



Climate Adaptation-Aware Flood Prediction for Coastal Cities Using Deep Learning

Bilal Hassan¹, Areg Karapetyan¹, Aaron Chung Hin Chow¹, and Samer Madanat¹

¹Division of Engineering, New York University Abu Dhabi, Abu Dhabi, United Arab Emirates

Correspondence: Bilal Hassan (bilal.hassan@nyu.edu)

Abstract. Climate change and sea-level rise (SLR) pose escalating threats to coastal cities, intensifying the need for efficient and accurate methods to predict potential flood hazards. Traditional physics-based hydrodynamic simulators, although precise, are computationally prohibitive and impractical for city-scale coastal planning applications. Deep Learning (DL) techniques offer promising alternatives, however, they are often constrained by challenges such as data scarcity and high-dimensional output requirements. Leveraging a recently proposed vision-based, low-resource DL framework, we develop a novel, lightweight Convolutional Neural Network (CNN)-based model designed to predict coastal flooding under variable SLR projections and shoreline adaptation scenarios. Furthermore, we demonstrate the ability of the model to generalize across diverse geographical contexts by utilizing datasets from two distinct regions: Abu Dhabi and San Francisco. Our findings demonstrate that the proposed model significantly outperforms state-of-the-art methods, reducing the mean absolute error (MAE) in predicted flood depth maps on average by nearly 20%. These results highlight the potential of our approach to serve as a scalable and practical tool for coastal flood management, empowering decision-makers to develop effective mitigation strategies in response to the growing impacts of climate change. **Project Page:** <https://caspiannet.github.io/>

1 Introduction

Of the world's 34 megacities (i.e., those with more than 10 million inhabitants), approximately 70%¹ are situated on or near the coast. Coastal cities, including megacities, host nearly 10% of the world's population, are 2.6 times denser populated as compared to inland areas, and are powerhouses of global trade and business activities (the legacy of maritime trade) (Pal et al., 2023). Yet, these cities, especially those located in low-elevation coastal zones (i.e., hydrologically connected areas along the shore below 10 m of elevation (van de Wal et al., 2024)), such as Venice, Italy; Hamburg, Germany (van de Wal et al., 2024); and in the United States: Annapolis, MD, Miami, FL; Houston, TX; San Diego, CA; as well as Manila, Philippines; and Dakar, Senegal (Griggs and Reguero, 2021), are hotspots of water-related disasters and climate change-induced risks. In fact, according to a 2021 article by UN-Habitat², 90% of megacities are vulnerable to sea level rise (SLR) and, as analyzed in (Hallegatte et al., 2013), the flood risk to coastal cities is expected to rise nine-fold by 2050. Land subsidence (i.e., the gradual sinking of land), which results mainly from anthropogenic activities and is especially pronounced in rapidly developing coastal cities (e.g.,

¹Calculated from the data reported in United Nations' World Cities Report 2024.

²UN-Habitat Press Release



Dubai or Jakarta (Cao et al., 2021; Ardha et al., 2024)), contributes to *relative* SLR and associated hazards, confronting coastal communities with the problem of managing multiple sources of risks (Barnard et al., 2024).

To safeguard the lives and livelihoods of their residents, coastal cities have been actively armoring their shorelines with engineered structures such as seawalls, levees, and storm barriers. Notable examples with already existing fortifications include the communities of San Francisco (SF), Sausalito and Oakland, located within the San Francisco Bay (Beagle et al., 2019); and the Greater New Orleans area, where many parts of the coastline have been extensively reinforced with multiple levees and seawalls that cover most of its perimeters (Papacharalambous et al., 2013). A more recent example involves New York City, where probability-based hydrodynamic modeling has identified areas of flood risk and has led to the protection of a two-and-a-half mile stretch of Lower Manhattan's shoreline with seawalls (at a total cost of \$1.8 billion) – an intervention largely driven by the aftermath of Hurricane Sandy in 2012, which flooded 17 % of New York and inflicted \$19 billion in damages (Lewis, 2023).

Construction of these coastal defense structures, however, significantly alters the shoreline geometry, subsequently influencing the local hydrodynamics, which in turn can translate into regional impacts. In particular, despite the conferred *local* flood-reduction benefits, such hold-the-line strategies that rely on engineered structures to maintain the present-day shoreline position tend to cause spillover effects, amplifying water levels outside their perimeters and spreading floodwaters to adjacent (otherwise unaffected) regions (Hummel et al., 2021; Wang et al., 2018; Haigh et al., 2020). For instance, as demonstrated in (Hummel et al., 2021), protection of certain individual stretches of the coast in SF Bay can increase the flooding in other zones by up to 36 million cubic meters owing to the effects of the modified shoreline geometry.

Therefore, for effective and responsible flood protection planning of coastal cities, it is essential to account for plausible hydrodynamic changes due to shoreline modifications and the potential negative ramifications thereof. To this end, physics-based high-fidelity simulators, such as Delft3D (Deltares), can be employed to resolve the detailed hydrodynamics around seawalls under SLR. While these tools can simulate complex coastal processes with detailed accuracy, providing valuable insights into flood dynamics, they are computationally prohibitive and resource-intensive, often requiring days to simulate a single shoreline protection scenario (Kyprioti et al., 2021; Rohmer et al., 2023; Karapetyan et al., 2024). This computational burden limits the practicality of such simulators coastal planning routines, where multiple simulations are needed to evaluate different adaptation strategies or SLR projections (Chen et al., 2024; Du et al., 2024).

In response to these limitations, data-driven methods, including artificial intelligence (AI), machine learning (ML), and deep learning (DL) techniques, have emerged as promising alternatives for rapid flood modelling and prediction (Bentivoglio et al., 2022; Mosavi et al., 2018; Muñoz et al., 2024; Zuhairi et al., 2022; Zhou et al., 2023; Nevo et al., 2022; Zhao et al., 2021). These surrogate models aim to emulate high-fidelity simulators by learning complex input-output relationships without explicitly modeling the underlying physical processes (Karapetyan et al., 2024).

However, despite these advancements, developing accurate DL-based models for high-resolution coastal flood prediction poses several challenges such as *data scarcity* and the *high dimensionality of outputs* (Kyprioti et al., 2021; Rohmer et al., 2023; Karapetyan et al., 2024). Previously, a deep vision-based framework was introduced to address these challenges by recasting flood prediction as a computer vision task (Karapetyan et al., 2024). This approach leverages the strengths of DL



models in handling high-dimensional outputs and deploying data augmentation, thereby enabling high-performance surrogate models even with limited training data. Taking a step further, this work extends and improves the DL model, and validates the proposed framework across two geographically distinct locations with different SLR values. More concretely, the key contributions of this study are as follows:

1. We develop a novel DL model (*CASPIAN-v2*) that accurately predicts high-resolution coastal flooding under various SLR scenarios and shoreline adaptation strategies. The proposed model employs a lightweight convolutional neural network (CNN) designed for fast and scalable flood prediction, significantly reducing computational time compared to traditional high-fidelity hydrodynamic models while maintaining high prediction accuracy.
2. We compile and provide two comprehensive datasets from vulnerable coastal cities, Abu Dhabi (AD) and San Francisco, covering different sea-level rise scenarios and shoreline adaptations. These datasets aim to facilitate future research and contribute to enhanced flood risk management strategies.
3. We rigorously evaluate the proposed framework against state-of-the-art (SOTA) ML and DL models, demonstrating superior performance with error reductions of up to 19.96%. Additionally, through fine-tuning, we tested the model on different SLR datasets, showing strong generalization capabilities across diverse scenarios.
4. We further validate *CASPIAN-v2* through explainable AI techniques, revealing that the predictions by the model closely align with physics-based, high-fidelity hydrodynamic simulations. These insights confirm the ability of the model to accurately identify inundated regions while offering interpretability that supports informed decision-making in flood risk assessment.

Put together, these contributions can assist urban policymakers in designing more effective and reliable coastal protection programs. Additionally, we open-source the code and datasets in the hope of facilitating further research and attracting greater attention to this problem within the machine learning community. Furthermore, we provide insights into the impact of different shoreline adaptation strategies on flood dynamics, supporting decision-makers in developing effective flood mitigation measures in response to climate change.

2 Related Works

Coastal regions are increasingly susceptible to flood-related hazards as a consequence of climate change and SLR (Archetti et al., 2011; Han and Mozumder, 2022; Liu et al., 2024). The compounded effects of urbanization, pluvial flooding (induced by rainfall), fluvial flooding (originating from river sources), and coastal flooding pose multifaceted challenges for disaster prevention and urban resilience planning (Liu et al., 2024; Archetti et al., 2011). Recent significant hurricanes, including Idalia (2023), Harvey (2017) and Sandy (2012), have starkly illustrated the catastrophic humanitarian, ecological, and economic consequences of coastal flooding (Han and Mozumder, 2022; Muñoz et al., 2024). Moreover, coastal municipalities are sig-



nificantly impacted by land subsidence; therefore, it is essential to consider all variables that contribute to flooding, including SLR, storm surges, and tidal inundation (Ezer, 2022; Barnard et al., 2024; Muñoz et al., 2024).

Traditional flood management techniques, which depend significantly on historical precipitation data and deterministic design principles, are largely insufficient in addressing the additional uncertainties introduced by climate change (Muñoz et al., 2024; Zhou et al., 2023). These methods significantly disregard the fact that weather patterns and growth contribute to the formation of impermeable surfaces and result in flooding of drainage systems (Aldiansyah and Wardani, 2023; Archetti et al., 2011). Given these constraints, researchers have begun to advocate for risk-based frameworks that prioritize resilient and flexible systems while encompassing a wider range of uncertainties (Bentivoglio et al., 2022; Nevo et al., 2022). Additionally, various decision-making systems for flood risk mapping, which integrate multiple aspects of vulnerability and hazards, have been created to improve urban flood resistance (Archetti et al., 2011; Liu et al., 2024).

