Secular Behavior of Solar Magnetic Activity: Nonstationary Time Series Analysis of the Sunspot Record

P. F. Verdes
Institute for Environmental Physics, University of Heidelberg, Im Neuenheimer Feld 229, D-69120 Heidelberg, Germany

P. M. Granitto and H. A. Ceccatto
Instituto de Física Rosario, Consejo Nacional de Investigaciones Científicas y Técnicas and Universidad Nacional de Rosario, Blvd. 27 de Febrero 210 Bis, S2000EZP Rosario, Argentina

February 11, 2004

Abstract. The sunspot record of solar magnetic activity is studied as a nonstationary time series by means of a previously developed algorithm for treating perturbed dynamical systems. This approach incorporates secular changes into the modeling process through an external driving parameter, whose temporal behavior is shown to correspond in this case to the long-term trend of the sunspot record. Our method is able to reduce in approximately 13% the prediction error of this series when compared to the standard stationary approach. Such reduction is remarkable in view of the benchmark status of the sunspot record in the statistical literature and, moreover, the fact that this gain is obtained on the performance of an already very competitive modeling technique based on ensembles of artificial neural networks.

Keywords: solar magnetic activity, sunspots, nonstationarity, time series analysis, artificial neural networks

1. Introduction

The global aspects of solar magnetic activity, as evidenced by the Wolf sunspot number (SSN), are well explained by dynamo theory (Priest, 1984). However, the nature of irregularities in the SSN time series is still under debate, being alternatively ascribed to chaos (Zeldovich and Ruzmaikin, 1990; Knobloch and Landsberg, 1996) or stochasticity (Choudhuri, 1992; Mininni et al., 2000; Mininni et al., 2002) as the underlying mechanism. A short cycle in solar magnetic activity was first discovered by Schwabe (1843), and later reported by Wolf (1852), who estimated its period in approximately 11 years. Of course, this cycle is not strictly periodic, with variations both in amplitude and length, and also periods of inactivity like the Maunder minimum (Maunder, 1922). Longer cycles have also been identified, like, for instance, the Gleissberg cycles (Gleissberg, 1971; García and Mouradian, 1998), although they are more difficult to study because of the limited time-span of observations.

There are both academical and practical reasons that justify the widespread interest in understanding and predicting the magnetic activity of the Sun (Christensen-Dalsgaard et al., 1996). Because of this, and also in part for historical reasons (Yule, 1927), the SSN time series has become a standard workbench in the statistical literature, where new methods and ideas are continuously developed and tested. Furthermore, many of the recent ideas developed within the theory of dynamical systems and later incorporated to nonlinear time-series analysis have been applied to this series. However, in most of these studies the SSN record is treated as stationary, disregarding strong indications of important changes in solar magnetic behavior in the last centuries —most notably, the Maunder minimum in the second half of the 17th century. In fact, there are also features in the last 300 years of the SSN record that can hardly be ascribed to a stationary process. Firstly, only three Gleissberg cycles are observed, and their irregular behavior makes them very difficult to account for. Secondly, and more important, there has been an increase in solar magnetic activity during the last century (Lockwood et al., 1999; Solanki et al., 2000), with a large rise in the SSN behavior (Usoskin et al., 2003) that most probably cannot be well captured by standard modeling techniques for stationary processes.

In this work we study the SSN record considering it as a nonstationary time series. This is achieved by incorporating into the analysis a previously developed methodology for the accurate reconstruction of driving forces in perturbed dynamical systems (Verdes et al., 2001). According to it, secular changes in solar magnetic activity on timescales of the order of a century or larger are considered as perturbations on the "intrinsic" dynamics (short cycles). Here we obtain the profile of the effective "perturbing signal" during the three centuries covered by the available data, allowing us to model the SSN record more accurately. To avoid misinterpretations, we stress that in the present context there is clearly no force external to the Sun that is driving the dynamics. In Verdes et al. (2001) we already mentioned that the reconstructed nonstationary parameter called "driving force" might not represent the effect of a truly external perturbation but, as it is the case here, internal changes in the system or slow dynamical modes that are insufficiently recurrent within the observational period. Our approach incorporates this extra parameter only to model the fast dynamics (short cycles) more accurately. In fact, we show that our treatment is able to reduce in approximately 13% the one-step-ahead prediction error of ensembles of artificial neural networks (ANNs), considered a very competitive benchmark on this problem (Weigend and Gershenfeld, 1994; Calvo et al., 1995; Naftaly et al., 1997).

