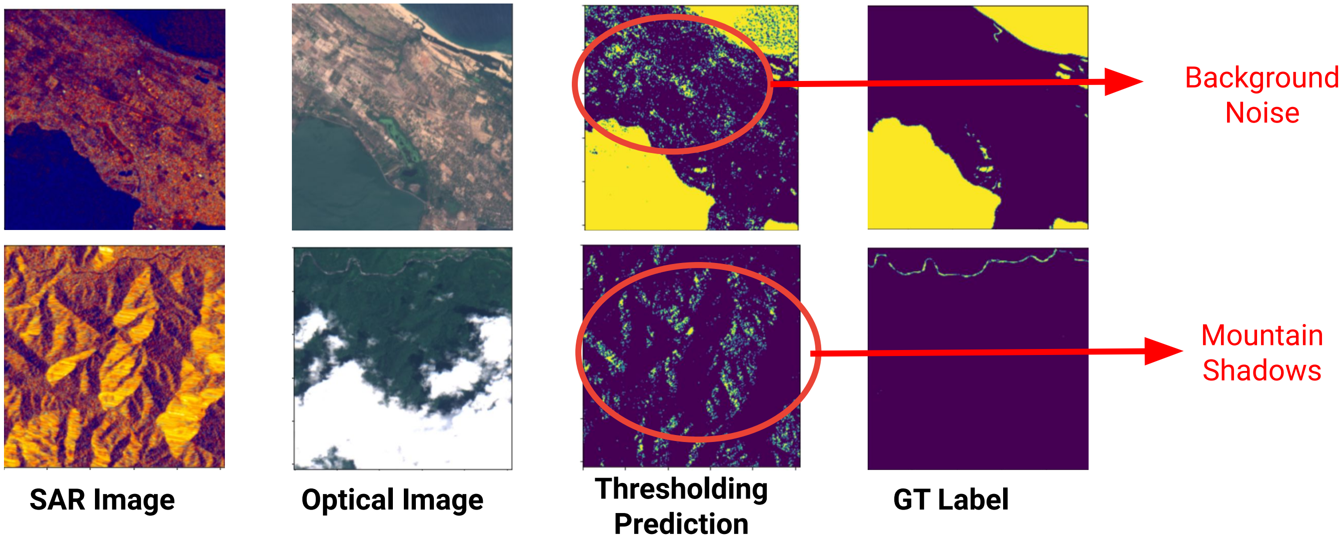

Dataset Details

Sen1Floods11 Dataset

- *Hand labeled data*
  - Water is hand labeled by experts
  - IID split:
    - 252 training images
    - 89 validation images
    - 90 test images
- *Weak labeled data*
  - Noisy labels created using MNDWI (Modified Normalized Difference Water Index) and NDVI (Normalized Difference Vegetation Index) band.
  - Total images: 4385 chips (only used for training)

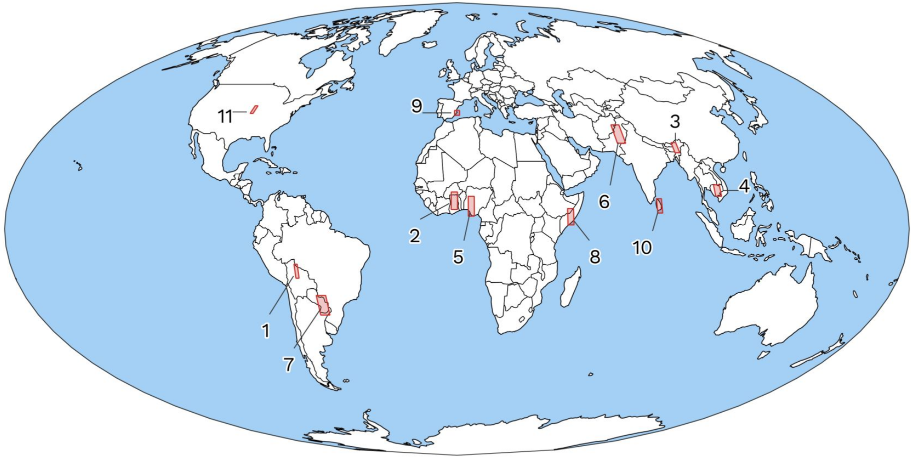

Locations from where flood event data was sampled.
Image taken from Sen1floods11 paper

Dataset Details

Additional Dataset

- Opportunistically sampled historical paired Sentinel-1 and Sentinel-2 data during flood events using Earth Engine**.
- Created weak label using NDWI (Normalized Difference Water Index).
- Total images: 23,260 chips (only used for training).

