

# Deep sequence-to-sequence neural networks for ionospheric activity map prediction

Noëlie Cherrier, Thibaut Castaings, and Alexandre Boulch

ONERA, The French Aerospace Lab,  
F-91761 Palaiseau, France

**Abstract.** The ability to predict the ionosphere activity is of interest for several applications such as satellite telecommunications or Global Navigation Satellite Systems (GNSS). A few studies have proposed models able to predict Total Electron Content (TEC) values of the ionosphere locally over measuring stations, but not worldwide for most of them. We propose a method using Deep Neural Networks (DNN) to predict a sequence of global TEC maps consecutive to an input sequence of past TEC maps, by combining Convolutional Neural Networks (CNNs) with convolutional Long Short-Term Memory (LSTM) networks. The numerical experiments show that the approach provides significant improvement over methods implemented for benchmarking and is competitive with state-of-the-art methods while providing global TEC predictions. The proposed architecture can be adapted to any sequence-to-sequence prediction problem.

**Keywords:** sequence prediction, neural network, forecasting, TEC, ionosphere, deep learning, CNN, LSTM




Fig. 1: Ground-truth (left) and predicted (right) TEC map example

## 1 Introduction

Ionospheric activity is mainly measured by the Total Electron Content (TEC), which is the total number of electrons in the ionosphere integrated along a vertical path above a given location. It is expressed in TEC Units (1 TECU =

$10^{16} \text{ el}/\text{m}^2$ ), usually ranging from a few units to one hundred TECU. During periods of high ionospheric TEC levels, the services provided by satellite telecommunication systems and Global Navigation Satellite Systems (GNSS) may be deteriorated due to changes in the paths of transionospheric radio waves inducing significant bitrate reduction and positioning errors [1, 2]. As a consequence, forecasting TEC globally (*i.e.* worldwide) in the short term (up to two hours) or in the longer term (a few days) increases the ability of the users of these services to evaluate, as one example, data loss or mispositioning risks in operations planning.

Several services relying on measurements provided by GNSS ground networks [3] exist to address this issue at a global scale, *i.e.* forecasting TEC values worldwide. CTIPe, an experimental tool implementing complex physics models [4] developed by the US Space Weather Prediction Center, produces global forecasts 30 minutes ahead of real-time. In Europe, the ESA Ionospheric Weather Expert Service Center combines products from different national services to provide global and regional 1-hour TEC forecasts. However, the records of the input data and forecasts are not published.

In the literature, several methods have been published to predict TEC a few minutes to several days ahead above specific stations using time series analysis and statistical methods [5–9]. Among these, a few provide a reconstruction of a small area [10, 11] with methods such as Bezier surface-fitting or Kriging. Several works are based on artificial neural networks [10, 12, 13], but they only focus on local stations. In order to make a prediction at a regional or global scale with these approaches, one model for each location must be adapted. In [14], a global analytical TEC model is proposed to address the global forecasting issue, using open source TEC data from the Center for Orbit Determination in Europe (CODE). This model is intended to apply to any temporal range, without relying on a record of TEC values.

In this paper, the objective is to predict global TEC maps 2 to 48 hours ahead of real-time. The proposed approach is entirely based on recurrent neural networks, taking advantage of convolutional and Long Short-Term Memory (LSTM) networks. Deep Neural Networks (DNN) have the advantage to enable complex modeling of large input data, such as global TEC maps in this case, with little or no prior knowledge.

The paper is organized as follows: Section 2 presents the approach used in this paper, including the data, the proposed network architecture and the training procedure; Section 3 discusses the results of our numerical experiments on TEC prediction; and finally the limits and perspectives of the proposed approach are discussed in Section 4.

## 2 Proposed approach

The objective is to design a neural network architecture able to be fed with an input sequence of a number of maps and to output the next 48 hours of TEC maps.