2. Nonstationary Time Series Analysis of the Sunspot Record

In Verdes et al. (2001) we proposed an algorithm for the accurate reconstruction of slowly-changing external forces acting on nonlinear dynamical systems. The method traces the evolution of the external driving force by locally linearizing the map dependency with the shifting parameter. It essentially relies on developing local models on (possibly overlapping) windows of M points and predicting the time series behavior on neighboring intervals. In this Section we briefly review the related theory and apply it to the sunspot record. As stressed above, in this case the reconstructed nonstationary parameter —denoted α in the following— does not represent the effect of a truly external perturbation to the Sun. However, beyond any putative physical interpretation of this parameter, α is the driver of the *model* we develop, since it is external to it (see below).

2.1. Theoretical Aspects

Consider an observational record $\{x_t\}_{t=1}^{N}$ generated by a deterministic dynamical system, where time t is measured in units of the lag τ between observations. We model this process in a d-dimensional pseudo-phase space according to

$$x_{t+1} = f(\boldsymbol{x}_t, \alpha_t) + \varepsilon_t, \tag{1}$$

with $\boldsymbol{x}_t = (x_t, x_{t-1}, ..., x_{t-d+1})$. As stated above, here α_t accounts for the effects of internal degrees of freedom varying on large time scales $T \gg \tau$ not modeled by f, and ε_t is some residual (Gaussian) noise. Then, if we split the data into $N_{\rm int}$ (possibly overlapping) intervals of M points each, for $M\tau \ll T$ we can write

$$x_{t+1}^{(m)} \simeq f(\mathbf{x}_t^{(m)}, \alpha^{(m)}) + \varepsilon_t, \tag{2}$$

where $\alpha^{(m)}$ is the mean value of the parameter in the m-th interval considered, and now t runs only on points of this interval. For a smooth dependence of f with α, we can employ a first-order expansion on its second argument centered in $\alpha^{(k)}$. Rearranging terms, we obtain:

$$x_{t+1}^{(m)} - f(\mathbf{x}_t^{(m)}, \alpha^{(k)}) \simeq \frac{\partial f}{\partial \alpha}(\mathbf{x}_t^{(m)}, \alpha^{(k)})[\alpha^{(m)} - \alpha^{(k)}] + \varepsilon_t, \tag{3}$$

which should be valid for intervals k and m close enough (in practice they can be taken with a substantial overlap to fulfil this condition). Notice that, to the same order of approximation, we can replace

$\frac{\partial f}{\partial \alpha}(\mathbf{x}_t^{(m)}, \alpha^{(k)})$ by $\frac{\partial f}{\partial \alpha}(\mathbf{x}_t^{(m)}, \alpha^{(m)})$. Averaging over all t in interval m we arrive to

$$E_k^m = A^m \Delta\alpha_k^m, \tag{4}$$

where we have introduced the notation $E_k^m = \left\langle x_{t+1}^{(m)} - f(\mathbf{x}_t^{(m)}, \alpha^{(k)}) \right\rangle_t$, $A^m = \left\langle \frac{\partial f}{\partial \alpha}(\mathbf{x}_t^{(m)}, \alpha^{(m)}) \right\rangle_t$, and $\Delta\alpha_k^m = \alpha^{(m)} - \alpha^{(k)}$. Here the sub-index k is fixed and denotes the interval used as reference, while the running super-index m describes the actual temporal variation of the driving parameter.

