Pergamon

Phys. Chem .Earth (C), Vol. 25, No. 4, pp.267-273, 2000

1464-1917/00/$ - see front matter

PII: S1464-1917(00)00015-5

Ionospheric foF2 Storm Forecasting using Neural Networks

P. Wintoft[1] and Lj. R. Cander[2]

[1]Swedish Institute of Space Physics, Solar-Terrestrial Physics Division, Lund, Sweden
[2]CLRC Rutherford Appleton Laboratory, Chilton, Didcot, Oxon, OX11 0QX, England

Received 21 May 1999; revised 17 June 1999; accepted 15 September 1999

Abstract. The ionosphere shows a large degree of variability on time scales from hours to the solar cycle length. This variation is associated with magnetospheric storms, the Earth's rotation, the season, and the level of solar activity. To make accurate predictions of key ionospheric parameters all these variations must be considered. Neural networks, which are data driven non-linear models, are very useful for such tasks. In this work we examine if the F2 layer plasma frequency, foF2, at a single ionospheric station can be predicted 1 to 24 hours in advance by using information of past foF2 observations, magnetospheric activity, and time as inputs to neural networks. Particular attention has been paid to periods when great geomagnetic storms were in progress with the aim to develop a successful ionospheric storm forecasting tool.

1 Introduction

This work focuses on the prediction of the ionospheric critical plasma frequency, foF2, on time scales from 1 hour to 24 hours ahead for a mid-latitude ionosonde in Europe (Slough). We are especially interested in predicting the development of foF2 during ionospheric storms and, therefore, precursors for ionospheric storms are studied. As long time series of both ionospheric and magnetospheric data exist, and because any successful system describing the time evolution of the ionosphere should be a nonlinear model, we have selected an artificial neural network for the study.

The short term (hourly) variation of foF2 has been studied extensively using neural networks. Models for one hour ahead predictions have been developed by a number of people (Altinay et al., 1997; Cander and Lamming, 1997; Cander et al., 1998; Wintoft and Cander, 1999; Kumluca et al., 1999). Other studies concentrate on predicting the noon foF2 value (Williscroft and Poole, 1996; Francis et al., 1998). Finally, on these time scales, efforts have been made to predict foF2 up to 24-30 hours ahead (Wintoft and Cander, 1998; Francis et al., 1999a,b). There is also a model for the prediction of the monthly median foF2 several months ahead (Lamming and Cander, 1998). These models are summarised in Table 1. All the models share the common architecture in that they use past values of foF2 as input. In addition, some models also contain information about solar activity (sunspot number or 10.7cm radio flux), geomagnetic activity (AE, Dst, Kp, Ap), or time information (season, local time). In a sense, much of this additional information is already present in the foF2 input time series. E.g., during solar maximum, the general level of foF2 is higher than during solar minimum and thus, based only on previous foF2 values, we already have this information without the sunspot number. In the same way, we can also use previous values of foF2 to determine the quiet ionosphere for, say, the next 24 hours. However, during disturbed conditions matters become more complicated, and the ionosphere is often in a disturbed state. If the disturbance is determined by the internal dynamics of the ionosphere it should again be possible to predict future foF2 values from past foF2 values. However, any disturbance that is driven, by e.g. magnetospheric activity, must be modelled with some external inputs other than foF2 itself.

It has been long known that ionospheric disturbances are related to magnetospheric activity that depends on the local time (Kirby et al., 1936) and season (Appleton et al., 1937; Kirby et al., 1937). A review on ionospheric F-region storms and their relation to magnetospheric activity is given by Prölss (1995). Although many different magnetospheric indices have been used in the above mentioned neural networks, their effects on the predictions have been minor. However, the reason is not the absence of a magnetospheric-ionospheric relation, nor that the neural networks are incapable of finding a re-

Correspondence to: P. Wintoft

Table 1. Different models for the prediction of foF2. The models in the table are: A97, Altinay et al. (1997); C98, Cander et al. (1998); W98, Wintoft and Cander (1998); W96, Williscroft and Poole (1996); F98, Francis et al. (1998); L98, Lamming and Cander (1998).