## 2.1 Datasets

Open source TEC data from the CODE is used in this study. Two datasets of TEC maps are available: Rapid TEC maps, which are accessible quicker, and Final TEC maps, more precise. In this study, final TEC maps are used with a  $5^\circ \times 2.5^\circ$  resolution on longitude and latitude and 2 hour temporal resolution, covering all latitudes and longitudes. One pixel in these maps represents the vertical TEC at this point.

Data from 1/1/2014 to 5/31/2016 is used for training, and data from 7/1/2016 to 12/31/2016 for testing. Splitting the data into sequences corresponding to five consecutive days allows for convincing results in terms of convergence and in the test procedure (see Section 3). Data is loaded as a sequence of 60 maps (one map every two hours): the first 36 maps (*i.e.* 3 days) are fed to the network, the last 24 maps (*i.e.* 48 hours) being the prediction targets. This amount of input is sufficient to understand the context, and more information would lead to an overuse of the network’s neurons (several neurons being dedicated to process out-of-date data).

## 2.2 Preprocessing

In order to fit the data to the input of the network and scale the dynamics of the TEC maps, the data undergoes feature normalization and is resized to  $72 \times 80$  pixels. To prevent the algorithm from focusing on deterministic phenomena that can be easily isolated, the frame of reference is also changed from Fixed-Earth to Heliocentric, so that the effect of Earth’s rotation is no longer visible to the DNN architecture. The transformed data is then the input of the network.

However, the resulting TEC maps appear to have a residual low spatial frequency, 24h-periodicity due to the Earth’s magnetic poles, distinct from the geographic poles. In order to avoid learning the residual periodicity, we choose the target  $T$  of the neural network to be the difference between the true TEC map  $I$  and a gaussian filtered TEC map at  $t-48h$   $I_{t-48h}^{blurred}$ , so that the subtracted component can always be computed from the input sequence. The gaussian filter is applied in order to remove high frequency variations from the past map.

$$T_t = I_t - I_{t-48h}^{blurred} \quad (1)$$

## 2.3 Network architecture

Deep Neural Networks (DNN) taking as input a sequence of images to output *e.g.* audio description or video labels have become a common topic in the past years [15]. However, sequence prediction problems in which both the inputs and targets are a sequence of images have barely been handled. In [16], the authors propose an architecture for precipitation nowcasting and introduce a convolutional LSTM structure.

To give an outline, the proposed approach uses Convolutional Neural Network (CNN) layers [17] to extract spatial features from an input TEC map. The

resulting features are then fed to a convolutional LSTM [16], derived from the fully-convolutional LSTM [18]. This enables to handle spatially structured data. Deconvolutional layers are then fed with the output of the convolutional LSTM in order to generate the residual TEC map corresponding to the next time step.

Fig. 2 presents the network architecture. To go into the details, we first explain the architecture designed to predict one TEC map (2 hours ahead of real-time), represented as the dotted black rectangle in Fig. 2. The network is built as the repetition of a column composed of three modules: the **encoder** (four  $3 \times 3$  convolutions with Rectified Linear Unit (ReLU) in between), a **convolutional LSTM cell** (with a  $3 \times 3$  convolution operation and ReLU) and a **decoder** (four  $3 \times 3$  deconvolutions and ReLU) in order to produce an output of the same size as the input. The columns are cascaded to enable the temporal information to go through the network, each column handling a single TEC map. The "many-to-one" architecture (dotted box in Fig. 2) is the starting point of the final architecture.

In order to generate a sequence of TEC maps, we derive the "many-to-one" architecture to obtain a "many-to-many" network able to output 48 hours of TEC maps. The same columns as in the "many-to-one" part are reused. The prediction process is achieved by recursively feeding the next column of the network with the sum of the last prediction and the filtered TEC map at  $t - 48h$  (see Subsection 2.2). The cost function is finally summed over the differences produced by the successive prediction columns.

