



# HIDRA3: a robust deep-learning model for multi-point ensemble sea level forecasting

Marko Rus<sup>1,2</sup>, Hrvoje Mihanović<sup>4</sup>, Matjaž Ličer<sup>1,3,★</sup>, and Matej Kristan<sup>2,★</sup>

<sup>1</sup>Slovenian Environment Agency, Office for Meteorology, Hydrology and Oceanography, Ljubljana, Slovenia

<sup>2</sup>Faculty of Computer and Information Science, Visual Cognitive Systems Lab, University of Ljubljana, Ljubljana, Slovenia

<sup>3</sup>National Institute of Biology, Marine Biology Station, Piran, Slovenia

<sup>4</sup>Institute of Oceanography and Fisheries, Split, Croatia

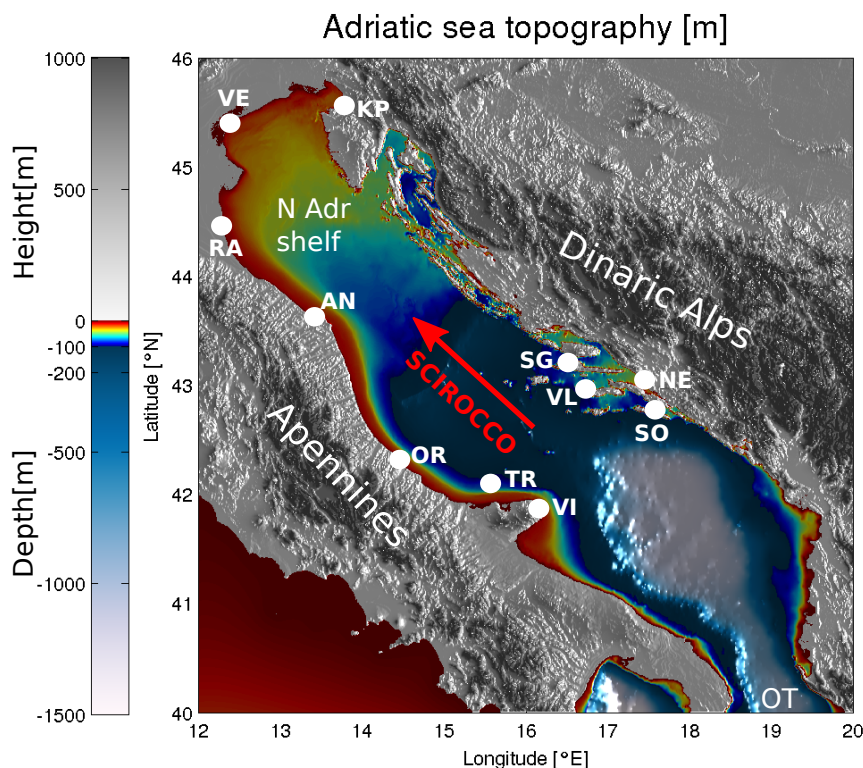
★These authors contributed equally to this research.

**Correspondence:** Marko Rus (marko.rus@fri.uni-lj.si)

**Abstract.** Accurate modeling of sea level and storm surge dynamics with several day-long temporal horizons is essential for effective coastal flood response and the protection of coastal communities and economies. The classical approach to this challenge involves computationally intensive ocean models that typically calculate sea levels relative to the geoid, which must then be correlated with local tide gauge observations of sea surface heights (SSH). A recently proposed deep learning model, HIDRA2, avoids numerical simulations while delivering competitive forecasts. Its forecast accuracy depends on the availability of a sufficiently large history of recorded SSH observations used in training. This makes HIDRA2 less reliable for locations with less abundant SSH training data. Furthermore, since the inference requires immediate past SSH measurements at input, forecasts cannot be made during temporary tide gauge failures. We address the aforementioned issues with a new architecture, HIDRA3, that considers observations from multiple locations, shares the geophysical encoder across the locations, and constructs a joint latent state, which is decoded into forecasts at individual locations. The new architecture brings several benefits: (i) it improves training at locations with scarce historical SSH data, (ii) it enables predictions even at locations with sensor failures, and (iii) it reliably estimates prediction uncertainties. HIDRA3 is evaluated by jointly training on eleven tide gauge locations along the Adriatic. Results show that HIDRA3 outperforms HIDRA2 and the standard numerical model NEMO by  $\sim 15\%$  and  $\sim 13\%$  MAE reduction at high SSH values, respectively, setting a solid new state-of-the-art. Furthermore, HIDRA3 shows remarkable performance at substantially smaller amounts of training data compared with HIDRA2, making it appropriate for sea level forecasting in basins with large regional variability in the available tide gauge data.

## 1 Introduction

Sea surface height (SSH) forecasting is a critical modeling task for two primary practical reasons. Elevated sea levels can result in severe flooding of densely populated coastal towns, while lower sea levels may restrict marine cargo traffic from navigating shallow coastal regions like northern Adriatic, where depths often fall below 15 meters (see Fig. 1). In the Adriatic basin, the challenges posed by high and low sea levels are dynamically distinct. Elevated sea levels are associated with intense pressure



**Figure 1.** Topography and bathymetry of the Adriatic region. White dots denote tide gauges used in this study. Abbreviations used on the map are as follows: KP - Koper, VE - Venice, RA - Ravenna, AN - Ancona, N Adr Shelf - Northern Adriatic Shelf, OR - Ortona, TR - Tremiti, VI - Vieste, SO - Sobra, SG - Stari Grad, VL - Vela Luka, OT - Otranto strait. The direction of Scirocco is marked with the red arrow. The image was created by the authors based on General Bathymetric Chart of the Oceans (GEBCO) bathymetry and elevation data, available at <https://download.gebco.net/> (last access: 14 June 2024).

lows and strong winds typical of cyclonic activity in the basin, whereas negative sea level extremes arise during periods of spring tides, combined with high atmospheric pressure due to anticyclonic presence.

Numerical general circulation models are commonly used for SSH prediction and provide basin-scale 4-dimensional simulations of the full sea state (temperatures, salinities, circulation, sea levels), whose evolution is nontrivial and numerically demanding (Umgiesser et al., 2022; Ferrarin et al., 2020; Madec, 2016). Furthermore, sea level information in these models typically corresponds to a minor part of the total numerical cost of these simulations, which often involve baroclinic and other processes that have very limited relevance for coastal flood predictions. A further complication arises with the fact that semi-enclosed basins exhibit meteorologically induced standing oscillations or seiches which get superimposed on the tidal sea level signal. This leaves the total sea level highly non-linearly dependent on the time lag between peak tide and peak seiche. A possible remedy to this sensitivity is to employ probabilistic or ensemble modeling (Rus et al., 2023; Bernier and Thompson,



2015; Mel and Lionello, 2014) with the hope of efficiently reproducing the envelope of possible sea levels. This approach, however, comes with a high numerical cost, making it unfeasible to run ensemble predictions multiple times a day at national meteocean services.