We assume now that the M data points in interval k cover the attractor well enough to allow modeling[1] $f(\bullet, \alpha^{(k)})$, obtaining a predictor $\widehat{f}(\bullet, \alpha^{(k)})$ and estimates $\widehat{E}_k^m$. In this case, and for contiguous intervals ($m = k \pm 1$), the system of equations (4) reads

$$\begin{cases} \widehat{E}_1^2 \simeq A^1 \cdot [\alpha^{(2)} - \alpha^{(1)}] \\ \widehat{E}_2^1 \simeq -A^2 \cdot [\alpha^{(2)} - \alpha^{(1)}] \\ \widehat{E}_2^3 \simeq A^2 \cdot [\alpha^{(3)} - \alpha^{(2)}] \\ \widehat{E}_3^2 \simeq -A^3 \cdot [\alpha^{(3)} - \alpha^{(2)}] \\ \widehat{E}_3^4 \simeq A^3 \cdot [\alpha^{(4)} - \alpha^{(3)}] \\ \ldots \end{cases} \tag{5}$$

This system can be solved up to a scale factor. In particular, taking A^1 as the parameter of the unknown scale transformation, the solution is

$$\Delta\alpha_k^{k+1} = (-1)^{k+1} \frac{\widehat{E}_{k+1}^k}{A^1} \prod_{r=1}^{k-1} \left(\frac{\widehat{E}_{r+1}^r}{\widehat{E}_r^{r+1}} \right). \tag{6}$$

According to the above discussion, we proposed the following procedure for reconstructing α: Choose a data segment k containing M iterates and use them to construct a model[2] $\widehat{f}(\bullet, \alpha^{(k)})$. Then, compute the average prediction errors $\widehat{E}_k^m$ for $m = k \pm 1$. Repeat for all possible k and reconstruct the driving parameter up to the arbitrary scale A^1 and shift $\alpha^{(1)}$ according to Eq. (6). The α_t profile can be obtained following this procedure —at least in principle. However, for intervals k and $m = k \pm 1$ with large overlaps, the quantities $E_k^{k\pm1}$ are in general small and their estimates $\widehat{E}_k^{k\pm1}$ tend to have large relative errors. The solution

[1] This can be done by using any data-driven modeling approach. That is, proposing a general model —e.g., linear ARMA model, ANNs, radial basis functions or polynomials— and optimizing its internal parameters by a maximum-likelihood or least-squares procedure using the observational data from interval k.

[2] In the concrete application to SSN we will use ANNs to do this, as described in the following two subsections.

(6) is very sensitive to these modeling errors because of the large productories involved both in the numerator and denominator, something which becomes particularly accute near stationary points of the driving parameter where $\widehat{E}_k^{k\pm1} \sim 0$. A way to cope with this problem is to solve the system (5) in a mean-squared-error sense, implementing a simple gradient-descent minimization of the error function

$$\begin{aligned} Err &= \sum_{k=1}^{N_{\rm int}-1} (\widehat{E}_k^{k+1} - A^{k+1}\Delta\alpha_k^{k+1})^2 + \sum_{k=2}^{N_{\rm int}} (\widehat{E}_k^{k-1} + A^{k-1}\Delta\alpha_{k-1}^{k})^2 \\ &+ \lambda \sum_{k=1}^{N_{\rm int}-1} (A^{k+1} - A^k)^2, \end{aligned} \qquad (7)$$

where the λ-term forces the method to seek for smooth solutions, as is usually done for ill-defined problems. The minimization of (7) with respect to its arguments A^k and $\Delta\alpha_k^{k+1}$ produces more regular results than the exact solution (6), because the effects of accidentally large modeling errors in $\widehat{E}_k^{k\pm1}$ will be moderated by the λ-term. Notice that this way of solving the problem can deal with stationary regions in the driving parameter, since in those regions the method will find solutions of (5) with $\widehat{E}_k^{k+1} \simeq 0 \simeq \Delta\alpha_k^{k+1}$ and $A^{k+1} \simeq A^k$. In Verdes et al. (2001) we showed that this is indeed the case. For the gradient-descent minimization of (7) we started from suitable random values of A^k and $\Delta\alpha_k^{k+1}$, and checked that different initializations do not sensibly change the results. For more details of this methodology please refer to Verdes et al. (2001).