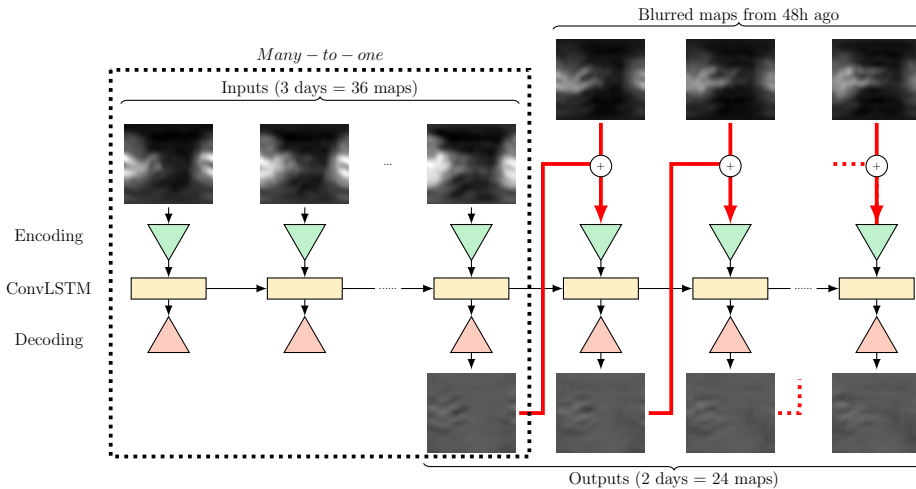


Fig. 2: Network architecture: encoders (*green*), convolutional LSTM cells (*yellow*), and decoders (*red*)

This architecture is inspired from the Convolutional Autoencoders [19] with a LSTM inserted in between. It is preferred to other well-known architectures

such as Fully Convolutional Autoencoders or Deep Belief Networks [20] which cannot handle 2D data. This approach is also preferred to three simpler methods for performance reasons:

- *Many-to-one*: the loss is computed only over one map and the sequence is recursively forecast from the previously predicted map. It leads to competitive results for 2h prediction but diverges quickly (see the results in Subsection 3.1).
- *Direct TEC map prediction*: this approach performs poorly. As a comparison, taking the map corresponding to  $t - 48h$  as a prediction leads to better performance (average RMS of 2.810 against 2.353 for the selected architecture).
- *Difference maps as input*: this method does not perform better than the selected approach (average RMS of 2.918 against 2.353 for the selected architecture).

**Cost function** The cost function for the DNN training is the pixel-wise  $\ell_1$ -norm between the predicted maps and the targets, summed with a measure of relative error as a regularizing factor in order to improve the performance in areas where the TEC level is low. The  $\ell_1$ -norm is chosen instead of  $\ell_2$ : the  $\ell_1$ -norm is known for capturing more details whereas the  $\ell_2$ -norm tends to smooth the results. In this study, we obtained an average RMS of 2.524 with  $\ell_2$ -norm and 2.353 with  $\ell_1$ -norm for the selected architecture. The cost function  $\mathcal{L}$  is given as:

$$\mathcal{L} = \sum_{t \in \mathcal{S}} w_t \sum_{i \in \mathcal{M}^t} \left( |P_i^t - T_i^t| + \frac{|P_i^t - T_i^t|}{T_i^t} \right) \quad (2)$$

with  $\mathcal{S}$  the sequence of TEC maps,  $w_t$  the weight at time  $t$ ,  $\mathcal{M}^t$  the TEC map at  $t$ ,  $P$  the predicted map and  $T$  the ground-truth map, where  $t$  indexes time and  $i$  is the map pixel index.

Different weight arrays  $[w_t]$  are considered, corresponding to different learning strategies: *e.g.* putting more weight on the first predicted TEC maps, arguing that the accuracy of the following predictions depend on them. In this study three weight profiles are taken into account: uniform weights over the 48-hour sequence; uniform weights on a 24h-window; linear decreasing weights over the 48-hour sequence. These weight arrays enable a better control over the temporal error and a possible improvement of performance at several prediction horizons.