Model	Network	Input	Output	Stations
A97	MLFF	foF2, Kp, LT	foF2 1 hour ahead	Poitiers
C98	MLFF	foF2, SSN, Ap	Hourly foF2, 1 to 5 hours ahead	Slough
W98	MLFF	foF2, Ap	Hourly foF2, 1 to 24 hours ahead	Slough
W96	MLFF	Season, SSN, a_k	Noon foF2	Grahamstown
F98	RBF	foF2	Hourly & monthly foF2, 1–30 steps ahead	Slough
L98	MLFF	Month, LT, SSN	Monthly median foF2	Poitiers

lation, but rather it can be attributed to the choice of training data and how the inputs are combined. For only a fraction of the time series there is major magnetospheric activity, and if insufficient care is taken when selecting the training data, these periods will be outnumbered by periods characterising more settled conditions.

Based on the ideas by Prölss (1995), it should be possible to predict ionospheric F-region positive and negative storms for mid-latitude stations several hours in advance using AE as an input to a neural network. Negative ionospheric storms typically develop after the onset of a geomagnetic storm during the preceding night. A possible explanation for this effect is that the ionosonde station moves with the rotation of the Earth under the ionospheric disturbance zone. As the auroral zone starts its expansion during the night, the ionosonde moves out of the disturbance zone into the daytime side while the magnetospheric storm is developing. Then, during the following night, the ionosonde reaches the disturbance zone, and the ionospheric negative storm starts. As the disturbed ionosphere is also convected into the daytime, the negative storm persists for the whole day. On the other hand, positive ionospheric storms are mainly observed in the daytime and are seen to follow the magnetospheric storm a couple of hours later. The hypothesis is that a traveling atmospheric disturbance (TAD) is generated by the magnetospheric storm, which travels from high latitudes towards the equator. The positive ionospheric storms should thus be predictable a couple of hours in advance, while the negative storms should be predictable up to a day in advance using magnetospheric activity indices as precursors.

In this work, we will take account of all the known variations that are seen in foF2 with a linear filter, and then use a neural network to find the more complex behaviour during disturbed ionospheric conditions. The foF2 data are taken from the ionosonde at Slough. As precursor for the ionospheric storm, we use the auroral electrojet index AE, as it has been used in many studies relating magnetospheric activity to ionospheric activity. Real time predictions of AE are available at `http://www.astro.lu.se/~henrik/`.

2 A model for foF2 storm predictions

2.1 Quiet time foF2 variations

We assume that the ionospheric F2 layer will be in a quiet condition when the daily average AE is below a certain threshold. If we select this threshold to be 200 nT then, on average, 1 day in every 2.6 days will be a quiet day. Clearly, even during days with a daily average AE of 200 nT there will be some level of activity. This can, however, be handled by the network that we will introduce in the next section. To develop a model that actually describes the quiet ionosphere is a difficult problem. Wrenn et al. (1987) developed a method for describing the quiet ionosphere to study the evolution of ionospheric storms. They concluded that simply using the monthly median foF2, which is updated once a month, is too coarse. Instead, they used a running monthly median model that only included magnetically quiet days. In this work we use a similar procedure and calculate the average foF2 for the past 7 days that are magnetically quiet and use this as a first prediction of foF2 one to two days ahead. Then, the difference between the predicted foF2 and the observed foF2 becomes

$$\delta(d+D,h) = \mathrm{foF2}(d+D,h) - \frac{1}{7}\sum_{i=1}^{7}\mathrm{foF2}(d-D_i,h), \quad (1)$$

where $\mathrm{foF2}(d,h)$ is the observed foF2 at day d and local time h, D is one or two days ahead, and D_i are the previous seven days with daily average AE less than 200 nT. This model will basically remove the obvious variations in foF2 that are related to the solar cycle, the season, and the diurnal variations. The remaining δ will be composed of intrinsic variations of foF2 (linear and/or nonlinear), geomagnetically driven variations, and other miscellaneous variations (e.g. associated with changes in solar EUV flux). An example of the observed foF2 and the quiet time reference is given in Fig. 1.