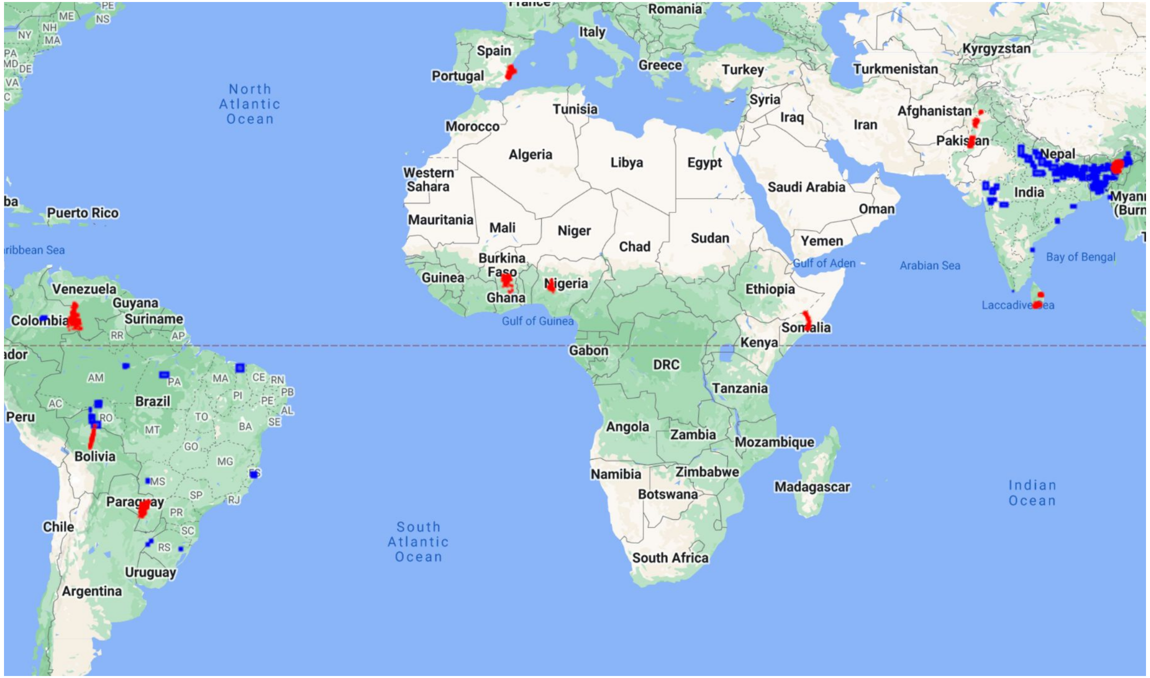

Red points: Sen1Floods11 data points region sampling
Blue points: External Dataset data points region sampling

** Though cloud free Optical (Sentinel-2) data is rarely available during inference, enough samples from previous flooding years are available with low cloud percentage for training.