In short, the proposed methodology can be summarized as follows:

1. choose a data segment k containing M iterates and employ any preferred data-driven modeling tool to build a local model of the dynamics,

2. estimate the cross-prediction errors $\widehat{E}_k^{k\pm1}$ of this model over the first-neighboring intervals,

3. repeat for all possible k,

4. with all the $\widehat{E}_k^{k\pm1}$ at hand, minimize Err (eq. 7) with respect to its arguments A^k and $\Delta\alpha_k^{k+1}$, and finally

5. reconstruct α_t by accumulation of the $\Delta\alpha$'s found in the previous step.

2.2. Nonstationary Modeling of the Sunspot Record

In this subsection we apply the methodology sketched above to model the Wolf SSN record of solar magnetic activity. This series is known to have a complex behavior with several time scales[3]. In particular, we considered the annual mean SSN from 1700 to 2001 ($N = 302$ data points), and used this record to build a set of $N - d$ input-output patterns $(\boldsymbol{x}_t, x_{t+1})$. We split this set in overlapping intervals of $M = 33$ contiguous patterns (approximately 3 short cycles), with neighboring intervals m and $m + 1$ differing only in one point (this maximum overlap ensures small changes between the average values of the driving parameter $\alpha^{(m)}$ and $\alpha^{(m+1)}$). For modeling purposes we used ANNs, which are known to perform very well on this data set (Weigend and Gershenfeld, 1994; Calvo et al., 1995; Naftaly et al., 1997). In particular, on each interval we trained a ANN on 22 randomly-chosen patterns, validating their learning process (to avoid overfitting) on the remaining 11 patterns. We considered feedforward ANNs with architectures $d : h : 1$; the number of input neurons d was varied between 2 and 9, the number of hidden units was fixed to $h = 10$, and the (linear) output neuron was used to predict the mean SSN one year ahead. In this way we obtained local models $\widehat{f}(\bullet, \alpha^{(m)})$ for all intervals, which allow to predict sunspots in the neighboring intervals $m \pm 1$ as required by the algorithm proposed in Verdes et al. (2001). The window size of $M = 33$ points used is a reasonable compromise between the need for having enough data for training and validation purposes while keeping the resulting model as local as possible (*i.e.*, without incorporating major nonstationary effects). Our choice corresponds to considering approximately two short cycles for modeling the local dynamics and one cycle for validation, which seem to be minimum requirements for reliably learning from the data. All the ANNs were trained until the minimum of the corresponding validation set error was achieved ("early-stopping" criterion). For simplicity, we have not changed the number of hidden units, which should not be important because of the early-stopping validation.

After obtaining these models, we applied the driving-parameter reconstruction algorithm for $d = 2, 4, 6$, and 9. In Fig. 1 we present the average results for α_t in 10 independent reconstructions of this parameter. We see that for $d = 2, 4$ and 6 we obtain very similar profiles, but this changes for $d = 9$. In all cases the driving parameter shows a cyclic

[3] The recently constructed Group SSN series (Hoyt and Schatten, 1998) contains 80% more raw data than the Wolf SSN, and is proven to be more homogeneous and reliable for the times prior to 1850. However, in this work we will consider the latter since it has been commonly used in previous statistical studies.

behavior with a period slightly smaller than 100 years, corresponding to the well-known Gleissberg cycle. More interestingly, for the smallest embedding dimensions d there is also an important rise in the perturbation amplitude in the last century, in agreement with the observed increase of SSN activity (Usoskin et al., 2003). In particular, the α profile for $d = 2$ can be shown to be practically equivalent to a 22-year low-pass FFT smoothing of the SSN record[4] (see Fig. 2). The observed behavior of α is also coincident with the increase in the Sun's coronal magnetic field since 1901, recently estimated from measurements of the near-Earth interplanetary magnetic field (Lockwood et al., 1999). This effect has been attributed to chaotic changes in the dynamo that generates the Sun's magnetic properties. For $d = 9$, however, we see a different trend in the secular change. Although one is tempted to trust the consistent results for the smaller dimensions, we discuss next a different criterion to select the best reconstruction.