Recent advancements in ML techniques have enhanced flood prediction models, allowing for real-time forecasting and significantly improving prediction accuracy (Chen et al., 2024; Ali et al., 2022). ML methods shorten computation time while maintaining acceptable accuracy levels. Numerous studies have assessed the risks and occurrences of floods using ML methods, including decision trees, random forest, support vector machines, and k -nearest neighbors (Mosavi et al., 2018; Ali et al., 2022; Syeed et al., 2022). Binary logistic regression has proven to be the most successful strategy for predicting flooding occurrences based on historical rainfall data when employing ML algorithms (Syeed et al., 2022).

Hybrid models that combine physical hydrodynamic simulations with ML techniques are increasingly favored to address the shortcomings of standalone methodologies (Chen et al., 2024; Du et al., 2024). The precision of urban flood predictions was markedly improved by integrating hydrodynamic models with ensemble stacking learning algorithms, illustrating the advantages of merging physical models with ML techniques (Du et al., 2024). Further, ensemble models integrating several data sources and ML techniques significantly improve flood prediction in complex flood scenarios (Mosavi et al., 2018; Adnan et al., 2023).

Compared to traditional ML methods, DL algorithms have greatly enhanced flood prediction capabilities. One-dimensional (1D) DL structures, such as long short-term memory (LSTM) networks and recurrent neural networks (RNNs), are widely applied in various fields, particularly in sequential forecasting. Additionally, hybrid models that combine sliding window techniques with exponential smoothing in LSTM greatly improve the accuracy of flood predictions. Moreover, CNNs and other two-dimensional (2D) DL methods have become increasingly popular for modeling floods and identifying spatial patterns. These methods can capture complex spatial features critical for understanding flood inundation by analyzing 2D data, such as flood extent maps, topographic grids, and other geographically distributed model outputs. They take into account spatial heterogeneity and local topographic factors to improve the understanding and prediction of floodwater movement in specific areas.

Despite these breakthroughs, evaluating coastal flooding due to climate change remains challenging, particularly when various shoreline adaptation strategies are taken into account. Current models frequently struggle with computational efficiency and the integration of dynamic factors such as land subsidence and urban morphology (Xu and Gao, 2024). While some research has employed hybrid surrogate models that combine ML approaches with traditional hydrodynamic models for real-time urban



coastal flooding forecasts (Xu and Gao, 2024), many flood prediction efforts focus on types of flooding or singular triggering
 125 variables. Additionally, numerous studies focusing solely on short-term extreme events overlook the implications of coastal
 adaptation techniques and the rising sea levels resulting from climate change (Jia et al., 2016; Guo et al., 2021).

A prior study (Karapetyan et al., 2024) provided a 2D DL framework for coastal flood prediction, achieving promising
 results, but was limited to a single location and a particular SLR scenario. This paper introduces a novel DL model designed to
 generalize across two distinct coastal regions and multiple SLR scenarios. This research seeks to enhance predictive accuracy
 130 and adaptability by integrating various flood-inducing factors into the framework, as illustrated in Figure 1. In doing this, we
 fill a gap in the literature for comprehensive models that can reliably predict floods under diverse climate change and shoreline
 protection scenarios, thereby improving advanced coastal flood risk management.




Figure 1. An overview of the proposed framework for coastal flood prediction. It begins with hydrodynamic simulations based on SLR data and coastal protection scenarios to generate raw flood data, which is then processed into spatial flood maps. The CASPIAN-v2 model, trained on these maps, predicts inundation patterns and flood extent. The framework can be fine-tuned with new data for improved adaptability. The different colored paths represent training (red), inference (green), and fine-tuning (blue) stages.

3 Study Area and Data Description

In this research, we examine two vulnerable metropolitan coastal areas (Abu Dhabi and the San Francisco Bay Area) to predict
 135 coastal flooding under various SLR and shoreline protection. Both locations feature low-lying topographies and significant
 urbanization, making them particularly susceptible to environmental effects. Our aim is to evaluate the effectiveness and appli-
 cability of the DL-based solution for forecasting inundation in these regions.



3.1 Study Area Description

3.1.1 Abu Dhabi

140 Abu Dhabi, located along the southern coast of the Arabian Gulf, faces rising flood risks from climate change-induced SLR, tidal flooding, and storm surges driven by extreme winds such as Shamal (Langodan et al., 2023). Projections estimate that a 0.5 m SLR, expected by 2050–2100 based on IPCC AR6 (IPCC, 2021), could inundate critical ecosystems like mangroves and artificial islands, potentially doubling flood zones when accounting for wind and wave action (Melville-Rea et al., 2021). The shallow bathymetry of the region amplifies these risks, with even minor sea-level increases threatening key infrastructure
 145 and densely populated areas, where over 85% of the population and 90% of the infrastructure lie just meters above sea level (al Kabban, 2019; Melville-Rea et al., 2021).

To assess flood risks, we divided the AD coastline into 17 operational landscape units (OLU), as outlined in (Chow and Sun, 2022). This framework captures unique features of both natural ecosystems and urban zones, enabling detailed flood vulnerability analyses under various shoreline adaptations. Figure 2 illustrates the AD coastline, the OLU divisions, and the
 150 inundation points for the 0.5 m SLR scenario.

3.1.2 San Francisco Bay Area

The San Francisco Bay Area also faces significant flood risks from SLR and tidal variability, which are exacerbated by climate change (Commission, 2018). Although storm surge impacts are lower than in AD due to its inland location, the low-lying urban zones, transportation networks, and hydrological systems (Napa River Basin) remain highly vulnerable to flooding (Commis-
 155 sion, 2018; Wang et al., 2018). Even moderate sea-level increases, combined with tidal events, threaten critical infrastructure and could lead to substantial economic losses.

For this research, we segmented the coastline of the Bay Area into 30 OLUs based on shoreline morphology, hydrology, and urban infrastructure, as adopted from previous studies (Hummel et al., 2021; Sun et al., 2020). Figure 3 illustrates the SF Bay Area, including OLU divisions and inundation points for the 1.0 m SLR scenario.

160 3.2 Data Sources and Hydrodynamic Simulations

To generate the necessary flood data for both regions, we conducted hydrodynamic simulations using the Delft3D model, which integrates key physical processes including SLR and tidal dynamics (see Supplementary Material, Section S.1). High-resolution bathymetry and digital elevation models (DEM) (with data sources such as TanDEM-X, Landsat-8, and Nautical Charts) were used for both regions to ensure accurate modeling of coastal topography that transitions smoothly between sea
 165 and the land. Additionally, for AD, we applied a SWAN model to simulate wind and wave effects, particularly the impact of Shamal winds, which can significantly intensify tidal flooding risks. Both the SWAN model and Delft3D models were forced with ERA5 meteorological data in the Arabian Gulf.




Figure 2. AD study area shown on the map of the United Arab Emirates. (a) All areas susceptible to flooding under a 0.5 m SLR scenario without any shoreline protections (b) The 17 OLUs defined along the AD shoreline where protections are to be tested for their effectiveness.




Figure 3. SF Bay study area shown on the map of the United States (a) All areas susceptible to flooding under a 0.5 m SLR scenario without any shoreline protections (b) The 30 OLU defined along the SF Bay shoreline where protections are tested for their effectiveness.



In both geographic locations we aimed to generate data that corresponds to a hypothetical future extreme flooding scenario. For AD, simulations were based on a 0.5 m SLR scenario, consistent with regional projections for mid-century SLR (as described above) (IPCC, 2021). The 0.5 m SLR scenario was then coupled with storm surges resulting from a sample Shamal event. In contrast, flood simulations for the SF Bay Area were conducted under three SLR scenarios: 0.5 m, 1.0 m, and 1.5 m. The data from the 1.0 m SLR scenario served as the primary training dataset for our model, while the 0.5 m and 1.5 m scenarios were used to evaluate the model's generalization capabilities under different SLR conditions. Table 1 shows the dataset details for both geographical regions, including SLR depths, OLU descriptions, and the number of unique shoreline protection scenarios.

Table 1. Dataset details for AD and SF regions, including OLUs, SLR depths, and the number of unique shoreline protection scenarios in the main, holdout, and generalizability sets.

Region	OLUs	SLR	Protection Scenarios
AD	17 OLUs: 1 (<i>Mussafah</i>), 2 (<i>Bain Al Jesrain</i>), 3 (<i>Grand Mosque District</i>), 4 (<i>AD Island West</i>), 5 (<i>Marina, CBD, Al Mina</i>), 6 (<i>AD Island East</i>), 7 (<i>Al Reem Island</i>), 8 (<i>Saadiyat Island</i>), 9 (<i>Yas Island</i>), 10 (<i>Al Raha Island</i>), 11 (<i>Al Shahama</i>), 12 (<i>Al Rahba</i>), 13 (<i>New Port City</i>), 14 (<i>Ghantoot</i>), 15 (<i>Lulu Island</i>), 16 (<i>Hudayriat Island</i>), 17 (<i>Inner Islands</i>)	0.5 m	142 (Main Set) 32 (Holdout Set)
SF	30 OLUs: 1 (<i>Richardson</i>), 2 (<i>Corte Madera</i>), 3 (<i>San Rafael</i>), 4 (<i>Gallinas</i>), 5 (<i>Novato</i>), 6 (<i>Petaluma</i>), 7 (<i>Napa - Sonoma</i>), 8 (<i>Carquinez North</i>), 9 (<i>Suisun Slough</i>), 10 (<i>Montezuma Slough</i>), 11 (<i>Bay Point</i>), 12 (<i>Walnut</i>), 13 (<i>Carquinez South</i>), 14 (<i>Pinole</i>), 15 (<i>Wildcat</i>), 16 (<i>Point Richmond</i>), 17 (<i>East Bay Crescent</i>), 18 (<i>San Leandro</i>), 19 (<i>San Lorenzo</i>), 20 (<i>Alameda Creek</i>), 21 (<i>Mowry</i>), 22 (<i>Santa Clara Valley</i>), 23 (<i>Stevens</i>), 24 (<i>San Francisquito</i>), 25 (<i>Belmont - Redwood</i>), 26 (<i>San Mateo</i>), 27 (<i>Colma - San Bruno</i>), 28 (<i>Yosemite - Visitacion</i>), 29 (<i>Mission - Islais</i>), 30 (<i>Golden Gate</i>)	1.0 m	285 (Main Set) 46 (Holdout Set)
		0.5 m	32 (Generalizability Set)
		1.5 m	32 (Generalizability Set)

We ran Delft3D simulations for three months to collect hourly inland inundation data under different coastal protection scenarios to create a dataset for training and validating our DL model. Our findings highlight the importance of holistic regional flood control measures, especially given the intricate interplay between protected and unprotected zones. Further, the datasets from two regions allowed us to assess the applicability and reliability of the DL model in different vulnerable coastal settings.