Dataset Details

Weak labels quality issues

Complete
river
absent

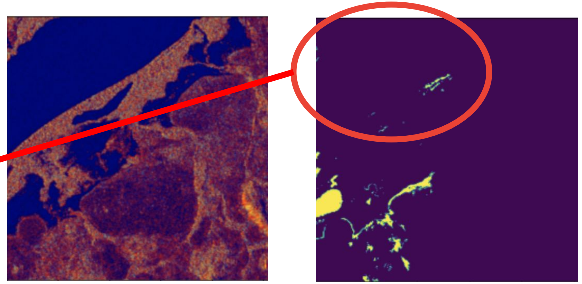

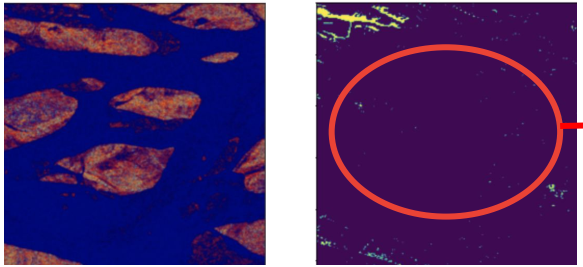

Complete
river
absent

Dataset Details

Weak labels quality issues

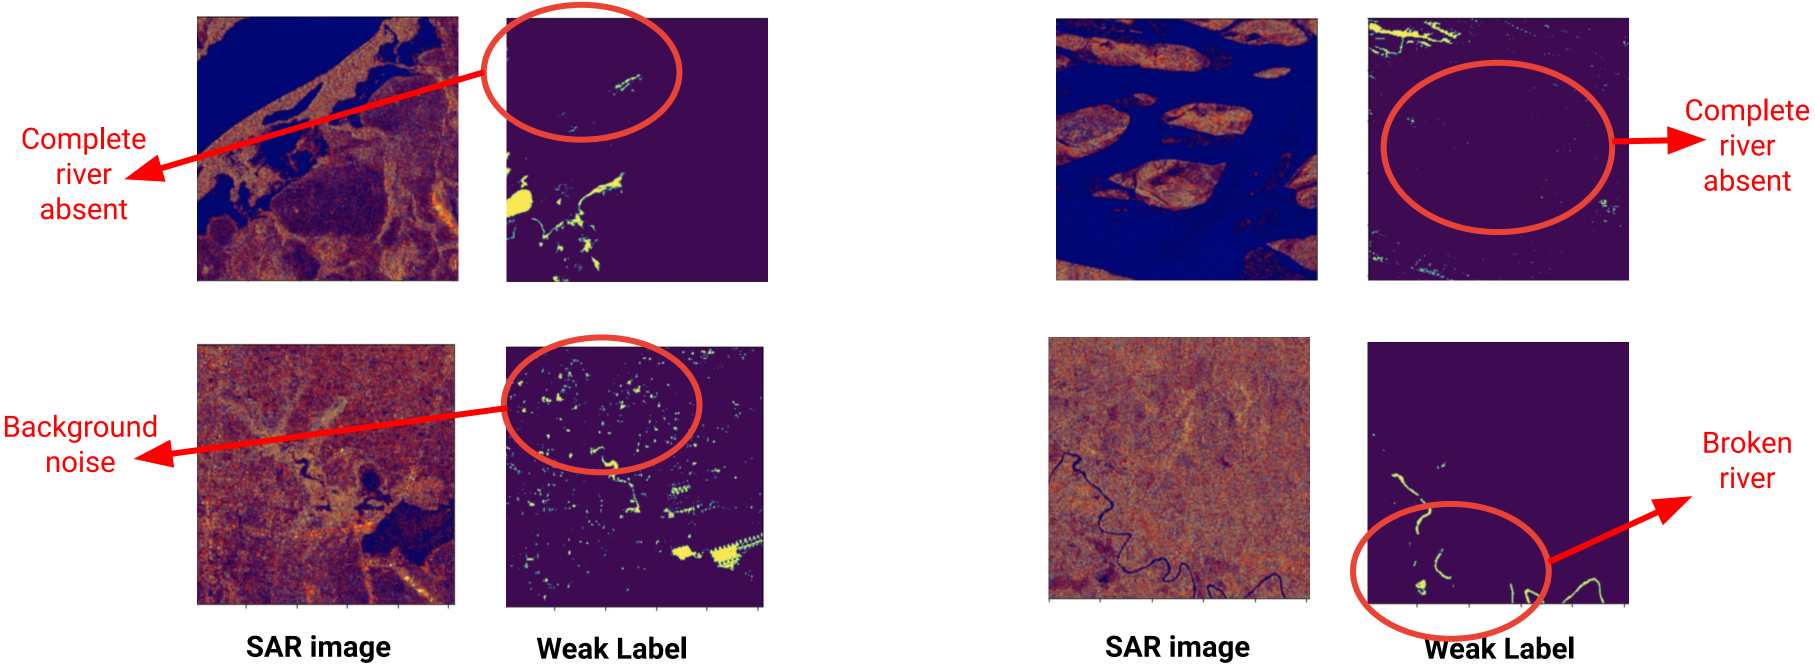

Methodology

Supervised Baseline

- Architecture: Deeplab v3+ with Xception65 backbone
- Loss: Cross entropy with additional weight given to edges

Baseline 1

- Input: Sentinel-1 image
- Dataset: Sen1Floods11 hand label train split
- Dataset size: 252 images

Baseline 2

- Input: Sentinel-1 image
- Dataset: Sen1Floods11 weak label
- Dataset size: 4385 images

