

## Anonymous Referee # 1

*The manuscript proposes a novel deep learning framework for predicting coastal flooding (only tidal). The framework has been applied to the urban areas of Abu Dhabi and San Francisco. The proposed model performs better in comparison to the existing DL methods. The manuscript is free from grammatical errors and is interesting to read. However, the following serious concerns need to be addressed,*

### **Response:**

We thank the reviewer for their positive assessment of our manuscript and for recognizing the novelty and potential of the proposed deep learning framework for coastal flood prediction. We appreciate the encouraging comments regarding the manuscript's clarity and relevance. We also acknowledge the important concerns raised and will address each of them in detail in the revised manuscript. Below, we outline the actions we plan to take to respond to each point and improve the manuscript accordingly.

*The developed model is only tuned for tidal flooding, and it does not account for the storm surge, which is a major source of coastal flooding. Especially, San Francisco is vulnerable to flooding by hurricanes. The paper doesn't mention the idea behind the omission of flooding due to storm surge, and why modelling only tidal flooding is crucial in these example cases. The motivation for setting up this kind of prediction system is very unclear.*

### **Response:**

We thank the reviewer for the comment. We acknowledge that the San Francisco area does indeed experience storm surges especially along the Pacific coastline and to a lesser extent within San Francisco Bay. In the Supplementary materials section S1, line 16, we justified not using a storm surge model for San Francisco Bay due to significant wave heights typically of about 0.07-0.2 m, an order of magnitude smaller than the 2.0-3.0m found on the California Coast outside San Francisco Bay. We will update our manuscript to clarify our focus area which is San Francisco Bay and its shoreline urban communities, as opposed to the coastal communities facing the Pacific Ocean. We have focused here because this is an area with a larger urban infrastructure such as highways that could suffer flood induced interruptions that affect the entire Bay Area transport network. Additionally, independent of storms, the addition of coastal protections (in the form of sea walls or levees) to one part of the Bay may affect other parts of the Bay; Hollemann and Stacey, 2014 showed that protecting the southern portion of the Bay may increase maximum water levels by 0.2 in the north of the Bay; Wang et al., 2018 showed that protection of the southern portions of the West Bay adversely affects the water levels of the East Bay, and vice-versa.

Meanwhile, the Abu Dhabi coastline experiences storm surges that are caused by sustained northwesterly winds that generally occur in the winter (called the Shamal winds). They are incorporated into the model by the coupling of Delft3d with SWAN, a spectral wave model, and storm runoff calculations were performed along the Abu Dhabi coastline in the same manner described in Chow and Sun, 2022. Details of the Abu Dhabi hydrodynamic models were provided in the Supplementary material, but in the next revision of the manuscript we will incorporate more

detailed descriptions of the Abu Dhabi Delft3d and SWAN coupled models into the main body of the manuscript..

*Also, what is the level of damage that can occur due to tidal flooding (as the storm surge is not included here) needs to be discussed so that the readers can understand the significance or the need for this framework.*

**Response:**

Thank you for the question. We will include in our revised manuscript that the motivations of studying the changes to tidal flooding within San Francisco Bay due to different protection scenarios are based on protection observations from previous studies that adding coastal protections to one part of the Bay may affect other parts of the Bay (Holleman and Stacey, 2014; Wang et al., 2018).

Meanwhile storm surges and tidal flooding have both been modeled for the Abu Dhabi coastline since there are effects both from storm surges with the presence of sustained onshore winds and tidal interactions between the multiple mangrove islands in the Abu Dhabi area (as shown in Figure 2).

