



HIDRA2: deep-learning ensemble storm surge forecasting in the presence of seiches – the case of Northern Adriatic

Marko Rus¹, Anja Fettich², Matej Kristan¹, and Matjaž Ličer^{2,3}

¹Faculty of Computer and Information Science, Visual Cognitive Systems Lab, University of Ljubljana, Ljubljana, Slovenia

²Slovenian Environment Agency, Office for Meteorology, Hidrology and Oceanography, Ljubljana, Slovenia

³National Institute of Biology, Marine Biology Station, Piran, Slovenia

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

Abstract. We propose a new deep-learning architecture HIDRA2 for sea level and storm surge modeling, which is extremely fast to train and apply, and outperforms both our previous network design HIDRA1 and the state-of-the-art numerical ocean model (a NEMO engine with sea level data assimilation), over all sea level bins and all forecast lead times. The architecture of HIDRA2 employs novel atmospheric, tidal and SSH feature encoders, as well as a novel feature fusion and SSH regression block. HIDRA2 was trained on surface wind and pressure fields from a single member of ECMWF atmospheric ensemble and on Koper tide gauge observations during years 2006–2018, and evaluated on the data from June 2019 to December 2020. Compared to HIDRA1, the overall mean absolute forecast error is reduced by 13.9 %, while on storm surge events it is lower by even a larger margin of 25.1 %. Consistent superior performance over HIDRA1 as well as NEMO is observed in both tails of the sea level distribution. Power spectrum analysis indicates that HIDRA2 most accurately represents the energy density peaks centered on the two lowest Adriatic wind-induced free oscillation eigenmodes (seiches) among all tested models. To assign model errors to specific frequency bands covering diurnal and semi-diurnal tides and the lowest two basin seiches, sea level band-pass filtering of several historic storm surge events is applied. HIDRA2 performs well across all frequency bands and most accurately predicts amplitudes and temporal phases of the Adriatic basin seiches. This is shown to be an important forecasting benefit due to the high sensitivity of total Adriatic storm surge sea level to the phase lag between peak tide and peak seiche.

1 Introduction

Global mean sea level rise, related to anthropogenic climate change, is causing a worldwide increase in coastal flooding frequency and is leading to a myriad of negative consequences for coastal communities, civil safety and economies. Shallow semi-enclosed coastal regional basins like Northern Adriatic (North Central Mediterranean Sea) are facing growing threats of coastal inundation and erosion (Ferrarin et al., 2020), sea water intrusions in freshwater reservoirs and are worsening the conditions for marine traffic. Northern Adriatic ports like Venice, Koper and Trieste, but also other cultural landmark towns like Chioggia or Piran, have been – or will be – forced to take expensive preventive measures to mitigate their exposure.

The problem of sea level forecasting on the Northern Adriatic Shelf (see Fig. 1 for the shelf location and depth) is two-fold: (i) high sea levels lead to severe flooding of densely populated coastal towns, while (ii) low sea levels may effectively inhibit



25 large marine cargo due to very shallow depths (often below 15 meters) of marine waterways on the shelf and especially in the Gulf of Trieste. Reliable forecasting of both tails, high and low, of sea level distribution is therefore imperative for services like civil safety and cargo scheduling activities in local ports.

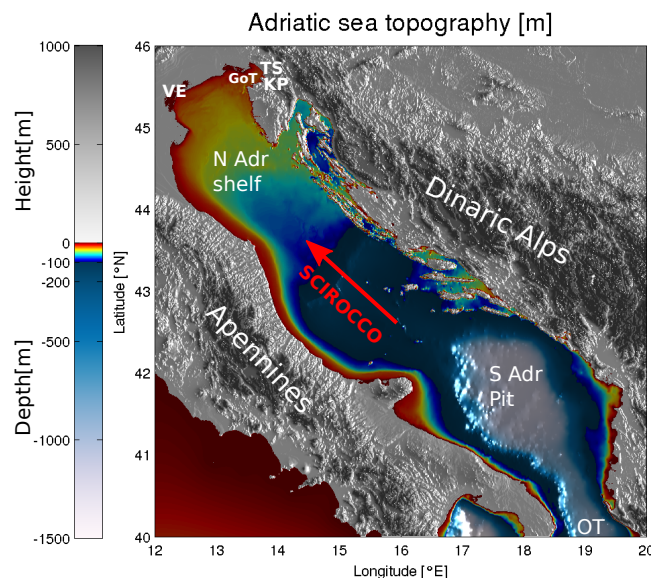


Figure 1. Topography and bathymetry of the Adriatic region. Abbreviations used on the map are as follows: TS - Trieste, KP - Koper, GoT - Gulf of Trieste, VE - Venice, N Adr Shelf - Northern Adriatic Shelf, S Adr Pit - Southern Adriatic Pit, OT - Otranto Strait. Direction of Scirocco is marked with a red arrow. The image was created by the authors based on EMODnet bathymetry data, available at <https://portal.emodnet-bathymetry.eu/> (last access: 8 June 2022) and Copernicus European Digital Elevation Model, available at <https://land.copernicus.eu/imagery-in-situ/eu-dem/eu-dem-v1-0-and-derived-products/eu-dem-v1.0> (last access: 8 June 2022).

The two distribution tails, however, represent two dynamically separate problems. High sea levels always occur due to intense pressure lows and corresponding strong winds during cyclonic activity in the basin, while extremely low sea levels typically occur through a combination of prolonged periods of high atmospheric pressure and spring tides. Equilibrium ocean response to slow changes in air pressure is captured by the inverse barometer effect, while the wind set-up of the sea level occurs through the vertical momentum flux across the air-sea interface. Dominant winds in the Adriatic basin are southeasterly Scirocco, blowing along the major axis of the basin (see Fig. 1), and the north-easterly cross-basin Bora. Strong Scirocco events lead to severe storm surges, excitation of basin seiches (Bajo et al., 2019) and potentially severe flooding in the Northern Adriatic. Adriatic along-basin seiches have eigenperiods of 21.5 hours (fundamental eigenmode) and 10.9 hours (first excited eigenmode), see e.g. Medvedev et al. (2020), and decay on the timescale of days, mostly due to radiation losses through Otranto (Cerrovecki et al., 1997).

The key difficulty of sea level forecasting in the Adriatic basin arises from high sensitivity of total sea level to the phase lag between the gravitationally generated tides (independent from meteorological forcing) and meteorologically generated



40 basin seiches (independent from gravitational forcing). This sensitivity can translate reliable atmospheric forecasts with very limited errors in timing and trajectory of the cyclone into substantial errors in the sea level forecast. Probabilistic ensemble forecasting of sea level envelopes with error variance estimation (Žust et al., 2021; Ferrarin et al., 2020; Bernier and Thompson, 2015; Mel and Lionello, 2014) was therefore explored to tackle this drawback. However, the ensemble sea level forecasting is numerically expensive, requires specialized expensive hardware and introduces delays in prediction. To avoid the high numeric
45 cost of ensemble sea level forecasting, computationally efficient machine-learning-based ensemble models have recently been explored (Žust et al., 2021). While these models require a substantial amount of training data to learn the complex relations for reliable predictions, the inference is numerically cheap – for example, the neural network HIDRA1 ensemble (Žust et al., 2021) is a million times faster than the operational numerical ocean model ensemble based on NEMO engine (Madec, 2016) at Slovenian Environment Agency.

50 Machine learning has thus been explored by several research groups for single-point sea level forecasting. The early approaches (Imani et al., 2018) were based on classic machine learning models such as support vector machines (Vapnik, 1999) with radial basis function kernels. In their work, Pashova and Popova (2011) and Karimi et al. (2013) utilized shallow fully connected neural networks, but due to simplistic network architectures that did not utilize the numerical atmospheric forecast, they could only report the desired accuracy for short temporal horizons. Ishida et al. (2020) used long short-term memory
55 (LSTM) (Hochreiter and Schmidhuber, 1997) networks together with several atmospheric variables to improve one-hour prediction into the future, but did not expand the prediction horizon. Braakmann-Folgmann et al. (2017) predicted further in time by applying a combination of LSTMs and convolutional neural networks, but at a very coarse level. Autoregressive neural networks were considered in (Hieronymus et al., 2019) to increase the temporal resolution and the prediction horizon. Most recently, a convolutional neural network HIDRA1 (Žust et al., 2021) with a specialized architecture to utilize atmospheric data,
60 sea surface heights and astronomic tides was proposed. To the best of our knowledge, HIDRA1 is currently the most accurate machine-learning sea surface height prediction model with a several days long prediction horizon at an hourly resolution. But while HIDRA1 performed favorably in comparison to the NEMO model used in that study, it failed to beat the NEMO setup at very high and very low ends of sea level distributions. In other words, extreme sea levels (coastal floods on one end and very low sea levels on the other), which interest us the most, were not yet captured with sufficient reliability and present an open
65 challenge for machine learning methods.