35 To limit numerical cost, one effective option is to use 2-dimensional barotropic models, which have demonstrated efficacy in sea level and ensemble modeling (Bajo et al., 2023; Ferrarin et al., 2023, 2020). Recently, machine learning has emerged as another promising option for reducing numerical costs while enhancing performance. Early models (Imani et al., 2018) used traditional techniques such as Support Vector Machines (SVMs) (Vapnik, 1999), while Ishida et al. (2020) employed Long Short-Term Memory networks (LSTMs) (Hochreiter and Schmidhuber, 1997) with atmospheric variables for one-hour  
40 predictions. Braakmann-Folgmann et al. (2017) extended prediction horizon using a combination of LSTMs and Convolutional Neural Networks (CNNs). Autoregressive neural networks were explored in Hieronymus et al. (2019), increasing temporal resolution and the prediction horizon. While these methods were orders of magnitude faster than their numerical counterparts, their accuracy was generally inferior.

A decisive improvement of accuracy was possible through the employment of deep convolutional architectures from com-  
45 puter vision, such as models from the HIDRA modeling family (Žust et al., 2021; Rus et al., 2023). HIDRA1 and HIDRA2 networks avoid the need for basin-scale simulations of full ocean state and focus the network capacity to only predict sea levels at a single geographic location. HIDRA2 (Rus et al., 2023) has been shown to substantially outperform several numerical models in speed as well as in accuracy of sea level predictions in Koper (northern Adriatic) and is currently implemented also in the Baltic basin on Estonian coasts (Barzandeh et al., 2024).

50 The HIDRA2 model currently provides SSH forecasts with a 6-day temporal horizon. Its inputs consist of atmospheric winds and pressures as well as SSH observations in the past 24 hours at the prediction location. This means that a sufficiently large history of measured SSH and atmospheric values is required to train HIDRA2 at each geographic location separately. While the atmospheric values are available from sources such as ECMWF (Leutbecher and Palmer, 2007), the SSH values can only be obtained from in-situ tide gauges. This brings two main drawbacks. First, for some locations, a long history of SSH  
55 measurements is not available, which limits the training capability of HIDRA2, which has to learn not only the SSH information encoding but also the relevant atmospheric features, the input data fusion process and the prediction heads. Second, in case of sensor failure, immediate past SSH values are unavailable, and predictions cannot be made.

To address these limitations, we introduce HIDRA3, a new deep-learning architecture designed to simultaneously learn from and predict SSH across multiple tide gauge locations. This new multi-location formulation is the main contribution of  
60 this paper. HIDRA3 can further handle missing values in observation datasets at input locations and implicitly reconstruct them in latent space from other available locations. In addition to SSH values, HIDRA3 also predicts SSH forecast uncertainties, which enhances the interpretability of the output.

HIDRA3 thus brings forward several benefits. Firstly, the model utilizes information from multiple locations to improve prediction accuracy at each individual location. Since a single HIDRA3 is trained for all locations, many parameters are  
65 shared between the locations so that locations with scarce historical training data benefit from locations with abundant data. Additionally, predictions can be made even during sensor failures, as HIDRA3 exploits all available observed information at



other locations to obtain the predictions at the locations with the missing SSH values. Lastly, due to its numerical efficiency, HIDRA3 enables fast and reliable ensemble sea level forecasting multiple times a day. The remainder of the paper is organized as follows. Section 2 details the new HIDRA3 architecture, Sect. 3 compares HIDRA3 with state-of-the-art numerical and machine-learning models, while conclusions are drawn in Sect. 4.

## 2 HIDRA3 deep learning sea level model

### 2.1 HIDRA3 training and testing datasets

Our objective is to forecast hourly SSH values for  $N = 11$  tide gauges located along the Adriatic coast (Fig. 1) over a three-day period. HIDRA3 achieves this by leveraging a comprehensive set of ocean state parameters. This includes the past 72 hours of available sea level observations from stations shown in Fig. 1, with data availability for each station detailed in Table 1. Additionally, it considers both past and future astronomic tides at these stations, and past and future 72 hours of gridded geophysical variables from atmospheric and ocean numerical models.

Location	SSH Availability in 2000–2022	Thresholds [cm]
Koper	95.0%	-69.3, 65.7
Venice	64.6%	-64.3, 61.3
Ancona	50.4%	-39.9, 44.6
Vieste	48.9%	-33.3, 36.5
Ortona	48.3%	-34.0, 39.6
Ravenna	43.4%	-56.3, 57.2
Neretva	38.9%	-32.6, 37.8
Sobra	24.1%	-33.4, 37.0
Stari Grad	23.9%	-34.0, 38.7
Tremiti	20.5%	-32.4, 37.0
Vela Luka	16.6%	-31.9, 38.6

**Table 1.** Availability of SSH measurements between 2000 and 2022 for 11 tide gauge locations used in training and evaluating HIDRA3, and defined thresholds [1<sup>st</sup>, 99<sup>th</sup> percentile] for low and high SSH values used in this study. See Fig. 1 for station locations.


The training time window spans two intervals: from 2000 to 2018 and from 2021 to 2022. The testing time window covers the period from June 2019 to the end of 2020. This specific testing interval was selected due to the occurrence of numerous high SSH and coastal flood events in the northern Adriatic. During training, we assume that the past 72 h of SSH observations are available for at least one tide gauge in the training set, while they may be missing at several or all other locations. Note that very different amounts of tide gauge data are available at different stations; particularly, all stations situated south of Ancona have less than 50% data availability throughout the training interval (refer to Table 1).



The data collected from the tide gauges were subjected to detailed manual quality checks. Astronomic tides were calculated from tide gauge data in one-year intervals using the UTIDE Tidal Analysis package for Python (Codiga, 2011).

In the context of this analysis, a dataset of sea-level extremes was constructed for each station. SSH readings are categorized as *low* if they are below the 1<sup>st</sup> percentile, and as *high* if they surpass the 99<sup>th</sup> percentile of the observed values at the respective location. The thresholds determined are listed in Table 1.

Tide gauges in the northern Adriatic (KP, VE, RA) and those in the middle Adriatic (OR, TR, VI, SO, VL, NE, SG) form two separate clusters that fall within the barotropic Rossby radius of each other and thus exhibit similar SSH phases of high and low sea levels. This is illustrated in Fig. 2, which depicts the SSH at all stations, and in Fig. 3, which shows mean absolute differences between all stations.