*The manuscript needs to contain the details of how the hydrodynamic model, whose results have been used for training the DL model, is validated. What kind of events were used? How closely they matched the observed flooding also needs to be included in the paper.*

**Response:**

Thank you for the comment. Previous work has performed validation of the hydrodynamic models used in this paper with the current tidal gauges so that the observed water levels (without sea level rise) are well reproduced by the simulation – see Wang et al., 2017 and Chow and Sun, 2022. The simulations were performed with a 0.5m SLR rise which reflects a possible future scenario for San Francisco Bay in the year (somewhere between 2050-2100 depending on the climate change scenario pathway (between SSP2-4.5 and SSP5-8.5) from IPCC AR6 report (referenced in Section 3.1.1 of the Supplementary Materials). As such the simulated flooding without SLR serves as a predicted flooding extent that is not currently observed.

*The paper argues that the developed model is computationally efficient, but it fails to elucidate the details on how fast it is in comparison to the hydrodynamic model. This is very crucial to establish the significance of this work.*

**Response:**

We appreciate the reviewer's comment regarding the need for a computational efficiency analysis. In the revised manuscript, we will include a detailed comparison of training time, inference time, and GPU memory usage across all evaluated models. Specifically, our CASPIAN-v2 model requires approximately 23 hours to train on a local machine equipped with an NVIDIA RTX 4090 GPU. Its inference time is around 0.227 seconds per scenario, and the model is lightweight, comprising only 0.38 million parameters.

In contrast, simulating a single scenario using traditional hydrodynamic models (Delft3D for San Francisco and Delft3D coupled with SWAN for Abu Dhabi) takes approximately 14 hours and 19

hours, respectively, on high-performance computing infrastructure. This means that running all 72 scenarios in the test set (36 for Abu Dhabi and 36 for San Francisco) would take around 49 days using these simulators. In comparison, CASPIAN-v2 can process the same 72 scenarios in roughly 17 seconds.

We will include these comparisons in a new table and accompanying discussion to clearly demonstrate the significant time and resource savings achieved by our surrogate model, directly addressing the computational motivation highlighted in the manuscript.

*There is no spatial validation of the over-/under-/correctly predicted flooded areas. For example, see Nithila Devi et al. (2024), where a spatial fitness index has been used to assess the spatial accuracy.*

**Response:**

We thank the reviewer for the suggestion and for referencing Nithila Devi et al. (2024). We agree that spatial validation adds important insight into model performance. In the revised manuscript, we will reference past tidal gauge validation performed for the San Francisco model (Wang et al., 2017) and the Abu Dhabi model (Chow and Sun, 2022) for vertical water level validation with observations. While no flooding extent has been observed in the case of zero sea level rise, our model attempts to predict the extent of flooding in a future scenario of 0.5m sea level rise.

*In San Francisco, there is also riverine and pluvial flooding present. How only modelling the tidal flooding sufficient in this case for assessing the flood protection capabilities?*

**Response:**

Thank you for the comment. Although our San Francisco model does include riverine input from the Sacramento and San Joaquin Rivers, the inflow rates into the Bay were baseline values rather than for extreme fluvial flood events. While we acknowledge that incorporating more refined hydrodynamic forcing conditions to include pluvial and riverine floods, as well as extreme storm events, can refine the hydrodynamic model to reflect more extreme flooding, our overall scope in this paper is in the use of machine learning to be able to act as a surrogate for a hydrodynamic model running at a more average condition but under a Sea Level Rise of 0.5m and in the presence of a coastal protection over many different stretches of the shoreline. Our focus on tidal flooding within San Francisco Bay was to highlight the unique tidal behavior within the Bay where the construction of sea walls along certain portions of the shoreline Bay may in fact exacerbate the sea level within the Bay to increase by up to 1 m (Holleman and Stacey, 2014).

*Flood protection measures on the coast are very vaguely discussed. It is very difficult to understand what they are, how they function, and how they are incorporated into the hydrodynamic and DL models. It will be nice to discuss with figures showing their placement, functionality and model representation explicitly.*

**Response:**

Thank you for the comment. In the revised manuscript, we will include and add more information contained in Supplemental Materials to the main body of the paper, namely: Sections 3.1.2 for

San Francisco, which references previous work that provides more details of the coastal protection assumptions; and Section 3.1.1 that describes the flood protection measures taken for Abu Dhabi, which also references previous work that provide more details.

