Deep Learning-Based Prediction of Coastal Storm Tide Flood Maps

Molly McKenna
J.C. Dietrich, T.A. Cuevas López, & Dylan Anderson

High-Definition Maps for Communicating Flood Risk

Maps shape risk perception

➔ People understand flooding through spatial visuals, not numbers.

Maps make risk clearer and more actionable

➔ Communities interpret risk faster and at local scales

ADCIRC needs downscaling

➔ Mesh → high resolution map that communities and emergency managers can use

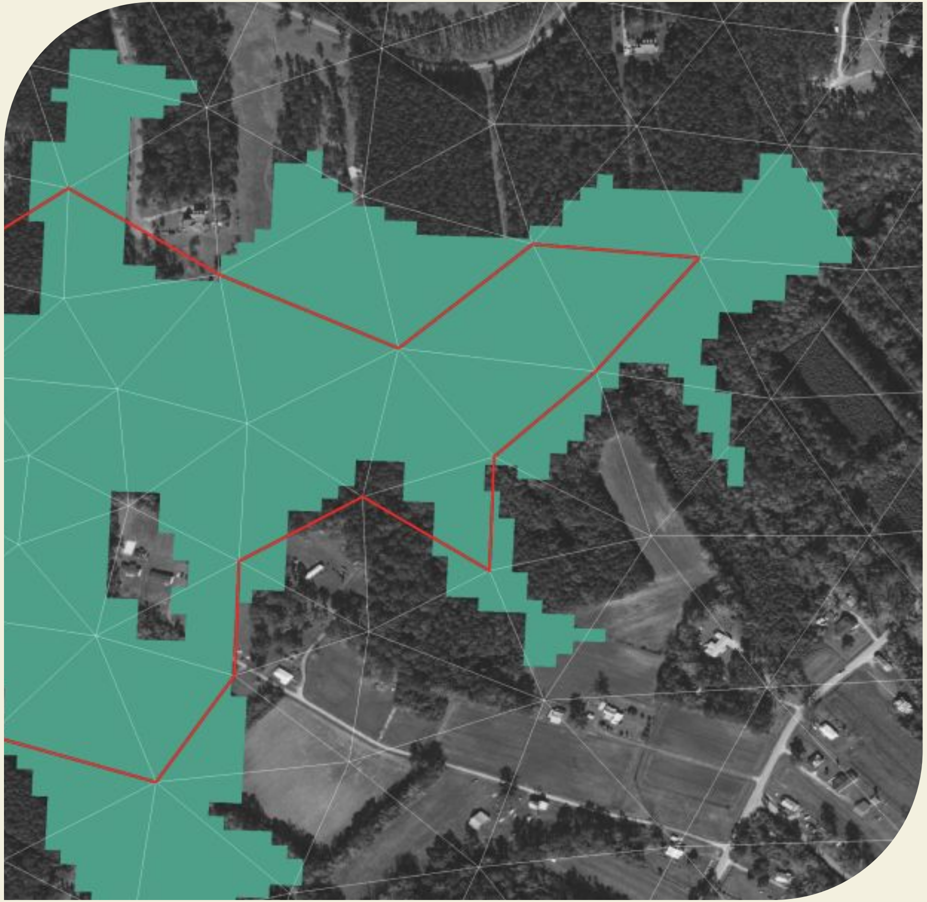

Example of Kalpana downscaling.

NN Models for Coastal Flooding

NN Models for Coastal Flooding

1. **Early Models: point-based surge predictions**
2. Later Models: hybrid models, hundred/thousand point surge predictions
3. Storm Surge predictions of coarse spatial maps
4. First storm-tide Surrogate model: point-based

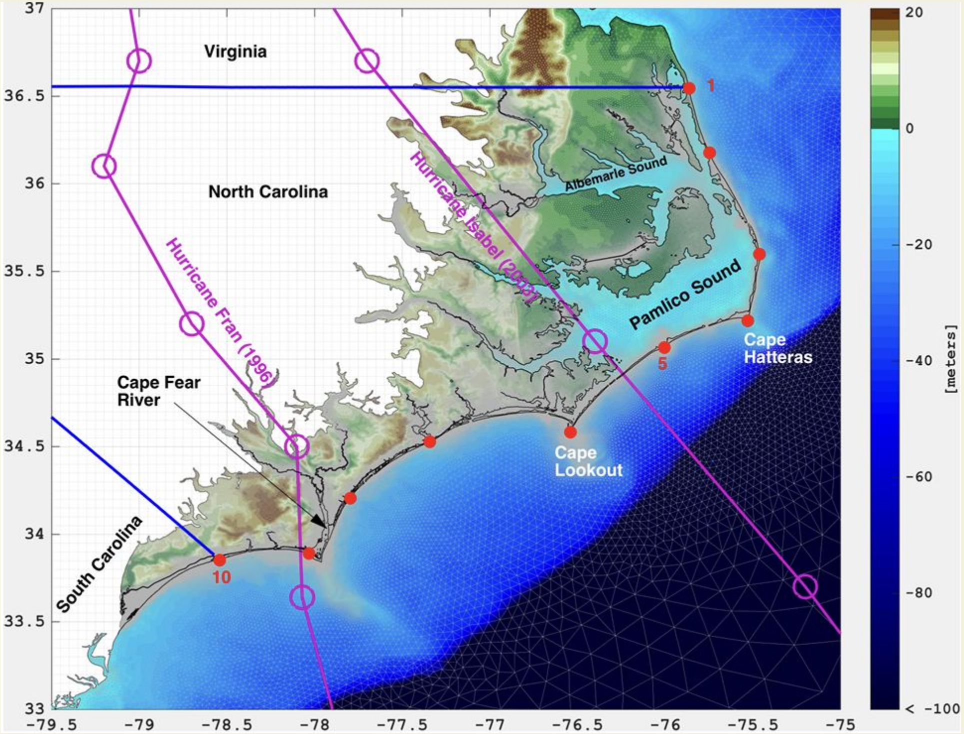

ADCIRC mesh of NC. Ten output locations as red points. The paths of Hurricane Fran and Isabel are shown. (Bezuglov et al., 2016, Figure 1)

NN Models for Coastal Flooding

1. Early Models: point-based surge predictions

2. **Later Models: hybrid models, hundred/thousand point surge predictions**

3. Storm Surge predictions of coarse spatial maps

4. First storm-tide Surrogate model: point-based

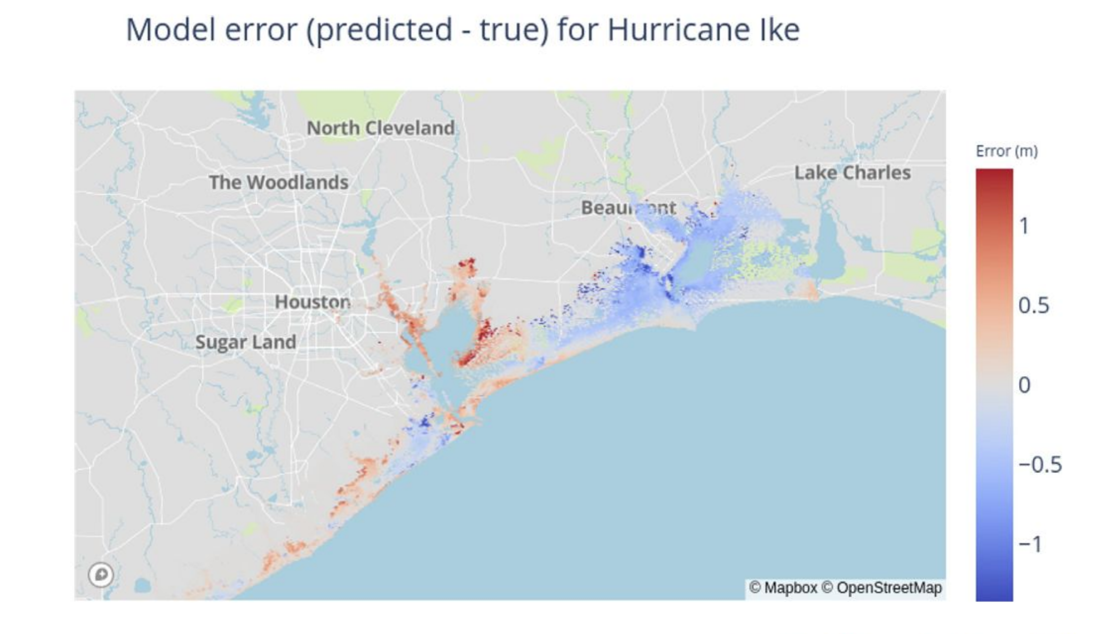

Fig. 10. Model error for Ike. Map data courtesy of OpenStreetMap (OpenStreetMap contributors, 2022).

Model Error of the Neural Network model in TX for Hurricane Ike [Pachev et al. 2023]

NN Models for Coastal Flooding

1. Early Models: point-based surge predictions
2. Later Models: hybrid models, hundred/thousand point surge predictions
3. **Storm Surge predictions of coarse spatial maps**
4. First storm-tide Surrogate model: point-based

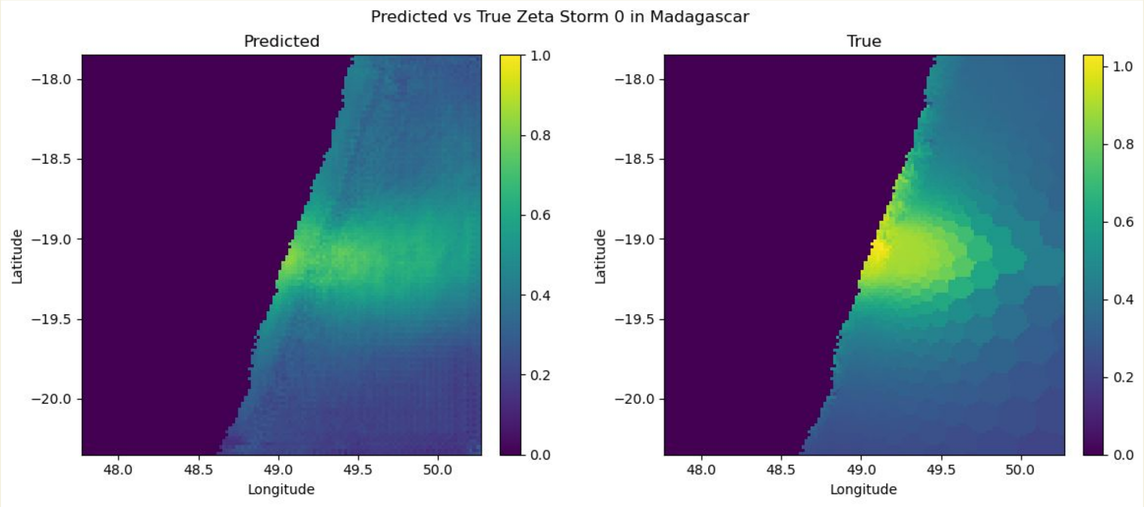

Model Error of the Global Neural Network model in Madagascar [Pachev et al. 2026]

Storm Tide is the combined effect of storm surge along with tide.