**Figure 2.** A representative example period of SSH from all 11 tide gauges used in training and evaluation of HIDRA3.

ERA5 reanalysis data (Hersbach et al., 2018) was employed for training purposes, while ECMWF Ensemble Prediction System data (Leutbecher and Palmer, 2007) was employed for evaluation to reflect the practical forecasting setup in which future reanalysis does not exist and forecasts are used. The metocean model training dataset used in this study consists of 10-meter winds, mean sea level air pressure, sea surface temperature (SST), mean wave direction, mean wave period and significant height of combined wind waves and swell. All atmospheric model input fields were spatially cropped to the Adriatic basin and subsampled to a  $9 \times 12$  spatial grid, following our previous work (Rus et al., 2023).

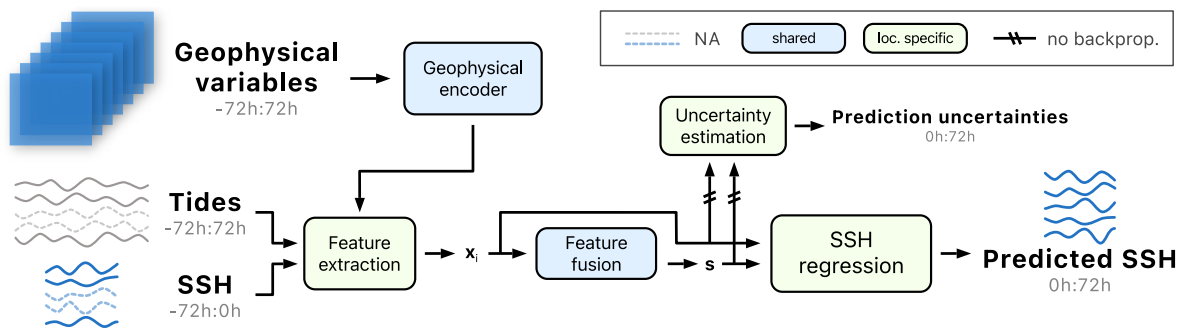
## 2.2 HIDRA3 model architecture

HIDRA3 proceeds as follows (see Fig. 4 for the architecture). Geophysical variables (wind, air pressure, SST, waves) are encoded once for all locations in the Geophysical encoder module (Sect. 2.2.1), while the sea level data (full SSH and astronomic tide) and geophysical features are processed separately for each location in the Feature extraction module (Sect. 2.2.2). Features are fused in the Feature fusion module (Sect. 2.2.3), which takes sea-level features from locations with available data and



**Figure 3.** Mean absolute differences [cm] of SSH measurements between different tide gauge locations. Abbreviations used here are: KP - Koper, VE - Venice, RA - Ravenna, AN - Ancona, OR - Ortona, TR - Tremiti, VI - Vieste, SO - Sobra, VL - Vela Luka, NE - Neretva, and SG - Stari Grad.

encodes them into a single joint state feature vector  $\mathbf{s}$ , which is decoded by the SSH regression head (Sect. 2.2.4) into the SSH predictions at the individual geographic locations. Finally, the Uncertainty estimation module predicts forecast uncertainties (Sect. 2.2.5).




**Figure 4.** The HIDRA3 architecture. The Geophysical encoder encodes pressure, wind, SST and waves. The Feature extraction module extracts features separately for each location, creating feature vectors  $\mathbf{x}_i$ , which are then fused in the Feature fusion module. The SSH regression module produces SSH forecasts, while the Uncertainty estimation module predicts their standard deviations. The location-specific modules and the modules shared among the locations are depicted in green and blue, respectively. Missing SSH data is denoted by a dashed curve. The notation  $a:b$  indicates hourly data points from the interval  $(a, b]$ , while the prediction point is at the index 0.



## 2.2.1 Geophysical encoder module

The Geophysical encoder takes past and future, i.e. -72:72 h in hourly steps, of geophysical variables as the input: wind ( $\mathbf{v} = (u_{10}, v_{10}) \in \mathbb{R}^{2 \times W \times H}$ ), pressure ( $\mathbf{p} \in \mathbb{R}^{W \times H}$ ), sea surface temperature ( $\mathbf{T} \in \mathbb{R}^{W \times H}$ ), and waves ( $\mathbf{w} \in \mathbb{R}^{4 \times W \times H}$ ), where  $W = 9$  and  $H = 12$  are the spatial dimensions of the resampled input fields. The waves tensor  $\mathbf{w}$  is composed of the following four components (hence the dimension  $4 \times W \times H$ ): the first two are sine and cosine encodings of the mean wave direction, while the remaining two components are the mean wave period and significant height of combined wind waves and swell.


The input variables are processed in two steps. In the first step, variables (wind, pressure, SST and waves) are encoded separately, and then in the second step, encodings from different variables are fused together. The first step is shown in Fig. 5. The encoder processes variable of size  $c \times 144 \times 9 \times 12$  by a 3D convolution block with  $64 \times 2 \times 3 \times 3$  kernels with stride  $2 \times 2 \times 2$  (the first dimension is temporal, while the second two are spatial), followed by a ReLU, a 3D dropout (Tompson et al., 2014) and another 3D convolution with  $512 \times 2 \times 4 \times 5$  kernels with stride  $2 \times 1 \times 1$ . Note that the strides and the kernel sizes are adjusted to reduce the dimensions to  $512 \times 36 \times 1 \times 1$ . For the sea surface temperature and wave variables, the number of kernels in the last convolutional layer is 64 instead of 512, forming the output of size  $64 \times 36 \times 1 \times 1$ . This modification has led to a marginal enhancement in performance, presumably due to the larger number of features representing wind and pressure compared to SST and waves, thereby assigning greater significance to the former variables, which hold more relevance in the context of SSH forecasting.



**Figure 5.** The first step of the Geophysical encoder module involves encoding each of the four variables separately. Note that wind has two input channels, while the waves data has four. The encoder consists of two 3D convolutions, reducing spatial dimension to  $1 \times 1$ . The variables  $k$ ,  $s$  and  $n$  denote the kernel size, the stride and the number of output channels, respectively. The number of channels is indicated in gray, the size of the temporal dimension is in red, and the spatial dimensions are in blue.

In the second step (see Fig. 6), encodings from all variables are concatenated and fused by a 1D convolution with 256 kernels of temporal size 5 and two 1D convolutions with 256 kernels of temporal size 1. This is followed by a SELU activation (Klambauer et al., 2017), 1D dropout and a residual connection (see Fig. 6). After flattening, the output is a single vector of 8192 geophysical features.






**Figure 6.** The second step of the Geophysical encoder module. Geophysical variables that were encoded independently in the first step are fused to form a single vector. 1D convolutions with residual connections and nonlinearities are applied to produce  $256 \times 32$  output, which is then flattened into a single vector of all geophysical features. The variables  $k$  and  $n$  denote the kernel size and the number of output channels, respectively. The number of channels is indicated in gray, and the size of the temporal dimension is in red.