2.2 The neural network inputs

The next step is to develop a neural network to predict the δ from Eq. 1. The inputs to the neural network that we consider are derived from the time series of δ, AE,

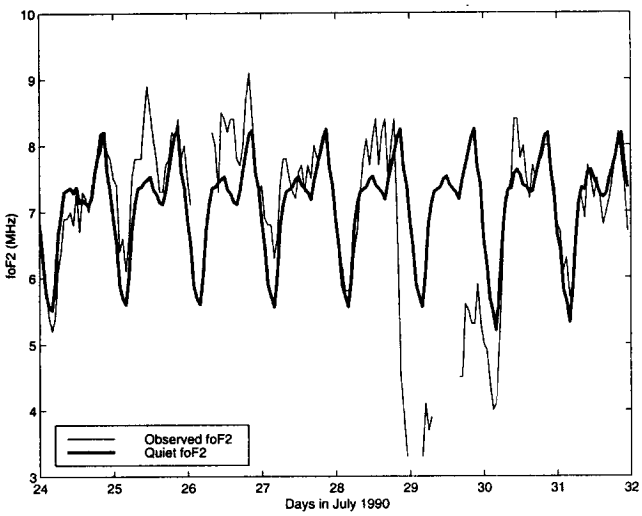

Fig. 1. Observed (thin line) and quiet (thick line) foF2 for a few days in July 1990. There is a clear negative ionospheric storm during the 29th.

local time, and season. The motivations for the choices of these parameters are described below.

The evolution of an ionospheric storm, following a geomagnetic storm, depends on when, in the local time sector, the geomagnetic storm starts. Severe geomagnetic storms that begin before sunrise are followed by negative ionospheric storm conditions during the whole day (Prölss, 1993) (and references therein). On the other hand, geomagnetic storms that begin in the local daytime sector often lead to positive ionospheric storms a couple of hours later.

The time of the year is another important factor that determines the evolution of ionospheric storms associated with geomagnetic storms. In the summer the negative ionospheric storms extend all the way from the polar region to the subtropics, while in the winter the negative storm effects are restricted to higher latitudes. The positive storms are also mainly observed in the winter.

The different input parameters to the neural network are arranged into a vector $\boldsymbol{x}$ with the elements x_i, where i denotes different input units. The first elements of $\boldsymbol{x}$ are

$$x_i = \delta(t - T_i),\ 1 \le i \le N, \tag{2}$$

where the time t is in hours, $T_1 = 0$ hours, $T_2 = 1$ hours, etc. to $T_N = N - 1$ hours. The following elements are

$$x_i = \mathrm{AE}(t - T_{i-N}),\ N + 1 \le i \le 2N. \tag{3}$$

The local time information is the sine of the local time, i.e.

$$x_i = \sin\left(2\pi \frac{h - T_{i-2N}}{24}\right),\ 2N + 1 \le i \le 3N, \tag{4}$$

where h is the local time in hours. Finally, the seasonal information is coded as

$$x_i = \cos\left(2\pi \frac{H - T_{i-3N}}{24 \cdot 365}\right),\ 3N + 1 \le i \le 4N. \tag{5}$$

where H is the number of hours from hour 0, January 1^{st}. With this choice of representing time, we see that Eq. 4 for the local hours 23 and 00, which are close in time, have values that are close. As the sine function is double valued in any one cycle, two values are needed to describe time uniquely. One approach is to use $\sin(t)$ and $\cos(t)$. Here, we instead use a time window of sine values and N in Eq. 4 must thus be larger or equal to 2. If N is larger than 2 there will be redundancy in the input, however, for any local time the input vector will be unique. The redundant units at the input can improve on the learning, but that will not be further examined here. The seasonal information is coded as a time series of cosine values. During, e.g., a 24 hour period, i.e. $N = 24$, the input values will all be similar, as Eq. 5 recycles once every 8760 hours. Thus, during winter the inputs will be close to $+1$ and during summer the inputs will be close to -1. Spring and autumn will have values close to 0. This coding will work if we assume that the ionospheric storms develop differently in winter and summer, and that the spring and autumn storms show similar behaviour.

It should be noted that the δ and AE used for the input have been normalized by subtracting the mean and dividing by the standard deviation. Thus, both the δ and AE used for input have zero mean unit variance.

As we are using linear output units we do not normalize the δ for the output. The desired output is

$$y = \delta(t + \tau), \tag{6}$$

where τ is 1 to 24 hours.

If we denote the function that a neural network, for a τ hours ahead prediction, implements as F_τ, then the relation becomes

$$\hat{y} = F_\tau(\boldsymbol{x}), \tag{7}$$

where $\hat{y}$ is a prediction of y.

2.3 Training, validation, and test sets

Three independent data sets are extracted for the training (training set), optimization (validation set), and testing (test set) of the neural networks (Haykin, 1994). During training, the weights of the network are found from the error backpropagation algorithm. Several different networks are trained where the type of inputs and the number of hidden units are varied. Then the validation set is used to determine the optimal network. Finally, the optimal network is tested on the test set to determine how well it will work for new data. All the following figures that show network predictions are cases from the test set.