NN Models for Coastal Flooding

1. Early Models: point-based surge predictions
2. Later Models: hybrid models, hundred/thousand point surge predictions
3. Storm Surge predictions of coarse spatial maps
4. **First storm-tide Surrogate model: point-based**

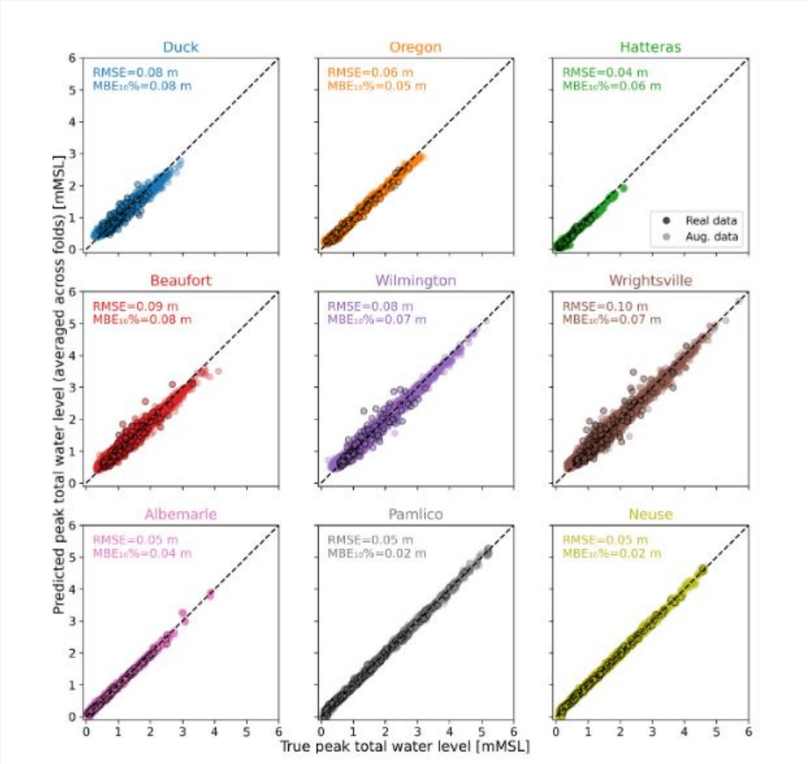

Model Error of storm tide at 9 NC locations [Cuevas López 2025]

Although these studies show that neural networks can generalize coastal processes, currently, none of these models are using NN/ML to predict **storm tide maps.**

How can neural network models learn the relationships between tides, coastal storms, and coastal landscapes to predict maps of storm tide flooding?

Objective: A neural network model that produces high resolution storm tide maps.

Model Overview

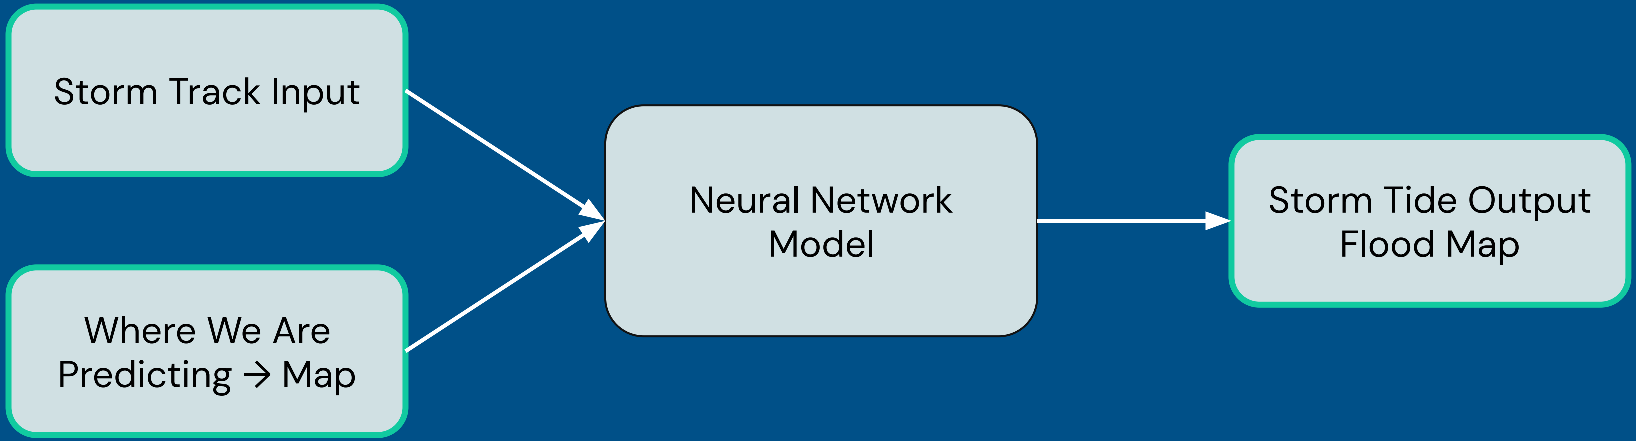

Model Basis:
(Cuevas López 2025)

1,813 synthetic Atlantic storms selected (From STORM, a statistical extension of historical storms)

Storms were chosen that specifically impacted NC

Each storm run in ADCIRC with simulated tides + atmospheric forcing

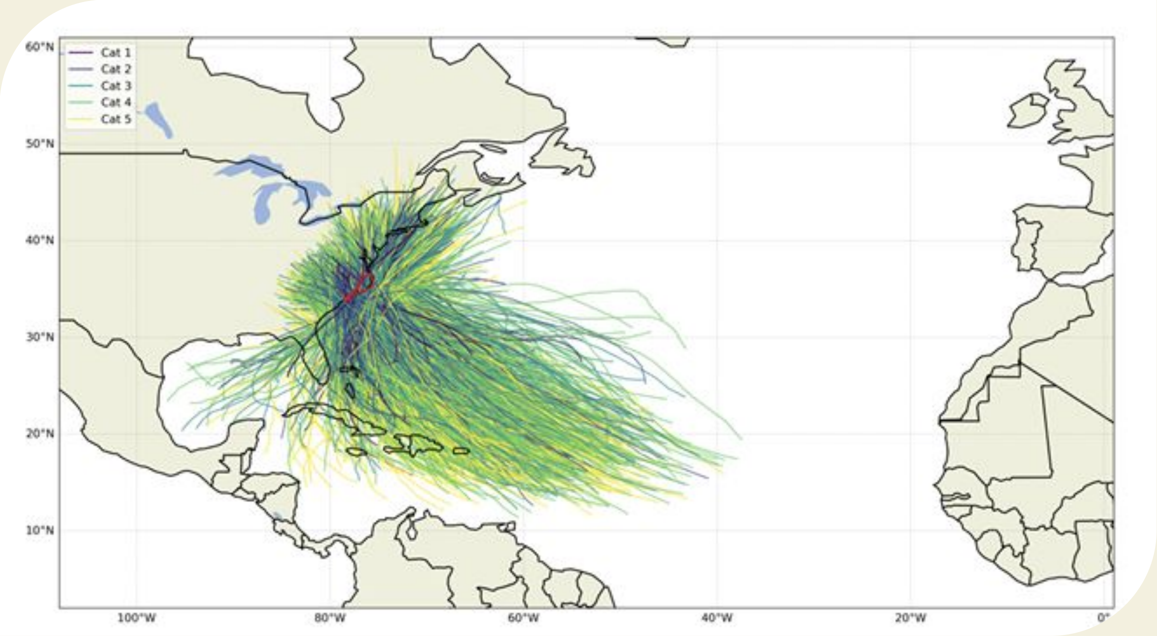

Selection of Storms that impact NC [Cuevas López et al. 2025]

Selected Track Inputs

This is a time series of the most important 25 hours.

- Storm longitude and latitude
- Wind Speed
- Pressure
- Radius to Max Wind Speed
- Heading Direction
- Forward Speed, u, v
- Tidal locations (Duck, Hatteras, Beaufort, Albemarle, Pamlico, Neuse, Oregon Inlet)
- Storm dist u,v to tidal location
- Distance to image center
- Distance u, v to image center

