

SEA LEVEL ANOMALY PREDICTION USING RECURRENT NEURAL NETWORKS

Anne Braakmann-Folgmann, Ribana Roscher, Susanne Wenzel, Bernd Uebbing and Jürgen Kusche

Institute of Geodesy and Geoinformation, University of Bonn, Nussallee 15, D-53115 Bonn.
 Contact: {abraakmann / rroscher / susanne.wenzel / bernd.uebbing / kusche}@uni-bonn.de

ABSTRACT

Sea level change, one of the most dire impacts of anthropogenic global warming, will affect a large amount of the world’s population. However, sea level change is not uniform in time and space, and the skill of conventional prediction methods is limited due to the ocean’s internal variability on timescales from weeks to decades. Here we study the potential of neural network methods which have been used successfully in other applications, but rarely been applied for this task. We develop a combination of a convolutional neural network (CNN) and a recurrent neural network (RNN) to analyse both the spatial and the temporal evolution of sea level and to suggest an independent, accurate method to predict interannual sea level anomalies (SLA). We test our method for the northern and equatorial Pacific Ocean, using gridded altimeter-derived SLA data. We show that the used network designs outperform a simple regression and that adding a CNN improves the skill significantly. The predictions are stable over several years.

Index Terms— sea level, neural networks, CNN, RNN, deep learning, climate change, altimetry, time series analysis

1. INTRODUCTION

Modelling and predicting sea level anomalies (SLA) is currently a relevant topic, as sea level responds to global warming directly by thermal expansion of the ocean and indirectly by mass increase through melting ice sheets and glaciers. And with 60% of the world’s population living in coastal areas an accurate prediction of SLAs at a high spatial resolution is required. Sea level rise is both spatially and temporally highly variable and therefore especially hard to model. Predictions by ocean models rely on incomplete representation of physical processes, limited spatial resolution, uncertain initial conditions, and scenario-based boundary conditions.

Complementing physical ocean modelling, we develop a new approach using neural networks. Neural networks are a powerful mean to solve tasks such as classification, object detection and speech recognition, where they show promising results reaching accuracies superior to classical and shallow state-of-art machine learning algorithms [10].

So far, they have barely been explored in the context of sea level prediction. Makarynskyy et al. [6] use a single tide gauge to predict the local sea level in a harbour, and Wenzel and Schröter [13] use global tide gauge data, to derive average trends and amplitudes for eight different regions. However, they lack a high spatial resolution and the used one-layered fully connected network architectures are simple compared to current networks applied in the computer vision community.

Here, in contrast mostly convolutional neural networks (CNNs) and very deep networks are used, as they reach higher accuracies [10]. CNNs are particularly suitable for gridded data exploiting the spatial correlations. On the other hand recurrent neural networks (RNNs) and their enhancement Long Short Term Memory networks (LSTM, [3]) are specifically designed to model time series data. A combination of CNN and RNN has for example been used to predict one frame ahead in video sequences [8]. With sequence-to-sequence LSTMs [12] it is possible to predict several time steps ahead.

Our aim is first to design a sophisticated network composed of a CNN and an RNN to capture both spatial and temporal relations of radar-altimetric SLA fields in the Pacific Ocean and to predict the SLA map of the next month. Then we extend the prediction length to several years.

2. DATA

We use time series of SLAs from ESA (European Space Agency) CCI (Climate Change Initiative) [7] as input to our network. SLAs are the difference between actual sea surface height (SSH) and mean sea surface height. We use the Level 4 product of ESACCI, where all altimetry mission measurements have been merged into monthly grids with a spatial resolution of 1/4 degree. The data cover 23 years from January 1993 to December 2015. To validate the predictive skill, we split the data into training- (16 years), validation- (4 years) and test data (3 years).

The northern and equatorial Pacific make a perfect test region, since parts of it undergo rapid sea level rise while being affected by strong interannual and decadal variability [9]. Here we examine the region between 110° and 250° longitude and 15°S to 60°N latitude. This region holds 170,240 grid points, which give us almost 33 million values for training of the network and 47 million values overall.




Fig. 1. Different network designs. The input \dot{x} is all SLAs at one time step t . In a) and c) they are reshaped to a vector, in b) the grid structure is preserved. The network predicts the vector/grid y at the next one month (in a) and b)) or nine months (in c))

3. METHODOLOGY

SLAs reflect various complex non-linear interactions in the ocean and hence contain several deterministic and stochastic modes that are difficult to predict. Yet, SSH maps contain many reoccurring spatial patterns (e.g. eddy fields, gyres, ENSO) that can potentially be learned by neural networks.