In this paper we propose HIDRA2, our latest attempt at sea level forecasting using deep learning. In contrast to the previous version, HIDRA2 presents a novel architecture with new atmospheric, tidal and SSH feature encoders, as well as a novel feature fusion and SSH regression block. Additional conceptual novelty is that HIDRA2 predicts the full SSH rather than the residual (i.e., the difference between SSH and astronomic tide) as is the case for HIDRA1. The new model extracts relevant information
70 from different spatial locations in the atmosphere signal and predicts the SSH with a three-days horizon at an unprecedented accuracy, outperforming HIDRA1 as well as the state-of-the-art NEMO model.

The paper is organized as follows. Section 2 introduces sea level and atmospheric model data and performance measures used in our analysis. Section 3 details the new HIDRA2 architecture and the numerical ocean model setup, used as the performance



baseline. Section 4 reports the analysis of the HIDRA2 architecture, and provides detailed quantitative as well as qualitative
75 comparison with the state-of-the-art. Conclusions and outlook are drawn in Sect. 5.

2 Training and evaluation datasets

2.1 Sea level training data

Sea surface height (SSH) observations during the period 2006–2018 were retrieved from the Koper Mareographic Station
(45° 33' N, 13° 44' E; see Fig. 2 for location), which is maintained by the Slovenian Environment Agency (ARSO) and is part
80 of the European Sea Level Network (Pérez Gómez et al., 2022). Observations are obtained in 10-minute intervals, using both a
float-type sensor and an additional radar sea level sensor, and they undergo subsequent quality control at ARSO (Pérez Gómez
et al., 2022). Hourly data points are extracted to get the signal used in HIDRA2 training and evaluation.

The tidal signal in the sea level is independent of atmospheric processes and can be computed by tidal analysis and prediction
models. The tidal contribution to Koper SSH considered in this study is estimated from hourly instantaneous SSH values in
85 one-year segments using the UTIDE Tidal Analysis package for Python (Codiga, 2011). The total sea level series is then
decomposed into a tidal and a residual part, where we define the sea level residual as the arithmetic difference between the
total sea level and the tidal sea level. According to the ARSO operational protocol, the SSH is classified as a flood if it is higher
than $\Theta_{floods} = 300$ cm. Floods thus constitute 0.43 % of all training data.

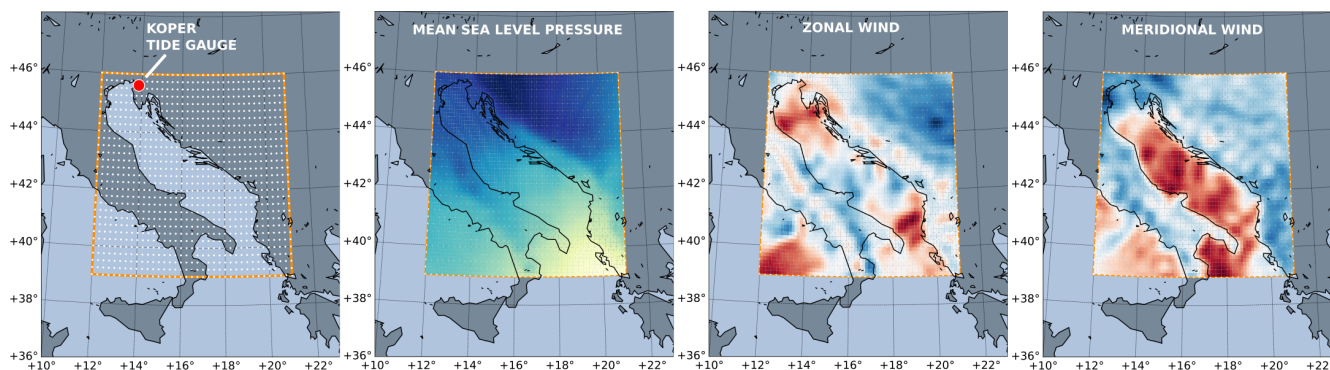


Figure 2. HIDRA2 input domain and dataset. The leftmost panel depicts ECMWF grid (white dots) and Koper tide gauge location (red circle). Three panels on the right depict snapshots of ECMWF atmospheric fields used during training.

2.2 Atmospheric training data

90 The atmospheric input used for HIDRA2 training was retrieved from the European Centre for Medium-Range Weather Forecasts (ECMWF) Ensemble Prediction System (Leutbecher and Palmer, 2007). The ECMWF ensemble forecasts come as a global atmospheric ensemble of fifty members with a 0.125° arc degree spatial resolution and a 3-hour temporal resolution.



The training dataset used in this study consists of (i) 10-meter winds and (ii) mean sea level air pressure from a single atmospheric ensemble member during the period 2006–2018. All ECMWF input fields were standardized and cropped to the Adriatic basin, represented by a 57×73 spatial grid (see Fig. 2). The forecasts were linearly interpolated to hourly timesteps to match the SSH temporal resolution. To simplify the training protocol, a single atmospheric sequence is constructed by concatenating the first 24 hours of daily consecutive ECMWF forecasts into the final atmospheric sequence used in training. HIDRA2 does not require explicit annotation of whether a location point belongs to land or sea, thus land masks are not generated.

2.3 Evaluation data

The evaluation input dataset for both HIDRAs and NEMO is disjoint from the training dataset and consists of ECMWF atmospheric predictions and Koper sea levels between 01 June 2019 and 31 December 2020. This period was chosen due to challenging conditions and unusually high incidence of floods. We use the ECMWF daily predictions, each containing fifty ensemble members with three-days prediction lead time. The data are standardized and the dimensionality of the atmospheric data is reduced in the same fashion as described in the Sect. 2.2, except that for inference, the full ECMWF three-day forecast is presented to the model. The floods represent 1.2 % of the test dataset.

2.4 Performance measures

Standard measures, i.e., the mean absolute error (MAE), the root mean square error (RMSE) and the model bias are used to evaluate prediction performance. To reflect the practical suitability, we additionally calculate the prediction accuracy as a ratio between the predictions which are within $\Theta_{acc} = 10$ cm from the ground truth and all predictions. The threshold Θ_{acc} reflects an acceptable deviation from the ground truth and was determined through discussion with the operational forecasting service at ARSO. The metrics are calculated globally by considering all prediction points, as well as separately only on floods to reflect the prediction performance at these critical rare events.

To further probe the flood event prediction capabilities, we make use of the standard performance measures from detection literature: precision Pr , recall Re and the F1 measure $F1$. Firstly, we need to define the flood event and then define the notion of the event being detected. Both of these have been defined in discussion with operational forecasters at ARSO. The anchor (i.e., temporal point) of a flood event is defined as a local maximum in an SSH sequence above $\Theta_{floods} = 300$ cm. If the predicted flood event anchor is within a 3 h margin (before or after) from the nearest ground truth flood event anchor, it is considered a true positive TP , otherwise it is a false positive FP . A flood event in the ground truth is considered a false negative FN if there is no matching flood event anchor in the predicted SSH. Like in the accuracy definition, the tolerance of $\Theta_{acc} = 10$ cm is applied, meaning that predictions below Θ_{floods} are also considered as TP when they appear in the margin of Θ_{acc} , and that false positives with ground truth within Θ_{acc} are ignored. The precision and recall are then calculated as