Methodology

Cross modal distillation training: Stage 1

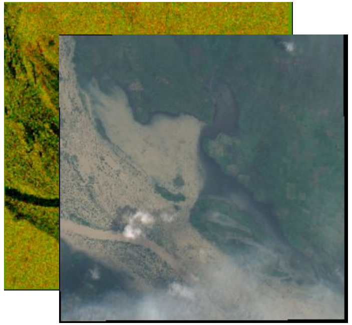

Teacher input:
Stacked Sentinel-1 +
Sentinel-2 image X_{s1+s2}

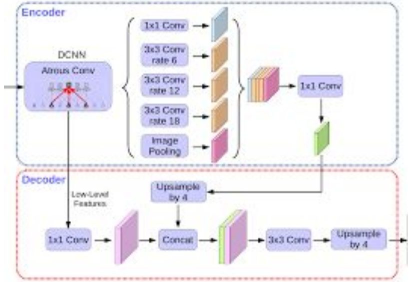

Teacher network f_t
(Deeplab v3+)

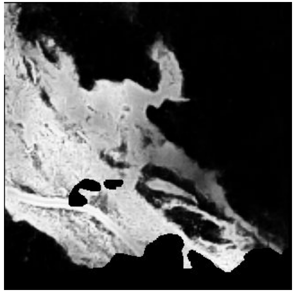

Teacher probability p_t

CE loss

Hand label

Methodology

Cross modal distillation training: Stage 2

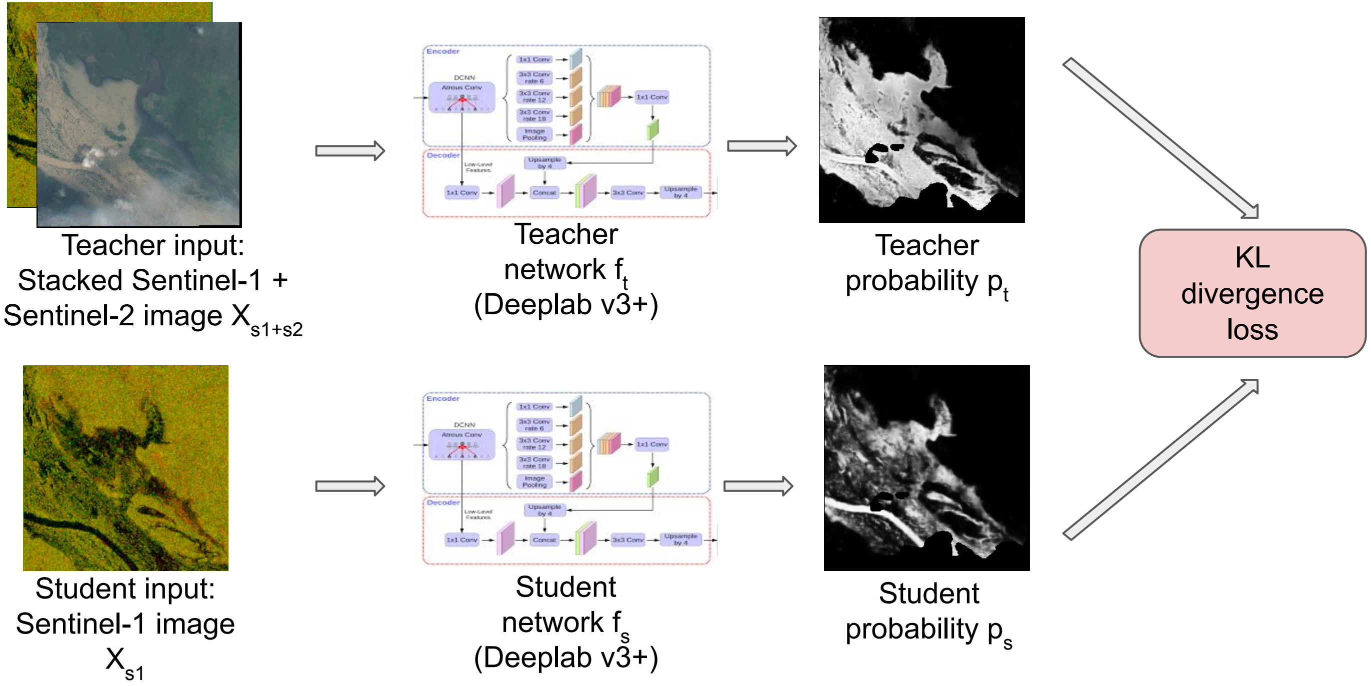

Methodology

Cross modal distillation inference stage

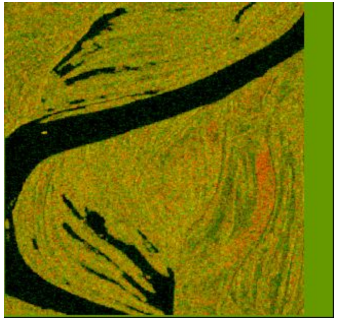

Input: Sentinel-1 image X_{s1}

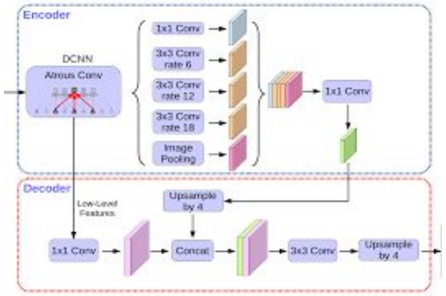

Student network f_s (Deeplab v3+)

Output

Methodology

Cross modal distillation

Motivation:

- Extract information from a more informative modality (Sentinel-2) to supervise paired Sentinel-1 SAR images
- Use a small hand labeled set and the large weak labeled as unlabeled set to train a teacher student network.

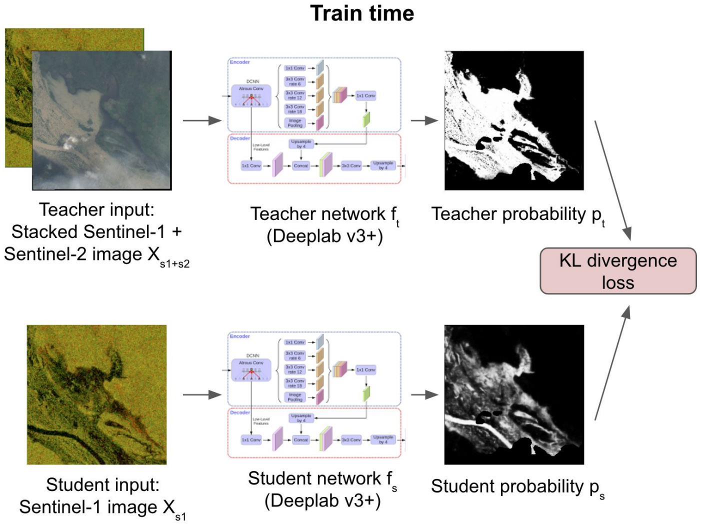

Cross Modal distillation framework

Quantitative Results

Method	mIoU water class
Otsu thresholding	54.58
Hand labeled Sentinel-1 supervised	67.63 ± 0.45

Result comparison on Sen1Floods11 handalbel test split