*Regarding the writing style, I find that the paper is more from a computer science background, rather than explaining clearly the application and relevance to the hydraulic modelling application. For example, the manuscript lacks details on training data used, accuracy of the training data, spatial prediction accuracy, specifics of protection measures, etc., as mentioned in the earlier comments.*

**Response:**

We appreciate the observation of the reviewer regarding the writing style and the need to better align the manuscript with the expectations of the hydraulic modeling community. In the revised version, we will move some of the material presented in the Supplementary Materials section into the main body of the manuscript, we will revisit relevant sections to better highlight the application relevance and hydrodynamic context of our work (with some of the model development and validation material taken from Wang et al., 2018 and Chow and Sun, 2022). Finally, we will improve the explanation of the training data (including its source and accuracy), provide further clarification on spatial prediction accuracy, and include additional details on the definitions of the Operational Landscape Units (OLUs) and the protection scenarios of OLU that were used. This will ensure that the manuscript clearly communicates both the modeling and application aspects of the study.

*The paper is very wordy with several irrelevant information and lacks crucial information necessary for the understanding of the readers. For example, information on kind of damages due to tidal flooding, validation etc. Please discuss the important information briefly rather than just citing other works.*

**Response:**

Thank you for the feedback. In the revised manuscript, we will reduce wordiness and remove less relevant content. We will also briefly include key information, such as tidal flood impacts and data validation, to improve clarity and minimize reliance on external citations.

*The following are the specific comments listed for each section,*

*Introduction and Related Works*

*The introduction is very wordy and unorganised. A separate section of “related works” is unwarranted and has a lot of redundant information. Therefore, sections 1 and 2 (introduction and related works) need to be shortened and merged into one, ensuring that there is a proper flow of content in them.*

**Response:**

Thank you for this suggestion. In the revised manuscript, we will merge the Introduction and Related Works sections to create a more concise and cohesive narrative. Redundant content will be removed, and we will ensure a clearer flow that integrates background, motivation, and relevant literature more effectively.

*Lines 17-20 can be shortened.*

**Response:**

Thank you for pointing this out. In the revised manuscript, we will shorten lines 17–20 to convey the key message more concisely while maintaining clarity and relevance.

*Lines 26 – 34 need to be shortened, retaining information relevant to the kind of study the paper deals with.*

**Response:**

Thank you for the suggestion. We will revise lines 26–34 to remove general content and retain only information directly relevant to the scope and objectives of this study.

*Please use terms like computationally expensive/intensive instead of the term “computationally prohibitive” if possible.*

**Response:**

Thank you for the suggestion. As suggested by the reviewer, we will replace the term “computationally prohibitive” with “computationally expensive” in the revised manuscript, where applicable.

*Mention literature related to computationally efficient modelling approaches (subgrid, parallelisation, simplified models, etc.), such as De Almeida et al. (2013), Neal et al. (2012), Li and Hodges (2019), Sanders and Schubert (2019), Nithila Devi et al. (2024), etc. Discuss such methodologies and bring out the importance of DL/ML/AI in flood forecasting and modelling.*

**Response:**

Thank you for this helpful recommendation regarding various methods of employing spatially nested models. In the revised manuscript, we will revise the introduction to include a brief discussion of computationally efficient hydrodynamic modeling approaches, as suggested by the reviewer. We wish to note that nesting has not been necessary for the DELFT3D portion of the model, as the model grid has been downscaled to about 30 m horizontal resolution for areas of interest such as the city of Abu Dhabi; and that the model is capable of modeling dry and wet grid cells. But we did employ a nesting method to incorporate the SWAN to model storm surges and waves for Abu Dhabi. We will highlight the limitations of traditional approaches and clearly motivate the role and advantages of AI-based methods in improving computational efficiency and scalability for flood forecasting and modeling.