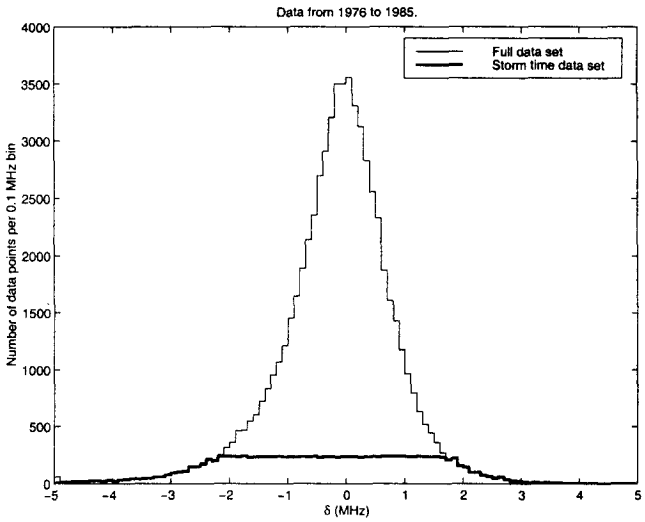

Fig. 2. The distribution of δ in 0.1 MHz bins. The thin line represents the full data set over the years 1976 to 1985, while the thick line is the sum of the training and validation sets.

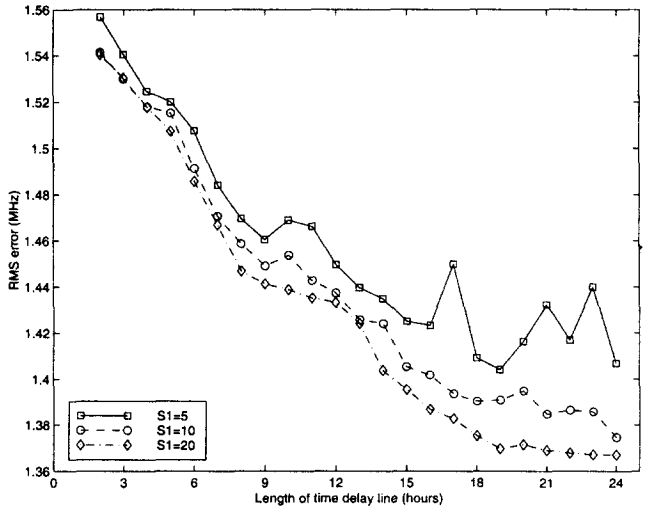

Fig. 3. The change in RMS error when the number of hidden units (S1) and the length of the time delay line are changed.

The training and validation sets are chosen from the years 1976 to 1985, which cover solar cycle 21. The distribution of δ for this period is shown by the thin line in Fig. 2. We see that small values of δ dominates the distribution. As we are interested in storm periods, and if we assume that large values of δ represents storm conditions, then a network that is trained on the full data set will not be able to learn the prediction of storm periods. To overcome this problem we select a subset for the training and validation sets. Our approach is to limit the number of data points to 250 in each 0.1 MHz bin. This will ensure that the number of storm hours will increase as compared to the quiet hours. The thick line in Fig. 2 shows the distribution of the training and validation sets. Finally, the test set is chosen from the years 1986 to 1994, which cover solar cycle 22.

3 Results

Here we describe the one hour ahead predictions and then summarize the results for 6, 12, and 24 hour ahead predictions.

3.1 One hour ahead predictions

For the 1 hour ahead predictions we try different combinations of δ, AE, local time, and season as inputs and also vary the length of the time delay line. The smallest root mean square (RMS) error of the validation set is obtained when all parameters are used as inputs. However, from examining several ionospheric storms, it is seen that the predictions of storm onset are actually lagging by approximately one hour. The information for ionospheric storm onset that exists in AE seems to be discarded in this type of model and it is not obvious why this is the case. However, it can be noted, that the goal of the training algorithm is to minimize the RMS error of the predictions of all the data points in the training set and this is apparently achieved from the past values of δ. Similar results are obtained, using the same network architecture, in the prediction of the 10 minute average local magnetic field data from solar wind data. Again, using past local magnetic field data together with solar wind data, leads to a shift in the predicted local magnetic field at storm periods (Hans Gleisner, Private communication). The problem is related to the fact that the correlation between $\delta(t)$ and $\delta(t+1)$ is much higher than AE(t) and $\delta(t+1)$. The RMS error used on the validation set is thus not a measure that should be used alone to determine the optimal model. More elaborate methods to examine this could also be employed as e.g. discussed by Smith (1999). As we are mainly interested to examine AE as a precursor of ionospheric storms we will not include δ in the input. This will lead to higher RMS errors but improve the prediction at storm onset.