$$Pr = \frac{TP}{TP + FP}, Re = \frac{TP}{TP + FN}, \quad (1)$$

while the F1 measure that summarizes the detection performance, i.e.,

$$F1 = 2 \frac{Pr \cdot Re}{Pr + Re}, \quad (2)$$

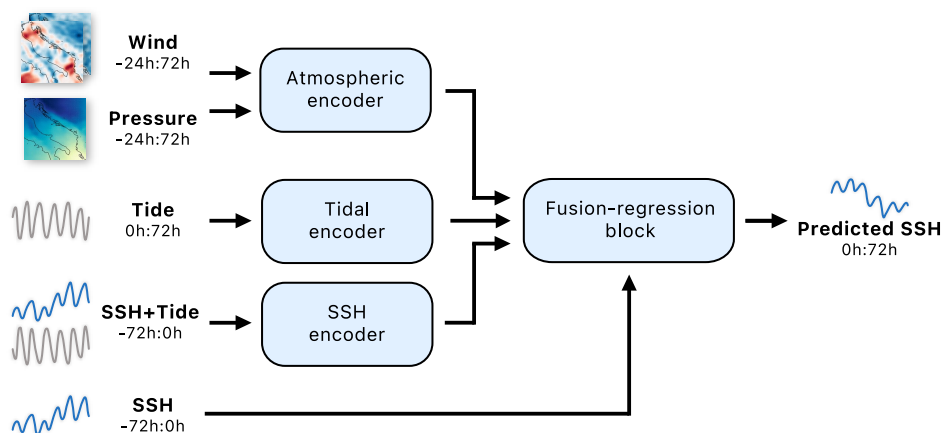


Figure 3. The HIDRA2 architecture. The Atmospheric encoder embeds the wind and pressure sequences with learnable temporal subsampling and pattern prototype matching to extract relevant features from different geographic locations and fuse them temporally into a single feature embedding. The Tidal and SSH encoders encode the future tide evolution and the past SSH and tide observations, respectively. All features are re-calibrated, fused with the past 72 h SSH and regressed into the final SSH predictions by the Fusion-regression block. Notation $a:b$ indicates hourly data points from the interval $(a, b]$, while the prediction point is at the index 0.

125 is defined as the harmonic mean between precision and recall.

3 Numerical models

3.1 HIDRA2

The proposed HIDRA2 is the second generation of a deep neural model for sea surface height prediction, with HIDRA1 (Žust et al., 2021) being the first. The new architecture is shown in Fig. 3. The input data is encoded by three encoders: the wind and
130 pressure sequences for the past 72 h are processed and merged by the Atmospheric encoder (Sect. 3.1.1), the tidal signal for the future 72 h is encoded by the Tidal encoder, and the sea surface height measurements coupled with the tide for the past 24 h are encoded by the SSH encoder (Sect. 3.1.2). The outputs of all three encoders are re-calibrated, fused with the past 72 h SSH and regressed into the final SSH hourly predictions for the future 72 h by the Fusion-regression block (Sect. 3.1.3). The following subsections detail the individual blocks.

135 3.1.1 Atmospheric encoder

The atmospheric data for the Adriatic basin at a given time-step is represented by a 57×73 spatial grid, i.e., an image. HIDRA2 assumes that coarse spatial resolution of atmospheric data contains enough information to provide satisfactory results, so it first downsamples the atmospheric data to 9×12 grid by an average pooling operation.




Figure 4. The first stage of the Atmospheric encoder. The input are four consecutive hours with two wind channels (this case), or pressure. 3×3 convolution is applied, followed by a prototype matching layer, outputting a single vector of size 512. Note that 24 independent passes are performed in parallel for the entire atmospheric input sequence. The variables k and n denote the kernel size and the number of output channels, respectively.

The Atmospheric encoder is composed of two stages. In the first stage (shown in Fig. 4), the sequences of the wind and pressure images are independently processed by their respective encoding blocks, which use the same architecture. The wind image sequence of 96 h (the past 24 h and future 72 h) is divided into 24 groups of four consecutive hours, which are processed independently. The spatial and temporal dimension of each group is reduced by a learnable 2D convolutional layer with a 3×3 kernel, stride 2 and 64 output channels¹. A ReLU activation and Dropout layers are applied, followed by a prototype matching layer with 512 learnable 4×5 prototypes (i.e. convolutional kernels of input map size). They extract features from different spatial positions, thus producing a 512-dimensional feature vector per group, i.e., 24 temporal vectors of size 512. The same processing architecture is applied to the pressure image sequence to produce 24 vectors of size 512. The two outputs are then concatenated to form a mixed set of 24 wind-pressure features of size 1024.

The second stage of the Atmospheric encoder (Fig. 5) extracts the temporal atmospheric features by considering the consecutive wind-pressure features extracted by the first stage. A 1D convolutional layer with the kernel temporal dimension size 5 and with 256 output channels² is applied, entangling the information from temporal segments equivalent to 20 h in length. Note that because we are using a convolutional layer instead of the fully connected layer, the number of learnable parameters of the entire Atmospheric encoder is independent of the forecast horizon. Each of the obtained 20 features³ is then independently processed by a network containing two blocks of residual connections, each involving 1D convolution with kernel temporal dimension size 1 (i.e., 1×256 kernels), a SELU activation (Klambauer et al., 2017) and a dropout layer. Finally, each of the obtained twenty 256-dimensional features are convolved by 32 1×256 kernels to reduce their dimensionality to 20×32 .

3.1.2 Tidal and SSH encoders

Both the tidal and SSH encoders use the same architecture, the only difference is in the size of the encoders' input. Fig. 6 depicts the SSH encoder, which takes as the input the past 72 h SSH measurements and the tide concatenated into a 72×2

¹Note that the number of output channels is equal to the number of different kernels used in the layer.

²Note that the sizes of the kernels in the 1D convolutional layer are 1024×5 , but with the first dimension matching the size of the input features, the convolution displacements are only along the second dimension, hence the 1D convolutional direction implied by the layer's name.

³The fact that the number of output segments is equal to the 20 h timespan is coincidental.




Figure 5. The second stage of the Atmospheric encoder. Features from all time points and both wind and pressure are processed with a 1D convolution, followed by two blocks with residual connections. The last convolution reduces feature dimensionality. The variables k and n denote the kernel size and the number of output channels, respectively.




Figure 6. The SSH encoder encodes a concatenation of the past SSH and tide by a 1D convolution, followed by two blocks with residual connections, max-pooling temporal reduction and convolution-based feature reduction. The variables k and n denote the kernel size and the number of output channels, respectively.

input tensor and processes them in a similar fashion to the second stage of the Atmospheric encoder: the input is processed
 160 by convolution with $256 \ 3 \times 2$ kernels, which is followed by two consecutive residual blocks, a subsampling max-pooling
 layer and a final convolutional layer to reduce the dimensionality of features to 17×16 . The Tidal encoder follows the same
 architecture, the only difference is that the input is the tide forecast for the next 72 h.

3.1.3 Fusion-regression block

Atmospheric, Tide and SSH encoders produce temporal features of different importances and sizes. To account for that, the
 165 features are re-calibrated by normalization with means and variances of the features calculated during training, and then denor-
 malized with learned weights and biases. The form of normalization follows the batch normalization layer (Ioffe and Szegedy,
 2015), which applies a 0.9 momentum for updating the running means and variances during training. The normalized features
 are then concatenated and mixed by a fully connected layer, reducing their final dimension from 1184 to 512 (left part of the




Figure 7. The Fusion-regression block firstly re-calibrates the features (the C symbol), then concatenated features are passed to a dense layer, which fuses features and reduces their dimensionality. Undistorted SSH is appended and processed with two residual blocks. The final dense layer outputs the predictions. The variable n denotes the number of output channels.

Fig. 7). The obtained 512-dimensional domain context feature vector thus contains rich atmospheric and sea-surface height information from all time points and all parts of the input domain.

While the encoding and mixing operations extract the domain context, the explicit surface height information might not well be retained in the extracted feature vector. To re-inject this information, the obtained domain context feature vector is concatenated with the timeseries of past observed SSH before passing to the final regression block. The latter is composed of two fully connected layers with 584 units, SELU activations and residual connections, followed by a fully connected layer with 72 outputs for the 72 h prediction horizon (see Fig. 7).

3.1.4 The network training

HIDRA2 is trained end-to-end using mean squared error (MSE) loss between the predictions and the ground truth. We train the model using the AdamW optimizer (Loshchilov and Hutter, 2017) with standard parameter values (learning rate $1e-4$, and the running average damping parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$), and apply the cosine annealing (Loshchilov and Hutter, 2016) learning schedule to gradually reduce the learning rate during training to $1e-5$. The training batch size is set to 512 data samples, and the model is trained for 60 epochs, with 1000 batches in each epoch. Training takes approximately 1.5 hours on a single computer with NVIDIA Geforce RTX 2080 TI graphics card, while the inference of a single 72 h prediction for one member of the atmospheric ensemble takes only 4 ms.