3.3 Data Preprocessing

The raw data generated using the Delft3D software was not directly compatible with our proposed DL model. Therefore, we first standardized and mapped the inundation coordinates onto a 1024×1024 grid, ensuring that each cell was uniquely assigned an inundation point. Subsequently, we classified each point according to its proximity to the nearest protected or unprotected OLU. We then encoded shoreline protection scenarios as binary strings, where '0' indicates unprotected OLUs and '1' denotes protected OLUs. Finally, we constructed the input and output matrices for training by combining these classification values with the corresponding peak water level (PWL) measurements (see Supplementary Material Section S.2 for details).



4 Method

This section presents the proposed DL architecture to predict inundation at various SLR depths and coastal protection scenarios. Additionally, we introduce a novel hybrid loss function that amalgamates multiple key metrics to improve robustness and adaptability against diverse flooding patterns.




Figure 4. The CASPIAN-v2 model architecture. The encoder extracts hierarchical features using FE blocks; the bottleneck employs MARX blocks to capture high-level representations; and the decoder reconstructs outputs using FR and SEE blocks. Different colors show separate layer operations.

4.1 CASPIAN-v2 Architecture

The CASPIAN-v2 model improves and extends the capabilities of the previously developed CNN architecture to predict coastal flooding (Karapetyan et al., 2024). Unlike the previous version, CASPIAN-v2 integrates SLR data and has a more robust yet minimalistic architecture that generalizes across various geographical regions. Figure 4 illustrates the CASPIAN-v2 architecture, which consists of three main stages: encoder, bottleneck, and decoder. The following subsections provide a detailed explanation of the CASPIAN-v2 architecture.



4.1.1 Encoder Stage

The encoder stage of the CASPIAN-v2 model is designed to extract hierarchical features by progressively reducing the spatial dimensions of the input grid while increasing the depth of the feature maps. This process enables the network to capture multi-scale patterns essential for accurate flood prediction.

The model accepts two inputs: an input grid $\mathbf{I} \in \mathbb{R}^{H \times W \times 1}$, where H and W are the spatial dimensions, and a scalar SLR value denoted as S . Since the SLR input S contains global information affecting the entire spatial domain, it is integrated directly into the decoder part of the network. The encoder stage contains a series of *feature extraction* (FE) blocks, allowing the model to capture both local features, such as specific inundation points and their immediate surroundings, and global patterns, including the overall spatial distribution of protected and unprotected areas.

First, the preprocessed input grid $\mathbf{I}$ is processed through a series of depthwise separable convolutional layers to reduce the spatial dimensions and extract complex features. At each depth level $k = 1$ to K , where K is the total depth of the encoder, the feature map $\mathbf{X}_k$ undergoes several transformations. First, a 2×2 depthwise convolution with stride 2 is applied to the input feature map $\mathbf{X}_k$, capturing spatial correlations within each channel independently while significantly reducing computational cost compared to standard convolutions. Following that, a single stride 1×1 pointwise convolution is used to combine the outputs across channels, allowing for feature interaction and increasing the depth of the feature map. In addition, 2×2 pooling operations with stride 2 are applied at each depth level k . The pooled features from the previous layer are also concatenated with the output of the pointwise convolution, enhancing the feature representation by merging hierarchical features from different scales. This enables the network to capture intricate patterns by combining information in different resolutions, which is essential to interpret how local features contribute to the overall risk of flooding.

We used residual connections to maintain key spatial data and increase network depth. Incorporating a modified $\mathbf{X}_k$ into the concatenated output mitigates gradient vanishing and improves identity mapping learning. These connections preserve crucial features of the early layer and streamline network training.

This process is repeated for each FE block in the encoder, leading to a gradual reduction in the spatial dimensions of the feature maps. At each FE block k , the spatial dimensions are reduced by factor 2, so the resulting feature maps have dimensions $\frac{H}{2^k} \times \frac{W}{2^k} \times F$, where F is the number of channels after concatenation. Such a progressive decrease in the spatial dimensions allows the network to capture more significant receptive fields, collecting information from more extensive regions of the input grid, which is critical for the simulated spread of inundation under various SLR situations. The input grid $\mathbf{I}$, by the end of the encoder stage, transforms into dense feature maps $\mathbf{X}_{\text{enc}} \in \mathbb{R}^{H' \times W' \times F}$ that capture both local and global data patterns, serving as input to the next stage.

4.1.2 Bottleneck Stage

The bottleneck stage of the CASPIAN-v2 comprises a novel block called the *multi-attention ResNeXt* (MARX) block to enhance the ability of the model to focus on the most informative parts of the data. It integrates ResNeXt blocks (Xie et al., 2017) with the *convolutional block attention module* (CBAM) (Woo et al., 2018). The output feature maps from the encoder



230 $(\mathbf{X}_{\text{enc}} \in \mathbb{R}^{H' \times W' \times F})$ serve as the input to the bottleneck stage. The MARX Blocks process feature maps through a sequence of operations, starting with a ResNeXt block, followed by a CBAM module, and concluding with an additional ResNeXt block. In the first ResNeXt block, the input feature map is divided into G groups, and group-specific convolutions are applied using 1×1 and 3×3 kernels, as expressed in Eq. (1):

$$\mathbf{X}_{R1} = \sigma \left(\mathbf{X}_{\text{enc}} + \sigma \left(\sum_{g=1}^G \mathbf{W}_g \circledast \mathbf{X}_{\text{enc},g} + \mathbf{b}_g \right) \right), \quad (1)$$

235 where $\mathbf{X}_{\text{enc}}$ is the input feature map, $\mathbf{W}_g$ and $\mathbf{b}_g$ are the weights and biases for the g -th group, $\circledast$ denotes the grouped convolution operation, and σ is the activation function. $\mathbf{X}_{R1} \in \frac{H}{2^K} \times \frac{W}{2^K} \times F_g$ is the output feature map of the first ResNeXt block, where K is the total depth of the encoder stage, and F_g is the number of channels determined by multiplying the cardinality $\mathcal{C}$ with the bottleneck width $\mathcal{B}$ ($F_g = \mathcal{C} \times \mathcal{B}$).

Next, to refine the feature maps and enable the model to focus on the most informative aspects of the data, we integrated
240 the CBAM principle within the MARX block. The CBAM enhances the representational power of the model by sequentially applying attention mechanisms along both the channel and spatial dimensions. In the channel attention module, inter-channel relationships are captured by computing a channel attention map $\mathbf{M}_c \in \mathbb{R}^{1 \times 1 \times F}$, which reweights the channels of the feature map, as expressed in Eq. (2):

$$\mathbf{M}_c = \delta(\mathbf{W}_{c2} \cdot \sigma(\mathbf{W}_{c1} \cdot \mathbf{z})), \quad (2)$$

245 where $\mathbf{W}_{c1} \in \mathbb{R}^{F^r \times F}$ and $\mathbf{W}_{c2} \in \mathbb{R}^{F \times F^r}$ are the weights of the fully connected layers, F^r is a reduction ratio parameter, σ denotes the activation function, and δ is the sigmoid function. The aggregated channel descriptor $\mathbf{z} \in \mathbb{R}^F$ is obtained by applying global average pooling over the spatial dimensions.

$$\mathbf{z}_f = \frac{1}{H'W'} \sum_{i=1}^{H'} \sum_{j=1}^{W'} \mathbf{X}_{R1,f}(i,j), \quad (3)$$

where $\mathbf{X}_{1,f}$ is the f -th channel of the feature map $\mathbf{X}_{R1}$. The channel attention map $\mathbf{M}_c$ is then applied to the feature map via
250 element-wise multiplication ($\mathbf{X}'_{R1} = \mathbf{M}_c \odot \mathbf{X}_{R1}$). This operation emphasizes channels that are more informative for predicting inundation patterns influenced by SLR and protection measures. Following that, a spatial attention map $\mathbf{M}_s \in \mathbb{R}^{H' \times W' \times 1}$ is computed by initially aggregating the feature map across the channel dimension using the average pooling, as expressed in Eq. (4):

$$\mathbf{q}(i,j) = \frac{1}{F} \sum_{f=1}^F \mathbf{X}'_{R1}(i,j,f), \quad (4)$$

255 where $\mathbf{q} \in \mathbb{R}^{H' \times W'}$ are the aggregated feature maps. Afterward, a 7×7 convolution is performed to extract intricate inundation patterns. Next, the spatial attention map $\mathbf{M}_s$ is applied to the refined channel feature map ($\mathbf{X}_C = \mathbf{M}_s \odot \mathbf{X}'_{RX1}$), allowing the model to concentrate on the spatial regions that are most pertinent for predicting flood inundation, such as areas with high



vulnerability due to low elevation or insufficient protection. Finally, the output of the CBAM module, $\mathbf{X}_C$, is processed through a second ResNeXt block to better capture the representations of complex features.

$$\mathbf{X}_{R2} = \sigma \left(\mathbf{X}_C + \sigma \left(\sum_{g=1}^G \mathbf{W}'_g * \mathbf{X}_{C,g} + \mathbf{b}'_g \right) \right), \quad (5)$$

where $\mathbf{X}_{R2}$ is the output feature map of the MARX Block, and $\mathbf{W}'_g$, $\mathbf{b}'_g$ are the weights and biases for the g -th group in the second ResNeXt block. The MARX blocks allow the CASPIAN-v2 model to generalize across complex datasets by adaptively concentrating on the most informative features in both channel and spatial dimensions. By the end of the bottleneck stage, the feature maps ($\mathbf{X}_{bn} \in \mathbb{R}^{H' \times W' \times F}$) are transformed into rich, high-level representations that capture key information about the candidate input scenario. These refined features serve as a strong foundation for the decoder stage, where they are progressively upsampled and combined with the SLR scalar $\mathcal{S}$ to reconstruct the spatial resolution of the input grid.