### 2.2.2 Feature extraction module

In the Feature extraction module (see Fig. 7), location-specific features  $\mathbf{x}_i$  are extracted independently for each location. The SSH observations (72 values) and the astronomic tide (144 values) are concatenated and encoded with a dense layer (i.e. fully-connected layer) with 512 output channels. Geophysical features are also passed through a dense layer with 512 output channels. Outputs of both layers are concatenated and processed with four dense layers with 1024 output channels. A SELU, a dropout and residual connections are applied as shown in Fig. 7, yielding the enriched SSH latent representation of size 1024.



**Figure 7.** The Feature extraction module receives geophysical features, observed SSH measurements and astronomic tide forecast, which are processed independently for each location to extract location-specific features  $\mathbf{x}_i$ . The variable  $n$  denotes the number of output channels. The data dimensions are indicated in gray.





### 2.2.3 Feature fusion module

As indicated in Fig. 4, the feature fusion module combines the location-specific features  $\mathbf{x}_i \in \mathbb{R}^{1024}$ , into a unified state vector  $\mathbf{s} \in \mathbb{R}^{8192}$ . A critical design requirement for the module is robustness to missing data. Specifically, the architecture must guarantee that different tide gauges do not produce much different state vector  $\mathbf{s}$  so that subsequent modules receive consistent features, regardless of the number of operational input tide gauges.

A partial reconstruction  $\mathbf{s}_i$  of the state is computed from each location-specific feature vector  $\mathbf{x}_i$  by applying a fully-connected layer. In addition, a weight vector  $\mathbf{w}_i$  is computed by applying another fully-connected layer to  $\mathbf{x}_i$ . Each coordinate in  $\mathbf{w}_i$  reflects the extent by which the particular location contributes to the respective coordinate in the joint state vector. The joint state vector is thus computed as

$$\mathbf{s} = \sum_{i \in V} \hat{\mathbf{w}}_i \odot \mathbf{s}_i, \quad (1)$$

where  $\odot$  denotes element-wise array multiplication and  $\hat{\mathbf{w}}_i$  are the coordinate-normalized weight vectors. The components  $\hat{w}_{i,j}$  are defined computed using softmax function:

$$\hat{w}_{i,j} = \frac{e^{w_{i,j}}}{\sum_{k \in V} e^{w_{k,j}}}, \quad (2)$$

where the softmax ensures that the coordinate-wise weights sum to one across all locations.

### 2.2.4 SSH regression module

The SSH regression module receives two inputs: the joint state vector  $\mathbf{s}$  and the location-specific feature vector  $\mathbf{x}_i$ . A transformed feature vector  $\hat{\mathbf{x}}_i$  is computed by applying a dense layer with 1024 output features to  $\mathbf{x}_i$ . Note that  $\mathbf{x}_i$  does not exist in case of tide gauge failure at location  $i$ . In this case,  $\hat{\mathbf{x}}_i$  is calculated by applying a dense layer with 1024 output features to the joint state vector  $\mathbf{s}$ . Finally, the state  $\mathbf{s}$  and the transformed feature vectors  $\hat{\mathbf{x}}_i$  are concatenated and processed by a dense layer with 72 output features to produce the 72 hourly SSH predictions  $\hat{\boldsymbol{\mu}}$ .

### 2.2.5 Uncertainty estimation module

Finally, the Uncertainty estimation module predicts the hourly standard deviations corresponding to the SSH estimate  $\hat{\boldsymbol{\mu}}$ . This layer has the same architecture as the SSH regression module (Sect. 2.2.4) with a slight distinction: it applies a soft-plus function (Glorot et al., 2011) on the output to ensure the positivity of the estimated standard deviation vector  $\hat{\boldsymbol{\sigma}}$ .

## 2.3 Training and implementation details

Training is split into two phases. In the first phase, SSH mean value prediction is trained using the MSE loss between the SSH predictions and the ground truth. We use AdamW (Loshchilov and Hutter, 2017) optimizer with learning rate  $1e-5$  and weight decay of 0.001. We apply the cosine annealing (Loshchilov and Hutter, 2016) learning schedule to gradually reduce the learning rate by a factor of 100. To simulate tide gauge failures during training, a random number of tide gauges are turned



off with a probability of 0.5. The batch size is set to 128 data samples, the model is trained for 20 epochs. All input data is standardized. Training takes approximately one hour on a computer with an NVIDIA A100 Tensor Core GPU graphics card.

In the second phase, uncertainty prediction is trained. The second phase needs to have a different training dataset, which is necessary to prevent the underestimation of the uncertainty. Hence, only the year 2020 was used to train the second phase. In the second phase, the negative log-likelihood is optimized:

$$\mathcal{L} = -\frac{1}{N} \sum_{i=1}^N \sum_{t=1}^{72} \log \frac{1}{\hat{\sigma}_{i,t}} \exp \left[ -\frac{(\mu_{i,t} - \hat{\mu}_{i,t})^2}{\hat{\sigma}_{i,t}^2} \right], \quad (3)$$

where  $\mu_{i,t}$  is the ground truth SSH value for the  $i$ -th tide gauge at prediction lead time  $t$ ,  $\hat{\mu}_{i,t}$  is the predicted SSH value, and  $\hat{\sigma}_{i,t}$  is the predicted standard deviation. The second phase applies a learning rate of  $1e-4$  for 40 epochs; batch size, weight decay, and learning rate schedule are the same as in the first stage.

In operational forecasting and for the evaluation, an ensemble of  $n_{\text{ens}} = 50$  atmospheric members is utilized. HIDRA3 processes each member individually, yielding 50 SSH forecasts and associated standard deviations at each time step. These are merged into a single probabilistic forecast through Gaussian moment matching (Kristan and Leonardis, 2013). The prediction is the average of all ensemble forecasts, while the variance is computed as follows:

$$\bar{\sigma}_t^2 = \frac{1}{n_{\text{ens}}} \sum_{i=1}^{n_{\text{ens}}} (\hat{\sigma}_{t,i}^2 + \hat{\mu}_{t,i}^2) - \left( \frac{1}{n_{\text{ens}}} \sum_{i=1}^{n_{\text{ens}}} \hat{\mu}_{t,i} \right)^2, \quad (4)$$

where  $\hat{\mu}_{t,i}$  denotes the SSH forecast at time  $t$  for ensemble member  $i$ , and  $\hat{\sigma}_{t,i}$  represents the corresponding standard deviation estimate.