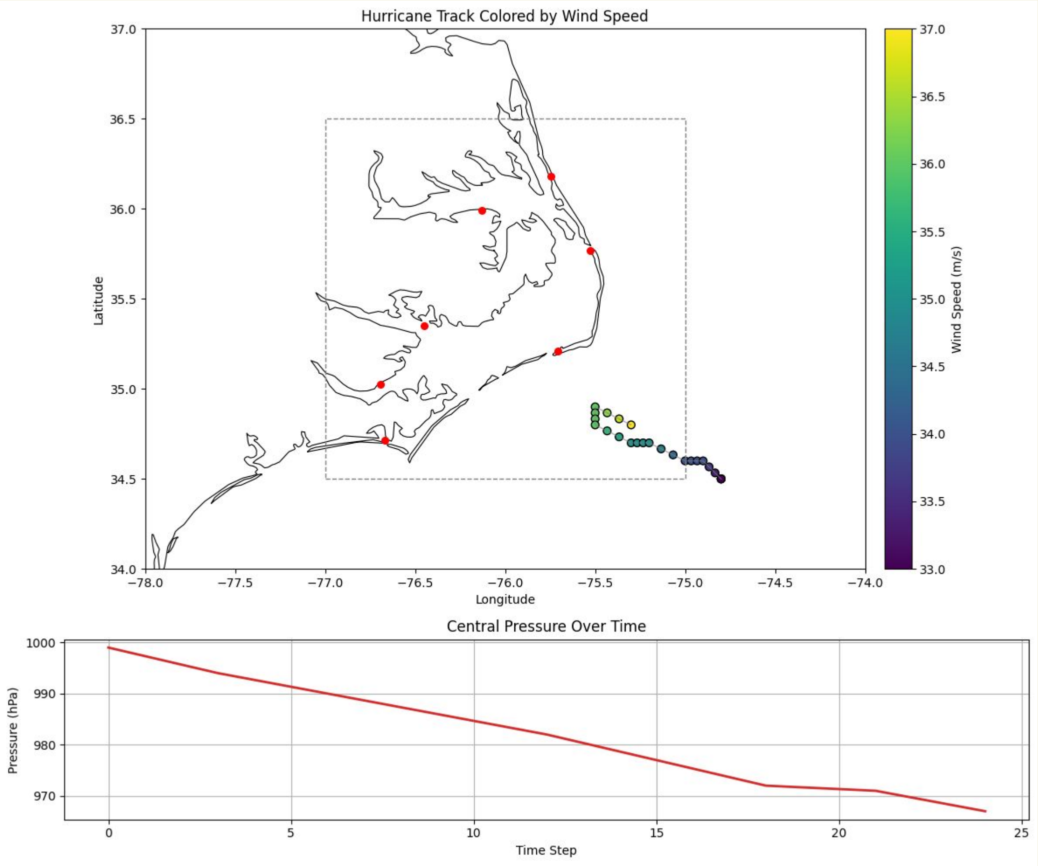

Land Input

Represents the land and coastal morphology of the prediction region

Provides the model with a reference frame for where flooding can occur

Gives the model the initial conditions of the landscape

This helps the model focus on the storm flooding and not the dry land elevations

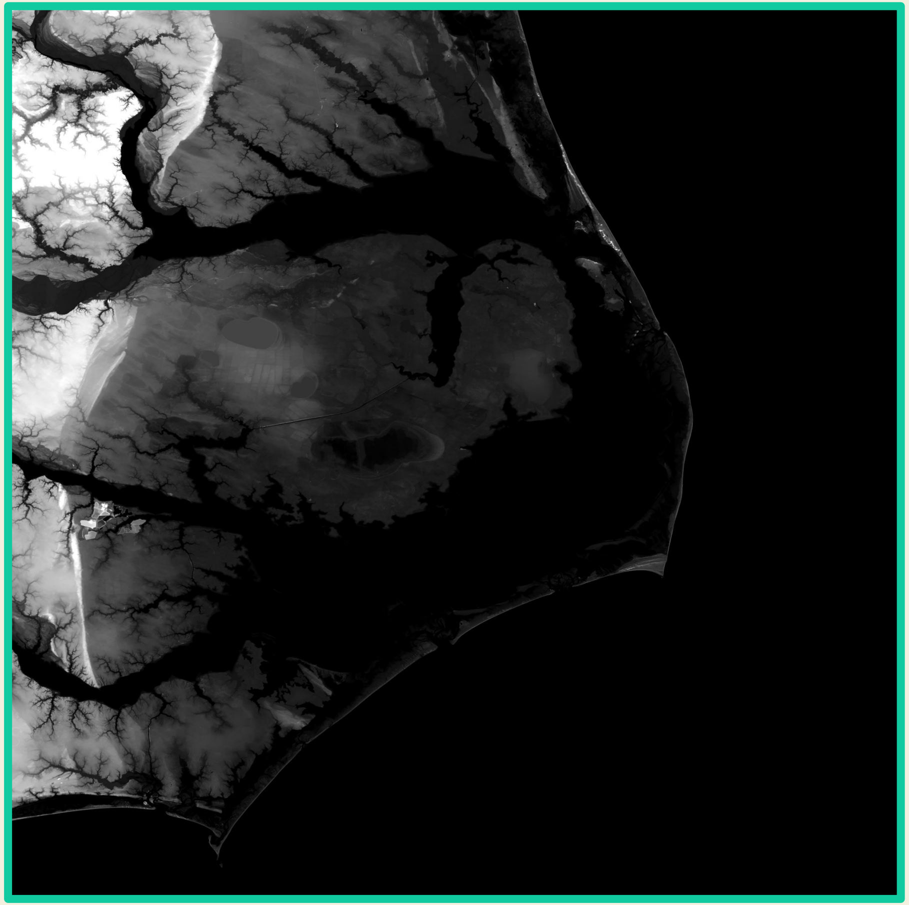

DEM of study area in grayscale (scaled from meters at 0.1 meter bins)