**Data loading into the network** The sequences are loaded chronologically from the training set, from 1/1/2014 to 5/31/2016, and reloaded again from the beginning of the dataset until convergence in training. This sequential sampling approach is more efficient than random sampling in this study (average RMS of 3.311 against 2.353 with the sequential sampling).

### 3 Results

Once it has been trained, the network can be fed with any 3-day sequence from the test set and produce 2-day forecasts consecutive to this sequence.

**Baselines** Two simple prediction methods are implemented for benchmarking:

- *The 3-day mean prediction* (or constant prediction): the mean map computed over the three days of input data is the prediction result for the next 48 hours of data.
- *The periodic prediction*: the predicted map at time  $t$  is exactly the map at  $t - 48\text{h}$ . This exploits the 24 hour-periodicity of TEC maps (see Section 2.2).

The three alternative architectures presented in Subsection 2.3 are also considered as baselines.

We finally compare our method to previous works [5, 8, 12, 14] in Subsection 3.3.

**Performance criterion** The RMS error is used to compare the performance of the different methods. Other performance criteria may be taken into account, such as the Mean Absolute Percentage Error (MAPE) which is used in several related works to evaluate the forecasts [6, 9]. However, MAPE is not suited to this study since we predict global TEC maps, which include low-level TEC areas (*e.g.* Earth poles) where a small estimation deviation leads to very high MAPE (see discussion in Section 4).

#### 3.1 Quantitative results

The RMS errors averaged for each prediction horizon for the three weighting profiles and for the baselines are summarized in Table 1 for the period from 10/20/2016 to 12/20/2016. In any case, the proposed approach performs clearly better than the 3-day mean prediction, equal or better than the periodic prediction at all prediction horizons, and is also significantly better than the three alternative architectures presented in Subsection 2.3.

The global generic model proposed in [14] achieves a mean RMS deviation of 7.5 TECU at any given time, whereas the approach proposed in this article performs significantly better for 48-hour ahead forecasting (2.4 TECU in average for the uniform weighting).

Decreasing weighting is the best for the first prediction horizons, most likely because it has the bigger weights at the beginning. The weights emphasizing only the first 24 hours provide good results until 24 hours as well, and then the uniform weighting takes over the other two, since no other weight arrays focus on the last part of the forecasting range. Combining weight arrays enables a good overall optimization of the error.

Direct TEC map prediction does not perform as well as the selected architecture. The network may have to dedicate most of its weights to learn the periodic

Table 1: Mean RMS of models at a given prediction horizon over period from 10/20/2016 to 12/20/2016

Method	2h	6h	12h	18h	24h	30h	36h	42h	48h
<b>Baselines:</b>									
3-day mean	3.072	3.102	3.121	3.140	3.173	3.217	3.241	3.260	3.291
Periodic (48h)	2.518	2.516	2.514	2.514	2.510	2.509	2.509	2.508	2.505
Direct TEC map	2.712	2.639	2.601	2.575	2.728	3.001	3.003	2.963	3.142
Difference as input	2.447	2.800	2.904	2.898	2.915	2.958	3.007	3.041	3.041
Many-to-one	1.996	3.055	5.356	5.658	7.240	7.586	8.671	8.926	9.427
<b>Selected architecture:</b>									
- Uniform weighting	2.050	2.265	2.254	2.256	2.331	<b>2.434</b>	<b>2.459</b>	<b>2.475</b>	<b>2.490</b>
- Decreasing weights	<b>1.987</b>	2.239	2.248	2.258	2.409	2.585	2.635	2.647	2.674
- Only first 24h	2.078	<b>2.195</b>	<b>2.197</b>	<b>2.214</b>	<b>2.297</b>	2.486	2.513	2.500	2.498

component and fails to catch transformations of interest. Similarly, giving the difference maps as input to the network does not improve the performance. The reason might be that a new temporal dependency is induced while subtracting the past maps, making it harder to extract the relevant phenomena.