*Highlight the need for high-resolution modelling, especially in the complex urban terrain and cite relevant literature.*

**Response:**

We appreciate the suggestion of the reviewer to emphasize the need for high-resolution modeling, particularly in complex urban terrain. Initially, data from hydrodynamic simulations was in tabular format, which, although containing peak water levels and coordinates, lacked explicit spatial structure and contextual relationships among points. This made it difficult for a learning algorithm to fully understand the spatial dependencies critical in urban flood prediction.

To address this, we transformed the tabular data into a high-dimensional spatial grid format (1024x1024) suitable for convolutional neural networks (CNNs). This transformation allowed us to encode not just the water level values but also the precise spatial positioning of each point within the urban landscape. By incorporating additional contextual information, we classified points by assigning two pixel values to each grid cell (one indicating proximity to protected OLU boundaries and the other to unprotected areas). This spatial encoding allowed the model to learn fine-grained flood patterns and localized behaviors more effectively within the CNN framework. This high-resolution spatial representation is essential in urban areas, where small variations in terrain or infrastructure can significantly influence inundation patterns. In future work, this framework also supports the integration of additional channels or modalities to further enrich spatial context.

*Line 56, what do you mean by high dimensionality of outputs?*

**Response:**

Thank you for pointing this out. By “*high dimensionality of outputs*”, we refer to the pixel-wise prediction of inundation values over large spatial grids (e.g., 1024×1024), where the model must produce a continuous value for each pixel. This increases the complexity of capturing spatial dependencies and patterns, particularly under varying coastal boundary and sea level rise conditions. In the revised manuscript, we will clarify the sentence accordingly to avoid ambiguity.

*Since this paper expands on a deep-vision based framework, explain briefly about this in the introduction so that the readers from diverse backgrounds can appreciate the work. (Line 57)*

**Response:**

Thank you for this suggestion. In the revised manuscript, we will include a brief explanation of the deep vision-based framework in the Introduction, as suggested by the reviewer. This addition will help readers from diverse backgrounds understand that our approach involves using convolutional neural networks (CNNs) to learn spatial patterns directly from gridded input data, enabling pixel-wise prediction of flood extent under varying coastal conditions.

*The paper lacks a dedicated and concrete objective description; rather, it mentions the conclusions from the study in the introduction. Please move the lines 61 – 81 to conclusions or discussions, if redundant, considering removing them.*

**Response:**

We thank the reviewer for the feedback. The content in lines 61–81 was intended to highlight the key contributions of the study, which is a common practice in scientific writing to help readers quickly grasp the novelty and scope of the work. However, we understand that it may have been interpreted as a set of concluding remarks. To address this, we will revise the wording to clearly frame these points as the main contributions and ensure they are presented as part of the objective and motivation in the Introduction. If any parts appear repetitive or better suited for the conclusion, we will move or remove them accordingly to improve the flow and clarity.

*Lines 83 – 90, redundant information. Please remove them while merging sections 1 and 2.*

**Response:**

Thank you for the suggestion. We will remove lines 83–90 while merging the sections, as they are redundant and do not contribute additional value to the streamlined narrative.

*Lines 91 – 98, irrelevant information.*

**Response:**

Thank you for the comment. As part of merging and streamlining the Introduction and Related Works sections, we will remove lines 91–98, as they contain information that is not directly relevant to the focus of this study.

*Mention what kind of flood protection strategies exist here.*

**Response:**

We thank the reviewer for the comment. While structural flood protection strategies, such as levees, seawalls, and revetments, are commonly used, our study does not focus on the specific type or engineering design of these structures. Instead, our primary concern is to evaluate the spatial importance of different shoreline segments in mitigating flood impacts.