ADCIRC Outputs → Raster Surface

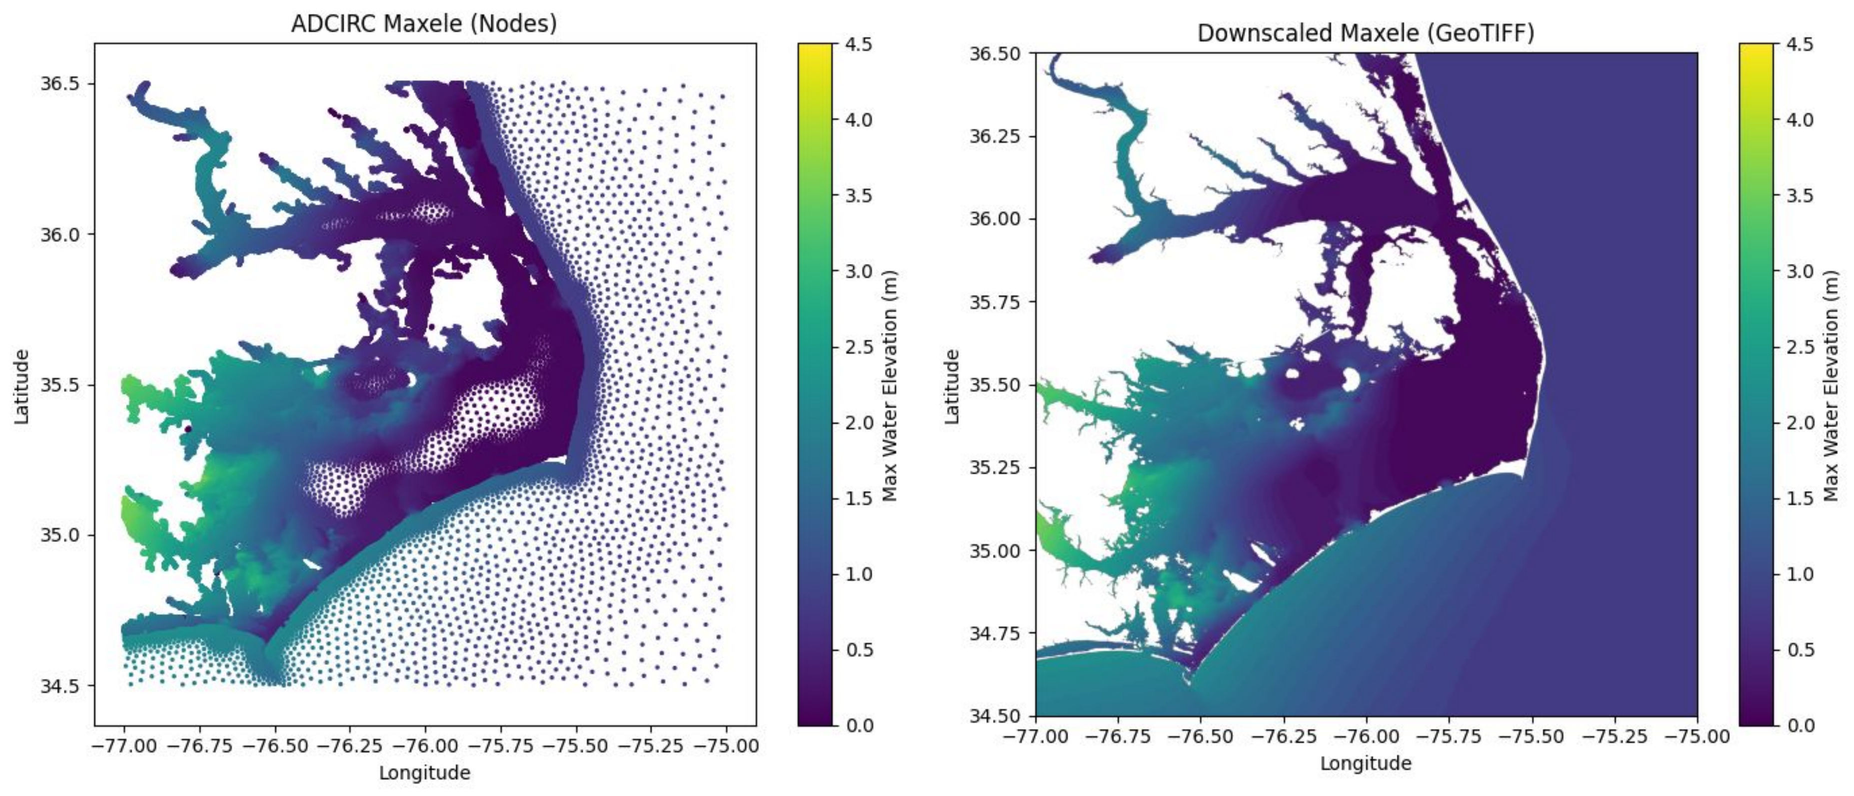

Raster Surface → Grayscale Image

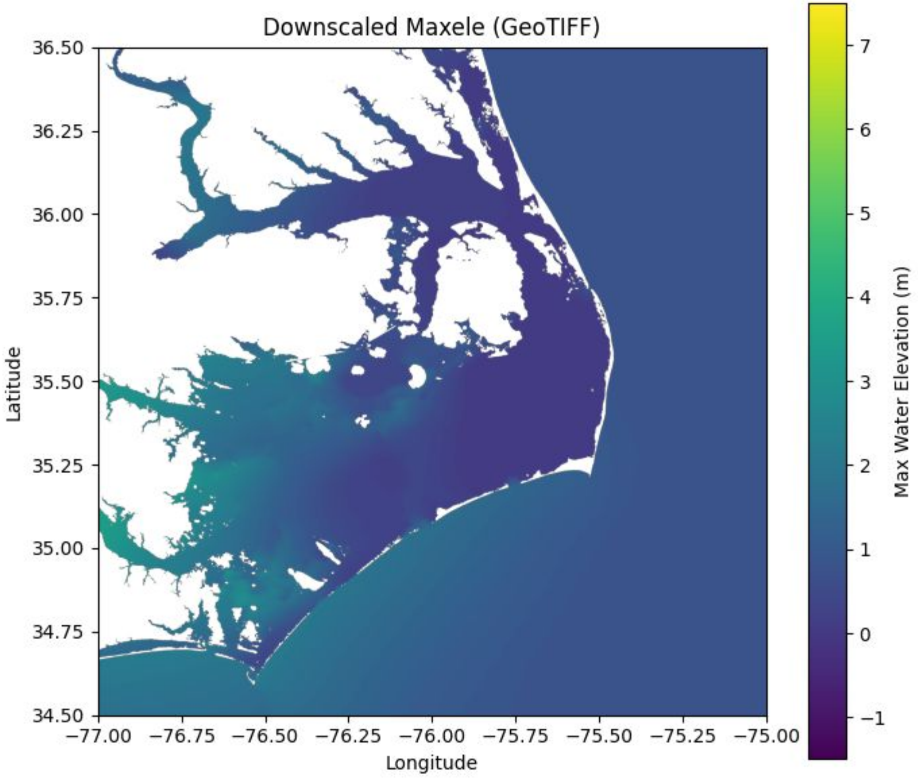

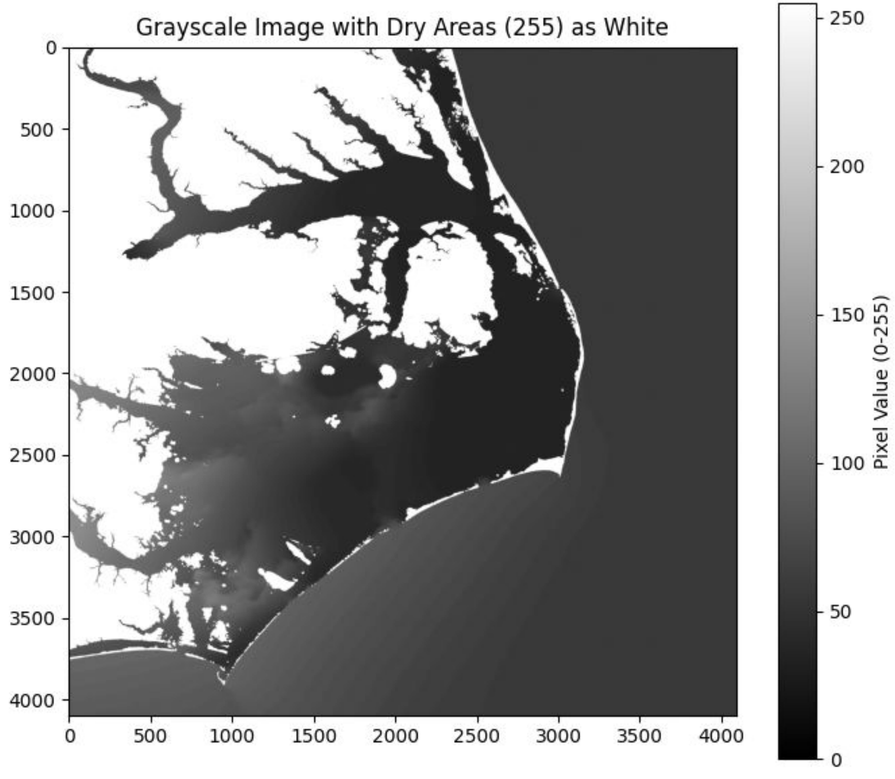

Grayscale Image → Continuous Surface

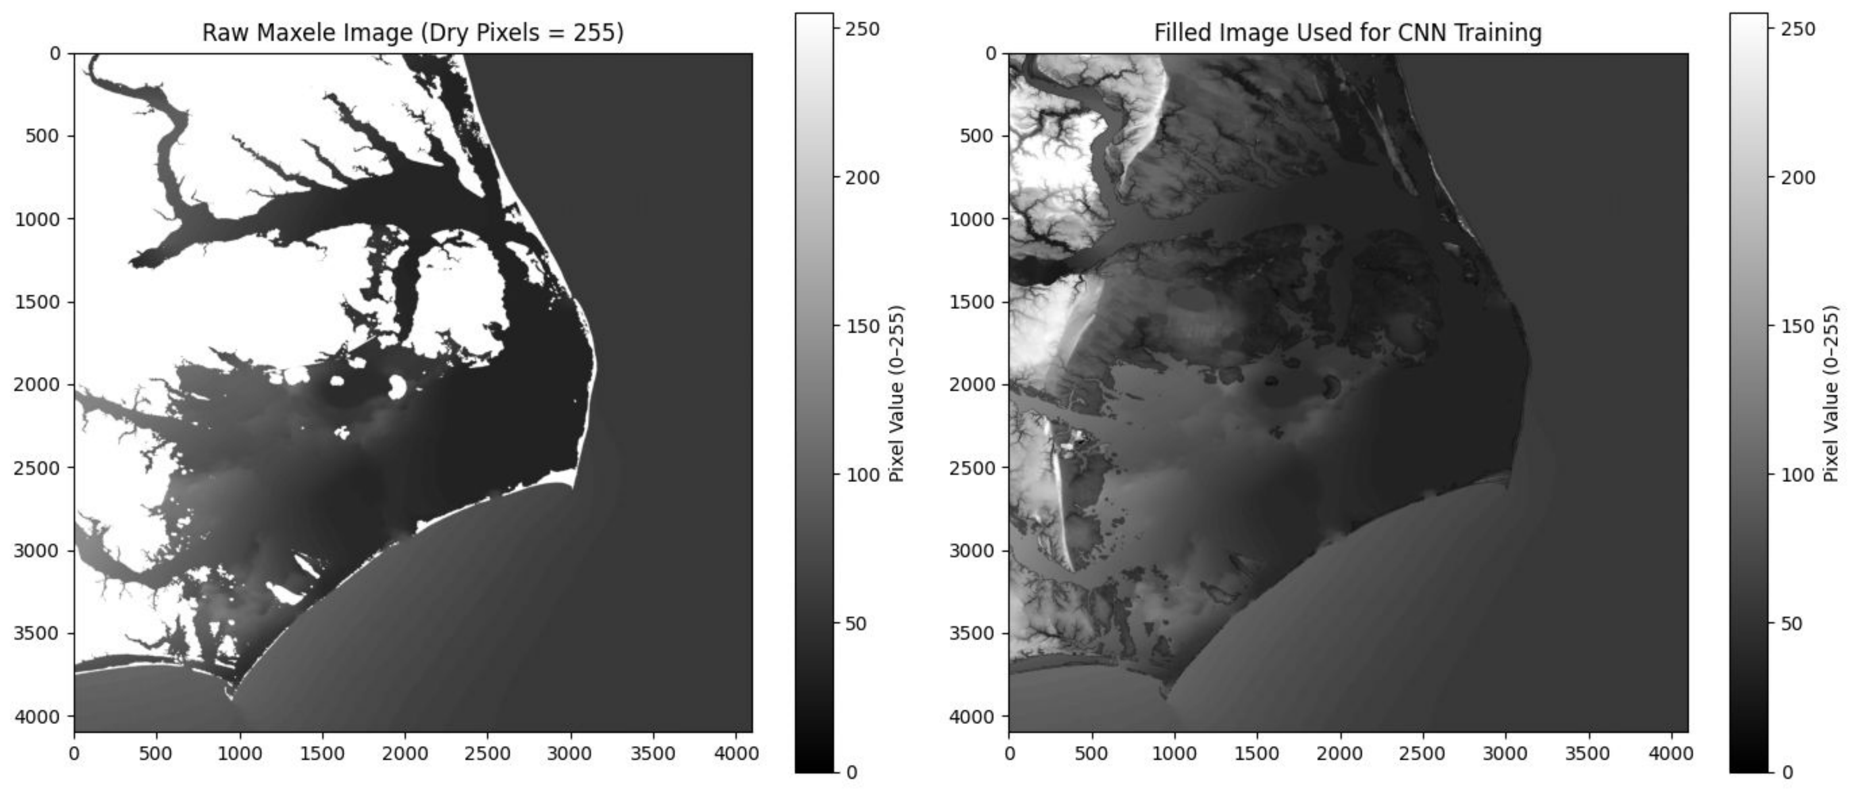

Data Augmentation

We divide the full 2°×2° domain into 64 tiles

Each tile corresponds to a 0.25° × 0.25° domain at 512 x 512 pixel resolution, or ~50 m x 50 m pixel resolution

Ensures full spatial coverage of the study area

Expands the dataset from 1,813 storms → 116,032 storm-tide tiles

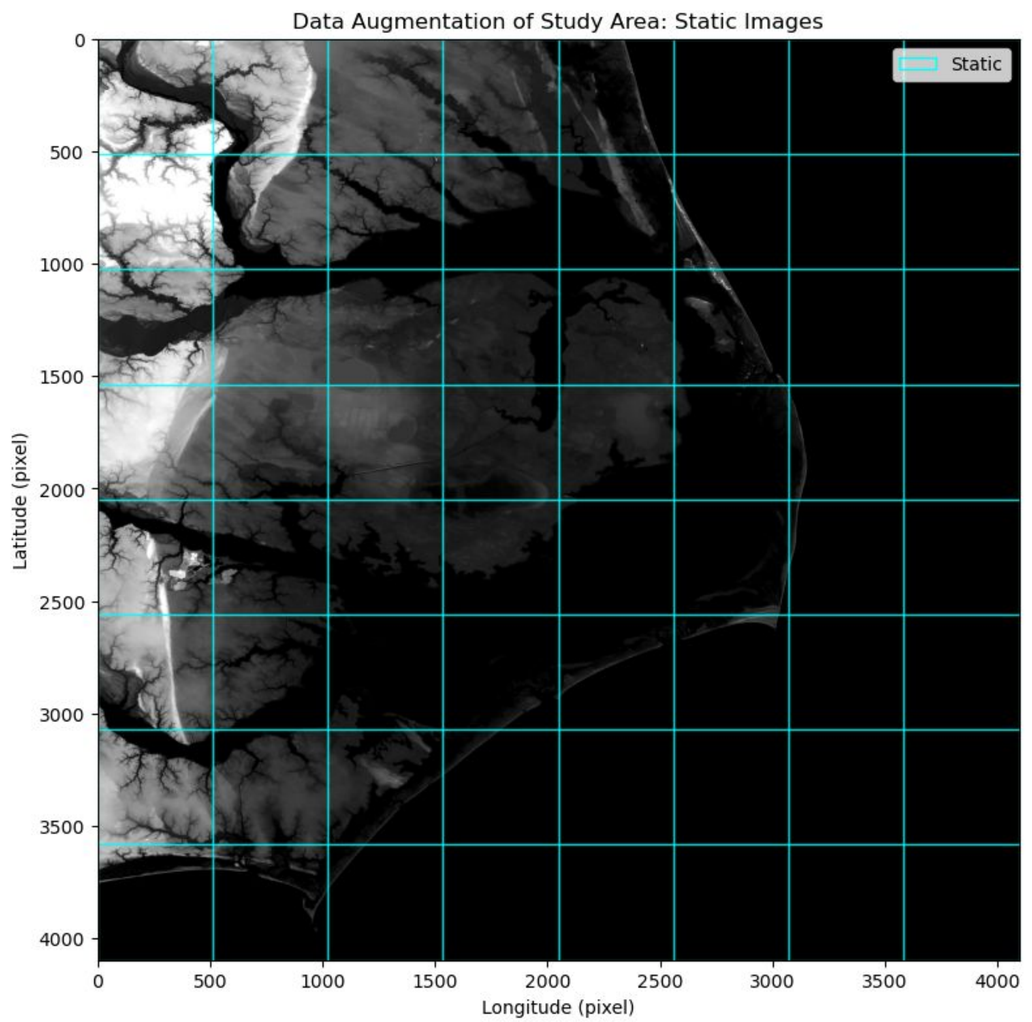

Model Library Summary

- Inputs: Storm intensity + tidal time series + geographic information of the storm to the prediction location.
- Inputs: land map for geospatial reference
- Outputs: Storm tide maps of max flooding for each storm for 0.25° longitude and latitude extents.