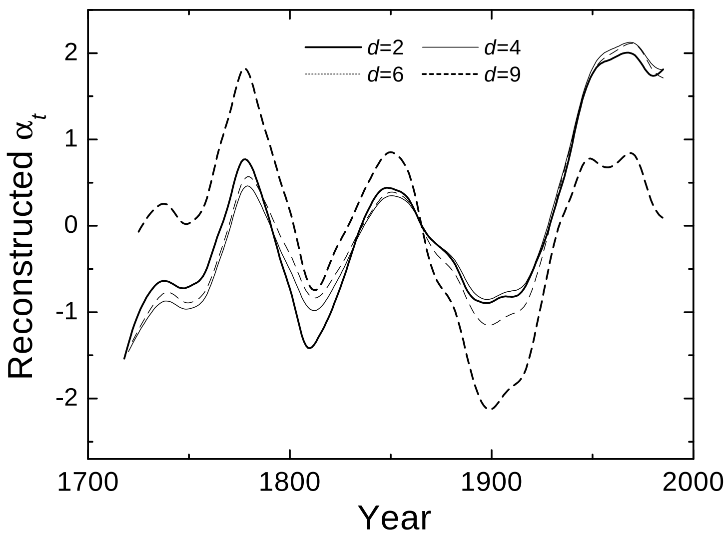

Figure 1. Average result of 10 independent reconstructions of the driving parameter α_t. The different curves correspond to local embedding dimension d equal 2 (full thick line), 4 (thin full line), 6 (thin dashed line), and 9 (thick dashed line).

Once the curve α_t has been determined, one can use it to build a global model of the sunspot dynamics that incorporates this secular (nonstationary) behavior into the process. For this we considered ensembles of ANNs, since they are known to perform better than single networks (Sharkey, 1999). We randomly selected 20% of the points in

[4] We thank one of the referees for pointing out this to us.

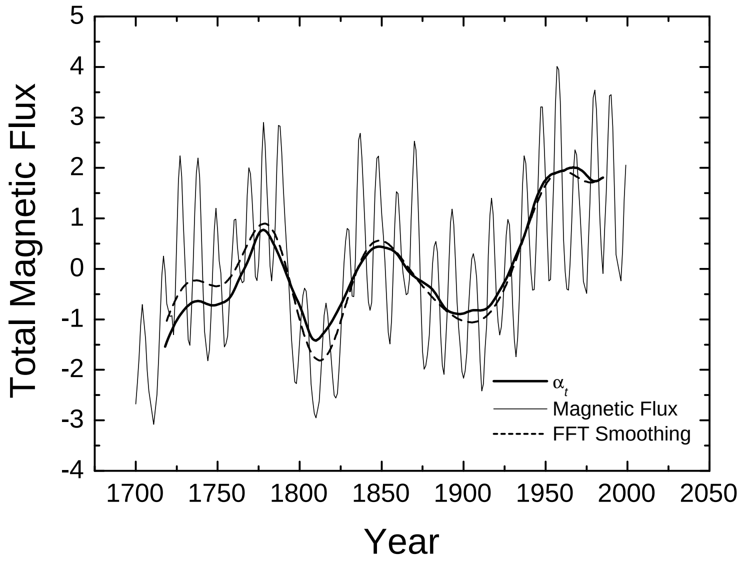

Figure 2. Comparison between the α_t profile for $d = 2$ (thick full line) and a suitably-rescaled 22-yr low-pass FFT smoothing of the SSN record (dashed line). Also plotted is the evolution of the Sun's total magnetic flux (thin full line) obtained by Solanki et al. (2002).

the whole sunspot record for testing purposes[5], and used the remaining 80% data points to train ensembles of 20 ANNs with architectures $(D+1) : 8 : 1$ $(D = 2, 9)$, following the ensemble construction technique proposed in Granitto et al. (2001). Notice that: i) in addition to the D delayed inputs there is an extra input unit that is fed with the α_t value, thus providing a local reference for the dynamics, and ii) the embedding dimension D for this global modeling needs not be the same as d.