The numbers show the aggregated mean and standard deviation of IoU from 5 runs.

Quantitative Results

Method	mIoU water class
Otsu thresholding	54.58
Hand labeled Sentinel-1 supervised	67.63 ± 0.45
Weak labeled supervised: Sen1Floods11 weak	67.76 ± 2.41

Result comparison on Sen1Floods11 handalbel test split
The numbers show the aggregated mean and standard deviation of IoU from 5 runs.

Quantitative Results

Method	mIoU water class
Otsu thresholding	54.58
Hand labeled Sentinel-1 supervised	67.63 ± 0.45
Weak labeled supervised: Sen1Floods11 weak	67.76 ± 2.41
Weak labeled supervised: Sen1Floods11 + Additional weak label data	68.94 ± 1.11

Result comparison on Sen1Floods11 handalbel test split
The numbers show the aggregated mean and standard deviation of IoU from 5 runs.

Quantitative Results

Method	mIoU water class
Otsu thresholding	54.58
Hand labeled Sentinel-1 supervised	67.63 ± 0.45
Weak labeled supervised: Sen1Floods11 weak	67.76 ± 2.41
Weak labeled supervised: Sen1Floods11 + Additional weak label data	68.94 ± 1.11
Cross modal distillation	**71.91 ± 0.41**

Result comparison on Sen1Floods11 handalbel test split
The numbers show the aggregated mean and standard deviation of IoU from 5 runs.