### 2.3.1 Network initialization

Considering the complex information flow from many variables that exist on different scales and are represented by latent descriptors of different sizes in the network, care is required at training-time parameters initialization. The weights of convolutional and dense layers are initialized using a standard Xavier initialization (Glorot and Bengio, 2010), while biases are initialized using a normal distribution with a standard deviation of 0.1. In the Feature extraction module, the deeper layers are given lower initial weights, assigning less significance to complex features during the initialization phase. Specifically, for each of the four recurrent dense layers in the Feature extraction module, the weight and bias are scaled by a factor of  $0.5^{i-1}$ , where  $i$  represents the layer's position in the sequence, ranging from 1 to 4.

In the SSH regression module and the Uncertainty estimation module, an additional step of weight scaling is implemented during initialization. In the module  $\mathbf{s}$  and  $\hat{\mathbf{x}}_i$ , which have different dimensions, are concatenated. To ensure that the contributions of  $\mathbf{s}$  and  $\hat{\mathbf{x}}_i$  to the final output are initially proportionally balanced, the ratio of their sizes is used to scale the weights of the final dense layer in the SSH regression module and the Uncertainty estimation module.



### 190 3 Results

We compare HIDRA3 with the state-of-the-art deep model HIDRA2 (Rus et al., 2023) and with the standard Copernicus Marine Environment Monitoring Service (CMEMS) product MEDSEA\_ANALYSISFORECAST\_PHY\_006\_013 (Clementi et al., 2021) numerical model NEMO (Madec, 2016).

195 Since NEMO predicts the full ocean state, including SSH, on a regular grid, the respective tide gauge locations are approximated by the nearest-neighbor locations in the grid. NEMO forecasts are post-processed by standard offset adjustment of the initial 12 hours forecast to obtain a zero-bias compared to the respective tide gauge, as in Rus et al. (2023).

#### 3.1 SSH forecast performance

200 The following performance measures (Rus et al., 2023) are employed: mean absolute error (MAE), root mean squared error (RMSE), accuracy (ACC), bias, recall (Re), precision (Pr) and F1 score. These performance metrics are reported in Table 2 separately for all SSH values (*overall*) and for low and high SSH values (see Sect. 2 for the definitions).

	Model	MAE (cm)	RMSE (cm)	ACC (%)	Bias (cm)	Re (%)	Pr (%)	F1 (%)
Overall	NEMO	2.65	3.56	97.76	-0.31	/	/	/
	HIDRA2	2.63	3.56	98.15	-0.17	/	/	/
	HIDRA3 (ours)	<b>2.42</b>	<b>3.28</b>	<b>98.60</b>	<b>-0.00</b>	/	/	/
Low SSH Values	NEMO	4.19	5.23	92.91	2.88	/	/	/
	HIDRA2	<b>3.27</b>	4.27	95.94	<b>1.02</b>	/	/	/
	HIDRA3 (ours)	3.30	<b>4.24</b>	<b>96.16</b>	1.33	/	/	/
High SSH Values	NEMO	4.68	6.19	89.14	-3.02	94.53	<b>99.40</b>	96.79
	HIDRA2	4.80	6.53	89.49	-2.35	96.62	97.82	97.18
	HIDRA3 (ours)	<b>4.06</b>	<b>5.61</b>	<b>91.63</b>	<b>-2.06</b>	<b>97.58</b>	98.67	<b>98.09</b>


**Table 2.** Performance calculated on all SSH values, low SSH values and high SSH values, averaged over all locations. The proposed HIDRA3 has the best performance overall and on high SSH values, and a comparable performance on low values to HIDRA2.

205 Table 2 shows that HIDRA3 outperforms both competing models across all metrics computed for all SSH values and for high SSH values. Notably, HIDRA3 achieves an 8.0% lower MAE compared with HIDRA2 and an 8.7% lower MAE than NEMO. On high SSH values, HIDRA3 achieves a 15.4% lower MAE than HIDRA2 and 13.2% lower MAE than NEMO. These trends are consistent across other evaluation metrics. Furthermore, HIDRA3 demonstrates a near-zero bias and achieves the highest recall among the models. While NEMO has slightly higher precision, this comes at the cost of a lower recall, implying missed flood events, while HIDRA3 exhibits a solid balance between precision and recall, as evidenced by its highest F1 score.

Figure 8 shows the MAE separately for each tide gauge location. Prediction is most challenging for the northern locations, which yield comparatively larger errors than other locations. This is likely due to shallow waters and topographic amplifications



in the northern Adriatic. In these conditions, however, HIDRA3 demonstrates the most significant improvement in MAE when compared with NEMO. In Koper, the MAE is lower by 22.6% and 27.7% on high SSH values, in Venice by 19.0% and 22.2% on high SSH values, and in Ravenna by 13.9% and 30.0% on high SSH values. On the other hand, southern locations show much smaller and comparable MAE errors with both NEMO and HIDRA3, which is also due to overall lower sea level variability in the deeper southern part of the basin with less topographic amplification.




**Figure 8.** The Overall MAE and MAE on high SSH values calculated for different models and tide gauge locations. HIDRA3 significantly outperforms HIDRA2 (Rus et al., 2023), and also outperforms NEMO (Madec, 2016) at northern locations.

Figure 9 and Fig. 10 visualize December 2019 flooding predictions for Venice and Stari Grad, respectively. The first 24 h of daily predictions were concatenated for each model to get a single time series. In Venice (Fig. 9), HIDRA3 predictions most accurately match the ground truth observations. For instance, NEMO fails to predict the December 14<sup>th</sup>–16<sup>th</sup> floods, while HIDRA3 captures them very well. Also, during the peaks in the second part of the month, HIDRA3 is most accurate.


In Stari Grad (Fig. 10), the first peak is missed by all models, while subsequent peaks above the high SSH threshold, as well as low water levels, are predicted most accurately by HIDRA3. NEMO and HIDRA2 often underestimate the range of sea level variability in Stari Grad, with maximums being too low and minimums being too high. HIDRA3, on the other hand, produces a solid forecasting result, even though the training data availability in Stari Grad is merely 23.9% (see Table 1).

### 3.2 Evaluation under tide gauge failures

To evaluate the robustness of HIDRA3 to tide gauge failures, we simulated the failures at each location by making the SSH measurements at a respective location unavailable at prediction time. In such a scenario, HIDRA2 (Rus et al., 2023) is unable to produce predictions. For NEMO, a tide gauge failure primarily means that the bias offset adjustment, which translates predicted SSH (over geoid) into the total water level, is not available. Consequently, we offset the NEMO SSH at the location of the failed tide gauge using the average bias of all the non-failed locations. We refer to this offset version of NEMO results as NEMO<sub>0</sub>.




**Figure 9.** Comparison of the HIDRA3, HIDRA2 and NEMO predictions with the ground truth (black line) on the December 2019 floods in Venice. The high SSH threshold is marked with a red line.