- **But how would a model relate these storm tracks to these flood images?**

Model Overview

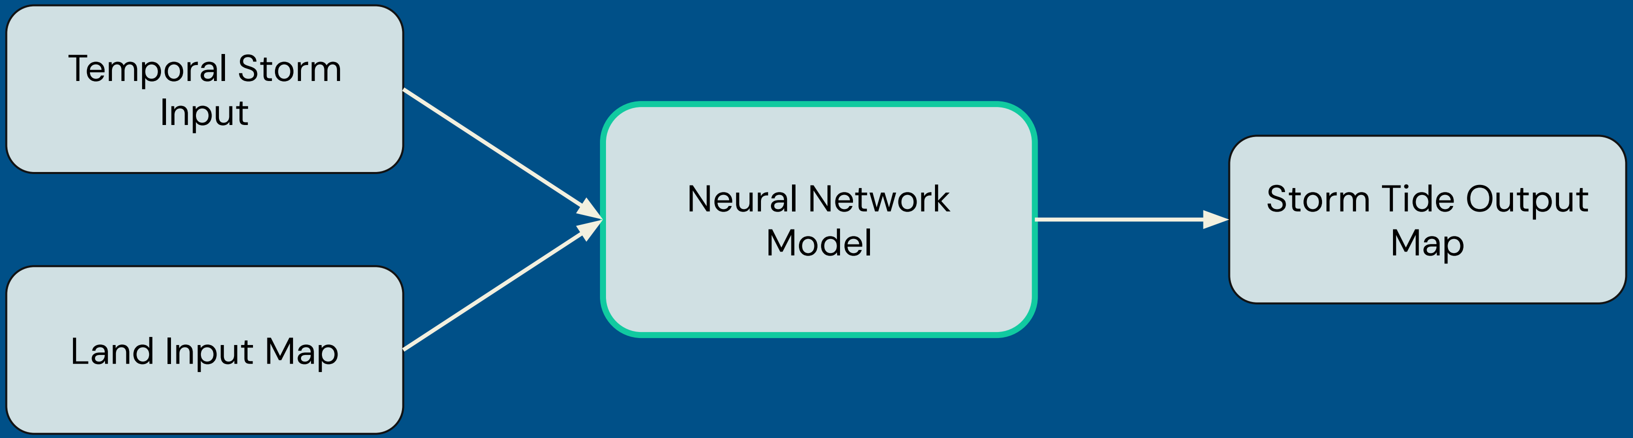

Spatial Input Stream:
helps the model focus on the dry land area and focus on the flooding rather than the initial land elevations.

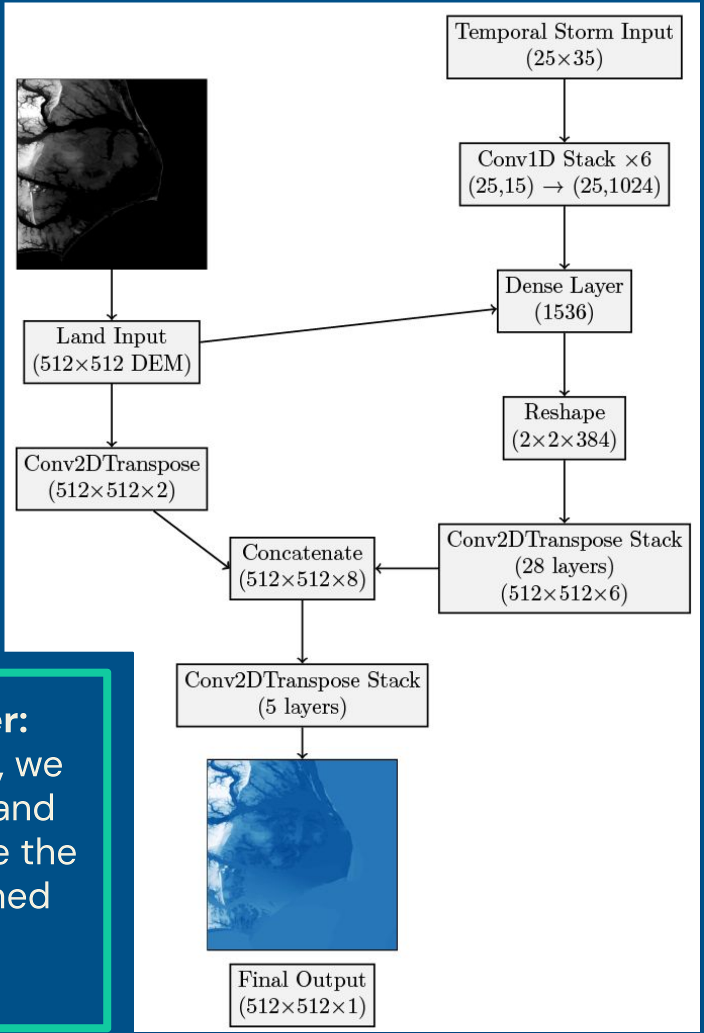

Temporal Input Stream:
helps the model learn how the storm evolves over time and how those changes influence flooding.

Spatial Encoder:
At high resolution, we reintroduce the land elevation to ensure the model stays aligned with the true geography.

Model Results

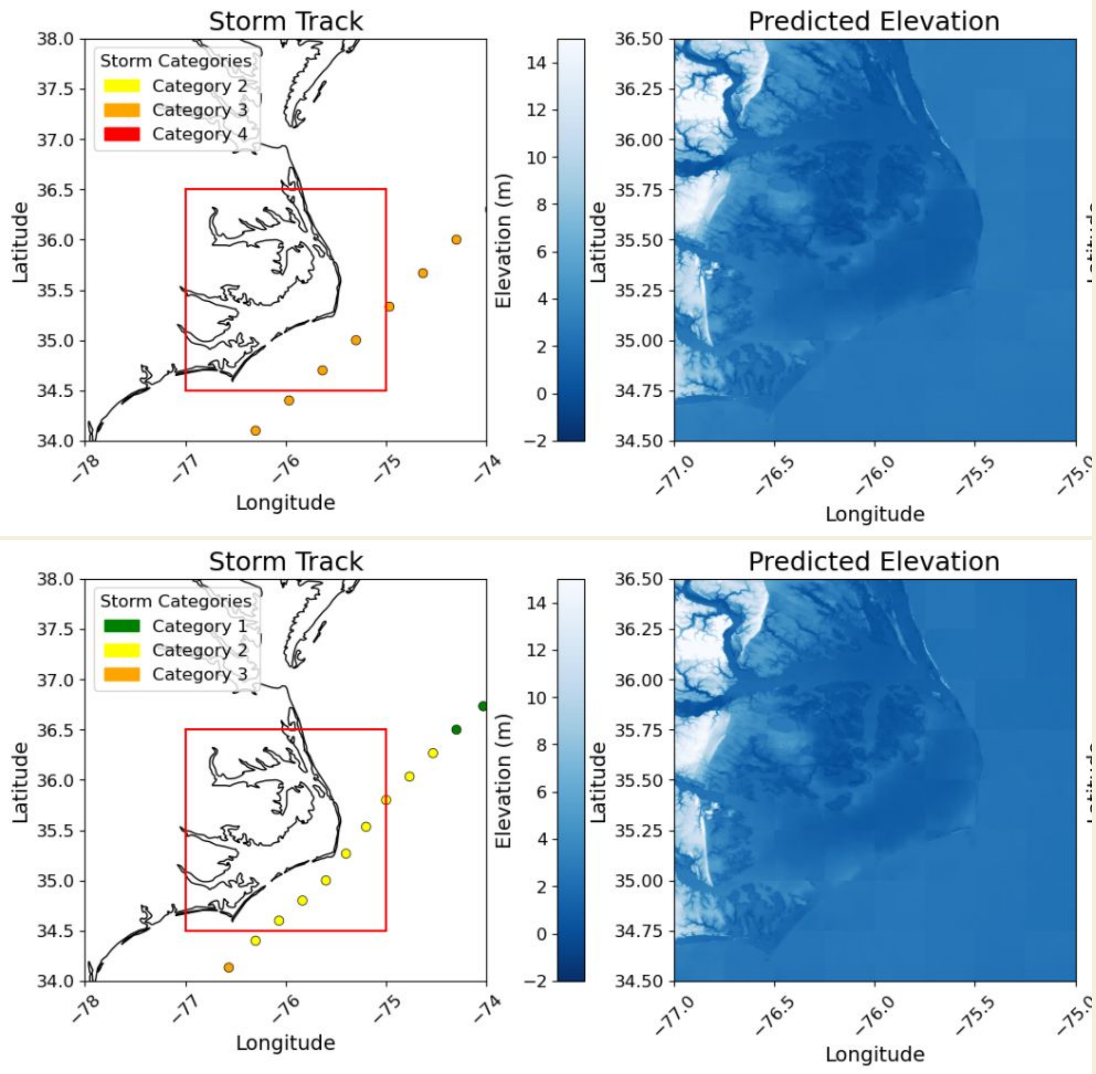

Takeaways

This model produces a spatially continuous map!

It can produce this map in 100 milliseconds

The spatial resolution of each pixel in this map (4096x4096 pixels) is about 50 m x 50 m per pixel

We are producing a high resolution storm tide map over NC in a tenth of a second!