Specifically, the study area in Abu Dhabi includes 17 defined shoreline protection segments, while San Francisco includes 30. Our analysis explores how different combinations of these segments, when protected or left unprotected, influence inland inundation patterns. The objective is to identify which segments are most critical in reducing flood extent and should be prioritized for protection, irrespective of the exact structural form eventually adopted.

It is also important to note that the Delft3D hydrodynamic simulator used in this study does not inherently distinguish between different types of flood protection structures in its core computations. Whether a barrier is a levee, seawall, or any other type, the model treats it generically as a solid obstruction defined by its geometry and crest elevation. This further justifies our focus on the presence or absence of protection along spatial segments rather than structural details.

This evaluation is performed using our trained deep learning model, which can rapidly assess the impacts of hundreds of protection patterns across scenarios. In contrast, performing such analysis with traditional hydrodynamic simulators, such as Delft3D or Delft3D+SWAN, would be highly time-consuming; a single simulation can take up to 19 hours. Our surrogate model reduces this computation to a matter of seconds per scenario, enabling efficient and scalable analysis of spatial protection strategies. We will clarify this focus more in the revised manuscript.

*Lines 123 – 125, not clear.*

**Response:**

We thank the reviewer for pointing this out. We agree that the lines could be made clearer. In the revised manuscript, we will rephrase this part as follows:

Despite progress in flood modeling, current approaches often face limitations in computational efficiency and in incorporating dynamic factors such as land subsidence and urban morphology



(Xu and Gao, 2024). Some recent studies have introduced hybrid surrogate models that combine machine learning with traditional hydrodynamic simulations for real-time coastal flood forecasting (Xu and Gao, 2024). However, many existing efforts remain focused on specific flood types or single triggering factors, without jointly considering long-term climate impacts such as sea level rise and evolving shoreline adaptation strategies.

This revised version restores the missing context and clarifies the challenges and gaps that motivate our study.

*How is the training done for different SLRs?*

**Response:**

We thank the reviewer for the question regarding the training process for different SLR scenarios. Our dataset includes two regions: Abu Dhabi and San Francisco. For Abu Dhabi, we had simulation data available for a single SLR level: 0.5 meters. For San Francisco, we had data corresponding to three SLR levels: 0.5 m, 1.0 m, and 1.5 m. Among these, the 1.0 m SLR scenario had sufficient simulation samples, while the other two had relatively limited data.

To ensure stable and effective training, we used all available 0.5 m data from Abu Dhabi and the 1.0 m data from San Francisco for the primary training phase. This dataset was split into training, validation, and test sets to train and initially evaluate the model.

After training, we performed fine-tuning using a subset of the limited samples available for the 0.5 m and 1.5 m SLR scenarios in San Francisco. The purpose of this step was to adapt the model and assess its ability to generalize across different SLR levels. The final evaluation for these additional SLR depths was then conducted on the remaining, unseen samples.

This approach allowed the model to effectively learn from the most data-rich scenarios while still extending its predictive capability to other SLR conditions through targeted fine-tuning.

*Figure 1 can be moved to Methodology and please describe the overall framework and the steps involved there.*

**Response:**

Thank you for the suggestion. In the revised manuscript, we will move Figure 1 to the Methodology section and enhance the accompanying description to clearly outline the overall framework and the key steps involved in the proposed approach.

*Study Area and Data Description*

*Line 136, what do you mean by environmental effects?*

**Response:**

Thank you for the suggestion. In the revised manuscript, we will clarify by stating that the environmental impacts are mostly focused on the effects on transportation links, and specifically whether important arterials such as shoreline highways or freeways will be flooded due to SLR.



*Briefly describe OLU discretization in a few sentences for the benefit of readers.*

**Response:**

Thank you for the question. In the revised manuscript, we will clarify that for San Francisco Bay, the discretization of the Operational Landscape Units (OLUs) was taken from SFEI, 2019. For Abu Dhabi, the urban coastline was also discretized into OLUs based on the Abu Dhabi Plan 2030 (Abu Dhabi Urban Planning Council, 2007). Maps of the OLU boundaries used for both San Francisco Bay and Abu Dhabi were shown in Section S1 in the supplementary Materials Section.