By varying the number of hidden units and varying the length of the time delay line we determine the optimal network that has the smallest RMS error on the validation set using the inputs AE, local time, and season. In Fig. 3, we see that the minimum error is obtained with a network with 20 hidden neurons and a time delay line of about 19 hours. The network thus has $3 \times 19 = 57$ inputs.

Fig. 4 shows a prediction using the optimal network for a period taken from the test set covering 7 days in 1990. The prediction of the negative ionospheric storm, starting on the 28^{th}, closely follows the observed values. In the lower panel the AE index shows major storm activity. We can also examine the model response (dashed curve) for geomagnetically quiet conditions by setting

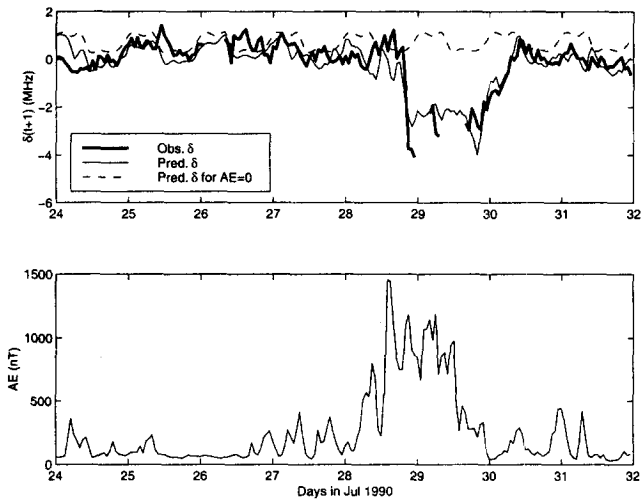

Fig. 4. Predicted and observed δ for one hour ahead predictions. When AE is set to zero the model output is the dashed line.

AE=0 for all inputs. We see that there is a diurnal variation and an offset towards positive values. If we assume that the other inputs are independent of AE, we can conclude that the quiet reference model does not remove the quiet diurnal variation completely and that our definition of "quiet" includes some activity.

Two other examples of ionospheric storms are given in Figs. 5 and 6. The time span is again 7 days and the axis limits are the same for easier comparison. The storm period in August 1990 (Fig. 5) contain several interesting features. A positive ionospheric storm begins in the daytime of the 21^{st} and is then followed by a negative storm. Except for the positive storm and for a few hours on the 23^{rd}, the prediction closely follows the observations. The model usually performs poorly for positive ionospheric storms. The other storm period in October 1990 shows severe negative ionospheric activity where δ reaches as low as -6 MHz. The onset of the storm and the general level of disturbance are predicted. However, on the 10^{th} and the 12^{th}, the observed values are much lower than the predicted values. Examining the geomagnetic activity, we see that the AE index is not extremely disturbed and is actually lower than for the previous two examples.

3.2 Predictions 6, 12, and 24 hours ahead

The same procedure used to train and optimize the networks for the one hour predictions is also adopted for the 6, 12, and 24 hour predictions. The optimal networks have 20 hidden units and time delay lines of 24 to 36 hours. We selected the October 1990 period to illustrate the accuracy of the predictions as shown in Figs. 7 to 9. We see that the 6 hour ahead prediction still accurately predicts the storm onset and then follows an evolution similar to the 1 hour ahead prediction. When