To evaluate the validity of the reconstructed parameters in Fig. 1, we alternatively input to the extra unit the profiles obtained with $d = 2$ and 9, and considered also for comparison feeding it with a linear ramp (to show that a mere reference to time is not enough to fully uncover the nonstationary dynamics). Finally, the generalization error E_α for the obtained ensembles were estimated using the test set. For the sake of comparison, we also performed the same calculations with an architecture $(6 + 1) : 8 : 1$, where the first 6 inputs correspond to $(x_t, x_{t-1}, x_{t-2}, x_{t-3}, x_{t-8}, x_{t-9})$, an optimal set to model the SSN

[5] Notice that this procedure tests one-step-ahead predictions only. However, they provide a clear indication of the model accuracy on the whole SSN record.

dynamics according to Pi and Peterson (1994). These 9 cases (three ANN architectures times three α profiles) were run on parallel for 100 different random splits of the data, in order to produce paired t-tests on different null hypotheses. We summarize the final results as follows: i) Irrespective of the input embedding employed, with 99.9% of confidence the $\alpha_{d=2}$ profile produces better modeling effects than the other two profiles; ii) Irrespective of the driving parameter profile used, with more than 99.9% of confidence the Pi-Peterson embedding produces a better modeling of the solar magnetic activity than $D = 2$ and $D = 9$; iii) Using the Pi-Peterson embedding, with 99% of confidence the $\alpha_{d=2}$ profile leads to a better modeling than the $\alpha_{d=9}$ and linear ramp profiles.

Could one further distinguish between $d = 2, 4, 6$ in reconstructing α? To answer this question we have run 100 experiments with different random splits of the data as above, all of them with Pi-Peterson embedding and feeding alternatively the α profiles corresponding to $d = 2, 4, 6$. For the sake of comparison, we also considered modeling the sunspot record without the driving parameter input. We concluded, with a confidence level above 99%, that the profile corresponding to $d = 2$ produces smaller modeling errors.

2.3. Nonstationary vs. Stationary Modeling

Next thing to consider is: How much can we gain by modeling the solar magnetic activity in this way, instead of treating it as a stationary time series? To answer this question we computed the relative performance differences $(E - E_\alpha)/E$ between the test set errors obtained using the $\alpha_{d=2}$ profile (E_α) and without this extra input (E). The results of the 100 runs are plotted as a histogram in Fig. 3. Most of the times the ensembles of ANNs trained with α_t outperform the usual (global) fitting. In particular, its predictions are, on average, 13.4% better than those of the stationary model; moreover, for some test sets one can have up to 35% of gain in performance. There is also a loss of less than 1.2% in a few cases, but this can be associated to peculiarities of the random test sets and/or statistical fluctuations in the ANN training processes.

2.4. Consistency Checks

To provide a test on the reconstruction algorithm that leads to the profiles in Fig. 1, we performed the same study on the monthly averages of the SSN and also on the raw daily values. In the first case we modeled locally the 1749-2002 time series using a $5 : 20 : 1$ feedforward ANN, following the same steps discussed for the annual mean. For the much noisier record of daily values (available from 1849 to the present), the architecture used was $8 : 10 : 1$. These nearly-optimal

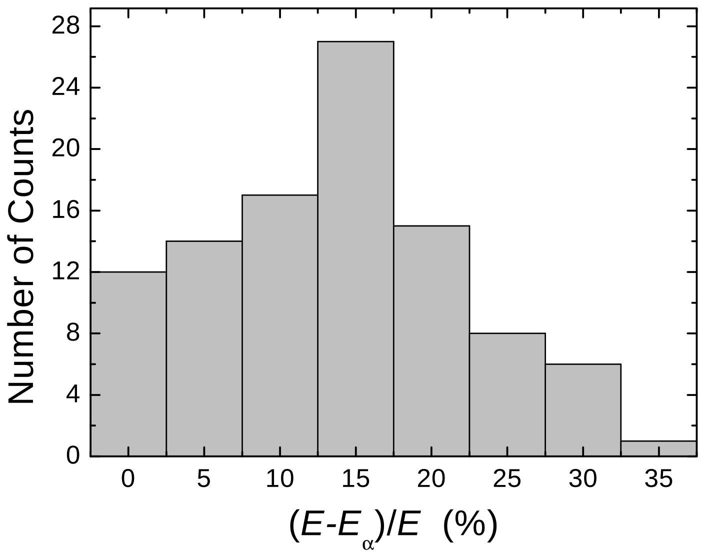

Figure 3. Relative performance difference $(E - E_\alpha)/E$ between the test set errors obtained using the $\alpha_{d=2}$ profile (E_α) and without this extra input to the ANNs (E). The results of 100 independent runs are plotted as a histogram.