3.1.5 Summary of differences to HIDRA1

While there are many differences between HIDRA2 and HIDRA1, we summarize only the major conceptual ones for a clearer exposition of the contributions. HIDRA1 uses wind, pressure and 2 m temperature from ECMWF predictions, while our preliminary study showed that the new HIDRA2 architecture does not benefit from the temperature, thus only wind and pressure are considered. HIDRA1 concatenates all atmospheric inputs at a timestep and encodes them by Resnet (He et al., 2016) blocks. While Resnet excels in computer vision tasks that rely on high-level semantic feature abstraction, we argue that tailored shal-



190 lower encoders are more appropriate for the extraction of meaningful atmospheric patterns. HIDRA2 thus separately encodes
the wind and pressure by shallow encoders, which apply spatial pattern features extraction, and then mixes the features from the
two atmospheric variables by extracting temporal patterns. While this allows HIDRA2 to extract multiple spatial patterns in the
data, only a single set of spatial weights is used to fuse the atmospheric features at a given time-step in HIDRA1, consequently
reducing its expressive power. HIDRA1 first averages four-hour atmospheric input data to temporally subsample the input,
195 while HIDRA2 considers per-hour inputs and learns the appropriate spatio-temporal subsampling to maximize its predictive
power. Another advantage of HIDRA2 is that it encodes the SSH input and mixes it with the atmospheric features early in the
network to create a domain context feature vector before the final regression, while HIDRA1 considers only the atmospheric
data for the context vector. Finally, HIDRA1 predicts the SSH residual (i.e., the difference between SSH and the astronomic
tide), while HIDRA2 directly predicts the full SSH.

200 3.2 NEMO ocean model

The Copernicus Marine Environment Monitoring Service (CMEMS) product MEDSEA_ANALYSISFORECAST_PHY_006-
_013 (see Clementi et al. (2021)), was used as a challenging numeric baseline for HIDRA2. This product is based on a
Mediterranean basin configuration of the NEMO ocean model (Madec, 2016), and provides daily ocean forecasts for sea
surface height above the geoid, temperature, salinity, circulation and mixed layer depth. The model domain spans the entire
205 Mediterranean basin with a $1^\circ/24$ resolution and has 141 unevenly spaced vertical levels. The model solutions are operationally
constrained to near-real-time observations using a 3D variational assimilation scheme of temperature, salinity and along-track
satellite sea level anomaly observations. The atmospheric forcing to the CMEMS model is provided by the ECMWF. Further
details about the modeling setup can be found in (Clementi et al., 2021). In this study, an SSH timeseries at the closest point to
Koper tide gauge was extracted from the Mediterranean ocean model forecast. Henceforth, we will be referring to this modeling
210 product simply as NEMO.

The NEMO sea level forecast, denoted as y^{nemo} , at any given location reflects departures from the local geoid and hence
does not represent the absolute local depth of the water. The latter is furthermore also driven by low-frequency processes on
the scales of many weeks or months which are often difficult to capture for regional basin models on synoptic timescales. The
NEMO SSH predictions are therefore offset-adjusted to obtain the Koper location total sea level (required by port authorities
215 and civil rescue) as follows. A time-averaged NEMO SSH offset ϵ_n on the n -th hour of the forecast day is defined as

$$\epsilon_n = n^{-1} \sum_{k=1}^n [y^{\text{nemo}}(t_k) - y^{\text{kp}}(t_k)], \quad (3)$$

where $y^{\text{kp}}(t)$ is the observed Koper sea level. The value of ϵ_{12} is then subtracted from $y^{\text{nemo}}(t)$ to ensure a zero bias for the
first 12 h of the forecast day. Note that, despite this correction, the complete forecast time-series still exhibits a non-zero bias.
A similar bias correction is not required in HIDRAs, since they predict the full SSH and learn to appropriately adjust for the
220 bias automatically.



4 Results and discussion

4.1 HIDRA2 architecture analysis

As noted in Sect. 2, HIDRA2 was trained on the period 2006–2018 and evaluated on the period between 01 June 2019 and 31 December 2020. For evaluation, a single prediction is obtained by averaging predictions over all fifty ECMWF ensemble members. In the following we analyze the architectural choices of HIDRA2. The prediction of full SSH is justified in Sect. 4.1.1, while Sect. 4.1.2 reports an ablation study. The results of all experiments are collected in Table 1.

4.1.1 Predicting the full SSH vs residual

A valid hypothesis can be made that predicting the residual (i.e., the difference between the full SSH and the tide) might be more beneficial than predicting the full SSH, since the network parameters might be better utilized by focusing only on the part of SSH not affected by the astronomic tide. In fact, HIDRA1 (Žust et al., 2021) does exactly this.

To explore this hypothesis, HIDRA2 was modified to predict only the residual ($\text{HIDRA2}_{\text{res}}$), by replacing the SSH input in the SSH encoder by the residual, as well as replacing the SSH with the residual in the Fusion-regression block. Results in Table 1 indicate a slight overall performance increase when only the residuals are considered in HIDRA2. However, considering only storm surges, we observe a substantial increase of the prediction error (+12.7 % MAE). This means that full SSH prediction is very beneficial for predicting floods, while incurring only a small drop in the overall performance. This is an acceptable trade-off that justifies the full SSH prediction.

4.1.2 Ablation study

Three HIDRA2 modifications were considered to evaluate the importance of individual encoders: HIDRA2 without the Atmospheric encoder ($\text{HIDRA2}_{\setminus\text{atmE}}$), HIDRA2 without the Tidal encoder ($\text{HIDRA2}_{\setminus\text{tidE}}$), and HIDRA2 without the SSH encoder ($\text{HIDRA2}_{\setminus\text{sshE}}$). The results in the Table 1 show that MAE increases with each modification, particularly during storm surge events. Removal of the Atmospheric encoder results in the most significant performance drop, indicating that the atmospheric features convey by far the most relevant predictive information. A significant performance drop is observed as well when removing the Tidal encoder. The SSH encoder has the smallest impact, yet still importantly contributes to the prediction accuracy during storm surge events.

Note that two types of input data (the SSH and the tide) are considered in the SSH encoder. Further HIDRA2 modifications were thus explored to evaluate their importance: a version with the tidal input omitted ($\text{HIDRA2}_{\setminus\text{tidI}}$) and a version without the SSH input ($\text{HIDRA2}_{\setminus\text{sshI}}$). The results in Table 1 show that the removal of each leads to a consistent but moderate increase of the errors overall. However, the errors increase substantially during storm surge events, indicating the importance of using both types of inputs.

We observe a similar situation when removing the atmospheric and SSH/tide feature re-calibration in the Fusion-regression block ($\text{HIDRA2}_{\setminus\text{norm}}$). Results in Table 1 indicate that feature normalization does not affect performance in normal conditions,



Table 1. Performance of ablated HIDRA2 designs evaluated over all sea level bins (the *Overall* columns) and only on storm surge events (*Storm surge events* column).

modification	Overall				Storm surge events						
	MAE [cm]	RMSE [cm]	Bias [cm]	Acc [%]	MAE [cm]	RMSE [cm]	Bias [cm]	Acc [%]	Re [%]	Pr [%]	F1 [%]
HIDRA2 _{res}	4.10	5.81	-0.43	92.89	10.94	15.29	-8.36	59.57	85.06	96.32	90.34
HIDRA2 _{atmE}	7.54	11.40	-0.67	75.55	27.81	34.25	-25.44	23.44	45.45	89.74	60.34
HIDRA2 _{tidE}	4.62	6.37	-0.25	90.43	10.86	14.76	-8.18	57.42	83.12	94.81	88.58
HIDRA2 _{sshE}	4.21	5.90	0.10	92.49	10.47	14.45	-7.66	59.57	85.71	96.35	90.72
HIDRA2 _{tidI}	4.22	5.93	-0.25	92.36	10.62	14.48	-7.92	59.14	84.42	96.30	89.97
HIDRA2 _{sshI}	4.21	5.89	-0.17	92.56	10.38	14.21	-7.70	60.00	86.36	96.38	91.10
HIDRA2 _{norm}	4.15	5.83	0.19	92.72	10.11	14.10	-7.37	61.29	86.36	95.68	90.78
HIDRA2	4.16	5.83	0.36	92.78	9.71	13.80	-6.72	63.44	89.61	95.83	92.62

but it substantially contributes to the prediction accuracy on storm surges. A closer inspection of HIDRA2_{norm} showed that the scale of the tidal features is four times larger than the scale of the atmospheric features. Inclusion of the re-calibration blocks, however, remedies this by making the scales of all features (atmospheric, SSH and tidal) approximately the same.