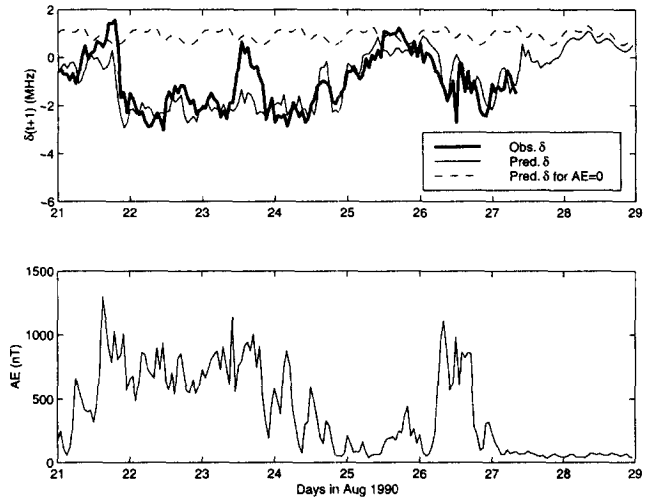

Fig. 5. Predicted and observed δ for one hour ahead predictions. When AE is set to zero the model output is the dashed line.

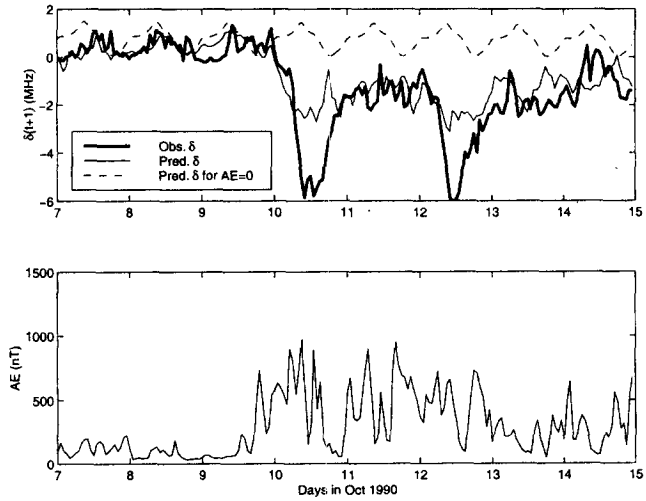

Fig. 6. Predicted and observed δ for one hour ahead predictions. When AE is set to zero the model output is the dashed line.

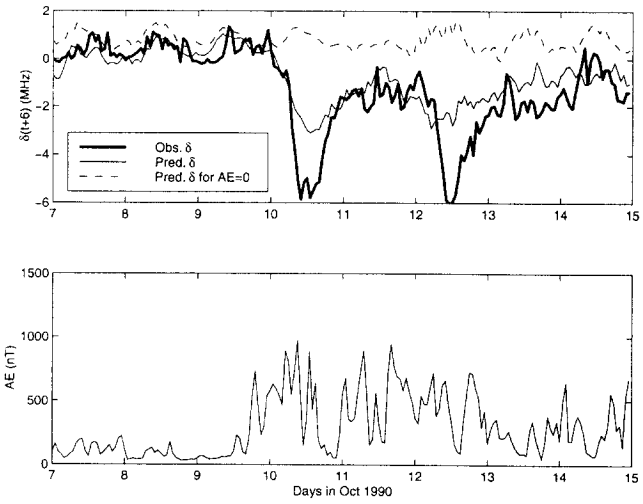

Fig. 7. Predicted and observed δ for six hour ahead predictions. When AE is set to zero the model output is the dashed line.

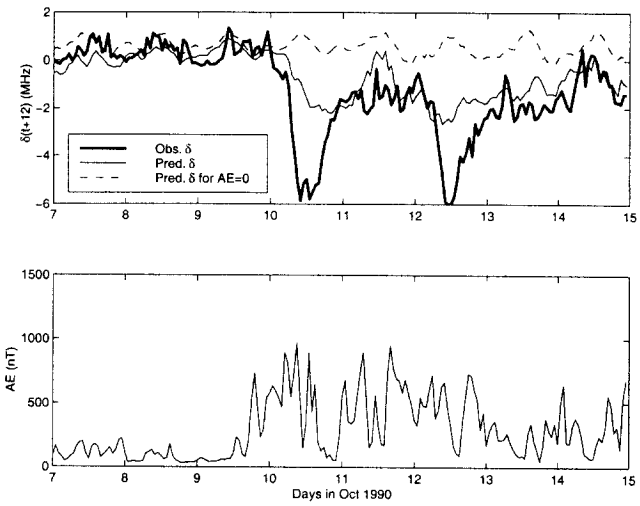

Fig. 8. Predicted and observed δ for twelve hour ahead predictions. When AE is set to zero the model output is the dashed line.

the prediction time is increased to 12 hours, we see a shift of several hours at the storm onset and for the 24 hour predictions the model now fails. A similar development can also be seen for the other periods that we showed for the 1 hour predictions.