Artificial neural networks are a means of machine learning inspired by the human brain to learn higher-order representations and perform diverse tasks. In contrast to other machine learning techniques, neural networks are able to extract relevant features and their weight in the model. CNNs have shown to be efficient networks which are especially designed to capture spatial dependencies by using gridded input data.

RNNs are neural networks that can deal with time series analysis, taking sequences as in- and output [12]. In contrast to feedforward networks they incorporate a self-loop, which enables the net to memorise the previous inputs (horizontal information flow in Fig. 1). LSTMs [3] are a special kind of RNN that are even capable of learning long-term dependencies. Convolutional LSTMs (ConvLSTM, [11]) preserve the input's spatial structure by replacing all matrix multiplications within the LSTM with convolutions. To learn higher-order representations, both RNNs and CNNs can be stacked on top of each other.

4. EXPERIMENTAL SETUP

We design different models using the framework Keras [1] with Tensorflow backend. We train them with the Adam optimizer [4] to minimise the MSE between their prediction and the truth in 150 epochs using decaying learning rates.

LSTM network: Here the input is a vector with all SLAs of one month. The network consists of three LSTM layers with 60 units each, followed by a fully connected layer that learns how to make a prediction of the next month for each

$1/4^\circ$ grid cell from the 60 features (Fig. 1a). This prediction is our output. The network consists of over 51 million parameters to train. We employ a hard sigmoid activation function for the hidden-to-hidden state translations and a tanh activation function in the LSTMs' hidden-to-output transformation, allowing the net to learn non-linear relations. To further improve the generalisation capability we employ LSTM-modified dropout [2] of 0.8. On an Intel i5-2400 CPU training takes 8 hours. After each prediction of one month we use the true values of the preceding month as input for the next prediction.

CNN+ConvLSTM network: Here the input SLAs are used in their natural grid structure, preserving spatial information. In this network we combine a CNN with a ConvLSTM as shown in Fig. 1b. The CNN consists of two convolutional layers with 32 filters, a kernel size of 3×3 and a ReLU activation function each, followed by a pooling layer with kernel size 4×4 and stride 4. We use the extracted feature maps as input to two ConvLSTM layers with 40 units each, a kernel size of 3×3 and the same activation functions as used in the LSTM. The last layer is a deconvolutional (DeConv, [14]) layer to regain the spatial extent of our input and to map the 40 feature maps to a single SLA value at each grid cell. This map of SLAs is our output prediction for the next month. Each predicted grid cell has a receptive field of 6° . This net has to learn only 229,413 parameters due to pooling. We apply batch normalisation to the input. Dropout is not used. On an Intel i5-2400 CPU training this net takes 18 hours.

Sequence LSTM network: This network is equal to the LSTM network, but takes sequences of nine months as in- and output (see Fig. 1c). While the former two networks only predict one month in advance, this network is able to predict the following nine months at once. After the prediction of a nine months long sequence, we use the true values of the preceding nine month as input for the next prediction. On the Intel i5-2400 CPU training this net takes 8 hours, too.

Sequence LSTM-P network: This network is trained just like the Sequence LSTM. However here we take the pre-

dicted nine months as input to predict nine months further.

5. RESULTS

In this section we assess the networks' performance and compare the predictions made by our networks to a regression approach. We use the sum of a trend, acceleration, annual and semi-annual sine and cosine as regression model. The regression parameters are estimated for each grid cell individually.

Figure 2 shows the measured time series of SLAs at one point in the North Pacific Gyre. The position is marked with a cross in Fig. 3. We visualize the prediction during training (blue), validation phase (green) and the actual test phase (red), compared to the known true data (black).




Fig. 2. Results for one grid cell (marked by a cross in Fig. 3): True (black) and predicted (blue = training, green = validation and red = test data) time series of SLAs

Table 1 shows the RMSEs between true SLAs and the pre-

RMSE averaged over	training data	validation data	test data
LSTM network	0.062 m	0.071 m	0.076 m
CNN+ConvLSTM	0.047 m	0.050 m	0.051 m
Sequence LSTM	0.059 m	0.079 m	0.077 m
Sequence LSTM-P	0.083 m	0.080 m	0.081 m
Regression	0.078 m	-	0.154 m

Table 1. RMSEs between the networks' or the regression's prediction and the true SLAs averaged over all grid cells

dictions made by our networks or the regression. The training error of the Sequence LSTM-P is higher, because we here feed the predictions 20 times back into the network.