255 4.2 Comparison with the state-of-the-art

HIDRA2 is compared to HIDRA1 (Žust et al., 2021), which is currently the state-of-the-art in machine-learning SSH prediction, and to NEMO (Madec, 2016), which presents a state-of-the-art geophysics-based SSH prediction method. The methods are evaluated with respect to different SSH values in Sect. 4.2.1, Sect. 4.2.2 reports performance with respect to the lead times, while spectral analysis is reported in Sect. 4.2.3. The last two sections discuss performance on historical storm surge events (Sect. 4.2.4) and the forecast spectral decomposition of these events (Sect. 4.2.5).

4.2.1 SSH forecast performance

The overall prediction performance and performance restricted to storm surge events are shown in Table 2. HIDRA2 outperforms both HIDRA1 and NEMO overall as well as on storm surges, yielding a lower MAE/RMSE and higher accuracy. While HIDRA1 achieves a lower bias, its RMSE/MAE are substantially higher – HIDRA2 outperforms HIDRA1 in MAE by 13.9 % overall, and by 25.1 % on the storm surges. NEMO achieves the highest precision of flood detection ($Pr = 100\%$), meaning that all detected floods are true positives. However, this comes at a cost of missing several floods, resulting in a low recall of $Re = 68.83\%$. A similar situation is observed for HIDRA1. The recalls for these two methods (NEMO: 68.83 % and HIDRA1: 79.22 %) are substantially lower than that of HIDRA2 ($Re = 95.83\%$), which detects nearly all of the floods with few false positives. The excellent tradeoff between the precision and recall of HIDRA2 is reflected in its F1 score (92.62 %), which is substantially higher than that of NEMO (81.45 %) and HIDRA1 (88.09 %).



Table 2. Performance of HIDRA1, HIDRA2 and NEMO over all sea level bins (the *Overall* columns) and only during storm surge events (*Storm surge events* columns). Performance of using the tide as SSH prediction is included for reference.

	Overall				Storm surge events						
	MAE [cm]	RMSE [cm]	Bias [cm]	Acc [%]	MAE [cm]	RMSE [cm]	Bias [cm]	Acc [%]	Re [%]	Pr [%]	F1 [%]
Tide	13.82	18.86	-5.13	47.45	55.75	59.45	-55.75	0.00	0.00	/	/
NEMO	6.54	8.52	-1.23	79.14	13.03	17.09	-11.24	49.68	68.83	100.00	81.54
HIDRA1	4.83	6.87	0.09	89.83	12.97	18.02	-10.53	54.62	79.22	99.19	88.09
HIDRA2	4.16	5.83	0.36	92.78	9.71	13.80	-6.72	63.44	89.61	95.83	92.62

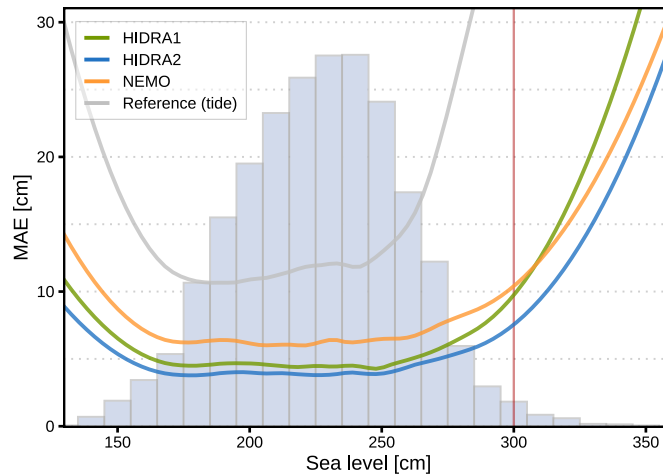


Figure 8. HIDRA and NEMO performances with regard to sea level bins (grey histogram in the bottom layer). The coastal flood threshold Θ_{floods} is marked with a vertical red line.

For detailed analysis, we visualize the MAE values of the tested methods with respect to the sea level heights in Figure 8. On storm surges, NEMO outperforms HIDRA1, while HIDRA2 consistently outperforms both HIDRA1 and NEMO by up to several cm. In the low sea level distribution tail, HIDRA1 exhibits a lower MAE (by 4 cm) than NEMO, while HIDRA2 again outperforms both HIDRA1 and NEMO by 2 cm and 7 cm, respectively. The solid HIDRA2 performance in the low end of the sea level distribution is particularly important to note because of its potentially high significance to cargo scheduling in the very shallow seas surrounding the Port of Koper, which is currently restricted only to the periods of high tides. In summary, HIDRA2 outperforms both state-of-the-art methods for all sea level heights, thus displaying its remarkable prediction skill in moderate as well as extreme values of the sea surface height.




Figure 9. MAE score of HIDRA and NEMO models with regard to prediction lead time (between 1 h and 72 h).

4.2.2 Performance with regard to forecast lead time

280 We next analyzed how the prediction lead time affects the prediction errors. Fig. 9 shows the MAE scores with respect to the prediction lead time for the values between 1 h and 72 h. The MAE of the prediction gradually increases with the lead time for all the tested methods. While overall being a solid performer with MAE well below 10 cm, NEMO exhibits the highest MAE and also the highest MAE variance. Clear signals are observed with 12 h and 24 h periods in the NEMO MAE. Since NEMO includes tides, we suspect this periodicity stems mostly from the errors in either amplitude or phase of the tidal part of the NEMO sea level signal but further research would be necessary to properly substantiate this claim. Interestingly, while HIDRA2 consistently outperforms HIDRA1 for all lead times, the shapes of the MAE curves show resemblance. While the 12 h period does not seem to be present in the MAE curves of these two models, their 24 h period is clearly present. Further research would, however, be required to substantiate and explain the observed MAE curve behavior.

4.2.3 Spectral analysis

290 To investigate the spectral properties of the modeled and observed SSH timeseries, spectral densities were computed on the HIDRA2, HIDRA1 and NEMO predictions. Unless otherwise stated, all time-series analyzed in this section were obtained by concatenating (in time) the first 24 hours of each daily HIDRA2, HIDRA1 and NEMO three-day forecast. Spectral densities (shown in Fig. 10) were then computed as absolute values of a 1D Fast Fourier Transform of the respective series over a fixed frequency domain of $(1 \text{ h})^{-1} - (72 \text{ h})^{-1}$.

295 Figure 10 indicates that all methods adequately represent the tidal dynamics in Koper. The energy content around the two lowest basin eigenmodes is, however, more discriminatory: NEMO clearly underestimates the spectral density both around the ground state seiche (at 21.5 h period) and around the first excited state (10.9 h period). Similar behavior was noticed in our

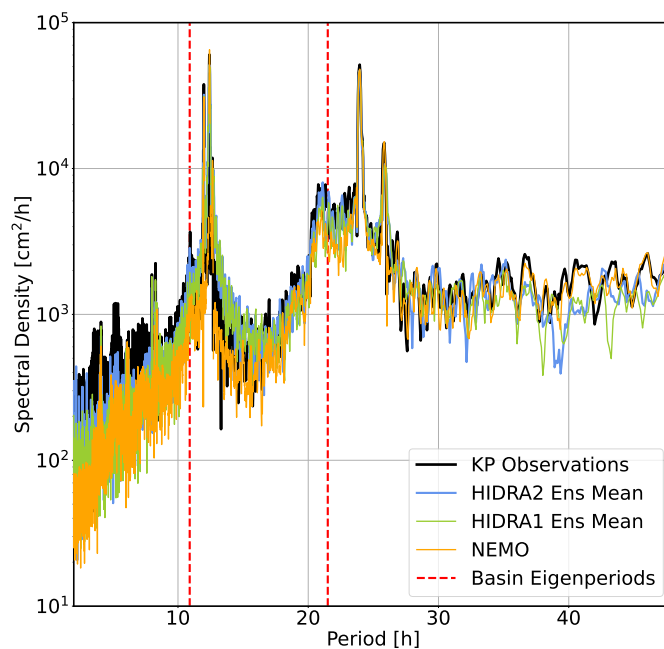


Figure 10. Spectral density of SSH timeseries from Koper tide gauge, HIDRA2, HIDRA1 and NEMO, during time window between 01 June 2019 and 31 December 2020. Sharp peaks at (roughly) 12 and 24 hours indicate the presence of tides, while the two dashed vertical red lines mark the periods of the two lowest Adriatic sea level eigenmodes.

previous work with an independent configuration of NEMO (Žust et al., 2021). HIDRA1 underestimates the energy of this part of the signal as well, but nevertheless does a bit better by packing more energy density in these two bands. Predictions of HIDRA2 are clearly the closest to the observations in both bands – this model does not underestimate the spectral density around the 21.5 h peak at all and it is the only method that exhibits a clear peak at the 10.9 h period.