**Figure 10.** Comparison of the HIDRA3, HIDRA2 and NEMO predictions with the ground truth (black line) on the December 2019 floods in Stari Grad, Hvar Island, Croatia. The high SSH threshold is marked with a red line.

The average performance over all locations with the failed tide gauges is shown in Table 3 and visualized in Fig. 11. HIDRA3 achieves a lower MAE than NEMO<sub>0</sub> across all locations, most notably in Koper. We emphasize that this is *not* equal to the



**Figure 11.** The Overall MAE and MAE on high SSH values calculated for locations with simulated sensor failures. Note that the increased NEMO<sub>0</sub> errors compared to NEMO errors in Figure 8 are due to a lower bias correction quality during periods of local sensor failures.


230 error of the NEMO model but rather a demonstration of an expected error in case of a sensor failure, where we are forced to infer offset adjustment from available stations.

Note that for the majority of the southern locations, the MAE on high SSH values is higher for HIDRA3 than NEMO<sub>0</sub>, with the differences being quite small. Nevertheless, in Koper, HIDRA3 outperforms NEMO<sub>0</sub> by a large margin and attains a lower mean MAE on high SSH values. This is also reflected in other metrics reported in Table 3.


	Model	MAE (cm)	RMSE (cm)	ACC (%)	Bias (cm)	Re (%)	Pr (%)	F1 (%)
Overall	NEMO <sub>0</sub>	3.26	4.15	95.81	<b>0.03</b>	/	/	/
	HIDRA3 (ours)	<b>2.63</b>	<b>3.52</b>	<b>98.35</b>	-0.07	/	/	/
Low SSH Values	NEMO <sub>0</sub>	4.00	4.95	<b>96.00</b>	3.08	/	/	/
	HIDRA3 (ours)	<b>3.30</b>	<b>4.26</b>	95.75	<b>1.02</b>	/	/	/
High SSH Values	NEMO <sub>0</sub>	5.12	6.48	86.82	-2.96	92.20	<b>99.81</b>	95.24
	HIDRA3 (ours)	<b>4.46</b>	<b>6.04</b>	<b>89.94</b>	<b>-2.32</b>	<b>97.38</b>	98.81	<b>98.06</b>

**Table 3.** Performance of HIDRA3 and NEMO<sub>0</sub> under the target location tide gauge failure.

235 We next simulated a pair of tide gauge failures: the one close to the river mouth of Neretva River, Croatia, and the Venice tide gauge. Figures 12 and 13 show the corresponding SSH observation time series along with the concatenation of the first 24 h of six daily forecasts. While the previous HIDRA2 cannot predict during tide gauge failures, HIDRA3 seamlessly predicts SSH values with errors lower than NEMO<sub>0</sub>.



**Figure 12.** Forecasts on December 2019 floods at the Neretva tide gauge with simulated sensor failure. The SSH measurements are shown in black, while the dotted lines indicate the failure simulations. HIDRA2 is unable to make predictions during sensor failure, while HIDRA3 continues to deliver reliable predictions. A two standard deviations band is drawn around HIDRA3 predictions to visualize the estimated uncertainty. Red lines indicate high SSH values.




**Figure 13.** Forecasts on December 2019 floods at the Venice tide gauge with simulated sensor failure. The SSH measurements are shown in black, while the dotted lines indicate the failure simulations. HIDRA2 is unable to make predictions during sensor failure, while HIDRA3 continues to deliver reliable predictions. A two standard deviations band is drawn around HIDRA3 predictions to visualize the estimated uncertainty. Red lines indicate high SSH values.





### 3.3 Influence of the training set size

240 One of the key hypothesized strengths of HIDRA3 with respect to HIDRA2 is its capability to leverage data from different  
 tide gauges, particularly when SSH data availability is limited and would severely inhibit HIDRA2 training capability. To  
 evaluate this aspect, we re-trained HIDRA3 and HIDRA2 for the Venice location with different historical window sizes while  
 keeping the data from other stations intact. The training set was gradually increased from 1 year worth of data to 10 years  
 worth of training data. The results are shown in Fig. 14, and indicate that the performance of HIDRA2 gradually improves with  
 245 increasing the training set size and converges if at least 8 years of data are available. On the other hand, HIDRA3 converges  
 rapidly and consistently outperforms HIDRA2 for all considered dataset sizes. Remarkably, even with a single year of training  
 data, HIDRA3 obtains performance comparable to when trained on the entire duration of 10 years.



**Figure 14.** MAE on high SSH values when restricting training data in Venice to a maximum of 10 years. HIDRA3 achieves substantially lower MAE than HIDRA2, particularly under limited historical training data.

### 3.4 Ablation study

#### 3.4.1 Importance of considering multiple locations


250 To further analyze the proposed architecture, we retrained HIDRA3 at each location separately (ignoring all others) to compare  
 it under the same inputs as HIDRA2 – we refer to this variant as HIDRA3<sub>1</sub>. Results in Table 4 and Fig. 15 show that, on average,  
 HIDRA3<sub>1</sub> achieves better performance than both HIDRA2 and NEMO in the overall metrics, with HIDRA3<sub>1</sub> particularly  
 excelling on the high SSH values. HIDRA3<sub>1</sub> consistently outperforms HIDRA2 across all locations, indicating the superiority  
 of the proposed architecture even when it cannot exploit the information from multiple locations.

255 Multiple locations nevertheless substantially benefit the training process, as can be seen by comparing HIDRA3<sub>1</sub> to HIDRA3.  
 The MAE of HIDRA3<sub>1</sub> is higher by 7.4% and 6.9% on high SSH values, meaning that observing multiple locations is bene-  
 ficial for the overall performance, and that HIDRA3 is able to leverage and combine the information from all stations into a  
 more accurate forecast.



	Model	MAE (cm)	RMSE (cm)	ACC (%)	Bias (cm)	Re (%)	Pr (%)	F1 (%)
Overall	NEMO	2.65	3.56	97.76	-0.31	/	/	/
	HIDRA2	2.63	3.56	98.15	-0.17	/	/	/
	HIDRA3 <sub>1</sub> (ours)	<b>2.60</b>	<b>3.47</b>	<b>98.40</b>	<b>0.02</b>	/	/	/
Low SSH Values	NEMO	4.19	5.23	92.91	2.88	/	/	/
	HIDRA2	<b>3.27</b>	<b>4.27</b>	<b>95.94</b>	<b>1.02</b>	/	/	/
	HIDRA3 <sub>1</sub> (ours)	3.52	4.47	95.58	1.79	/	/	/
High SSH Values	NEMO	4.68	6.19	89.14	-3.02	94.53	<b>99.40</b>	96.79
	HIDRA2	4.80	6.53	89.49	-2.35	96.62	97.82	97.18
	HIDRA3 <sub>1</sub> (ours)	<b>4.34</b>	<b>5.98</b>	<b>90.94</b>	<b>-1.72</b>	<b>96.82</b>	98.54	<b>97.65</b>

**Table 4.** Performance of NEMO, HIDRA2 and HIDRA3<sub>1</sub>, where HIDRA3<sub>1</sub> is the model trained separately on every single location.