architectures were chosen after some experimentation with several embedding dimensions and number of hidden units. In both cases, the local modeling was performed using again a window of 33 years, like in the calculations that led to Fig. 1. The corresponding α-profiles found reassuringly agree with the previous result obtained using the annual average, being hardly distinguishable from those in Fig. 1.

A different test on the algorithm was carried out by removing from the SSN time series features related to weak nonstationarities —linear trend and running variance. For this, we segmented the (short) solar cycles and normalized their amplitudes to 1, so that after reuniting these pieces the resulting time series has no average drift nor long-scale amplitude modulation. In spite of this arbitrary breaking into pieces and normalizing process performed on the SSN series, the overall profile obtained is, surprisingly, very similar to the results shown in Fig. 3. This is a consequence of the fact that the nonstationary changes are essentially present both in the amplitude *and* period of the short cycle, as indicated by the relation between these quantities found in studies of the solar cycle shape (Usoskin and Mursula, 2003). Furthermore, it proves the efficacy of our method in discovering hidden parameter drifts beyond simple modulation factors or linear running means.

3. Conclusions

In this work we have studied the sunspot record as a nonstationary time series. Our treatment incorporates secular changes in solar magnetic activity as an external input to the model, thus providing a local reference for the short-cycle dynamics. We found that this "hidden driving signal" input to our model must be simply the long-term trend of the SSN time series. We have shown that this procedure is able to reduce in approximately 13% the modeling error of this series when compared to the standard stationary approach. Remarkably, this performance gain is obtained using ensembles of ANNs, an already very competitive technique for this application within the standard approach.

The present work opens up different possibilities for further studies on solar magnetic activity. First, since the obtained driving parameter has a smooth variation on timescales of the order of a few short cycles, we can extrapolate this quantity to perform more stable and accurate long-term predictions of this activity. Secondly, by keeping α_t constant in multi-step predictions we may further analyze known connections between total magnetic flux and amplitude, length, and general shape of short cycles (Usoskin and Mursula, 2003). These and other studies are currently under way.

Acknowledgements

We are grateful to S. K. Solanki for useful comments and for providing us with his results of Solanki et al. (2002) plotted in Fig. 2. This work was supported by Fundación Antorchas (P.F.V.) and CONICET (P.M.G. and H.A.C.).

References

Calvo, R. A., H. A. Ceccatto, and R. D. Piacentini: 1995, 'Neural network prediction of solar activity'. *Astrophys. J.* **444**, 916.

Choudhuri, A. R.: 1992, 'Stochastic fluctuations of the solar dynamo'. *Astron. Astrophys.* **253**, 277.

Christensen-Dalsgaard, J., W. Dappen, S. V. Ajukov, E. R. Anderson, H. M. Antia, S. Basu, V. A. Baturin, G. Berthomieu, B. Chaboyer, S. M. Chitre, A. N. Cox, P. Demarque, J. Donatowicz, W. A. Dziembowski, M. Gabriel, D. O. Gough, D. B. Guenther, J. A. Guzik, J. W. Harvey, F. Hill, G. Houdek, C. A. Iglesias, A. G. Kosovichev, J. W. Leibacher, P. Morel, C. R. Proffitt, J. Provost, J. Reiter, E. J. Rhodes, F. J. Rogers, I. W. Roxburgh, M. J. Thompson, and R. K. Ulrich: 1996, 'The Current State of Solar Modeling'. *Science* **272**, 1286.

García, A. and Z. Mouradian: 1998, 'The Gleissberg cycle of minima'. *Solar Phys.* **180**, 495.

Gleissberg, W.: 1971, 'The probable behaviour of sunspot cycle'. *Solar Phys.* **21**, 240.