It appears that HIDRA2 is capable of generating a seiche-like behavior in its predictions. Spectral density, however, discards the temporal component of the signal, and adequate spectral density in the $(21.5 \text{ h})^{-1}$ and $(10.9 \text{ h})^{-1}$ frequency bands says little about whether Adriatic seiches are generated by HIDRA2 at the appropriate times, namely during the storm surge events. To inspect this aspect of HIDRA behavior, we now proceed to analyze the predictions during several historic storm surges.

4.2.4 Performance during historic storm surge events

Historic Adriatic storm surge events are used to qualitatively compare the HIDRA2 performance with the state-of-the-art. Storm surges in question occurred during November and December 2019 and were of historic proportions by any criterion. The Slovenian coast was flooded over ten times in a single month and sea levels in Venice were among the highest ever observed. Furthermore, the events in November 2019 turned out to be difficult to model due to the formation of a transient and very localized low pressure over the Gulf of Venice, which went unresolved in most models (Cavaleri et al., 2020). These



events, along with those from December 2019, therefore represent a highly challenging benchmark for any atmospheric model and even more for any downstream SSH prediction method.




Figure 11. Comparison of HIDRA2 ensemble (top), HIDRA1 ensemble (middle) and NEMO deterministic run (bottom) during the November 2019 flooding sequence in the Northern Adriatic. Semi-transparent regions in the top two plots depict the minimum-maximum envelope of each HIDRA ensemble.

Figure 11 shows HIDRA2, HIDRA1 and NEMO SSH forecasts for the Adriatic storm surge of November 2019. None of the
315 models successfully predicted the first and highest sea level peak on 12 November 2019, but HIDRA2 and NEMO both give a better forecast than HIDRA1. As noted in Cavalieri et al. (2020), this peak was difficult to forecast due to the delicate timing between the peak of winds and the peak of the full moon tide, combined with the formation of an unresolved local pressure disturbance over the west coast of Northern Adriatic. Relative timing of these influences turned out to be a *sine qua non* for a successful prediction – neither the winds nor pressure were, in themselves, in any way extraordinary. It is further shown in
320 Cavalieri et al. (2020) that this particular storm surge could have been up to 25 cm higher had this scenario evolved 12 hours earlier when tidal peaks were themselves higher.

The peak on 13 November is slightly better predicted by both HIDRAs than by NEMO, with HIDRA2 exhibiting a somewhat lower forecast spread than HIDRA1. Apart from this peak, all methods captured the sea level variability quite well, which is in itself an implicit testament to the high skills of ECMWF atmospheric products.




Figure 12. Comparison of HIDRA2 ensemble (top), HIDRA1 ensemble (middle) and NEMO deterministic run (bottom) during the December 2019 flooding sequence in the Northern Adriatic. Semi-transparent regions in the top two plots depict the minimum-maximum envelope of each HIDRA ensemble.

325 Floods of December 2019 are another example of HIDRA2 superior performance over HIDRA1 and NEMO. SSH obser-
330 vations and predictions in Koper during this period are depicted in Fig. 12. Several conclusions about HIDRA2 behavior may
be reached with regard to this particular flood. HIDRA2 ensemble appears to be closest to the observations and exhibits a
substantially lower forecasting spread than the HIDRA1 ensemble. Low forecasting spread is acceptable when in conjunction
with a well-behaved ensemble mean. In this case, the HIDRA2 ensemble mean is in excellent agreement with the observations.
335 The same could be said for HIDRA1, albeit to a lesser degree. NEMO, however, completely misses the first two peaks between
15 and 17 December, slightly underestimates (like HIDRA1) the highest peak on 23 December, and overall underestimates the
minimum-maximum range of the sea level variations, corresponding to poorly predicted ebb levels after 23 December. The
vertical sea level range is much better captured by both HIDRAs, especially by HIDRA2. This result is consistent with our
demonstration that HIDRA2 exhibits the lowest error in both the high and the low tail of sea level distributions (Fig. 10).

335 To inspect the behavior of the ensemble forecast spread, three timeseries were created from daily (72 h long) forecasts during
evaluation time window between 01 June 2019 and 31 December 2020. The first timeseries was constructed by concatenating
each first day (i.e., 1–24 h of forecast) from each of the daily forecasts, thus containing predictions with lead times of 1–24 h
on each respective day in the evaluation time window. The second and the third timeseries were constructed by concatenating

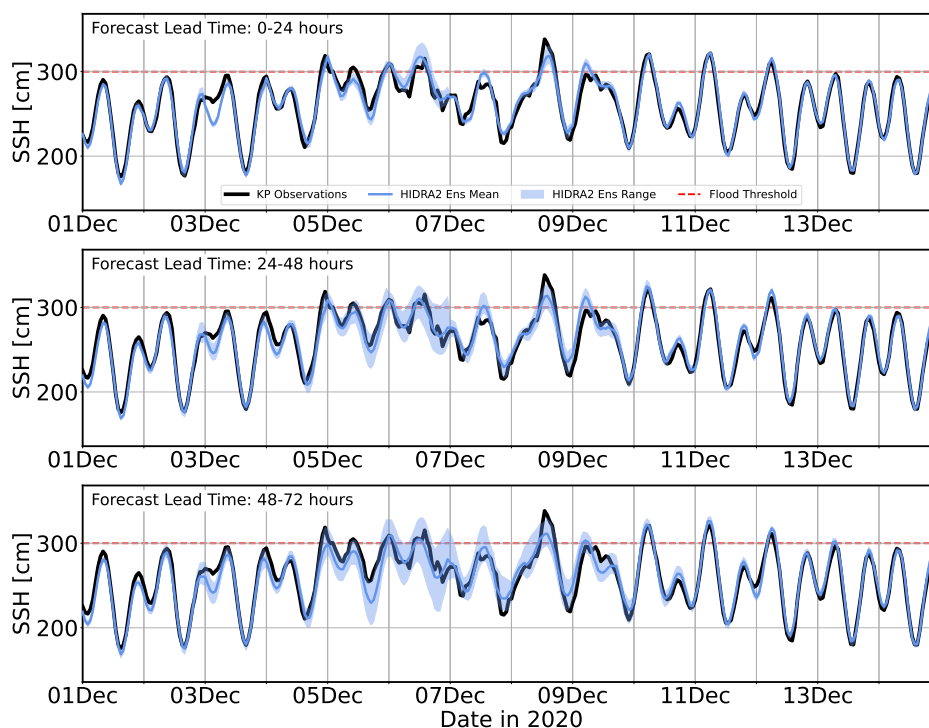


Figure 13. Comparison of HIDRA2 ensemble spread at forecast day 1 (top panel), day 2 (middle panes) and day 3 (middle panel), during December 2020 in Koper. Semi-transparent regions in the plots depict the minimum-maximum range of the HIDRA2 ensemble.

25–48 h (49–72 h) of forecast on each respective day in the evaluation time window. All three timeseries for the December
340 2020 floods are shown in Fig. 13. As expected, from the growing ensemble spread in the atmospheric forcing, HIDRA2 spread
is growing with forecast lead time as well. As we draw closer to a particular flooding event, the forecast spread drops, indicating
an increased prediction certainty.

4.2.5 Spectral decomposition of forecasts during storm surge events

To investigate the performance in geophysically relevant energy bands, we band-pass filtered the observed and the predicted
345 SSH signals in energy bands, centered around four important periods: semi-diurnal tide (12 h period), diurnal tide (24 h period),
fundamental basin along-axis eigenmode (21.5 h period) and first excited along-axis eigenmode (10.9 h period).

Although incomplete, this SSH decomposition allows qualitative estimation of the excitation intensity of the basin eigen-
modes during a particular storm surge, and also helps to qualitatively assign forecasting errors to specific frequency bands.
However, since the amplitudes of filtered signals in Fig. 14 directly depend on the filter bandwidths, they should not be inter-
350 preted as direct contributions to the sea level due to respective geophysical phenomena (i.e., two tidal signals, two eigenmodes).
They should rather be read strictly as an additional insight into the model performance within a specific band with reference to
filtered observations in the same band.