**Figure 15.** MAE comparison between locations, where HIDRA3<sub>1</sub> is the model trained separately on every single location.

### 3.4.2 Impact of sea temperature and waves

260 To evaluate the importance of using sea surface temperature and wave data to enhance our model, we removed the encoders for these variables from the HIDRA3 architecture and retrained it. We observed an average 0.8% and a 2.7% increase in the overall MAE and MAE on high SSHs, respectively. These results suggest that sea surface temperature and wave data are not the main sources of the HIDRA3’s excellent performance; however, they do bring non-negligible contribution.



### 3.4.3 Importance of the skip connection

265 Finally, we quantify the impact of passing the location-specific features  $\mathbf{x}_i$  directly to the SSH regression module in addition to  
the fused features from other locations, as shown in Fig. 4 in Sect. 2.2.3. For that, we change the SSH regression module so that  
it only accepts the fused features  $\mathbf{s}$  and train the model by limiting training data from Venice to 1 year. This setup was chosen  
to showcase the performance change in locations with less historical SSH measurements. In Sect. 3.3, we have observed that  
HIDRA3 achieves MAE on high SSH values of 6.24 cm, but with the changed architecture, MAE rises to 7.25 cm, which is an  
270 increase of error by 16.2%. This verifies our architectural design choice, which passes the location-specific features  $\mathbf{x}_i$  to the  
SSH regression module in addition to the joint state, thus improving the location-specific information flow.

## 4 Conclusions

We propose HIDRA3, a new deep-learning architecture for a multi-location sea level prediction with a temporal horizon span-  
ning several days. Machine learning models have recently emerged as a highly competitive alternative to general circulation  
275 models, but they are often challenged by the need for extensive periods of training data and operational reliance on real-time  
SSH observations, which may not always be available due to sensor failures. HIDRA3 addresses these challenges by consider-  
ing multiple tide gauge locations simultaneously and by being robust to individual tide gauge sensor failures. This approach not  
only improves training where historical data is scarce, but also enables predictions in the absence of real-time SSH data from  
a subset of training locations. HIDRA3 furthermore introduces a new module for the estimation of prediction uncertainties,  
280 enhancing the interpretability of the forecasts and their integration into downstream operational systems.

In a challenging experimental setup, HIDRA3 outperforms both the current state-of-the-art deep learning model HIDRA2  
and the CMEMS version of the NEMO general circulation model. Specifically, HIDRA3 achieves approximately 15.4% lower  
MAE on high SSH values compared with HIDRA2 and 13.2% lower MAE compared with NEMO. Results show excellent  
prediction capabilities even with limited training data, indicating a remarkable generalization capability of the proposed archi-  
285 tecture.

The ability to exploit data from multiple locations for improved individual predictions, robustness to sensor failures, and  
uncertainty estimation capabilities make HIDRA3 a powerful tool for coastal flood forecasting in regional basins with high  
variability in the availability of tide gauge data. Our future work will focus on densifying the prediction locations, ultimately  
leading to fully dense two-dimensional temporal predictions, which will further extend the application outreach of the HIDRA  
290 models.

*Code and data availability.* Implementation of HIDRA3 and the code to train and evaluate the model is available in the Git repository  
<https://github.com/rusmarko/HIDRA3> (last access: 27 June 2024). The persistent version of the HIDRA3 source code is available at <https://doi.org/10.5281/zenodo.12570449> (Rus et al., 2024a). HIDRA3 pretrained weights, predictions for all 50 ensembles, geophysical training and evaluation data and SSH observations from Koper (Slovenia) are available at <https://doi.org/10.5281/zenodo.12571170> (Rus et al.,



295 2024b). Sea level observations from Italian tide gauges are provided by The National Institute for the Environment Protection and Re-  
search (ISPRA) and are publicly available at the following address: <https://www.mareografico.it> (last access: 27 June 2024). Sea level  
observations from Neretva station are property of Croatian Meteorological and Hydrological Service (DHMZ) and are available upon re-  
quest at the following address: [https://meteo.hr/proizvodi\\_e.php?section=proizvodi\\_usluge&param=services](https://meteo.hr/proizvodi_e.php?section=proizvodi_usluge&param=services) (last access: 27 June 2024). Sea  
level observations from Sobra, Vela Luka and Stari Grad (Croatia) are provided by the Institute of Oceanography and Fisheries (IOR)  
300 and are publicly available at the Intergovernmental Oceanographic Commission Sea Level Station Monitoring Facility (IOC SLSMF;  
<http://www.ioc-sealevelmonitoring.org>, last access: 27 June 2024).

*Author contributions.* MR was the main designer of HIDRA3. MK led the machine-learning part of the research and contributed to the  
design of HIDRA3. ML provided the geophysical background and led the oceanographic part of this research. HM and MR performed  
quality control over sea level observations. MR, ML and MK wrote the draft of the paper. All authors contributed to the final version of the  
305 manuscript.

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

*Acknowledgements.* The authors would like to thank the Academic and Research Network of Slovenia - ARNES and the Slovenian National  
Supercomputing Network - SLING consortium for making the research possible using the ARNES computing cluster. The authors would  
like to acknowledge the efforts of all the technical staff at the Italian National Institute for the Environment Protection and Research (ISPRA),  
310 the Croatian Meteorological and Hydrological Service (DHMZ), The Institute of Oceanography and Fisheries (IOR) and at the Slovenian  
Environment Agency (ARSO) for maintaining an operational system of tide gauges.

*Financial support.* Matjaž Ličer acknowledges the financial support from the Slovenian Research and Innovation Agency ARIS (contract  
no. P1-0237). This research was supported in part by ARIS programme J2-2506 and project P2-0214 and by supercomputing network SLING  
(ARNES, EuroHPC Vega).