Overall Model Performance

All Test Results

Model RMSE: 0.2722 m

0–0.5 m RMSE = 0.3661 m

1–3 m RMSE = 0.3970 m

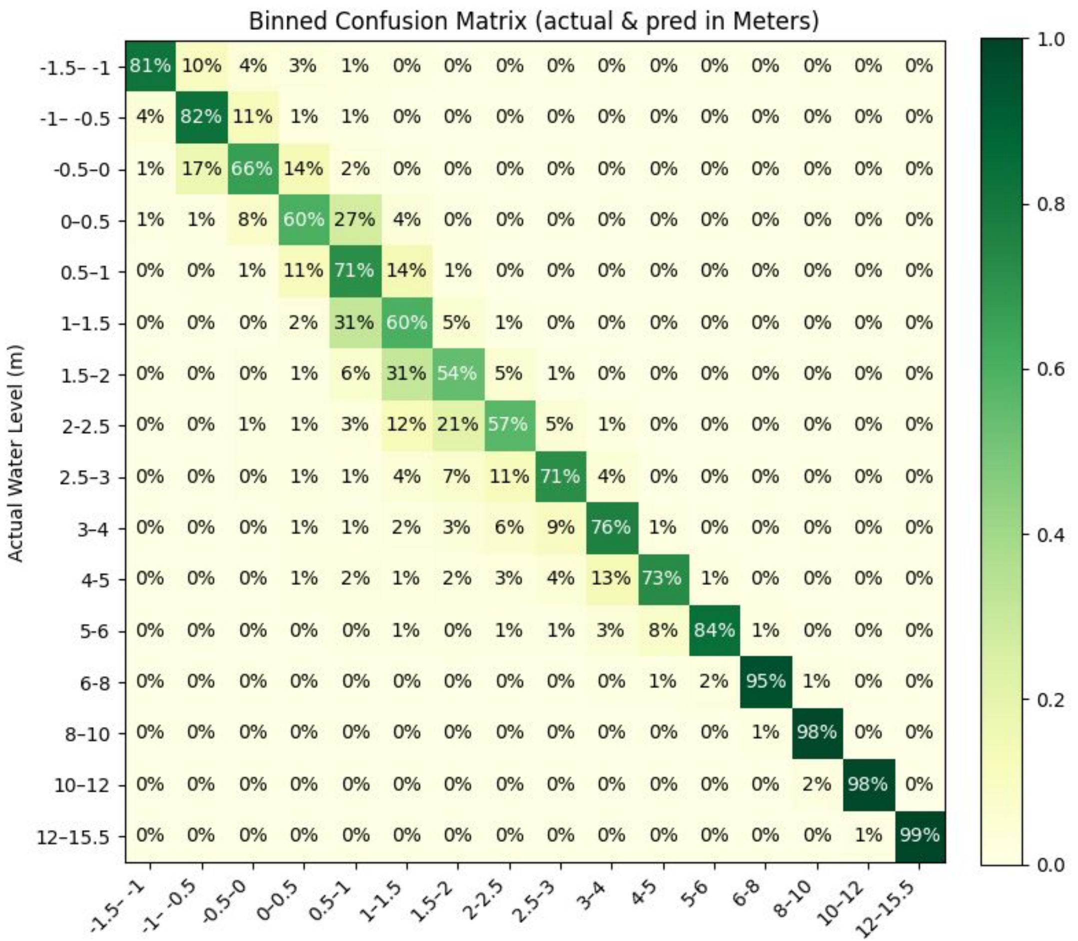

Key Takeaway: This model has high accuracy while learning the spatial structure of storm-tide flooding.

Domain Prediction Results

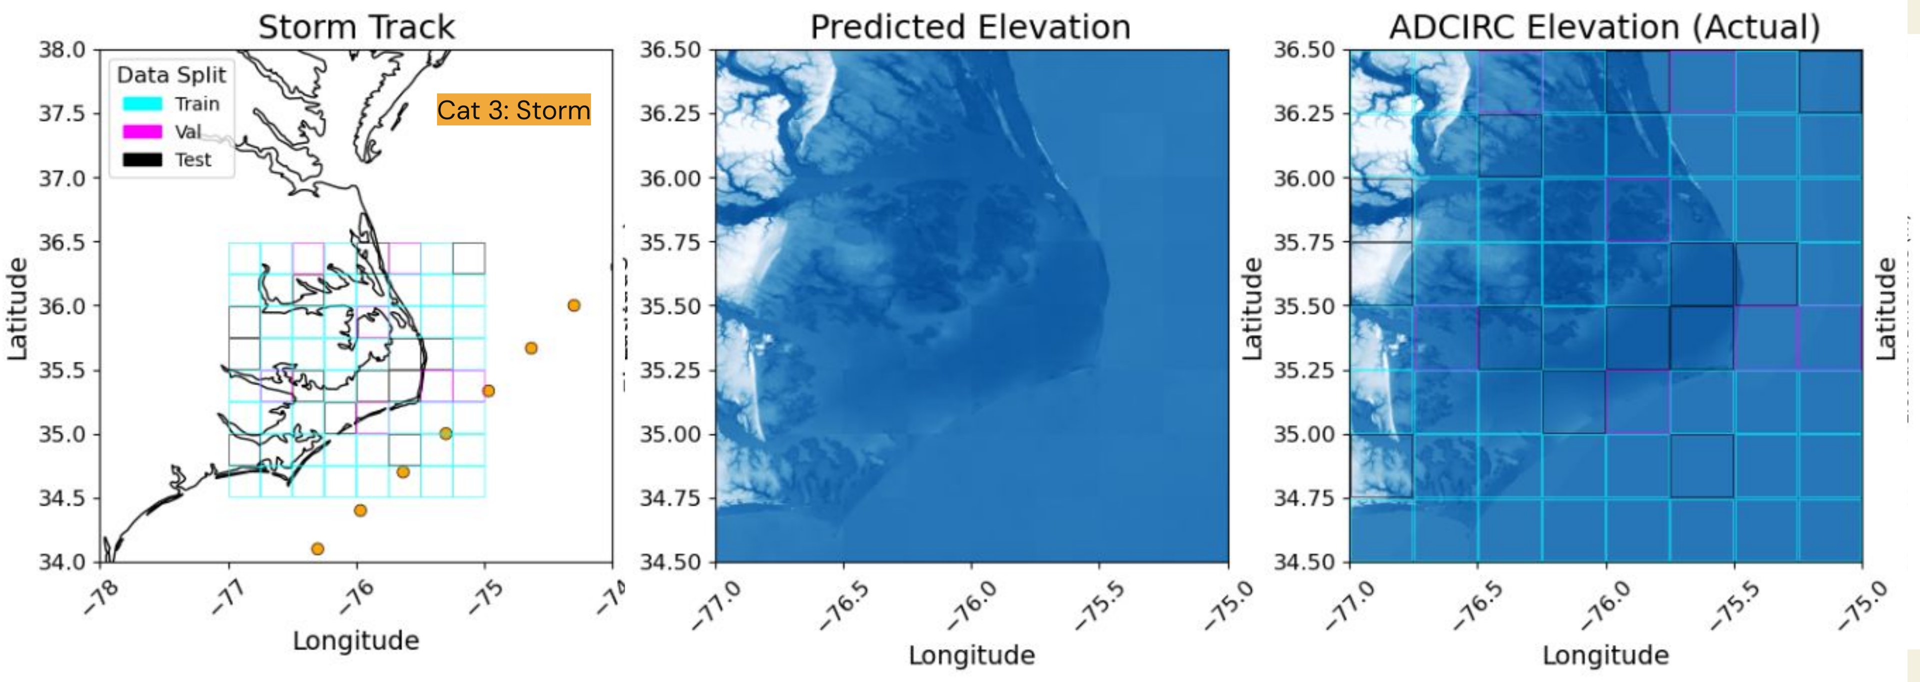

Key Takeaway: The data augmentation technique for physics aware strategy for storm-tide map tiles expands the effective training dataset while preserving the physical integrity. The model is learning how storms impact tiles.

Spatial Performance

Spatial Error Map

Most of the domain lies within 0.2 meters of the zero bias, indicating small errors and spatial consistency across the study area.

Under-estimation (–0.05 to –0.2 m) appears where steep storm-tide gradients make predictions highly sensitive to small elevation errors.

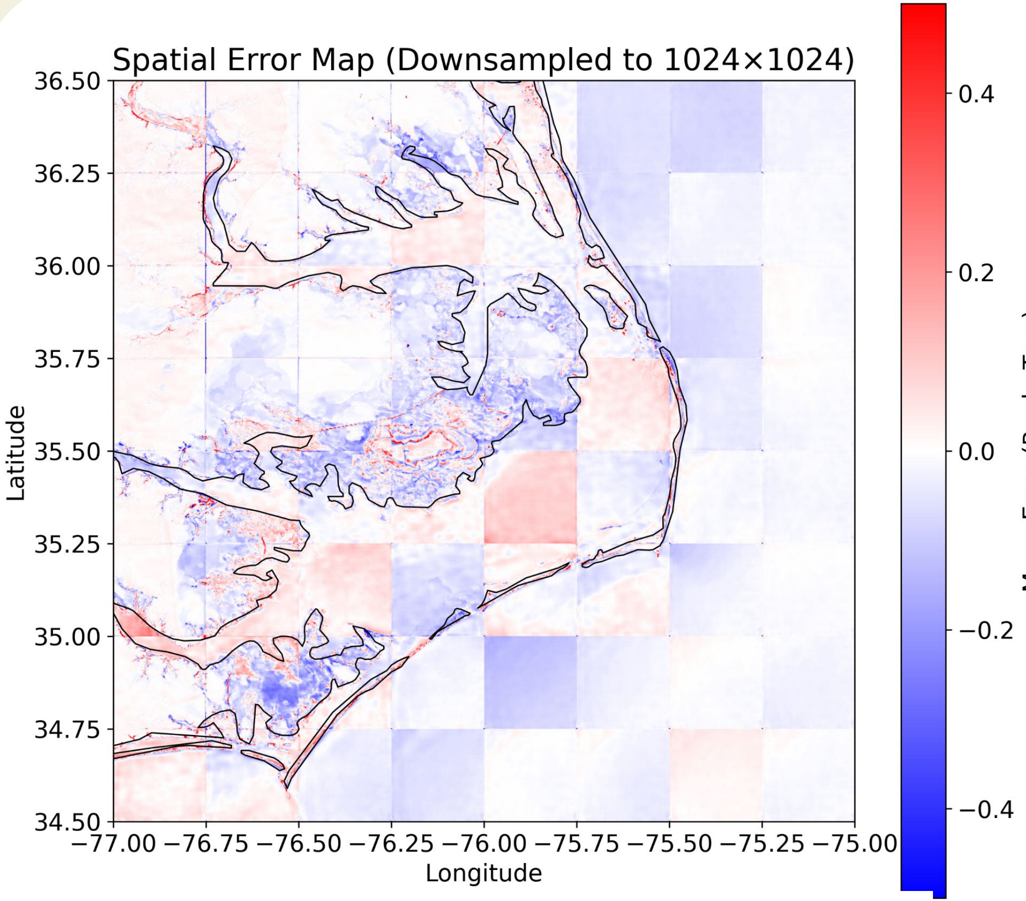

Key Takeaway: The model produces high resolution, spatially continuous storm-tide maps that describe behavior at local scales.