4.1.3 Decoder Stage

The decoder stage in the proposed CASPIAN-v2 model employs a series of *feature reconstruction* (FR) blocks to progressively upsample the feature maps. After the bottleneck, the refined feature maps $\mathbf{X}_{bn} \in \mathbb{R}^{H' \times W' \times F}$ serve as input to the decoder. The main goal is to restore the near-original spatial dimensions $H \times W$. At each depth level k , the decoder up-samples feature maps by a factor of 2 through transpose convolution, producing $\mathbf{X}_{dec,k} \in \mathbb{R}^{\frac{H'}{2^{k-1}} \times \frac{W'}{2^{k-1}} \times F}$. Upsampled maps are then concatenated to corresponding encoder outputs through skip connections, ensuring critical spatial details lost during downsampling are maintained.

To further strengthen the focus of the decoder on critical spatial regions and reflect SLR effects, we propose a novel *SLR-enhanced encoding* (SEE) block. It learns dynamic weighting from the SLR input to adjust decoder features. In the SEE block, the pooled feature maps of each encoder level are aggregated and passed through dense layers to generate spatial weighting coefficients, as expressed in Eq. (6).

$$\mathbf{w}_{sp,k} = \sigma \left(\mathbf{W}_{sp,k} \mathcal{F} \left(\mathcal{AP} \left(\sum_{f=1}^F \mathbf{P}^{(f)} \right) + \mathbf{b}_{sp,k} \right) \right), \quad (6)$$

where $\mathbf{P}^{(f)}$ is the f -th channel of the encoder feature map, $\mathcal{AP}$ denotes average pooling, $\mathcal{F}$ denotes the flattening operation, $\mathbf{W}_{sp,k}$ is weight matrix, and $\mathbf{b}_{sp,k}$ is bias term. Simultaneously, the SLR scalar input $\mathcal{S}$ is processed through a dense layer to produce $\mathbf{w}_{SLR}$. The spatial and SLR features are then concatenated to form $\mathbf{w}_{combined} = [\mathbf{w}_{sp,k}; \mathbf{w}_{SLR}]$, which is passed through another dense layer with sigmoid activation to produce the final weighting coefficients $\mathbf{w}_{SEE,k}$. These weighting coefficients are then reshaped and applied to the decoder feature map via element-wise multiplication $\mathbf{X}'_{dec,k} = \mathbf{X}_{dec,k} \odot \mathbf{w}_{SEE,k}$, where $\mathbf{X}_{dec,k}$ is the decoder feature map at the corresponding depth level, and $\odot$ denotes element-wise multiplication. This configuration of the SEE Block allows the model to adaptively weigh the decoder features based on both spatial information from the encoder and the global influence of SLR, enhancing the model's ability to predict flood inundation patterns under varying SLR scenarios.



At the output of the final FR block, a convolutional operation is applied to produce a preliminary output grid $\mathbf{O}_{\text{conv}} \in \mathbb{R}^{H \times W \times 1}$. To further enhance this output, the model computes the sum of the features across the channels of the last decoder feature map $\mathbf{X}'_{\text{dec},K} \in \mathbb{R}^{H \times W \times F}$. Moreover, the SLR input $\mathcal{S}$ is again incorporated at this stage by processing through a dense layer and then applying to the summed features via element-wise multiplication, as expressed in Eq. (7):

$$\mathbf{X}_{\text{sum}} = \left(\sum_{f=1}^F \mathbf{X}'_{\text{dec},K} \odot \mathbf{w}_{\text{SLR}} \right), \quad (7)$$

where $\mathbf{X}'_{\text{dec},K}$ is the f -th channel of the feature map. Finally, the enhanced summed features are added to the preliminary output grid before applying the activation function.

$$\mathbf{O} = \sigma(\mathbf{O}_{\text{conv}} + \mathbf{X}_{\text{sum}}), \quad (8)$$

where $\mathbf{O} \in \mathbb{R}^{H \times W \times 1}$ is the final output grid representing the predicted flood inundation map, and σ is the activation function. This allows extra information from the decoder feature maps by summing across the channels, thereby enriching the final output with more comprehensive spatial features. The grid $\mathbf{O}$ reflects the likelihood or extent of flooding at each spatial point, considering both the local features learned by the encoder and the broader SLR effects used in the decoder. This integrated design helps the CASPIAN-v2 model generate accurate and robust flood inundation maps, which are crucial for planning and mitigating coastal regions impacted by SLR.

4.2 Loss Function

Predicting PWL under different SLR scenarios is challenging due to outliers and the need to balance error sensitivity across multiple regions. To tackle these issues, we introduce a hybrid loss function that combines Huber (Huber, 1992), Log-Cosh (Saleh and Saleh, 2022), and Quantile (Koenker and Bassett Jr, 1978) losses in a weighted setup. The Huber loss L_h aims to robustly minimize small prediction errors while limiting the impact of outliers, and it uses a threshold δ to manage the sensitivity of the error. The L_h for each sample i is computed as expressed in Eq. (9):

$$L_{h,i} = \begin{cases} \frac{1}{2}(y_{p,i} - y_{t,i})^2 & \text{if } |y_{p,i} - y_{t,i}| \leq \delta, \\ \delta \cdot |y_{p,i} - y_{t,i}| - \frac{1}{2}\delta^2 & \text{otherwise} \end{cases} \quad (9)$$

where $y_{t,i}$ and $y_{p,i}$ represent the actual and estimated PWL values. We set δ within the range of 0.3 and 0.7, which is dynamically determined to balance sensitivity and robustness. Moreover, we integrate Log-Cosh loss (L_{cosh}) to smooth gradients in regions with large variations, helping to maintain prediction stability in different areas affected by SLR. The L_{cosh} is expressed as in Eq. (10):

$$L_{\text{cosh},i} = \log(\cosh(y_{p,i} - y_{t,i})), \quad (10)$$



In addition, the quantile loss L_q differentiates errors by assigning distinct penalties to underestimation and overestimation, dictated by a quantile parameter $\tau = 0.75$. This loss dynamically adjusts to minimize quantile-specific errors, calculated as in
 315 Eq. (11):

$$L_{q,i} = \begin{cases} \tau \cdot (y_{p,i} - y_{t,i}) & \text{if } y_{p,i} \geq y_{t,i}, \\ (1 - \tau) \cdot (y_{t,i} - y_{p,i}) & \text{otherwise} \end{cases} \quad (11)$$

To achieve an optimal balance, we linearly combine the three loss components into a comprehensive hybrid loss function L_{total} , weighted by empirically tuned coefficients $\alpha_h, \alpha_c, \alpha_q$. The final loss is expressed as in Eq. (12):

$$L_{\text{custom}} = \alpha_h \cdot L_h + \alpha_c \cdot L_{\text{cosh}} + \alpha_q \cdot L_q, \quad (12)$$

320 where $\alpha_h, \alpha_c, \alpha_q \geq 0$ and $\alpha_h + \alpha_c + \alpha_q = 1$. These weights are empirically determined to optimize predictive performance. By integrating these components, our custom hybrid loss function balances error sensitivity, maintains robustness to outliers, and addresses asymmetric error distributions, enhancing the model's predictive accuracy for PWL under varying SLR scenarios.

5 Experimental Setup

This section outlines the parameters employed to train, validate, and evaluate the proposed DL model. We detail the dataset
 325 splits, augmentation strategies, baseline models, and evaluation metrics to validate and compare the performance of the CASPIAN-v2 model.

5.1 Dataset Splits

We split both the AD and SF datasets into 3 subsets for training, validation and testing purposes, as shown in Table 2. To
 330 enhance the model's generalization ability and robustness under varying coastal conditions, we employed a systematic data augmentation strategy that generates multiple augmented versions of the training and validation subsets for both the AD and SF datasets.

Table 2. Dataset details for training the proposed model.

Region	Total	Train	Validation	Test
AD	142	96	10	36
SF	285	225	24	36

The augmentation process primarily involves a random remove function, which applies random spatial cutouts and scaling factors to the original samples. We create distinct yet related variants of the original dataset by systematically applying these



transformations multiple times ($24\times$ for AD and $10\times$ for SF). Compared to the original sparse dataset, the proposed augmen-
335 tation strategy produces a more balanced and richer dataset, comprising 2,304 training samples and 240 validation samples for
AD, along with 2,250 training samples and 240 validation samples for SF.

5.2 Model Optimization and Training Protocol

The CASPIAN-v2 model was implemented in Python 3.10 using TensorFlow 2.10.1 and was trained on a 64-bit Windows
operating system. We utilized an Intel Core i9-14900K (3.20 GHz) machine with 64 GB of RAM and an NVIDIA GeForce
340 RTX 4090 GPU. The CASPIAN-v2 architecture (shown in Figure 4) was refined through extensive ablation studies, which
evaluated the impact of loss functions, MARX blocks, the SEE block, and SLR integration (see Supplementary Material
Section S.3). The model was trained using the Adam optimizer and the proposed hybrid loss function. The training process
lasted for 200 epochs with a batch size of 2. The remaining hyperparameters were fine-tuned using Bayesian Optimization and
Random Search to ensure optimal performance.