4 Discussion and conclusions

From the above results it is clear that information concerning the magnetosphere is important for the prediction of ionospheric storms in the F2 layer. Here, we have used the hourly AE index to describe the magnetospheric storms. The index is convenient because real time predictions are available and its inclusion is sound from a physical point of view.

However, there are problems with the index as illustrated in Fig. 10. There is a strong UT variation of the number of hours with high levels of activity. The peak at UT16 could imply that the magnetic observatories at the corresponding location dominate the derivation of AE. This UT variation will affect the training of the neural networks in a negative way. This might partly explain the problem of the 1 hour prediction of positive ionospheric storms. Another problem is that during severe geomagnetic storms the auroral oval expands to lower latitudes and the AE index, which is derived from high latitude stations, shows a lower disturbance level than is actually the case.

As the observed δ are available in real time they should also be used in a model for the prediction of foF2. As we saw in the previous section it is not possible to make one hour ahead predictions when both δ and AE are used as inputs simultaneously. However, using only AE the ionospheric storms are predicted but the overall error increases. To overcome this problem, another network

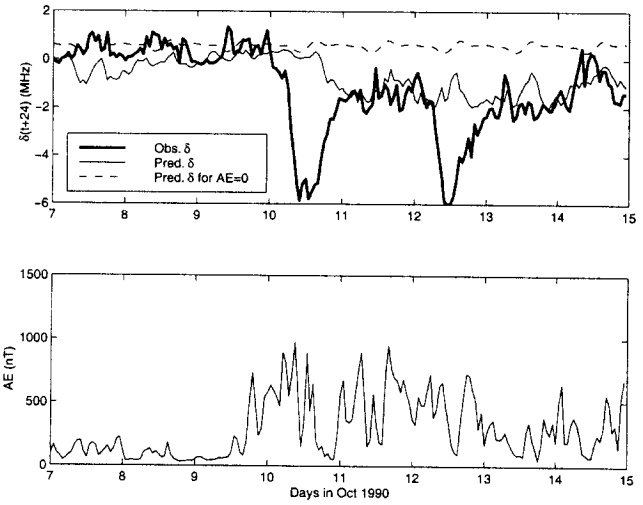

Fig. 9. Predicted and observed δ for 24 hour ahead predictions. When AE is set to zero the model output is the dashed line.

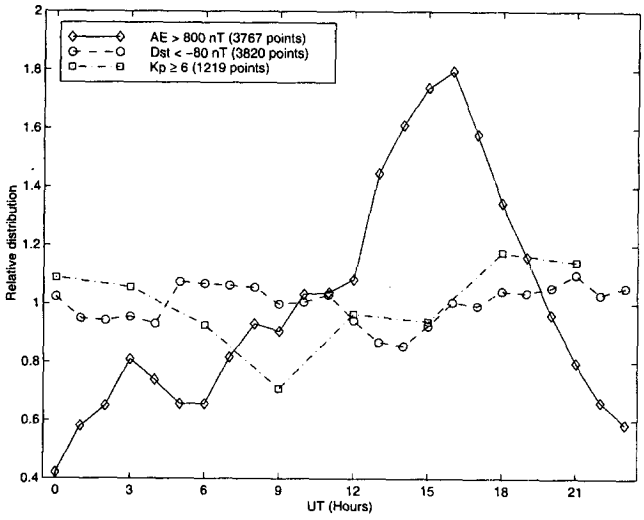

Fig. 10. The relative distribution of geomagnetic storms as a function of UT for the period 1976 to 1994. For AE there are 3767 points (hours) with AE>800 nT and the relative distribution of these points in UT is shown by the solid line. In the same way the relative distribution for Dst<-80 nT (3820 hours) and Kp≥6 (1219 three-hour intervals) are shown by the dashed and the dash-dotted curves, respectively.

could be used to predict the difference between the AE-network and the observed δ using past δ as input.

For future work it proposed that inputs other than AE are used. Other global magnetospheric indices could be used, such as AU and AL, as they more closely describe the actual physical processes in the ionosphere. Local magnetic field measurements could also be used, which have the advantage that they can easily be made available in real time. A further step could be to use solar wind measurements directly. Of course, this relies on the assumption that a solar wind monitor exists.

Acknowledgements. The authors are grateful to the two referees for their valuable comments.