Considering the record of TEC values in order to forecast 48 hour of TEC maps, with any weighting, helps improving over a generic model such as [14] which does not take advantage of the last TEC data.

### 3.2 Performance over the test period

The RMS value at 12 hours ahead of the input sequence is evaluated for sequences in the test period from 10/01/2016 to 12/31/2016 (see Fig. 3).

RMS errors are stationary after 10/20/2016. They reach an upper point at 7-8 TECU around 10/15/2016 (see Fig. 3). However, during the day 10/17/2016, the proposed algorithm is not outperformed by the periodic prediction.

The algorithm is designed to predict residual differences between TEC maps separated by 48 hours. In this case, there is a significant difference between the maps on 10/15/2016 and 10/13/2016. Such perturbations may be linked to external phenomena such as solar events or geomagnetic storms (discussed in Section 4).

### 3.3 Indicative comparison to local prediction methods

As an attempt to compare to previous works on TEC forecast, Table 2 presents a synthesis of the performance obtained in three related studies and an indicative comparison with the results presented in this paper. The comparison is only indicative since these works differ by their prediction horizons or by the considered areas and since several studies focus on one or a few specific measuring stations instead of producing a worldwide TEC prediction. The last column

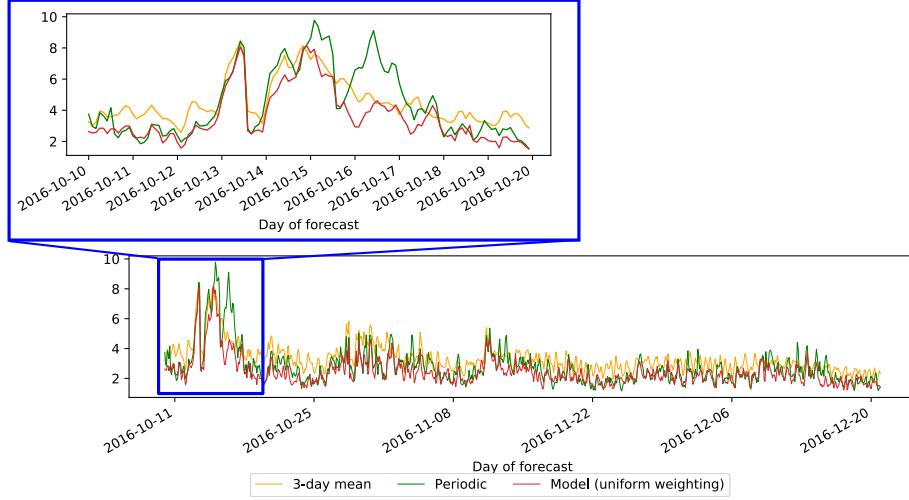


Fig. 3: RMS value for 12 hours ahead forecast over whole test period (bottom) and zoom on a disturbed period (top)

shows the RMS obtained with the proposed approach, computed as the mean of RMS values at the latitude(s) of the station(s) used by the cited papers.

Table 2: Results of previous works

Reference	Description	RMS (ref)	RMS (proposed)
[5] Chunli D., Jinsong P.	1 day forecast, 22°N in China region (2008 data)	1.45	2.049
[12] Huang, Z., Yuan, H.	24h forecast, three stations at latitudes 39.61°N, 30.53°N, 25.03°N (2011 data)	$\leq 2$	1.936
[8] Niu, R. <i>et al.</i>	Forecast of global mean TEC value (0 to 48h ahead (2012-2013 data))	3.1	0.800

The obtained results are competitive with state-of-the-art models. To put this in perspective, one should take into account the experimental periods of these studies. Periods from 2011 to 2016 are considered as an active period for the ionosphere due to solar activity, whereas 2008 is a very calm year. Good performance is easier to achieve in calm periods, as explained in [6]. The data used in this study is taken from 2014 to 2016, which makes the results from [12]



and [8] comparable. On the other hand, the better results of [5] are emphasized by the fact that the experiments are done during a calm year.