We applied a fifth-order Butterworth band-pass filter with the sampling rate of $(1\text{h})^{-1}$. Low and high cutoff frequencies, which define the semi-diurnal filtering band $\Delta\omega_{12}$, were set to $\Delta\omega_{12} = [(12.5\text{h})^{-1}, (11.5\text{h})^{-1}]$. Similarly, diurnal cutoff frequencies were set to $\Delta\omega_{24} = [(24.5\text{h})^{-1}, (23.5\text{h})^{-1}]$. Fundamental seiche filtering band was estimated from Fig. 10 to be $\Delta\omega_{21.5} = [(20\text{h})^{-1}, (24\text{h})^{-1}]$ which is also consistent with the seiche window used in Vilibić (2006). Finally, the first excited eigenmode band is defined as $\Delta\omega_{10.9} = [(11.4\text{h})^{-1}, (10.5\text{h})^{-1}]$. An example of this decomposition for November 2019 is shown in Fig. 14.

All models exhibit underestimation of the amplitude but are otherwise in phase with the observations in the $\Delta\omega_{12}$ band. In $\Delta\omega_{24}$, NEMO seems to be performing very well, with HIDRA2 slightly underestimating the range of the signal in this band. In the band $\Delta\omega_{21.5}$ NEMO is again closest to filtered observations while both HIDRA models overpredict the vertical range of the observed signal. Band $\Delta\omega_{10.9}$ is underpredicted in all models, but seems best (or rather least poorly) resolved by HIDRAS, with NEMO additionally exhibiting a substantial phase shift in the signal.

In any case, since both tidal bands and the ground state seiche are reliably predicted by all models, the reason for the forecasting errors must lie in the higher frequency bands with periods below 10.9 hours. This seems consistent with the occurrence of highly transient and localized low pressure over Venice mentioned in Cavaleri et al. (2020) and will be the subject of further research.

Similar remarks can be made regarding the December 2019 coastal flooding, depicted in Fig. 15. This event marked a sub-optimal performance of NEMO, which is systematically underestimating SSH peaks and the overall vertical range of the SSH variability during this time window (Fig. 15, top panel). This caused NEMO to miss four floods out of eight. HIDRA models perform better, with HIDRA2 most reliably predicting all the flood peaks, most notably those on 15, 16 and 23 December 2019.

The second and third panels in Fig. 15 demonstrate that all models are reliable in the diurnal tidal band $\Delta\omega_{24}$ but that HIDRA2 overestimates the signal in $\Delta\omega_{12}$. Since the overall performance of HIDRA2 is the best of all three models, it is unclear whether overshoots in $\Delta\omega_{12}$ could be interpreted as compensations for the underestimations in the nearby $\Delta\omega_{10.9}$ band. The bottom two panels in Fig. 15, however, indicate that part of the modeling errors stem from their underestimation of the basin seiches.

In the $\Delta\omega_{21.5}$ band, the HIDRA2 predictions most closely resemble the observations, followed by HIDRA1 and then NEMO (which is most severely underestimating this part of the signal). HIDRA2 is also the most reliable method in $\Delta\omega_{10.9}$ – but it nevertheless systematically underestimates the observations. HIDRA1 and NEMO performances are significantly worse, reaching one-half of the amplitude of HIDRA2 and one-third that of the observations. Poor performances of HIDRA1 and NEMO in $\Delta\omega_{21.5}$ and $\Delta\omega_{10.9}$ bands are simply another reflection of the fact depicted in Fig. 10, namely that both of these models struggle to generate an appropriate amount of energy in the bands around free oscillation eigenmodes.



5 Conclusions

385 This study presents a deep-learning based sea level model HIDRA2, suitable for operational sea level ensemble modeling due to its speed and accuracy. This work is a conceptual continuation of our previous attempt at sea level forecasting (Žust et al., 2021) and represents a substantial advancement over the first version (HIDRA1), setting a new state-of-the-art in machine learning SSH forecasting. The new architecture is validated by ablation studies. The performance is benchmarked against the current state-of-the-art Mediterranean setup of NEMO ocean model with sea level data assimilation, available as part of Copernicus Marine Service. We demonstrate that HIDRA2 outperforms both HIDRA1 and NEMO across all sea level bins. We further show that HIDRA2 most accurately represents the energy contents in the bands around relevant geophysical periods (diurnal and semi-diurnal tides, and the lowest two free oscillation basin eigenmodes).

Performance is analyzed over several historic storm surge events. Spectral decomposition of the total sea level signal into bands centered around tides and basin seiches is carried out to assign modeling errors to specific energy bands of the predicted sea levels. HIDRA2 consistently outperforms the other models by demonstrating the highest skill in exciting the ground state Adriatic basin seiche at the appropriate time and with the appropriate phase and amplitude.

HIDRA2 is a good example of how entanglement of deep-learning and geophysics may lead to reliable and numerically cheap models, which are able to mimic complex physical phenomena on the level of the best numerical physical models. Nevertheless, several extensions could be additionally explored. One possible extension is data ingestion from several tide gauges along the Adriatic coast and verification of whether the prediction accuracy at individual locations improves in such a multi-point prediction setup. Another extension is the inclusion of real-time in-situ measurements such as synoptic observations, satellite scatterometer and wind measurements. It would be interesting to migrate HIDRA2 to other semi-enclosed basins like the Baltic Sea or the Chesapeake Bay to validate generalization to other locations. These will be some of our future works.

Code and data availability. HIDRA2 code is available in the Git repository <https://github.com/rusmarko/HIDRA2> (last access: 30 June 2022). The persistent version of the HIDRA2 source code is available at <https://doi.org/10.5281/zenodo.6784842> (Rus et al., 2022a). ECMWF ensemble data are available through the Meteorological Archive and Retrieval System (MARS), but access is limited to member countries. Sea level datasets employed in this paper are available at <https://doi.org/10.5281/zenodo.6793489> (Rus et al., 2022b).

Appendix A: Additional forecasting examples

A1 December 2020

410 Coastal flood on 8 December 2020 is significantly underpredicted in HIDRA1 and even more so in NEMO. HIDRA2 demonstrates much better performance. Decompositions of the sea level signal into $\Delta\omega_i$ bands indicate that tidal bands are predicted well in all three models. The leading cause of NEMO model SSH underestimation clearly lies in the misrepresented amplitude of the fundamental basin seiche.



415 *Author contributions.* MR was the main designer of HIDRA2 and reimplemented HIDRA1 in PyTorch. MK led the machine-learning part of the research and contributed to the design of HIDRA2. ML provided the geophysical background relevant for the HIDRA2 design and led the geophysical part of the research. AF and ML prepared the atmospheric and sea level training and evaluation datasets. MR, MK and ML analyzed the results and wrote the paper. All authors contributed to the final version of the manuscript.

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



References

- 420 Bajo, M., Međugorac, I., Umgiesser, G., and Orlić, M.: Storm surge and seiche modelling in the Adriatic Sea and the impact of data assimilation, *Quarterly Journal of the Royal Meteorological Society*, 145, 2070–2084, <https://doi.org/https://doi.org/10.1002/qj.3544>, 2019.
- 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.
- 425 Cavaleri, L., Bajo, M., Barbariol, F., Bastianini, M., Benetazzo, A., Bertotti, L., Chiggiato, J., Ferrarin, C., and Umgiesser, G.: The 2019 Flooding of Venice and Its Implications for Future Predictions, *Oceanography*, 33, 42–49, <https://doi.org/10.5670/oceanog.2020.105>, 2020.
- Cerovecki, I., Orlic, M., and Hendershott, M. C.: Adriatic seiche decay and energy loss to the Mediterranean, *Deep-Sea Research Part I-Oceanographic Research Papers*, 44, 2007–2029, [https://doi.org/10.1016/s0967-0637\(97\)00056-3](https://doi.org/10.1016/s0967-0637(97)00056-3), n/a, 1997.
- 430 Clementi, E., Aydogdu, A., Goglio, A. C., Pistoia, J., Escudier, R., Drudi, M., Grandi, A., Mariani, A., Lyubartsev, V., Lecci, R., Cretí, S., Coppini, G., Masina, S., and Pinardi, N.: 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, 2021.
- 435 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.
- Ferrarin, C., Valentini, A., Vodopivec, M., Klaric, D., Massaro, G., Bajo, M., De Pascalis, F., Fadini, A., Ghezzi, M., Menegon, S., Bressan, L., Unguendoli, S., Fettich, A., Jerman, J., Ličer, M., Fustar, L., Papa, A., and Carraro, E.: 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-](https://doi.org/10.5194/nhess-20-73-2020)
- 440 2020, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J.: Identity mappings in deep residual networks, in: *European conference on computer vision*, pp. 630–645, Springer, 2016.
- 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.
- 445 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 relevance vector machine, *Global and planetary change*, 161, 211–221, 2018.
- Ioffe, S. and Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: *International conference on machine learning*, pp. 448–456, PMLR, 2015.
- 450 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.
- Karimi, S., Kisi, O., Shiri, J., and Makarynsky, O.: Neuro-fuzzy and neural network techniques for forecasting sea level in Darwin Harbor, Australia, *Computers & Geosciences*, 52, 50–59, 2013.
- Klambauer, G., Unterthiner, T., Mayr, A., and Hochreiter, S.: Self-normalizing neural networks, *Advances in neural information processing systems*, 30, 2017.
- 455