*Mention past flood damages caused by tides.*

**Response:**

Thank you for the question. Our focus on tidal flooding within San Francisco Bay was to highlight the unique tidal behavior within the Bay where the construction of sea walls along certain portions of the shoreline Bay may in fact exacerbate the sea level in other areas of the Bay. For example, Hollemann and Stacey, 2014 showed that protecting the southern portion of the Bay may increase maximum water levels by 0.2 in the north of the Bay; Wang et al., 2018 showed that protection of the southern portions of the West Bay adversely affects the water levels of the East Bay, and vice-versa.

*Briefly mention here about the training dataset.*

**Response:**

We appreciate the suggestion to provide more information about the training dataset. In the revised manuscript, we will briefly elaborate the dataset in the Study Area and Data Description section as suggested.

*How was the Delft 3D model validated? How accurate is it?*

**Response:**

For San Francisco Bay, the Delft3D model was adapted from the CoSMoS model originally developed by Barnard (2014) and adapted to San Francisco Bay by Wang et al (2017), and validated in the past using tidal gages at 9 tidal gage locations in and around San Francisco Bay. Pearson correlation coefficients ranged from 0.9862 to 0.9996, while the RMS ratios (the ratio of modelled versus measured RMS amplitudes) ranged from 0.973 to 1.027 (please refer to Wang et al., 2017).

For Abu Dhabi, the Delft3D model was validated using water level data from 196 tidal gage locations throughout the Persian Gulf (as the hydrodynamic model encompassed the entire Gulf in addition to the western portions of the Gulf of Oman). The water levels at these locations were compared with one month's worth of hydrodynamic simulation, and the resulting absolute RMSE values ranged from 0.0013 to 0.0043 m in the vicinity of Abu Dhabi. More validation details for Abu Dhabi can be found in Chow and Sun, 2022.

*Rather than points, use line or raster to areas susceptible to flooding. Figure 2a.*

**Response:**

Thank you for the suggestion. As recommended, we will revise Figure 2a to replace the point-based visualization with a line or raster representation that more clearly delineates areas

susceptible to flooding. This will improve readability and better convey spatial flood risk patterns across the region.

*Table 1 – What does main set, hold out set, etc., mean? Please explain when you have a table or figure in the manuscript for a general reader.*

**Response:**

Thank you for the comment. In the revised manuscript, we will add a clear explanation of the terms used in Table 1 for general readers. Specifically, we will clarify that the main set is used for training, validation, and testing of the model; the holdout set includes manually selected challenging scenarios to evaluate model performance under complex protection configurations; and the generalizability set consists of San Francisco data with two additional sea level rise depths, used to fine-tune and assess the model's ability to generalize to new conditions not originally trained on.

*What are Shamal winds? Describe in a sentence or two with relevant literature.*

**Response:**

Thank you for the question. While we have referenced Langodan et al., 2023 in Section 3.1.1 relating to the Shamal winds, we will place an explanation of these winds in the main body of the revised text. While the Persian Gulf does not typically experience tropical cyclones, it is known for its northwesterly winds generally occurring with winds at about 20 m/s with sudden onset and sustained over a period of up to 3-5 days. These are called the Shamal winds (meaning “North” in Arabic) and occur at least 10 times annually, mainly during the winter months (Senafi et al., 2015; Li et al., 2020).

*It is unclear what you mean by three months. Period of simulation or the computational time itself? What is the significance of choosing this?*

**Response:**

Thank you for the question. A three-month period of simulation was chosen for both the San Francisco and the Abu Dhabi simulations. This was chosen to balance the need to model a larger number of tidal cycles that are modeled per simulation, with the computational time and storage space used for the simulations.