5.3 Baseline Models

We assessed the performance of CASPIAN-v2 model for coastal flood prediction against several SOTA ML and DL techniques.
We considered conventional ML methods, including the Naïve model, which utilizes a dummy regressor to forecast the mean
value of the target variable to serve as a basic reference for assessing more advanced models. Additionally, we trained random
forest, linear regression, extreme gradient boosting, support vector regression, lasso regression with polynomial features, and
350 kriging with principal component analysis to establish an ML benchmark. The hyperparameters for training these models were
optimized through a combination of Bayesian optimization and random search methods, allowing for efficient exploration of
the parameter space while preventing overfitting on the validation set.

In addition to traditional ML baselines, we tested several DL models adapted to the flood prediction task. These include
a simple feed-forward neural network architecture, specifically a multi-layer perceptron (MLP), and compact convolutional
355 transformers (CCT) (Hassani et al., 2021), which serve as baseline 1D DL models. Furthermore, we evaluated several 2D
DL models, including Attention-Unet (Oktay et al., 2018), and Swin-Unet (Cao et al., 2022). To adapt these models for flood
prediction, we replaced their segmentation heads with a 1×1 convolution layer followed by activation to output real-valued
flood depth predictions. We evaluated two versions of Attention-Unet: one with randomly initialized weights and another
(denoted as Atten-Unet*) with an encoder pre-trained on ImageNet (Deng et al., 2009), leveraging transfer learning to improve
360 performance in low-data scenarios. The final DL baseline was CASPIAN, which we previously proposed in (Karapetyan et al.,
2024). All DL models were trained using the Adam optimizer and the proposed hybrid loss function (L_{custom}). Additionally,
each model was trained for 200 epochs with a batch size of 2, and early stopping based on validation loss. The remaining
training hyperparameters for each model were tuned using Bayesian Optimization and Random Search with the Keras Tuner
to ensure a fair comparison.



365 5.4 Evaluation Metrics

To evaluate the performance of our model in predicting PWL values, we employ several metrics that capture various aspects of predictive accuracy and robustness.

- Average relative total absolute error (ARTAE): It quantifies the relative error between the predicted y_p^k and true values y_t^k by measuring the normalized L_1 difference:

$$370 \quad \text{ARTAE} \triangleq \frac{1}{N} \sum_{k=1}^N \frac{\|y_t^k - y_p^k\|_1}{\|y_t^k\|_1} \quad (13)$$

where N denotes the total data samples.

- Average root mean square error (ARMSE): It captures the root mean square error for each sample, as expressed:

$$\text{ARMSE} \triangleq \frac{1}{N} \sum_{k=1}^N \sqrt{\frac{1}{d_y} \sum_{i=1}^{d_y} \frac{(y_{t,i}^k - y_{p,i}^k)^2}{d_y}} \quad (14)$$

where d_y indicates the dimensionality of each sample.

- 375 – Average mean absolute error (AMAE): It provides an average absolute error measure over samples. Unlike ARMSE, it does not severely penalize larger deviations. The AMAE is calculated as:

$$\text{AMAE} \triangleq \frac{1}{N} \sum_{k=1}^N \sum_{i=1}^{d_y} \frac{|y_{t,i}^k - y_{p,i}^k|}{d_y} \quad (15)$$

- Coefficient of determination (R^2): The R^2 measures the proportion of variance explained by the model, indicating how well the predictions approximate the true values. It is computed as:

$$380 \quad R^2 \triangleq \frac{1}{N} \sum_{k=1}^N \left(1 - \frac{\sum_{i=1}^{d_y} (y_{t,i}^k - y_{p,i}^k)^2}{\sum_{i=1}^{d_y} (y_{t,i}^k - \bar{y}_t^k)^2} \right) \quad (16)$$

where $\bar{y}_t^k$ is the mean of the true values for the k -th sample.

- Threshold exceedance metric ($\delta > \Delta$): This metric represents the fraction of cases where the prediction error exceeds a specified threshold Δ , which is crucial for applications sensitive to large errors. It is defined as:

$$\delta > \Delta \triangleq \frac{1}{N} \sum_{k=1}^N \frac{|\{i : |y_{t,i}^k - y_{p,i}^k| > \Delta, i \in [d_y]\}|}{d_y} \quad (17)$$

- 385 – Non-inundated prediction accuracy (Acc[0]): The Acc[0] measures the fraction of instances where the true values are zero (non-inundated points), and the predictions accurately reflect this. This is particularly relevant for sparse targets. It is computed as:

$$\text{Acc}[0] \triangleq \frac{1}{N} \sum_{k=1}^N \frac{|\{i : y_{t,i}^k = 0, i \in [d_y]\}|}{d_y} \quad (18)$$

6 Results

390 In this section, we evaluate the performance of CASPIAN-v2 model through quantitative and qualitative analyses.

6.1 Quantitative Results

6.1.1 Performance Evaluation on Test Set

We first report the performance of CASPIAN-v2 on the test set, as shown in Table 3. For AD data, the model achieves an AMAE of 0.0586, ARMSE of 0.4079, and a high average R^2 score of 0.9556, indicating excellent explanatory power. The ARTAE of 4.2793% and low error percentages ($\delta > 0.5\%$: 1.02% and $\delta > 0.1\%$: 4.37%) highlight higher precision in accurately predicting flood inundation levels. Similarly for SF, the model achieves an AMAE of 0.0320, ARMSE of 0.2094, and an average R^2 score of 0.9214. While the ARTAE is higher at 8.8129%, the model maintains high accuracy metrics with an Acc[0] of 99.76% compared to 99.04% in AD.

On the combined dataset, CASPIAN-v2 performs consistently well with an AMAE of 0.0453, ARMSE of 0.3087, and an average R^2 score of 0.9385. The combined ARTAE of 6.5461% and low error percentages ($\delta > 0.5$: 0.89% and $\delta > 0.1$: 3.55%) demonstrate balanced performance across regions. The high Acc[0] of 99.39% further underscores the reliability of the model in accurately predicting coastal inundation.

Table 3. Evaluation of CASPIAN-v2 on test set. ↓ indicates that lower values are better, and ↑ indicates that higher values are better.

Dataset	AMAE ↓	ARMSE ↓	ARTAE ↓	Avg. $\delta > 0.5$ ↓	Avg. $\delta > 0.1$ ↓	Avg. R^2 Score ↑	Avg. Acc[0] ↑
AD	0.0586	0.4079	4.2793	1.02%	4.37%	0.9556	99.04%
SF	0.0320	0.2094	8.8129	0.75%	2.72%	0.9214	99.76%
Combined	0.0453	0.3087	6.5461	0.89%	3.55%	0.9385	99.39%

6.1.2 Performance Evaluation on Holdout Set

In this section, we present CASPIAN-v2 performance on the holdout set, specifically curated to include challenging shoreline protection scenarios for both regions (see Supplementary Material Section S.4). The results are reported in Table 4, where it can be observed that the model achieves an AMAE of 0.0792, an ARMSE of 0.4871, and an average R^2 score of 0.9525 for AD. Furthermore, the small percentages of errors ($\delta > 0.5$: 1.29% and $\delta > 0.1$: 5.48%) underscore its accuracy in predicting flood inundation levels.

Similarly, for SF, CASPIAN-v2 achieves an AMAE of 0.0317, an ARMSE of 0.2259, and an average R^2 score of 0.9694. Compared to AD, the ARTAE of 4.0009% indicates slightly more predictions that have larger relative errors. However, with Acc[0] of 99.64%, the model achieves better non-inundated prediction accuracy compared to 99.07% in AD-Holdout.



Overall, CASPIAN-v2 achieves an AMAE of 0.0512, an ARMSE of 0.3331, and an average R^2 score of 0.9625 on the aggregated holdout dataset. The ARTAE of 3.7167% and small error percentages ($\delta > 0.5$: 1.04% and $\delta > 0.1$: 4.17%) signify consistent performance in both regions. The higher Acc[0] of 99.41% further confirms its reliability in predicting flood
 415 inundation across diverse and challenging shoreline scenarios.

Table 4. Evaluation of CASPIAN-v2 on holdout set.

Dataset	AMAE ↓	ARMSE ↓	ARTAE ↓	Avg. $\delta > 0.5$ ↓	Avg. $\delta > 0.1$ ↓	Avg. R^2 Score ↑	Avg. Acc[0] ↑
AD - Holdout	0.0792	0.4871	3.3081	1.29%	5.48%	0.9525	99.07%
SF - Holdout	0.0317	0.2259	4.0009	0.86%	3.26%	0.9694	99.64%
Combined	0.0512	0.3331	3.7167	1.04%	4.17%	0.9625	99.41%

6.1.3 Comparison with State-of-the-Art Methods

Comparison with Machine Learning Models:

In this section, we compare the performance of CASPIAN-v2 against various traditional ML models for flood prediction, as shown in Table 5. The Naïve model shows high errors with an AMAE of 1.5343, ARMSE of 3.5444, and an average R^2
 420 score of 0.5450. Among traditional approaches, linear regression reduces errors significantly, achieving an AMAE of 0.1272, ARMSE of 0.1946, and an average R^2 score of 0.9464. The lasso with polynomial model further improves performance, giving an AMAE of 0.0937, ARMSE of 0.1202, and the highest average R^2 score of 0.9618 among traditional ML models.

Compared to the best traditional model (lasso with polynomial), CASPIAN-v2 reduces the AMAE by 51.65% (from 0.0937 to 0.0453). However, CASPIAN-v2 has a higher ARMSE of 0.3087 compared to 0.1202, indicating it minimizes mean errors
 425 effectively but may experience larger individual prediction errors. Despite this, CASPIAN-v2 outperforms traditional models across multiple metrics, leveraging DL and multi-dimensional data integration to achieve superior accuracy in flood prediction.