- 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.
- Madec, G.: NEMO ocean engine, https://www.nemo-ocean.eu/wp-content/uploads/NEMO_book.pdf, 2016.
- 460 Medvedev, I. P., Vilibić, I., and Rabinovich, A. B.: Tidal resonance in the Adriatic Sea: Observational evidence, *Journal of Geophysical Research: Oceans*, 125, e2020JC016168, <https://doi.org/10.1029/2020JC016168>, 2020.
- 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.
- Pashova, L. and Popova, S.: Daily sea level forecast at tide gauge Burgas, Bulgaria using artificial neural networks, *Journal of sea research*,
465 66, 154–161, 2011.
- Pérez Gómez, B., Vilibić, I., Šepić, J., Medugorac, I., Ličer, M., Testut, L., Fraboul, C., Marcos, M., Abdellaoui, H., Álvarez Fanjul, E.,
Barbalić, D., Casas, B., Castaño Tierno, A., Čupić, S., Drago, A., Fraile, M. A., Galliano, D. A., Gauci, A., Gloginja, B., Martín Guijarro,
V., Jeromel, M., Larrad Revuelto, M., Lazar, A., Keskin, I. H., Medvedev, I., Menassri, A., Meslem, M. A., Mihanović, H., Morucci, S.,
Niculescu, D., Quijano de Benito, J. M., Pascual, J., Palazov, A., Picone, M., Raicich, F., Said, M., Salat, J., Sezen, E., Simav, M., Sylaios,
470 G., Tel, E., Tintoré, J., Zaimi, K., and Zodiatis, G.: Coastal Sea Level Monitoring in the Mediterranean and Black Seas, *Ocean Science
Discussions*, 2022, 1–80, <https://doi.org/10.5194/os-2021-125>, 2022.
- Rus, M., Fettich, A., Kristan, M., and Ličer, M.: Code for HIDRA2: Deep-Learning Ensemble Storm Surge Forecasting in the Presence of
Seiches – the Case of Northern Adriatic, <https://doi.org/10.5281/zenodo.6784842>, 2022a.
- Rus, M., Fettich, A., Kristan, M., and Ličer, M.: Sea Level Datasets for HIDRA2 Training and Evaluation,
475 <https://doi.org/10.5281/zenodo.6793489>, 2022b.
- Vapnik, V.: *The nature of statistical learning theory*, Springer science & business media, 1999.
- Vilibić, I.: The role of the fundamental seiche in the Adriatic coastal floods, *Continental Shelf Research*, 26, 206–216,
<https://doi.org/https://doi.org/10.1016/j.csr.2005.11.001>, 2006.
- Žust, L., Fettich, A., Kristan, M., and Ličer, M.: HIDRA 1.0: deep-learning-based ensemble sea level forecasting in the northern Adriatic,
480 *Geoscientific Model Development*, 14, 2057–2074, <https://doi.org/10.5194/gmd-14-2057-2021>, 2021.




Figure 14. Comparison of total Koper SSH observations and forecasts (top panel) and their Band-Pass Filtered signals (bottom four panels) over four bands, centered around four geophysically relevant periods (semi-diurnal and diurnal tides and two lowest along-axis basin eigenmodes). Time window of the SSH signal spans the 7 November 2019 to 19 November 2019. Note different vertical scale in the bottom $\Delta\omega_{10.9}$ panel.

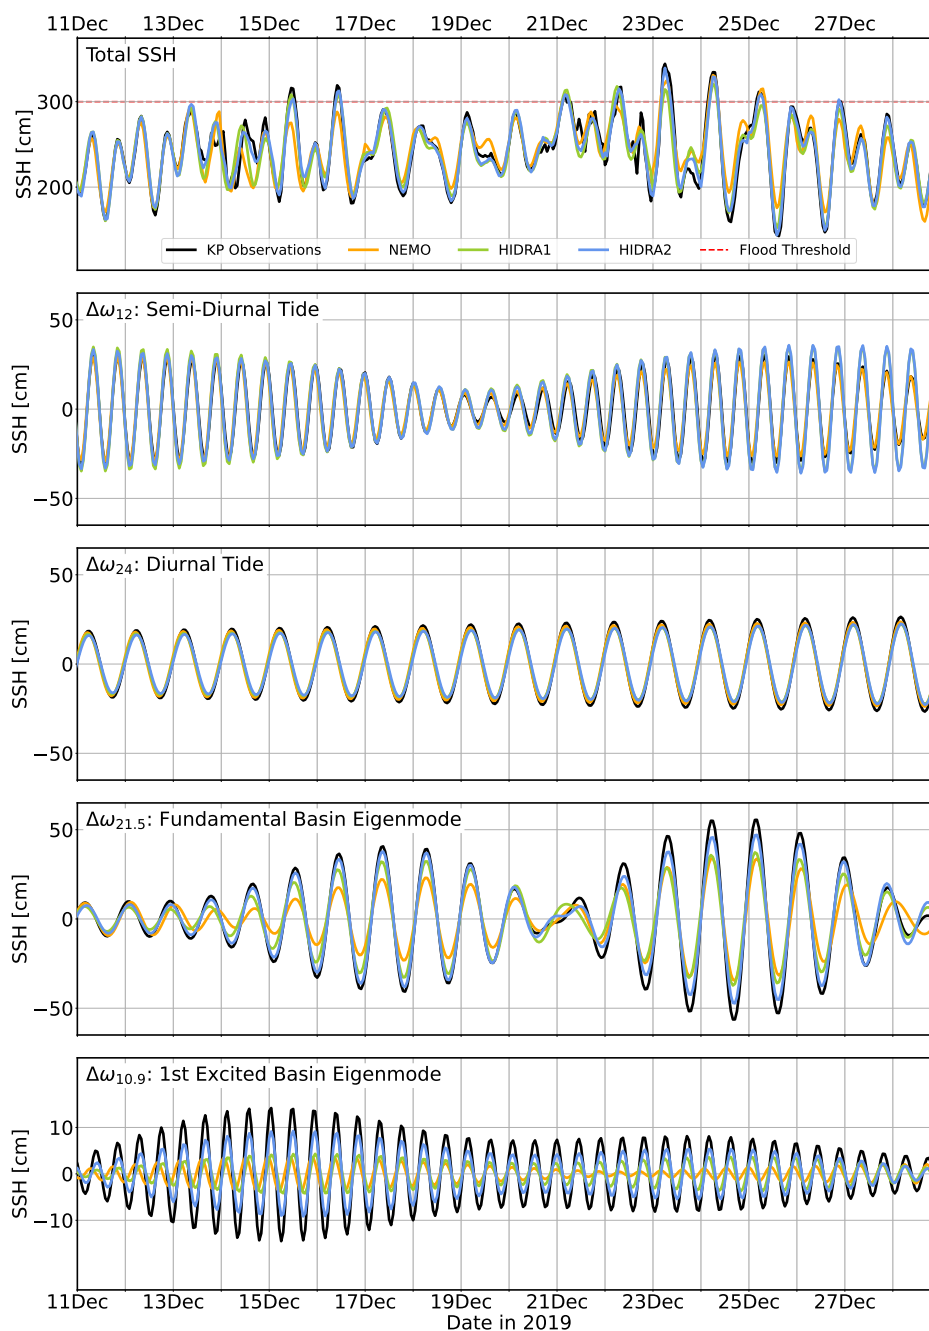


Figure 15. Same as Fig. 14, but for December 2019 coastal floods.




Figure A1. Same as Fig. 14, but for December 2020 coastal floods.