The proposed approach provides an accurate mean TEC value, significantly improving the results from [8]. In this study, giving more detailed data as input (*i.e.* a global TEC map) may indeed have helped to produce a better prediction than the one obtained by extrapolating from TEC mean values alone in [8].

## 4 Limits and perspectives

The presented method overcomes the baseline methods and is competitive with state-of-the-art local approaches measuring the RMS error.

MAPE (*i.e.* relative error) is mentioned in Subsection 3 as an alternative to RMS to evaluate the forecasts. Since global TEC maps are produced, the comparison with papers that use MAPE to evaluate local predictions is irrelevant as things stand. In this paper, we prioritized the achievement of a good global minimization with the  $\ell_1$ -loss rather than a good relative error. In order to investigate further the proposed approach, we could consider minimizing the MAPE instead of the  $\ell_1$ -loss. To avoid being over-influenced by ill-conditioned MAPE values, North and South poles (*i.e.* areas where TEC values are very small) could be cropped from TEC maps, so that the remaining TEC values are similar to the TEC values considered in [6] and [9], making fair comparison possible in future works.

During the implementation, trying to increase the complexity of the DNN by adding parameters and/or convolutional layers to the network did not improve the performance. Either another network is more appropriate to learn more complex dependencies, or there are dynamics in TEC evolution that can not be learned from TEC maps alone. For example magnetosphere or solar particles can significantly impact the ionosphere [21, 22] and its evolution. As a consequence, two possibilities are offered to try to improve performance: train another network architecture such as U-net architecture [23], or include solar information in the input to investigate the capacity of a DNN architecture to learn the dependency between TEC levels and solar activity. Both options require learning from a larger training set (data from CODE is available from 1998 to today, whereas only data from 2014 to 2016 is used in this study) in order to have more data from perturbation periods for the network to infer the right behavior.

## 5 Conclusion

A promising model able to forecast global TEC maps 2 to 48 hours ahead of real-time is proposed. This model takes advantage of Deep Neural Network architectures with CNNs and LSTM networks. The proposed approach is able to output a sequence of global TEC maps with comparable or smaller errors compared to single-station prediction models. In the specific case of [8], the global mean TEC value can also be significantly more accurately forecast.

First, the proposed approach forecasts global TEC maps as opposed to local TEC estimates proposed in other published works. Moreover, our method enables range forecasting from 2 to 48 hours ahead of real-time, which is a further horizon than the predictions provided by existing operational services.

This method opens new possibilities for using Deep Neural Networks to forecast global TEC maps: the proposed approach can be adapted to other needs such as forecasting a specific region or minimizing another criterion. The presented architecture can also be easily extended to other sequence prediction problems in which both the inputs and prediction targets are sequences of images.