Comparison with Deep Learning Models:

Existing 1D and 2D DL models show varied performance, as reported in Table 5. The CCT model achieves an AMAE of 0.9064, an ARMSE of 2.3292, and an average R^2 score of 0.6649, indicating moderate predictive capabilities. Atten-Unet and
 430 its variant Atten-Unet* improve performance with AMAE values of 0.1061 and 0.1032 and average R^2 scores of 0.9195 and 0.9210, respectively. Swin-Unet achieves further improvements, reducing the AMAE to 0.0629 and attaining an average R^2 score of 0.9514, reflecting its effectiveness in capturing spatial dependencies.

Compared to the second-best DL model, CASPIAN-v2 reduces the AMAE by 19.96% (from 0.0566 to 0.0453) and achieves an exceptional average Acc[0] of 99.39%, surpassing CASPIAN's 98.84%. These results highlight superior accuracy and
 435 robust generalization capabilities of CASPIAN-v2. The integration of advanced components such as MARX and SEE blocks in CASPIAN-v2, combined with an optimized Hybrid loss function, enables effective modeling of complex flood dynamics.



Table 5. Performance comparison of the proposed CASPIAN-v2 model with SOTA ML and DL models. The best and second-best results are highlighted in red and blue, respectively.

Type	Model	AMAE ↓	ARMSE ↓	ARTAE ↓	Avg. $\delta > 0.5$ ↓	Avg. $\delta > 0.1$ ↓	Avg. R^2 Score ↑	Avg. Acc[0] ↑
ML (1-D)	Naïve	1.5343	3.5444	1746.0693	74.92%	80.11%	0.5450	31.01%
	RF	0.5411	0.7310	264.9505	36.77%	72.20%	0.7962	34.19%
	Linear	0.1272	0.1946	64.9859	7.87%	14.03%	0.9464	59.28%
	XGBoost	0.2546	0.2446	164.1654	16.27%	49.88%	0.9347	44.10%
	SVR	0.2069	0.2423	72.3122	9.24%	41.17%	0.9298	45.46
	Lasso with Poly	0.0937	0.1202	28.1565	4.47%	15.04%	0.9618	55.78%
	Kriging	0.1098	0.2478	39.9073	5.22%	11.59%	0.9414	62.88%
DL (1-D)	MLP	0.6486	2.7247	524.1724	32.82%	41.94%	0.6572	36.91%
	CCT	0.9064	2.3292	843.5430	48.08%	64.63%	0.6649	34.01%
DL (2-D)	Atten-Unet	0.1061	0.3714	11.8245	3.14%	16.70%	0.9195	95.26%
	Atten-Unet*	0.1032	0.3627	11.6585	3.31%	15.62%	0.9210	94.99%
	Swin-Unet	0.0629	0.2788	6.7244	1.47%	12.94%	0.9514	98.10%
	CASPIAN	0.0566	0.3613	5.8573	1.01%	4.79%	0.9209	98.84%
	CASPIAN-v2 (Ours)	0.0453	0.3087	6.5461	0.89%	3.55%	0.9385	99.39%

*with pre-trained encoder on ImageNet (Deng et al., 2009).

6.1.4 Generalizability

This section reports the generalization performance of CASPIAN-v2 on unseen data. The model was fine-tuned using new SF data corresponding to 0.5 m and 1.5 m SLR depths, encompassing 30 protection scenarios where one OLU was protected at a time (more details in Supplementary Material Section S.5). For evaluation, 20% of the data (6 samples) was reserved, while the remaining 80% (24 samples) was used for fine-tuning and validation. Fine-tuning spanned 100 epochs with a progressive gradual recall approach, mixing the new data with the AD and SF holdout data in a 20:80 test/train ratio. The training set began with 70% of the AD and SF holdout set combined with 30% of the new data, gradually increasing to 70% by the end of training.

The results in Table 6 demonstrate strong generalization by CASPIAN-v2 across SLR scenarios. For SF 0.5 m data, the model achieved an AMAE of 0.0626, ARMSE of 0.2996, and average R^2 score of 0.9336. An ARTAE of 6.4240% and low error percentages ($\delta > 0.5$: 1.89% and $\delta > 0.1$: 7.79%) highlight its precision. For SF 1.5 m data, the model showed slightly suboptimal performance with an AMAE of 0.1005, ARMSE of 0.4565, and average R^2 score of 0.9196. The ARTAE of 4.3961% indicates balanced performance, with an average Acc[0] of 98.23% compared to 97.99% for 0.5 m data.



450 When retaining existing knowledge, CASPIAN-v2 achieved an AMAE of 0.0567 and ARMSE of 0.2274 on the AD holdout set for 0.5 m SLR, with an average R^2 score of 0.9901. The ARTAE of 2.5225% and low error percentages ($\delta > 0.5$: 0.53% and $\delta > 0.1$: 17.87%) emphasize its precision. For the SF holdout set at 1.0 m SLR, the model achieved an AMAE of 0.0433, ARMSE of 0.2318, and average R^2 score of 0.9685. The ARTAE of 4.6277% and error percentages ($\delta > 0.5$: 0.79% and $\delta > 0.1$: 9.61%) reflect its ability to balance low absolute and relative errors, with an Acc[0] of 99.34%.

455 Overall, the model achieved an AMAE of 0.0652, an ARMSE of 0.3040, and an average R^2 score of 0.9520, revealing robust generalization abilities of the model across various SLR settings. Further, the model achieved an ARTAE of 4.5871% and low error percentages ($\delta > 0.5$ ‰: 1.31% and $\delta > 0.1$ ‰: 12.07%), with a high Acc[0] of 98.69%. These findings highlight the ability of the CASPIAN-v2 model to effectively generalize to new and previously unseen scenarios with minor fine-tuning, making it a reliable tool for real-world inundation prediction.

Table 6. CASPIAN-v2 generalizability evaluation using different SLR data.

Dataset (SLR)	AMAE ↓	ARMSE ↓	ARTAE ↓	Avg. $\delta > 0.5$ ↓	Avg. $\delta > 0.1$ ↓	Avg. R^2 Score ↑	Avg. Acc[0] ↑
SF - Generalizability (0.5 m)	0.0626	0.2996	6.4240	1.89%	7.79%	0.9336	97.99%
SF- Generalizability (1.5 m)	0.1005	0.4565	4.3961	1.97%	14.51%	0.9196	98.23%
AD - Holdout (0.5 m)	0.0567	0.2274	2.5225	0.53%	17.87%	0.9901	99.18%
SF - Holdout (1.0 m)	0.0433	0.2318	4.6277	0.79%	9.61%	0.9685	99.34%
Overall	0.0652	0.3040	4.5871	1.31%	12.07%	0.9520	98.69%

460 6.2 Qualitative Results

6.2.1 Performance Evaluation on Test Set

In this section, we provide a qualitative assessment of the performance of CASPIAN-v2 on the test set. Figure 5 presents two randomly selected scenarios for the AD and SF regions, where it can be observed that the predicted inundation values of the proposed model closely align with the corresponding ground truth values. In single unprotected OLU scenarios (rows 1 and 3), the model accurately captures localized flooding effects, showing sensitivity to minor protection configuration changes.

465 Similarly, CASPIAN-v2 effectively handles the increased complexity of mixed OLU protection statuses (rows 2 and 4). These results highlight the robustness of the model in generalizing across diverse regions and protection patterns. Figure 5(c) shows the absolute error maps, where it can be observed that the CASPIAN-v2 model produced minimal errors, with deviations occurring mainly in areas with sharp transitions in flood depths. However, these small variations minimally affect the overall

470 prediction accuracy.




Figure 5. Evaluation of CASPIAN-v2 on the test datasets. (a) Ground truth inundation maps for representative AD and SF scenarios. (b) Predicted inundation values. (c) Absolute error distributions of predicted inundation values. Darker shades of blue indicate higher absolute errors, ranging from near 0% to greater than 25%.




Figure 6. Evaluation of CASPIAN-v2 on the holdout datasets. (a) Ground truth inundation maps for representative AD and SF scenarios. (b) Predicted inundation values. (c) Absolute error distributions of predicted inundation values. Darker shades of blue indicate higher absolute errors, ranging from near 0% to greater than 25%.



6.2.2 Performance Evaluation on Holdout Set

In this section, we demonstrate the performance of CASPIAN-v2 using a holdout set composed of particularly challenging coastal protection scenarios. Figure 6 illustrates two randomly selected configurations from both the AD and SF regions, each featuring a mix of protected and unprotected OLUs. Across these complex scenarios, CASPIAN-v2 continues to capture
475 inundation patterns with high fidelity, closely matching the ground truth and keeping errors at minimal levels. Notably, the model effectively handles abrupt changes in local flood behavior arising from varying protection measures. These results affirm the robustness and reliability of the CASPIAN-v2 model, even when encountering more challenging conditions beyond those in the main test set.

6.2.3 Comparison with State-of-the-Art Methods

480 We further evaluate CASPIAN-v2 predictions with SOTA models across various coastal protection scenarios, as shown in Figure 7. Here, it can be observed that CASPIAN-v2 consistently achieves lower absolute errors compared to other approaches in both regions. The Lasso with polynomial features baseline oversimplifies flood patterns, MLP struggles with spatial complexities due to its one-dimensional architecture, and Swin-Unet, while improved, still exhibits notable areas of higher error. The original CASPIAN model serves as a strong baseline, but CASPIAN-v2 outperforms it by delivering lower error levels and
485 more accurate representations of flood dynamics. These qualitative comparisons align with the quantitative results, highlighting the ability of the proposed model to achieve higher accuracy and visually superior predictions.