To examine the spatial patterns, we plot the true and predicted SLAs in November 2014 in Fig. 3. All networks are trained using data up to 2008. Starting the test phase in January 2013, the Sequence LSTM-P needs nine true months as input only once (i.e. till September 2013). The Sequence LSTM depends on true data every nine months, so June 2014 is the last true month needed to predict the sequence including November 2014. The LSTM and CNN+ConvLSTM depend on true inputs of the previous month (here October 2014).

Striking, overall all network architectures outperform the regression. Especially in Fig. 3f it can be seen that the regression fails at capturing the spatial structure. We observe that an additional CNN improves the accuracy significantly (compare Fig. 2a/b and 3b/c). In Fig. 3c the CNN+ConvLSTM network resolves nearly all spatial structures very well. However, we observe that in combination with a normal LSTM, the CNN brings no improvement (not shown here) - only if we also keep the spatial structure throughout the ConvLSTM.

Both extending the prediction length to nine months (compare Fig. 2a/c and 3b/d) and using the predicted values for the next prediction (compare Fig. 2c/d and 3d/e) leads to only slight degradation of predictive skill and the predictions stay close to the real values for a long time.

Figure 4 shows the spatial distribution of the RMSE averaged over the test phase for the CNN+ConvLSTM and the Sequence LSTM-P. In both cases the main errors occur around the Kuroshio current with its highly variable eddies.

6. CONCLUSION

In this work we develop a combined convolutional and recurrent neural network to analyse both the spatial and the temporal evolution of SLAs in the northern and central Pacific Ocean and to make accurate predictions for the future sea level. We validate the accuracy of our approach and compare the results to a regression. All the used network designs outperform the regression. Adding a convolutional neural network improves the accuracy significantly. Therefore we plan to extend this architecture to take sequences as in- and output. We are confident that in this way we will be able to improve predictions at longer timescales. The developed method could also be applied to other regions or even globally.




Fig. 3. Spatial structure of the SLAs in the selected region in November 2014 (within the test period)




Fig. 4. RMSE at each grid cell averaged over the test phase

7. REFERENCES

- [1] F. Chollet et al.: Keras. Published on GitHub (<https://github.com/fchollet/keras>), (2015)
- [2] Y. Gal and Z. Ghahramani: A Theoretically Grounded Application of Dropout in Recurrent Neural Networks. In NIPS, (2016)
- [3] S. Hochreiter and J. Schmidhuber: Long short-term memory. In Neural computation, (1997)
- [4] D. P. Kingma and J. Lei Ba: ADAM: A Method for Stochastic optimization. In ICLR, (2015)
- [5] P. Khorrami et al.: How Deep Neural Networks Can Improve Emotion Recognition on Video Data in Image Processing. In ICIP, (2016)
- [6] O. Makarynskyy et al.: Predicting sea level variations with artificial neural networks at Hillarys Boat Harbour, Western Australia. In Estuar. Coast. Shelf Sci., (2004)
- [7] G. D. Quartly et al.: A new phase in the production of quality-controlled sea level data. In Earth Syst. Sci. Data Discuss., (2017)
- [8] M. Ranzato et al.: Video (language) modeling: a baseline for generative models of natural videos. In CVPR, (2015)
- [9] R. Rietbroek et al.: Revisiting the contemporary sea-level budget on global and regional scales. In PNAS, (2016)
- [10] J. Schmidhuber: Deep Learning in Neural Networks: An Overview. In Neural Networks, (2015)
- [11] X. Shi, Z. Chen, H. Wang and D.-Y. Yeung: Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting. In NIPS, (2015)
- [12] I. Sutskever, O. Vinyals and Q. V. Le: Sequence to sequence learning with neural networks. In NIPS, (2014)
- [13] M. Wenzel and J. Schröter: Reconstruction of regional mean sea level anomalies from tide gauges using neural networks. In JGR, (2010)
- [14] M. D. Zeiler, D. Krishnan, G. W. Taylor and R. Fergus: Deconvolutional Networks. In CVPR, (2010)