Comparison to Other Models

Comparison to Process-Based and AI Models

Process-Based

SLOSH (Forbes et al. 2010)

- RMSE typically 0.5-1.0 m

ADCIRC (Pringle et al. 2023)

- RMSE = 0.15-0.35 m

AI Models

Early NN surge models (Bezuglov, Hashemi)

- RMSE = 0.2-0.6 m

Hybrid Surrogates (Naeini, Pachev 2023)

- RMSE = 0.26-0.33

Storm Tide NN (Cuevas Lopez 2025)

- RMSE = 0.08-0.19m

My Model

- **RMSE=0.2722 m**

Key Takeaway: This model is highly accurate relative to process-based and AI based alternatives.

Conclusions

1. This model has high accuracy while learning the spatial structure of storm-tide flooding.
2. The data augmentation technique for physics aware strategy for storm-tide map tiles expands the effective training dataset while preserving the physical integrity.
3. The model produces high resolution, spatially continuous storm-tide maps that describe behavior at local scales.
4. This model is highly accurate relative to process-based and AI based alternatives.
5. Machine-learning surrogates can preserve physical realism when trained on high-fidelity models.

Ongoing work: Addition of MDA tiles at frequently wet areas, and shuffling based on storms not randomly.

Thank you!
Questions?

Email: molly_mckenna11@triad.rr.com

Currently seeking full-time roles in coastal modeling

LinkedIn:

Link to Thesis:

1D CNNs

1D CNNs process the temporal storm-track sequence

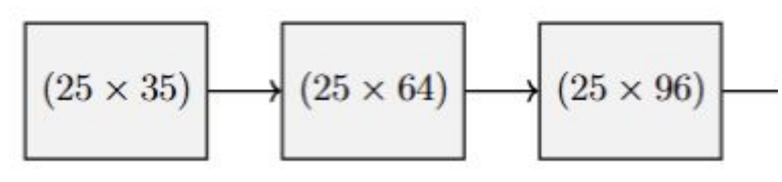

while extracting meaningful patterns

Upsampling layers expand the temporal features for fusion with spatial inputs

Helps the model learn how storm evolution shapes the final flood map.

The upsampling builds into a 2x2 seed image for the next stage of the model.

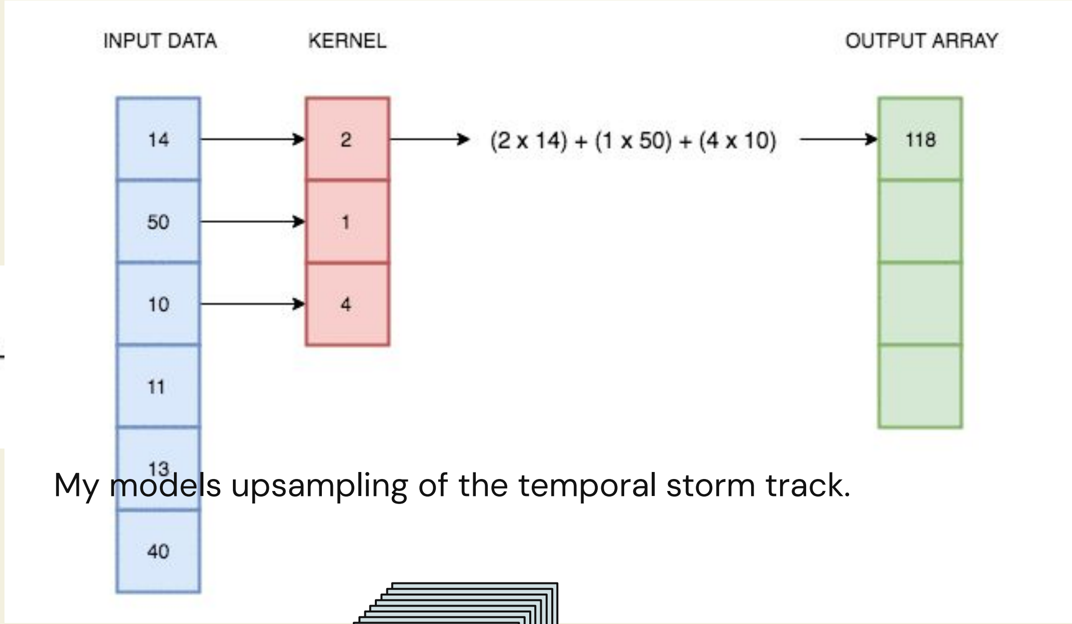

My models upsampling of the temporal storm track.

1D Convolut eep Learning Overview)

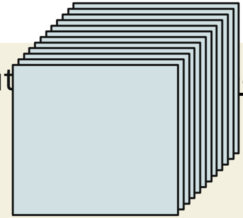

Example of the stack of seed images, the input to the next stage is 384, 2x2 images.

2D Convolutional Layers

Conv2DTranspose layers progressively upsample the 2×2×384 latent into the full 512×512 image

Each upsampling step increases spatial resolution and adds spatial complexity

Early layers learn coarse, low-resolution flood patterns; later layers refine fine-scale structure

Land elevation is re-introduced at high resolution to anchor the model to the correct geography

This guides the model to focus on the specific tile being predicted

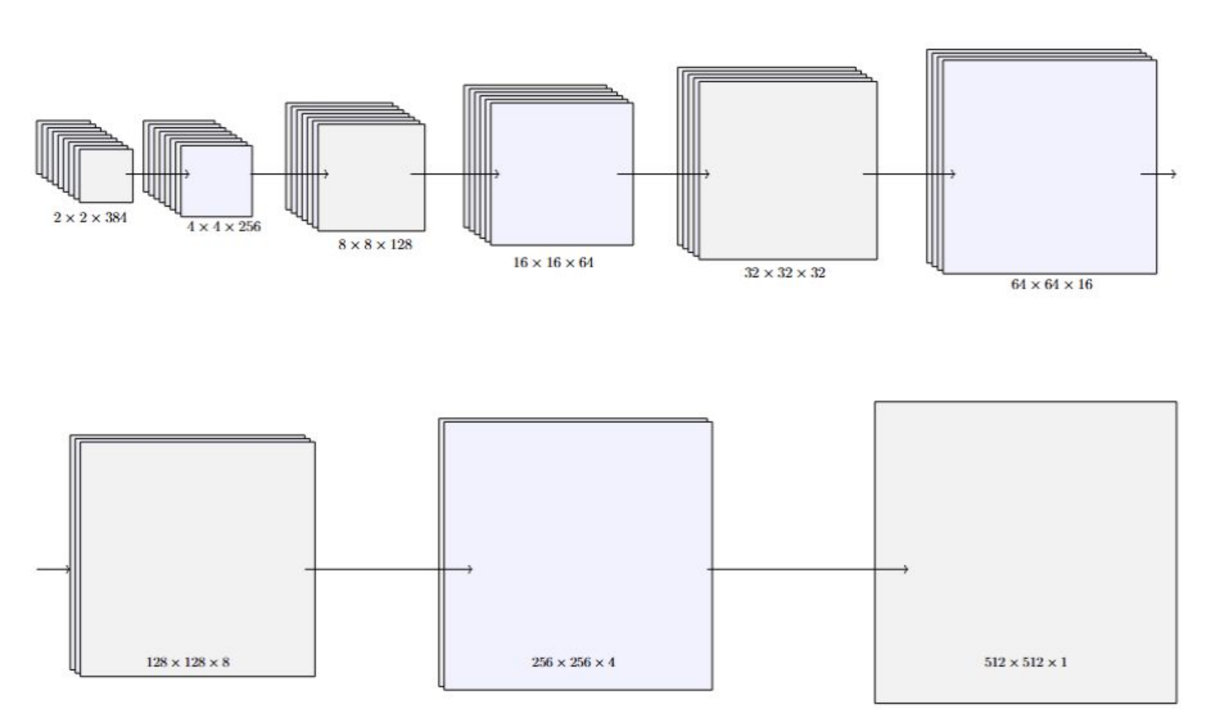

Training Loss in Stages

3 stages with a change in Huber loss to help the model perform more like Mean Absolute Error

Huber Loss = loss function combining Mean Squared Error & Mean Absolute Error

The model trained for 478 epochs : this took 370 hours

The reduction in δ makes the model more robust to outliers

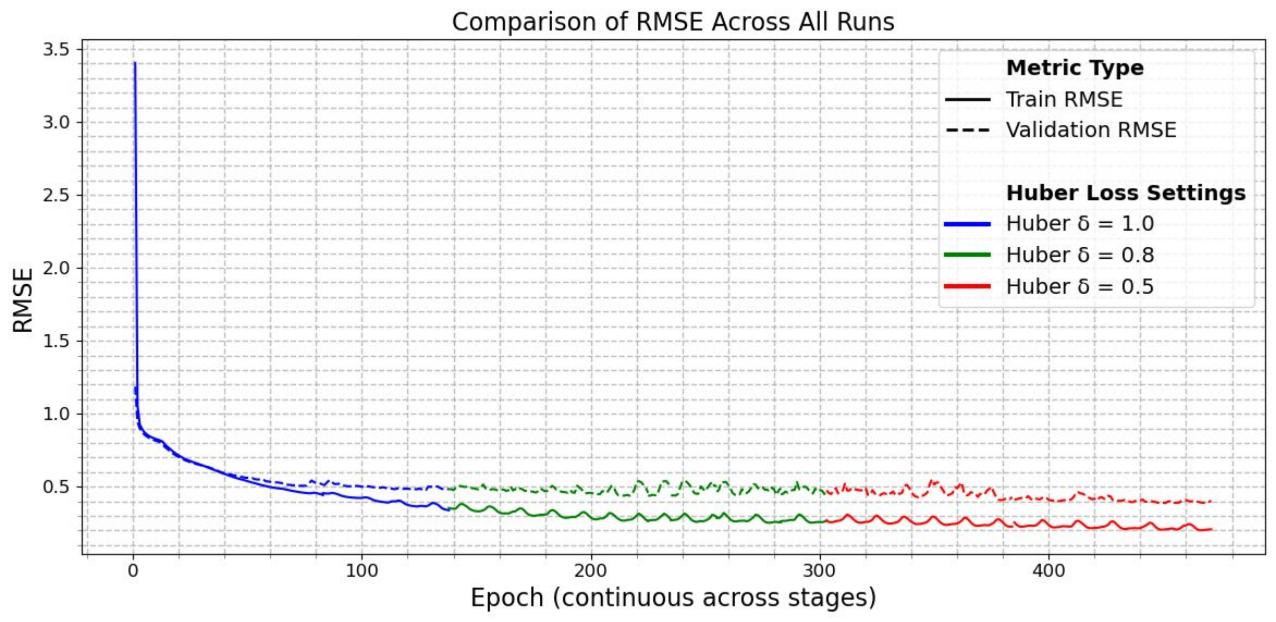

MDA Model

Wetness Variability

To determine the areas that are important for the model we computed wetness variability