6.2.4 Generalizability

We next evaluate the generalizability of CASPIAN-v2 under different environmental conditions by fine-tuning the model on two additional SLR data of 0.5 m and 1.5 m. Figure 8 shows the prediction results, illustrating that while the fine-tuned model
490 exhibits some localized discrepancies (Figure 8(c)), these deviations remain modest given the minimal training data and limited fine-tuning epochs. In the 0.5 m SLR scenario, the model yields relatively lower absolute errors in predicting flood extents. By contrast, the 1.5 m scenario exhibits slightly higher errors, likely due to the increased variability in PWL values. Nonetheless, the predictions generally align well with the ground truth inundation patterns.

Overall, these findings underscore adaptability of the proposed model to evolving coastal conditions, suggesting that with
495 sufficient training data and appropriately tuned hyperparameters, the model can maintain robust performance across a broad range of SLR scenarios.

7 Discussion and Conclusion

This research presents a novel DL model to predict coastal inundation across two geographical locations (AD and SF). The effectiveness of the proposed CASPIAN-v2 model is validated through extensive experiments, where it outperforms the exist-
500 ing SOTA methods, as shown in Table 5. Although traditional ML approaches are relatively fast to train, these methods lack




Figure 7. Qualitative comparison of CASPIAN-v2 with state-of-the-art approaches in predicting coastal flood inundation. (a) Ground truth inundation maps for representative AD and SF scenarios. (b–f) Corresponding absolute error distributions of predicted inundation values: (b) Lasso with polynomial features, (c) MLP, (d) Swin-Unet (Cao et al., 2022), (e) CASPIAN (Karapetyan et al., 2024), and (f) CASPIAN-v2 (Ours). Darker shades of blue indicate higher absolute errors, ranging from near 0% to greater than 25%.

the ability to capture complex spatial patterns in the data, thus producing less accurate results. Similarly, we found that 1D DL approaches do not scale effectively to large, spatially focused grids. Furthermore, jointly training these methods on AD and SF datasets was less successful and yielded poor results due to inconsistent input features, particularly the different number of OLU across regions and the need to address a broader array of shoreline adaptation scenarios. In comparison, the proposed 2D DL model can learn complex input patterns, enabling it to produce superior prediction results. Additionally, the data augmentation technique employed in this research exposes the model to various conditions, which enhances its resilience to




Figure 8. Generalizability evaluation of CASPIAN-v2 fine-tuned for 0.5 m and 1.5 m SLR scenarios. (a) Ground truth inundation maps. (b) Predicted inundation values. (c) Absolute error distributions of predicted inundation values. Darker shades of blue indicate higher absolute errors, ranging from near 0% to greater than 25%.

noise, missing data, and varying shoreline configurations. Moreover, CASPIAN-v2 demonstrates strong generalizability across different levels of SLR, which underlines its utility for future resilience planning.

A critical aspect influencing model performance is the underlying data distribution. Coastal flood scenarios tend to be highly imbalanced, with a predominance of non-inundated areas and relatively few points experiencing extreme water depths. Figure 9 shows the normalized distribution of inundation values across the entire dataset, revealing a strong skew toward zero and a disproportionately small fraction of cells with inundation levels. Despite this imbalance, CASPIAN-v2 maintains robust predictive capabilities. Its attention-based bottleneck components emphasize critical regions, such as those near unprotected OLU, where flooding is most likely to occur. The integration of a non-inundation accuracy metric (Acc[0]%) and threshold exceedance evaluations further confirms the model's skill in delineating safe areas and identifying regions at heightened risk.




Figure 9. Normalized inundation distribution in the complete dataset, depicting a significant class imbalance between non-inundated points (0s) and inundated points (non-0s).

To further demonstrate the interpretability and decision-making mechanisms of the CASPIAN-v2 model, the Grad-CAM technique (Selvaraju et al., 2020) was utilized, as shown in Figure 10. The regions highlighted in warm colors in Figure 10(c) indicate areas of heightened model attention during flood prediction. Notably, these regions correspond closely with unprotected OLU and other environmentally vulnerable zones. This spatial alignment substantiates the decision-making process of the model, offering empirical validation for its predictions. Moreover, such interpretability provides critical insights for stakeholders by clarifying why specific areas are identified as high-risk. This level of transparency is instrumental in fostering trust in DL-based solutions and aids in designing targeted resilience and adaptation strategies.

Beyond accuracy and interpretability, the CASPIAN-v2 architecture enables strong generalizability. The model was tested with different SLR levels in SF and demonstrated consistent predictive performance. This adaptability to changing boundary conditions and climate scenarios underscores the practicality of the model for future-facing resilience planning. Moreover, the use of data augmentation strategies, such as random spatial cutouts and scaling, proves instrumental in exposing the model to a broader range of conditions, enhancing its resilience to noise, missing data, and evolving shoreline configurations.

Although CASPIAN-v2 represents a significant advancement in surrogate modeling for flood prediction, it is essential to highlight certain limitations. The prediction accuracy of the model largely depends on the underlying hydrodynamic models, where any inaccuracies in land surface conditions, atmospheric influences, or bathymetric data can introduce biases that may compromise the model outputs. Future research could explore incorporating different geographical contexts, including detailed elevation data and hydro-connectivity information through additional input channels, thereby further enhancing inundation predictions and reliability. Moreover, domain adaptation techniques and incremental learning could accelerate the implementation




Figure 10. CASPIAN-v2 inundation prediction for AD (top) and SF (bottom): (a) Input images representing protected and unprotected OLU, (b) Predicted inundation with PWL intensity, (c) Grad-CAM visualizations highlighting model attention, where warmer colors indicate regions the model focused on most during prediction, aligning with unprotected and vulnerable areas.

of models in various domains. Additionally, improving overall operational utility would enhance scalability and interpretability through model compression, distributed training, or more sophisticated explainable artificial intelligence algorithms.

In conclusion, the CASPIAN-v2 model offers a robust, adaptable, and comprehensible approach to predicting coastal floods. The proposed model incorporates computer vision and a DL-inspired framework to address the complexities of diverse geographical regions, protection scenarios, and climate variability. The CASPIAN-v2 model effectively identifies critical inundation areas, handles uneven data distribution, and provides a clear rationale for its predictions. These strengths position CASPIAN-v2 as an essential tool for coastal resilience planning, helping decision makers, engineers, and legislators address current and future flood risks in the context of rapidly rising sea levels and changing coastal conditions.



Code and data availability. The code and data supporting the findings of this study are available on the project page: <https://caspiannet.github.io>

Author contributions. BH, AK, ACHC, and SM contributed to the conceptualization and validation of research. BH and AK performed the formal analysis, investigation, and visualization; BH also developed the software. ACHC managed data curation, contributed to the methodology, and provided essential resources. SM secured funding, oversaw project administration, and offered supervision. All authors participated in writing the original draft and in reviewing and editing the manuscript.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This research was supported by New York University Abu Dhabi.



References

- Adnan, M. S. G., Siam, Z. S., Kabir, I., Kabir, Z., Ahmed, M. R., Hassan, Q. K., Rahman, R. M., and Dewan, A.: A novel framework for addressing uncertainties in machine learning-based geospatial approaches for flood prediction, *Journal of Environmental Management*, 326, 116 813, 2023.
- al Kabban, M.: Sea level rise vulnerability assessment for Abu Dhabi, United Arab Emirates, Master Thesis in Geographical Information Science, 2019.
- Aldiansyah, S. and Wardani, F.: Evaluation of flood susceptibility prediction based on a resampling method using machine learning, *Journal of Water and Climate Change*, 14, 937–961, <https://doi.org/10.2166/wcc.2023.494>, 2023.
- Ali, M. H. M., Asmai, S. A., Abidin, Z. Z., Abas, Z. A., and Emran, N. A.: Flood prediction using deep learning models, *Int J Adv Comput Sci Appl*, 13, 972–81, 2022.
- Archetti, R., Bolognesi, A., Casadio, A., and Maglionico, M.: Development of flood probability charts for urban drainage network in coastal areas through a simplified joint assessment approach, *Hydrology and Earth System Sciences*, 15, 3115–3122, 2011.
- Ardha, M., Chulafak, G. A., Anggraini, N., Syetiawan, A., and Khomarudin, M. R.: Flood inundation prediction model related to land subsidence with Lidar in North Coastal Jakarta, in: Eighth Geoinformation Science Symposium 2023: Geoinformation Science for Sustainable Planet, vol. 12977, pp. 392–401, SPIE, 2024.
- Barnard, P. L., Befus, K. M., Danielson, J. J., Engelstad, A. C., Erikson, L. H., Foxgrover, A. C., Hayden, M. K., Hoover, D. J., Leijnse, T. W. B., Massey, C., McCall, R., Nadal-Caraballo, N. C., Nederhoff, K., O'Neill, A. C., Parker, K. A., Shirzaei, M., Ohenhen, L. O., Swarzenski, P. W., Thomas, J. A., van Ormondt, M., Vitousek, S., Vos, K., Wood, N. J., Jones, J. M., and Jones, J. L.: Projections of multiple climate-related coastal hazards for the US Southeast Atlantic, *Nature Climate Change*, 2024.
- Beagle, J., Lowe, J., McKnight, K., Safran, S., Tam, L., and Szambelan, S. J.: San Francisco Bay shoreline adaptation atlas: Working with nature to plan for sea level rise using operational landscape units, SFEI publication# 915, 2019.
- Bentivoglio, R., Isufi, E., Jonkman, S. N., and Taormina, R.: Deep learning methods for flood mapping: a review of existing applications and future research directions, *Hydrology and Earth System Sciences Discussions*, 2022, 1–50, 2022.
- Cao, A., Esteban, M., Valenzuela, V. P. B., Onuki, M., Takagi, H., Thao, N. D., and Tsuchiya, N.: Future of Asian Deltaic Megacities under sea level rise and land subsidence: current adaptation pathways for Tokyo, Jakarta, Manila, and Ho Chi Minh City, *Current Opinion in Environmental Sustainability*, 50, 87–97, slow Onset Events related to Climate Change, 2021.
- Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., and Wang, M.: Swin-unet: Unet-like pure transformer for medical image segmentation, in: European conference on computer vision, pp. 205–218, Springer, 2022.
- Chen, G., Hou, J., Liu, Y., Xue, S., Wu, H., Wang, T., Lv, J., Jing, J., and Yang, S.: Urban inundation rapid prediction method based on multi-machine learning algorithm and rain pattern analysis, *Journal of Hydrology*, 633, 131 059, 2024.
- Chow, A. C. and Sun, J.: Combining Sea level rise inundation impacts, tidal flooding and extreme wind events along the Abu Dhabi coastline, *Hydrology*, 9, 143, 2022.
- Commission, C. E.: San Francisco Bay Area Report, California's Fourth Climate Change Assessment, https://www.energy.ca.gov/sites/default/files/2019-11/Reg_Report-SUM-CCCA4-2018-005_SanFranciscoBayArea_ADA.pdf, 2018.
- Deltares: Delft3d, <https://oss.deltares.nl/web/delft3d>.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L.: Imagenet: A large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, pp. 248–255, Ieee, 2009.