*Method*

*Section 4.1, please shorten this to retain relevant general information. The rest can be moved to supplementary so that the readers don't lose interest.*

**Response:**

Thank you for the suggestion. As recommended, we will shorten Section 4.1 to retain only the most relevant general information about the model architecture. The more detailed architectural components and implementation specifics will be moved to the supplementary material to maintain reader engagement while keeping the core content concise.

*Describe a spatial fitness metric.*

**Response:**

Thank you for the comment. In the revised manuscript, we will include a spatial fitness metric to evaluate the alignment between predicted and ground truth flood extents, specifically assessing spatial accuracy in terms of overprediction, underprediction, and correctly matched inundated areas.

## *Results*

*Tabulate and discuss the values of the spatial fitness index for different DL/ML models.*

### **Response:**

Thank you for this valuable suggestion. In the revised manuscript, we plan to compute and tabulate the spatial fitness metric for our proposed CASPIAN-v2 model, along with any baseline DL/ML models evaluated in the study. We will include the resulting Table in the main text and briefly discuss the findings to highlight how CASPIAN-v2 performs in terms of spatial prediction accuracy relative to alternative approaches.

*Table 5, please describe the performance of the proposed method compared to the existing methods in the paper.*

### **Response:**

Thank you for the suggestion. In the revised manuscript, we will include a discussion describing the performance of the proposed method relative to existing approaches as summarized in Table 5. However, as most existing ML/DL-based flood prediction methods in the literature have not open-sourced their code and are often designed for point-based or aggregated outputs rather than 2D gridded flood maps, direct benchmarking remains challenging. Nonetheless, we will highlight key differences in methodology and performance trends to contextualize the advantages of CASPIAN-v2.

*Please use a zoomed-in figure to illustrate the effect of the representation of the protection measure in the manuscript.*

### **Response:**

Thank you for the suggestion. In the revised manuscript, we will include a zoomed-in figure, as part of the results section (e.g., Figures 5 and 6), to illustrate the effect of the representation of the protection measures. This will help highlight how different shoreline protection configurations influence localized inundation patterns and improve the interpretability of the results.

*Section 6.2.2, Lines 475 – 478, lacks a clear explanation.*

### **Response:**

Thank you for the comment. In the revised manuscript, we will clarify that the referenced abrupt changes in local flood behavior relate to spatial variations in inundation caused by diverse OLU protection scenarios, as illustrated in Figure 6. The figure presents results on a holdout set specifically curated to include challenging scenarios, such as configurations where one side of a Bay is protected while the other is not. These cases test the ability of the model to capture complex spatial dynamics. We will also highlight that the model demonstrates strong generalization to such unseen scenarios, even though the training data covers only a small subset of the possible  $2^{OLUs}$

combinations. This indicates that CASPIAN-v2 learns to infer inundation patterns based on the spatial protection configuration of OLU's.

*Section 6.2.4, a repeated heading, please be more specific and clearer.*

**Response:**

Thank you for pointing this out. In the revised manuscript, we will revise the headings within Section 6.1 and 6.2 to avoid repetition and improve clarity.

*Line 494, how is the ground truth information collected? What is the associated error?*

**Response:**

Thank you for this important question. In the revised manuscript, we will clarify that the ground truth inundation data used for training and evaluation were derived from physics-based hydrodynamic simulations. We will also provide information on the validation of these simulations and discuss the expected level of error or uncertainty in the simulation outputs, where available. This clarification will be included near Line 494 to improve transparency regarding data quality.

*Line 506, summarize in a few sentences about data augmentation.*

**Response:**

Thank you for the suggestion. In the revised manuscript, we will summarize the data augmentation strategy in a few concise sentences, as suggested.

*Figure 9, please move it to the supplementary.*

**Response:**

Thank you for the suggestion. In the revised manuscript, we will move Figure 9 to the supplementary material, as recommended.