## 315 References

- Bajo, M., Ferrarin, C., Umgiesser, G., Bonometto, A., and Coraci, E.: Modelling the barotropic sea level in the Mediterranean Sea using data assimilation, *Ocean Science*, 19, 559–579, <https://doi.org/10.5194/os-19-559-2023>, 2023.
- Barzandeh, A., Rus, M., Ličer, M., Maljutenko, I., Elken, J., Lagemaa, P., and Uiboupin, R.: Evaluating the application of deep-learning ensemble sea level and storm surge forecasting in the Baltic Sea, in: EGU General Assembly 2024, Vienna, Austria, 320 <https://doi.org/10.5194/egusphere-egu24-17233>, eGU24-17233, 2024.
- Bernier, N. B. and Thompson, K. R.: Deterministic and ensemble storm surge prediction for Atlantic Canada with lead times of hours to ten days, *Ocean Modelling*, 86, 114 – 127, <https://doi.org/https://doi.org/10.1016/j.ocemod.2014.12.002>, 2015.
- Braakmann-Folgmann, A., Roscher, R., Wenzel, S., Uebbing, B., and Kusche, J.: Sea level anomaly prediction using recurrent neural networks, arXiv preprint arXiv:1710.07099, 2017.
- 325 Clementi, E. et al.: Mediterranean Sea Analysis and Forecast (CMEMS MED-Currents, EAS6 system) (Version 1). Copernicus Monitoring Environment Marine Service (CMEMS), [https://doi.org/https://doi.org/10.25423/CMCC/MEDSEA-\\_ANALYSISFORECAST-\\_PHY-\\_006\\_013\\_EAS6](https://doi.org/https://doi.org/10.25423/CMCC/MEDSEA-_ANALYSISFORECAST-_PHY-_006_013_EAS6), 2021.
- Codiga, D.: Unified Tidal Analysis and Prediction Using the UTide Matlab Functions., Tech. rep., Graduate School of Oceanography, University of Rhode Island, Narragansett, RI, USA, <https://github.com/wesleybowman/UTide>, 2011.
- 330 Ferrarin, C., Orlić, M., Bajo, M., Davolio, S., Umgiesser, G., and Lionello, P.: The contribution of a mesoscale cyclone and associated meteotsunami to the exceptional flood in Venice on November 12, 2019, *Quarterly Journal of the Royal Meteorological Society*, 149, 2929–2942, <https://doi.org/https://doi.org/10.1002/qj.4539>, 2023.
- Ferrarin, C. et al.: Integrated sea storm management strategy: the 29 October 2018 event in the Adriatic Sea, *Natural Hazards and Earth System Sciences*, 20, 73–93, <https://doi.org/10.5194/nhess-20-73-2020>, 2020.
- 335 Glorot, X. and Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the thirteenth international conference on artificial intelligence and statistics, pp. 249–256, JMLR Workshop and Conference Proceedings, 2010.
- Glorot, X., Bordes, A., and Bengio, Y.: Deep sparse rectifier neural networks, in: Proceedings of the fourteenth international conference on artificial intelligence and statistics, pp. 315–323, JMLR Workshop and Conference Proceedings, 2011.
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., et al.: ERA5 340 hourly data on single levels from 1979 to present, Copernicus climate change service (C3S) climate data store (CDS), 10, 2018.
- Hieronymus, M., Hieronymus, J., and Hieronymus, F.: On the application of machine learning techniques to regression problems in sea level studies, *Journal of Atmospheric and Oceanic Technology*, 36, 1889–1902, 2019.
- Hochreiter, S. and Schmidhuber, J.: Long short-term memory, *Neural computation*, 9, 1735–1780, 1997.
- Imani, M., Kao, H.-C., Lan, W.-H., and Kuo, C.-Y.: Daily sea level prediction at Chiayi coast, Taiwan using extreme learning machine and 345 relevance vector machine, *Global and planetary change*, 161, 211–221, 2018.
- Ishida, K., Tsujimoto, G., Ercan, A., Tu, T., Kiyama, M., and Amagasaki, M.: Hourly-scale coastal sea level modeling in a changing climate using long short-term memory neural network, *Science of The Total Environment*, 720, 137 613, 2020.
- Klambauer, G., Unterthiner, T., Mayr, A., and Hochreiter, S.: Self-normalizing neural networks, *Advances in neural information processing systems*, 30, 2017.
- 350 Kristan, M. and Leonardis, A.: Online discriminative kernel density estimator with gaussian kernels, *IEEE transactions on cybernetics*, 44, 355–365, 2013.



- Leutbecher, M. and Palmer, T.: Ensemble forecasting, Tech. rep., <https://doi.org/10.21957/c0hq4yg78>, 2007.
- Loshchilov, I. and Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts, arXiv preprint arXiv:1608.03983, 2016.
- Loshchilov, I. and Hutter, F.: Decoupled weight decay regularization, arXiv preprint arXiv:1711.05101, 2017.
- 355 Madec, G.: NEMO ocean engine, [https://www.nemo-ocean.eu/wp-content/uploads/NEMO\\_book.pdf](https://www.nemo-ocean.eu/wp-content/uploads/NEMO_book.pdf), 2016.
- Mel, R. and Lionello, P.: Storm Surge Ensemble Prediction for the City of Venice, *Weather and Forecasting*, 29, 1044–1057, <https://doi.org/10.1175/WAF-D-13-00117.1>, 2014.
- Rus, M., Fettich, A., Kristan, M., and Ličer, M.: HIDRA2: deep-learning ensemble sea level and storm tide forecasting in the presence of seiches – the case of the northern Adriatic, *Geoscientific Model Development*, 16, 271–288, <https://doi.org/10.5194/gmd-16-271-2023>,  
360 2023.
- Rus, M., Mihanović, H., Ličer, M., and Kristan, M.: Code for HIDRA3: A Robust Deep-Learning Model for Multi-Point Sea-Surface Height Forecasting, <https://doi.org/10.5281/zenodo.12570449>, 2024a.
- Rus, M., Mihanović, H., Ličer, M., and Kristan, M.: Training and Test Datasets, Pretrained Weights and Predictions for HIDRA3, <https://doi.org/10.5281/zenodo.12571170>, 2024b.
- 365 Tompson, J., Goroshin, R., Jain, A., LeCun, Y., and Bregler, C.: Efficient object localization using convolutional networks. *CoRR* abs/1411.4280 (2014), arXiv preprint arXiv:1411.4280, 2014.
- Umgiesser, G., Ferrarin, C., Bajo, M., Bellafiore, D., Cucco, A., De Pascalis, F., Ghezzi, M., McKiver, W., and Arpaia, L.: Hydrodynamic modelling in marginal and coastal seas — The case of the Adriatic Sea as a permanent laboratory for numerical approach, *Ocean Modelling*, 179, 102–123, <https://doi.org/https://doi.org/10.1016/j.ocemod.2022.102123>, 2022.
- 370 Vapnik, V.: *The nature of statistical learning theory*, Springer science & business media, 1999.
- Žust, L., Fettich, A., Kristan, M., and Ličer, M.: HIDRA 1.0: deep-learning-based ensemble sea level forecasting in the northern Adriatic, *Geoscientific Model Development*, 14, 2057–2074, <https://doi.org/10.5194/gmd-14-2057-2021>, 2021.