Granitto, P. M., P. F. Verdes, H. D. Navone, and H. A. Ceccatto: 2001, 'A Late-Stopping Method for Optimal Aggregation of Neural Networks'. *Int. J. Neur. Sys.* **11**, 305.

Hoyt, D. V. and K. H. Schatten: 1998, 'Group Sunspot Numbers: A new Solar Activity Reconstruction'. *Solar Phys.* **181**, 491.

Knobloch, E. and A. S. Landsberg: 1996, 'A New Model of the Solar Cycle'. *Mon. Not. R. Astron. Soc.* **278**, 294.

Lockwood, M., R. Stamper, and M. N. Wild: 1999, 'A doubling of the Sun's coronal magnetic field during the past 100 years'. *Nature (London)* **399**, 437.

Maunder, E. W.: 1922, 'The Prolonged Sunspot Minimum 1675-1715'. *Brit. Astron. Assoc. J.* **32**, 140.

Mininni, P. D., D. O. Gómez, and G. B. Mindlin: 2000, 'Stochastic Relaxation Oscillator Model for the Solar Cycle'. *Phys. Rev. Lett.* **85**, 5476.

Mininni, P. D., D. O. Gómez, and G. B. Mindlin: 2002, 'Biorthogonal Decomposition Techniques Unveil the Nature of the Irregularities Observed in the Solar Cycle'. *Phys. Rev. Lett.* **89**, 061101.

Naftaly, U., N. Intrator, and D. Horn: 1997, 'Optimal ensemble averaging of neural networks'. *Network: Computation in Neural Systems* **8**, 283.

Pi, H. and C. Peterson: 1994, 'Finding the Embedding Dimension and Variable Dependencies in Time Series'. *Neur. Comput.* **6**, 509.

Priest, E. R.: 1984, *Solar Magnetohydrodynamics*. D. Reidel Publishing Company.

Schwabe, H.: 1843, 'Solar observations during 1843'. *Astron. Nachr.* **20**, 495.

Sharkey, A. J. C. (ed.): 1999, *Combining Artificial Neural Nets: Ensemble and Modular Multi-Net Systems*, Perspectives in Neural Computing. Springer-Verlag.

Solanki, S. K., M. Schussler, and M. Fligge: 2000, 'Evolution of the Sun's large-scale magnetic field since the Maunder Minimum'. *Nature (London)* **408**, 445.

Solanki, S. K., M. Schussler, and M. Fligge: 2002, 'Secular variation of the Sun's magnetic flux'. *Astron. Astrophys.* **383**, 706.

Usoskin, I. G. and K. Mursula: 2003, 'Long-Term Solar Cycle Evolution: Review of Recent Developments'. *Solar Phys.* **218**, 319.

Usoskin, I. G., S. K. Solanki, M. Schuessler, K. Mursula, and K. Alanko: 2003, 'Millennium-Scale Sunspot Number Reconstruction: Evidence for an Unusually Active Sun since the 1940s'. *Phys. Rev. Lett.* **91**, 211101.

Verdes, P. F., P. M. Granitto, H. D. Navone, and H. A. Ceccatto: 2001, 'Non-Stationary Time-Series Analysis: Accurate Reconstruction of Driving Forces'. *Phys. Rev. Lett.* **87**, 124101.

Weigend, A. and N. A. Gershenfeld (eds.): 1994, 'Time Series Prediction: Forecasting the Future and Understanding the Past', Vol. XV of *Santa Fe Institute Studies in the Science of Complexity*. Addison-Wesley.

Wolf, R.: 1852, 'Sunspot Epochs Since AD 1610: The Periodic Return of Sunspot Minima'. *Acad. Sci. Compt. Rend.* **35**, 704.

Yule, G. U.: 1927, 'On a Method of Investigating Periodicities in Disturbed Series with Special Reference to Wolfer's Sunspot Numbers'. *Philos. Trans. R. Soc. London A* **226**, 267.

Zeldovich, Y. B. and A. A. Ruzmaikin: 1990, *Magnetic Fields in Astrophysics*. Gordon and Breach.