References

Altinay, O., E. Tulunay, and Y. Tulunay, Forecasting of ionospheric critical frequency using neural network, *Geophys. Res. Lett.*, *24*, 1467–1470, 1997.

Appleton, E.V., R. Naismith, and L.J. Ingram, British radio observations during the second international polar year 1932–1933, *Philos. Trans., R. Soc.*, A236, 191–259, 1937.

Cander, Lj.R. and X. Lamming, Neural networks in ionospheric prediction and short-term forecasting, *10th International Conference on Antennas and Propagation, 14-17 April 1997, Edinburgh, IEE Conference Publication, 436*, 2.27–2.30, 1997.

Cander, Lj.R., S. Stankovic, and M. Milosavljevic, Dynamic ionospheric prediction by neural networks, *Proc. of the Second International Workshop on Artificial Intelligence Applications in Solar-terrestrial Physics*, ESA WPP-148 Proceedings, Eds. I. Sandahl and E. Jonsson, European Space Agency, Paris, 225–228, 1998.

Francis, N.M., A.G. Brown, A. Akram, P.S. Cannon, and D.S. Broomhead, Non-linear prediction of the ionospheric parameter foF2, *Proc. of the Second International Workshop on Artificial Intelligence Applications in Solar-terrestrial Physics*, ESA WPP-148 Proceedings, Eds. I. Sandahl and E. Jonsson, European Space Agency, Paris, 219–223, 1998.

Francis, N.M., A.G. Brown, P.S. Cannon, and D.S. Broomhead, Non-linear interpolation of missing points within the hourly foF2 time series, *Proc. of the Workshop on Space Weather*, ESA WPP-155 Proceedings, European Space Agency, The Netherlands, 419–422, 1999a.

Francis, N.M., A.G. Brown, P.S. Cannon, A. Akram, and D.S. Broomhead, Novel non-linear techniques for the prediction of noisy geophysical time series that contain a significant proportion of data drop outs, *Phys. Chem. Earth*, XXIV General Assembly of the European Geopysical Society, 1999b.

Haykin, S, *Neural networks, a comprehensive foundation*, Macmillan College Publishing Company, New York, 1994.

Kirby, S.S, T.R. Gilliland, N. Smith, and S.E. Reymer, The ionosphere, solar eclipse and magnetic storm, *Phys. Rev.*, 50, 258–259, 1936.

Kirby, S.S, N. Smith, T.R. Gilliland, and S.E. Reymer, The ionosphere and magnetic storms, *Phys. Rev.*, 51, 992–993, 1937.

Kumluca, A., E. Tulunay, I. Topalli, and Y. Tulunay, Temporal and spatial forecasting of ionospheric critical frequency using neural networks, Submitted to *Radio Sci.*, 1999.

Lamming, X., and Lj.R. Cander, Monthly median ionosphere frequencies prediction with neural networks, *Proc. of the Second International Workshop on Artificial Intelligence Applications in Solar-terrestrial Physics*, ESA WPP-148 Proceedings, Eds. I. Sandahl and E. Jonsson, European Space Agency, Paris, 229–234, 1998.

Prölss, G.W., On explaining the local time variation of ionospheric storm effects, *Ann. Geophys*, *11*, 1–9, 1993.

Prölss, G.W., Ionospheric F-Region Storms, *Handbook of atmospheric electrodynamics*, Ed. H. Volland, CRC Press, Inc., 195–248, 1995.

Smith, L.A., Sane, psychic and psychotic neural networks, *Phys. Chem. Earth*, XXIV General Assembly of the European Geopysical Society, 1999.

Williscroft, L.-A., and A.W.V. Poole, Neural networks, foF2, sunspot number and magnetic activity, *Geophys. Res. Lett.*, *23*, 3659–3662, 1996.

Wintoft, P., and Lj.R. Cander, Short-term prediction of foF2 using time-delay neural networks, *Phys. Chem. Earth* 24, 343-347, 1999.

Wintoft, P., and Lj.R. Cander, 24 hour predictions of foF2 using time-delay neural networks, Submitted to *Radio Sci.*, 1998.

Wrenn, G.L., A.S. Rodger, and H. Rishbeth, Geomagnetic storms in the Antartctic F-region. I. Diurnal and seasonal patterns for main phase effects, *J. Atmos. Terr. Phys.*, *49*, 901–913, 1987.