The areas of high wetness variability are where the areas change the most from wet to dry

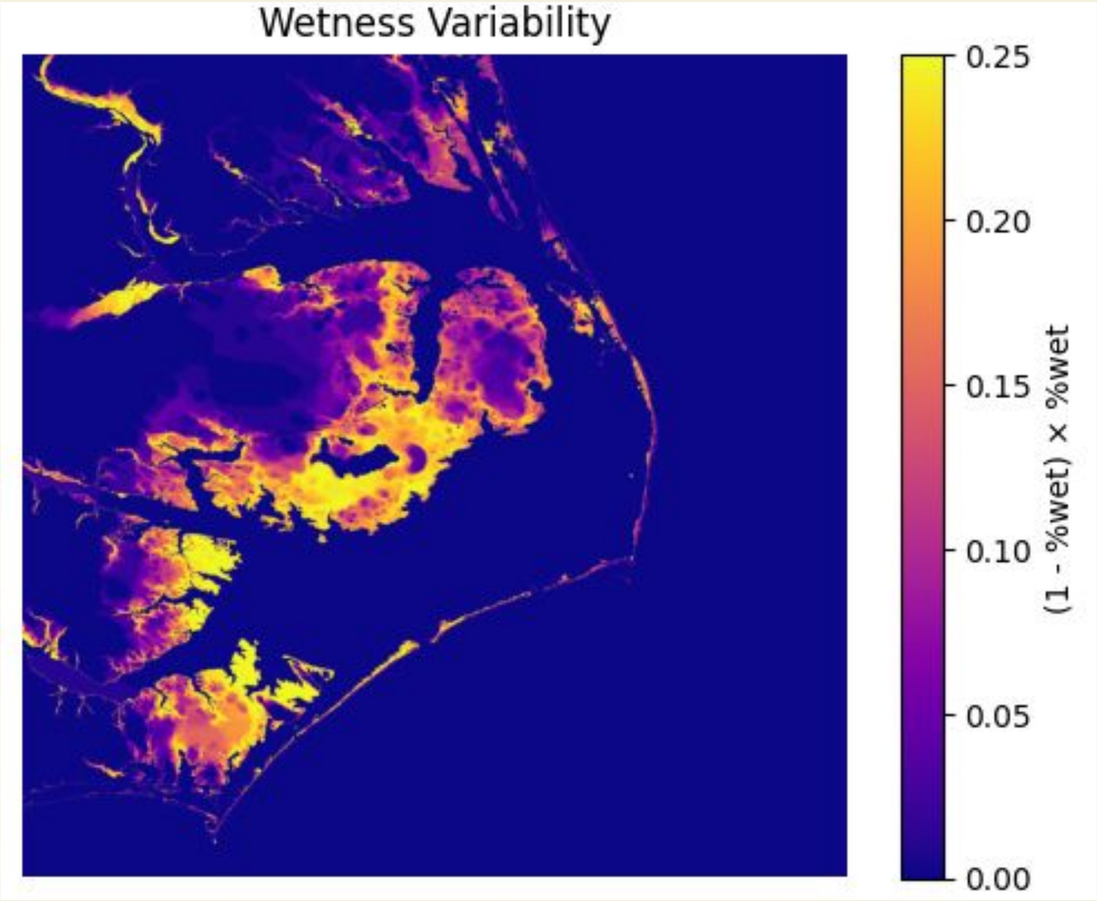

Maximum Dissimilarity Algorithm (MDA)

To select the most important a statistically significant points for the model to learn based on the wetness variability we implemented a MDA algorithm that selected image centerpoints to augment the dataset.

MDA Centerpoint Selection

Centerpoints for each storm were selected to be statistically relevant to the wetness variability region along with being distant from other points.

Each mda center point was checked for distance between other mda centerpoints → closest distance they could be was 250 pixels. (within one storm)

This increased the 1813 storm set into 152,292 storm maps

64 static points + 20 MDA points

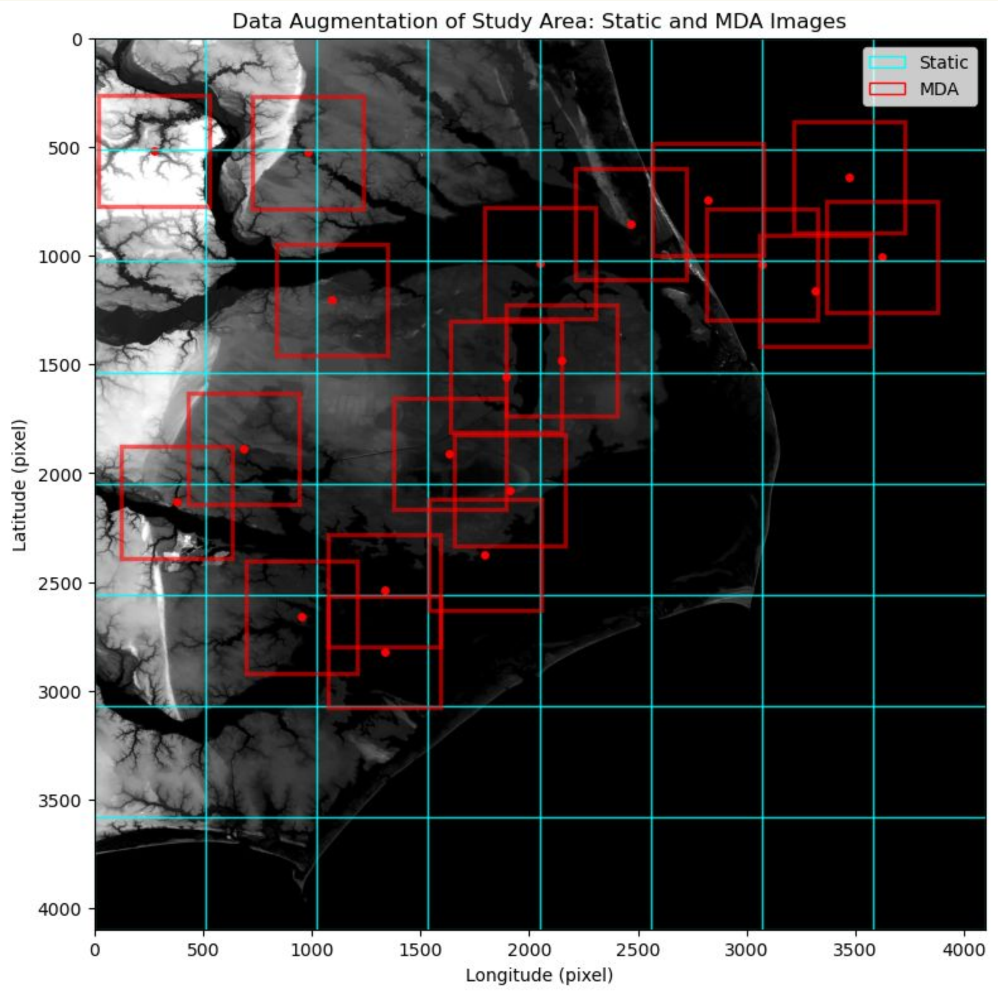

MDA Shuffled Storms

Shuffling Strategy

Two shuffling strategies:

- Fully randomized shuffling
- Shuffling storms and then shuffling within test, val, train

Storm shuffling gives the model the ability to see the entire storm in training or validation dataset.

This would also check the models ability to generalize on storms it has not seen before in the Test dataset

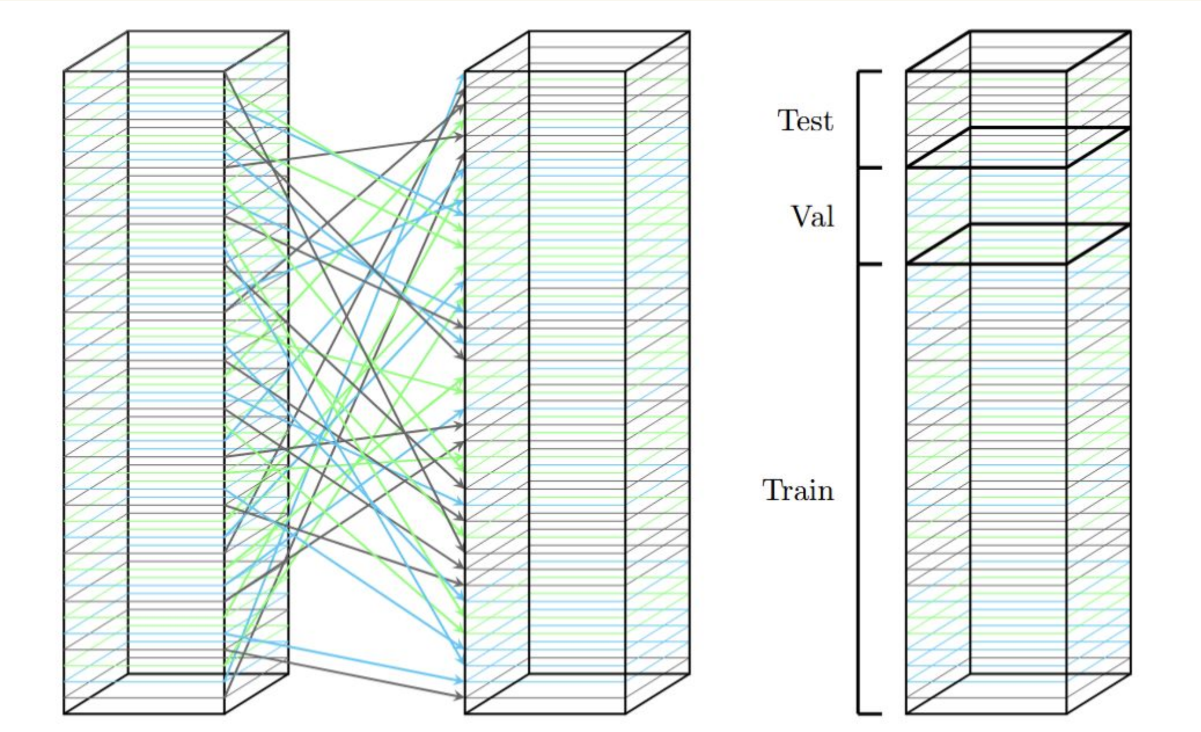

Model Performance

This mda shuffled storm model trained for 392 total epochs

RMSE performance was around 0.445 m

Similar to the other MDA model it saw higher RMSE per elevation bins.

Additionally the confusion matrix shows under predictions in the 0–2 m bins