## References

1. Datta-Barua, S., Lee, J., Pullen, S., Luo, M., Ene, A., Qiu, D., Zhang, G., Enge, P.: Ionospheric threat parameterization for local area GPS-based aircraft landing systems. *Journal of Aircraft* 47(4), 1141–1151 (2010)
2. Lee, J., Datta-Barua, S., Zhang, G., Pullen, S., Enge, P.: Observations of low-elevation ionospheric anomalies for ground-based augmentation of GNSS. *Radio Science* 46(6), 1–11 (2011)
3. Tulunay, E., Senalp, E.T., Cander, L.R., Tulunay, Y.K., Bilge, A.H., E. Mizrahi, S.S.K., Jakowski, N.: Development of Algorithms and Software for Forecasting, Nowcasting and Variability of TEC. *Annals of Geophysics, Supplement to Volume 47* 47(2/3), 1201–1214 (2004)
4. Millward, G.H., Muller-Wodarg, I.C.F., Aylward, A.D., Fuller-Rowell, T.J., Richmond, A.D., Moffett, R.J.: An investigation into the influence of tidal forcing on F region equatorial vertical ion drift using a global ionosphere-thermosphere model with coupled electrodynamics. *Journal of Geophysical Research: Space Physics* 106(A11), 24733–24744 (2001)
5. Chunli, D., Jinsong, P.: Modeling and prediction of TEC in China region for satellite navigation. In: 2009 15th Asia-Pacific Conference on Communications. pp. 310–313 (Oct 2009)
6. Elmunim, N.A., Abdullah, M., Hasbi, A.M.: Improving ionospheric forecasting using statistical method for accurate GPS positioning over Malaysia. In: 2016 International Conference on Advances in Electrical, Electronic and Systems Engineering (ICAEES). pp. 352–355 (Nov 2016)
7. Li, X., Guo, D.: Modeling and prediction of ionospheric total electron content by time series analysis. In: 2010 2nd International Conference on Advanced Computer Control. vol. 2, pp. 375–379 (March 2010)
8. Niu, R., Guo, C., Zhang, Y., He, L., Mao, Y.: Study of ionospheric TEC short-term forecast model based on combination method. In: 2014 12th International Conference on Signal Processing (ICSP). pp. 2426–2430 (Oct 2014)
9. Zhenzhong, X., Weimin, W., Bo, W.: Ionosphere TEC prediction based on Chaos. In: ISAPE2012. pp. 458–460 (Oct 2012)
10. Tulunay, E., Senalp, E.T., Radicella, S.M., Tulunay, Y.: Forecasting total electron content maps by neural network technique. *Radio Science* 41(4) (2006)
11. Wu, Y.W., Liu, R.Y., Jian-Ping, W., Wu, Z.S.: Ionospheric TEC short-term forecasting in CHINA. In: Proceedings of the 9th International Symposium on Antennas, Propagation and EM Theory. pp. 418–421 (Nov 2010)

12. Huang, Z., Yuan, H.: Ionospheric single-station TEC short-term forecast using RBF neural network. *Radio Science* 49(4), 283–292 (2014)
13. Senalp, E.T., Tulunay, E., Tulunay, Y.: Total electron content (TEC) forecasting by Cascade Modeling: A possible alternative to the IRI-2001. *Radio Science* 43(4) (2008)
14. Jakowski, N., Hoque, M.M., Mayer, C.: A new global TEC model for estimating transionospheric radio wave propagation errors. *Journal of Geodesy* 85(12), 965–974 (2011)
15. Donahue, J., Hendricks, L.A., Guadarrama, S., Rohrbach, M., Venugopalan, S., Saenko, K., Darrell, T.: Long-term Recurrent Convolutional Networks for Visual Recognition and Description. CoRR abs/1411.4389 (2014)
16. Xingjian, S., Chen, Z., Wang, H., Yeung, D.Y., Wong, W.K., Woo, W.c.: Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In: *Advances in NIPS*. pp. 802–810 (2015)
17. Jarrett, K., Kavukcuoglu, K., Ranzato, M., LeCun, Y.: What is the best multi-stage architecture for object recognition? In: *2009 IEEE 12th ICCV*. pp. 2146–2153 (Sept 2009)
18. Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural Comput.* 9(8), 1735–1780 (Nov 1997)
19. Masci, J., Meier, U., Cireşan, D., Schmidhuber, J.: Stacked Convolutional Auto-Encoders for Hierarchical Feature Extraction, pp. 52–59. Springer Berlin Heidelberg (2011)
20. Hinton, G.E., Salakhutdinov, R.R.: Reducing the Dimensionality of Data with Neural Networks. *Science* 313(5786), 504–507 (2006)
21. Webb, D.F.: Coronal mass ejections: origins, evolution, and role in space weather. *IEEE Transactions on Plasma Science* 28(6), 1795–1806 (Dec 2000)
22. Wells, H.W.: Effects of Solar Activity on the Ionosphere and Radio Communications. *Proceedings of the IRE* 31(4), 147–157 (April 1943)
23. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional Networks for Biomedical Image Segmentation. CoRR abs/1505.04597 (2015)