*References*

1. Bates, P. D., Horritt, M. S., & Fewtrell, T. J. (2010). A simple inertial formulation of the shallow water equations for efficient two-dimensional flood inundation modelling. *Journal of Hydrology*, 387(1–2), 33–45. <https://doi.org/10.1016/j.jhydrol.2010.03.027>
2. De Almeida, G. A. M., & Bates, P. (2013). Applicability of the local inertial approximation of the shallow water equations to flood modeling. *Water Resources Research*, 49(8), 4833–4844. <https://doi.org/10.1002/wrcr.20366>
3. Li, Z., & Hodges, B. R. (2019). Modeling subgrid-scale topographic effects on shallow marsh hydrodynamics and salinity transport. *Advances in Water Resources*, 129, 1–15. <https://doi.org/10.1016/j.advwatres.2019.05.004>
4. Neal, J., Schumann, G., & Bates, P. (2012). A subgrid channel model for simulating river hydraulics and floodplain inundation over large and data sparse areas. *Water Resources Research*, 48(11), 1–16. <https://doi.org/10.1029/2012WR012514>

5. Nithila Devi, N., & Kuiry, S. N. (2024). A novel local-inertial formulation representing subgrid scale topographic effects for urban flood simulation. *Water Resources Research*, 60(5), e2023WR035334.
6. Sanders, B. F., & Schubert, J. E. (2019). PRIMo: Parallel raster inundation model. *Advances in Water Resources*, 126, 79–95. <https://doi.org/10.1016/J.ADVWATRES.2019.02.007>

#### **Response:**

Thank you for sharing these references. We will incorporate the relevant literature, as suggested, into the revised manuscript to strengthen the discussion on computationally efficient flood modeling approaches and better contextualize the role of deep learning in this domain.

#### **Response References:**

1. Abu Dhabi Urban Planning Council. Plan Abu Dhabi 2030. Urban Structure Framework Plan. 2007. Available online: <https://www.ecouncil.ae/PublicationsEn/plan-abu-dhabi-full-version-EN.pdf>
2. Al Senafi, F.; Anis, A. Shamals and climate variability in the Northern Arabian/Persian Gulf from 1973 to 2012. *Int. J. Clim.* 2015, 35, 4509–4528.
3. Chow, A.C. and Sun, J., 2022. Combining Sea level rise inundation impacts, tidal flooding and extreme wind events along the Abu Dhabi coastline. *Hydrology*, 9(8), p.143.
4. Holleman, R.C. and Stacey, M.T., 2014. Coupling of sea level rise, tidal amplification, and inundation. *Journal of Physical Oceanography*, 44(5), pp.1439-1455.
5. IPCC. Summary for Policymakers. In *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*; Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S.L., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M.I., et al., Eds.; Cambridge University Press: Cambridge, UK, 2021.
6. Li, D.; Anis, A.; Al Senafi, F. Physical response of the Northern Arabian Gulf to winter Shamals. *J. Mar. Syst.* 2020, 203, 103280.
7. San Francisco Estuary Institute (SFEI), 2019 (Beagle, J., Lowe, J., McKnight, K., Safran, S., Tam, L. and Szambelan, S.J., authors). San Francisco Bay shoreline adaptation atlas: Working with nature to plan for sea level rise using operational landscape units (No. SFEI publication# 915).
8. Wang, R.Q., Herdman, L.M., Erikson, L., Barnard, P., Hummel, M. and Stacey, M.T., 2017. Interactions of estuarine shoreline infrastructure with multiscale sea level variability. *Journal of Geophysical Research: Oceans*, 122(12), pp.9962-9979.
9. Wang, R.Q., Stacey, M.T., Herdman, L.M.M., Barnard, P.L. and Erikson, L., 2018. The influence of sea level rise on the regional interdependence of coastal infrastructure. *Earth's Future*, 6(5), pp.677-688.