Qualitative Results

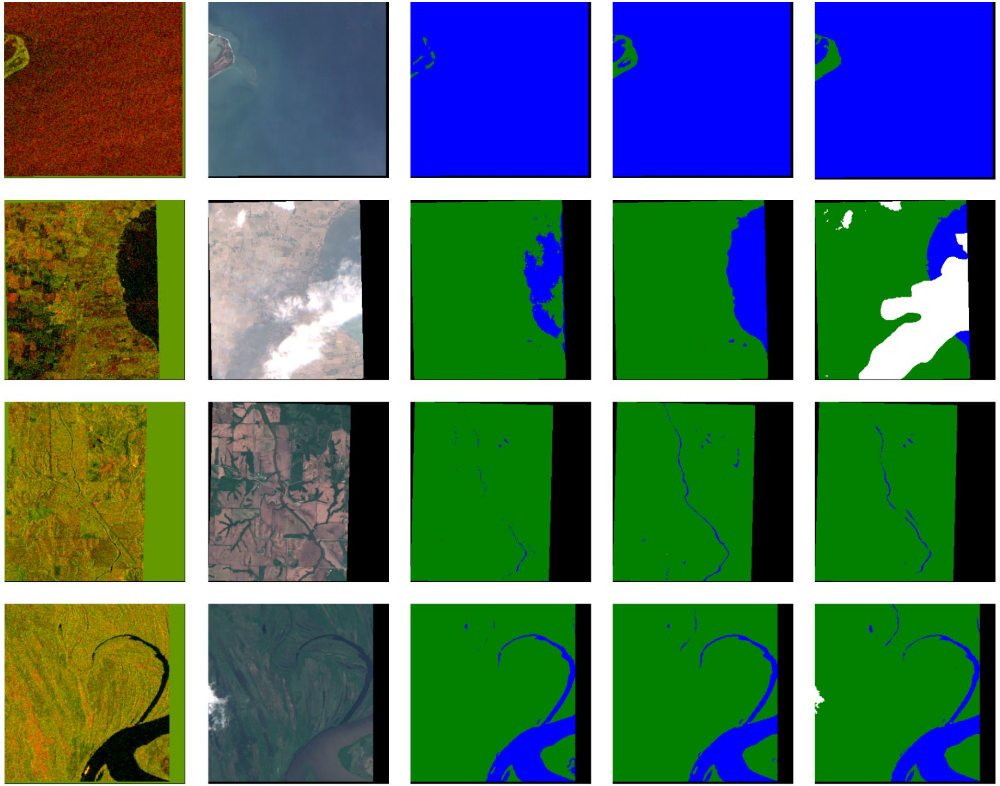

Sentinel-1 image | Sentinel-2 image | Weak supervised model output | Cross distillation model output | Ground truth

Failure Cases

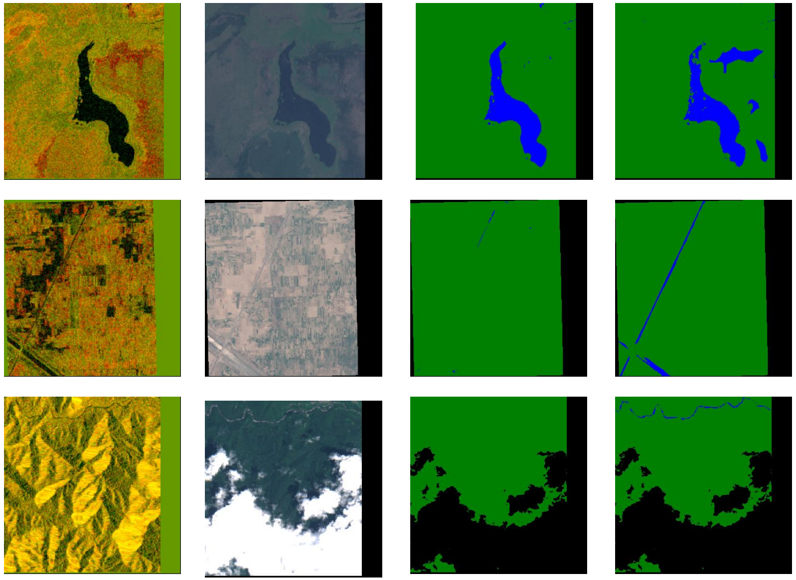

Sentinel-1 Image | Sentinel-2 Image | Cross distillation model output | Ground Truth

Thank You