- Du, B., Wang, M., Zhang, J., Chen, Y., and Wang, T.: Urban flood prediction based on PCSWMM and stacking integrated learning model, *Natural Hazards*, pp. 1–25, 2024.
- Ezer, T.: A demonstration of a simple methodology of flood prediction for a coastal city under threat of sea level rise: The case of Norfolk, VA, USA, *Earth's Future*, 10, e2022EF002 786, 2022.
- 590 Griggs, G. and Reguero, B. G.: Coastal adaptation to climate change and sea-level rise, *Water*, 13, 2151, 2021.
- Guo, Z., Leitao, J. P., Simões, N. E., and Moosavi, V.: Data-driven flood emulation: Speeding up urban flood predictions by deep convolutional neural networks, *Journal of Flood Risk Management*, 14, e12 684, 2021.
- Haigh, I. D., Pickering, M. D., Green, J. M., Arbic, B. K., Arns, A., Dangendorf, S., Hill, D. F., Horsburgh, K., Howard, T., Idier, D., et al.: The tides they are a-Changin': A comprehensive review of past and future nonastronomical changes in tides, their driving mechanisms, and future implications, *Reviews of Geophysics*, 58, e2018RG000 636, 2020.
- 595 Hallegatte, S., Green, C., Nicholls, R. J., and Corfee-Morlot, J.: Future flood losses in major coastal cities, *Nature Climate Change*, 3, 802–806, <https://doi.org/10.1038/nclimate1979>, 2013.
- Han, Y. and Mozumder, P.: Risk-based flood adaptation assessment for large-scale buildings in coastal cities using cloud computing, *Sustainable Cities and Society*, 76, <https://doi.org/10.1016/j.scs.2021.103415>, 2022.
- 600 Hassani, A., Walton, S., Shah, N., Abuduweili, A., Li, J., and Shi, H.: Escaping the big data paradigm with compact transformers, *arXiv preprint arXiv:2104.05704*, 2021.
- Huber, P. J.: Robust estimation of a location parameter, in: *Breakthroughs in statistics: Methodology and distribution*, pp. 492–518, Springer, 1992.
- Hummel, M. A., Griffin, R., Arkema, K., and Guerri, A. D.: Economic evaluation of sea-level rise adaptation strongly influenced by hydro-dynamic feedbacks, *Proceedings of the National Academy of Sciences*, 118, e2025961 118, 2021.
- 605 IPCC: Climate change 2021: the physical science basis, Contribution of working group I to the sixth assessment report of the intergovernmental panel on climate change, 2, 2391, 2021.
- Jia, G., Taflanidis, A. A., Nadal-Caraballo, N. C., Melby, J. A., Kennedy, A. B., and Smith, J. M.: Surrogate modeling for peak or time-dependent storm surge prediction over an extended coastal region using an existing database of synthetic storms, *Natural Hazards*, 81, 909–938, 2016.
- 610 Karapetyan, A., Chow, A. C. H., and Madanat, S.: Deep Vision-Based Framework for Coastal Flood Prediction Under Climate Change Impacts and Shoreline Adaptations, *arXiv preprint arXiv:2406.15451*, 2024.
- Koenker, R. and Bassett Jr, G.: Regression quantiles, *Econometrica: journal of the Econometric Society*, pp. 33–50, 1978.
- Kyprioti, A. P., Taflanidis, A. A., Nadal-Caraballo, N. C., and Campbell, M.: Storm hazard analysis over extended geospatial grids utilizing surrogate models, *Coastal Engineering*, 168, 103 855, 2021.
- 615 Langodan, S., Cavaleri, L., Benetazzo, A., Bertotti, L., Dasari, H. P., and Hoteit, I.: The peculiar wind and wave climatology of the Arabian Gulf, *Ocean Engineering*, 290, 116 158, 2023.
- Lewis, A.: After a Decade of Planning, New York City Is Raising Its Shoreline, *Yale School of the Environment*, <https://e360.yale.edu/features/new-york-city-climate-plan-sea-level-rise>, 2023.
- 620 Liu, Y., Zhang, T., Ding, Y., Kang, A., Lei, X., and Li, J.: Exploring the driving factors of compound flood severity in coastal cities: a comprehensive analytical approach, *Hydrology and Earth System Sciences Discussions*, 2024, 1–34, 2024.
- Melville-Rea, H., Eayrs, C., Anwahi, N., Burt, J. A., Holland, D., Samara, F., Paparella, F., Al Murshidi, A. H., Al-Shehhi, M. R., and Holland, D. M.: A roadmap for policy-relevant sea-level rise research in the United Arab Emirates, *Frontiers in Marine Science*, 8, 670 089, 2021.



- Mosavi, A., Ozturk, P., and Chau, K.-w.: Flood prediction using machine learning models: Literature review, *Water*, 10, 1536, 2018.
- 625 Muñoz, D. F., Moftakhari, H., and Moradkhani, H.: Quantifying cascading uncertainty in compound flood modeling with linked process-based and machine learning models, *Hydrology and Earth System Sciences*, 28, 2531–2553, 2024.
- Nevo, S., Morin, E., Gerzi Rosenthal, A., Metzger, A., Barshai, C., Weitzner, D., Voloshin, D., Kratzert, F., Elidan, G., Dror, G., et al.: Flood forecasting with machine learning models in an operational framework, *Hydrology and Earth System Sciences*, 26, 4013–4032, 2022.
- Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B., et al.: Attention u-net: Learning where to look for the pancreas, *arXiv preprint arXiv:1804.03999*, 2018.
- 630 Pal, I., Kumar, A., and Mukhopadhyay, A.: Risks to Coastal Critical Infrastructure from Climate Change, *Annual Review of Environment and Resources*, 48, 681–712, 2023.
- Papacharalambous, M., Davis, M., Marshall, W., Weems, P., and Rothenberg, R.: Greater New Orleans Urban Water Plan: Implementation, Waggoner & Ball Architects: New Orleans, LA, USA, 2013.
- 635 Rohmer, J., Sire, C., Lecacheux, S., Idier, D., and Pedreros, R.: Improved metamodels for predicting high-dimensional outputs by accounting for the dependence structure of the latent variables: application to marine flooding, *Stochastic Environmental Research and Risk Assessment*, 37, 2919–2941, 2023.
- Saleh, R. A. and Saleh, A.: Statistical properties of the log-cosh loss function used in machine learning, *arXiv preprint arXiv:2208.04564*, 2022.
- 640 Selvaraju, R. R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., and Batra, D.: Grad-CAM: visual explanations from deep networks via gradient-based localization, *International journal of computer vision*, 128, 336–359, 2020.
- Sun, J., Chow, A. C., and Madanat, S. M.: Multimodal transportation system protection against sea level rise, *Transportation Research Part D: Transport and Environment*, 88, 102 568, 2020.
- Syeed, M. M. A., Farzana, M., Namir, I., Ishrar, I., Nushra, M. H., and Rahman, T.: Flood prediction using machine learning models, in: 2022 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA), pp. 1–6, IEEE, 2022.
- 645 van de Wal, R., Melet, A., Bellafore, D., Camus, P., Ferrarin, C., Oude Essink, G., Haigh, I. D., Lionello, P., Luijendijk, A., Toimil, A., et al.: Sea Level Rise in Europe: Impacts and consequences, *State of the Planet*, 3, 1–33, 2024.
- Wang, R.-Q., Stacey, M. T., Herdman, L. M. M., Barnard, P. L., and Erikson, L.: The influence of sea level rise on the regional interdependence of coastal infrastructure, *Earth's Future*, 6, 677–688, 2018.
- 650 Woo, S., Park, J., Lee, J.-Y., and Kweon, I. S.: Cbam: Convolutional block attention module, in: *Proceedings of the European conference on computer vision (ECCV)*, pp. 3–19, 2018.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K.: Aggregated residual transformations for deep neural networks, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.
- Xu, L. and Gao, L.: A hybrid surrogate model for real-time coastal urban flood prediction: An application to Macao, *Journal of Hydrology*, 642, 131 863, 2024.
- 655 Zhao, G., Bates, P., Neal, J., and Pang, B.: Design flood estimation for global river networks based on machine learning models, *Hydrology and Earth System Sciences*, 25, 5981–5999, 2021.
- Zhou, Q., Teng, S., Situ, Z., Liao, X., Feng, J., Chen, G., Zhang, J., and Lu, Z.: A deep-learning-technique-based data-driven model for accurate and rapid flood predictions in temporal and spatial dimensions, *Hydrology and Earth System Sciences*, 27, 1791–1808, 2023.
- 660 Zuhairi, A. H., Yakub, F., Zaki, S. A., and Ali, M. S. M.: Review of flood prediction hybrid machine learning models using datasets, in: *IOP Conference Series: Earth and Environmental Science*, vol. 1091, p. 012040, IOP Publishing, 2022.