KENNETH E. WILLIAMS

PREDICTION OF SOLAR ACTIVITY WITH A NEURAL NETWORK AND ITS EFFECT ON ORBIT PREDICTION

The variability of solar activity over time demonstrates highly complex behavior. Neural networks provide useful models for predicting that behavior. Accurate predictions of solar activity are essential for satellite tracking and constellation management, orbit lifetime studies, and mission planning for long-term orbiting platforms, such as the proposed space station *Freedom*.

BACKGROUND

The variability of solar activity over time demonstrates highly complex behavior. Observations of solar activity as evidenced by the number of sunspots have been performed throughout history.[1] Regular observations have been recorded since 1610, shortly after the invention of the telescope; however, only observations made since about 1850 are considered highly reliable. In 1855, an international network of observers was organized by the Zurich Observatory in Switzerland to track daily sunspot activity. More recently, solar activity has been measured more precisely in terms of solar flux or received power per unit area at the Earth's surface at the 10.7-cm wavelength for which the Earth's atmosphere is more or less transparent. The flux, often referred to as the $F_{10.7}$ solar index or S number in units of 10^{-22} W/m^2 · Hz, has been measured daily since 1947 at the solar observatory in Ottawa, Canada. The Ottawa 10.7-cm S-number and Zurich sunspot number databases are strongly correlated (0.95 correlation coefficient) for the period from 1947 to the present. The relationship is given by the following equation:

$$F_{10.7} = 62.58 + 0.815n + 4.06 \times 10^{-4} n^2 ,$$

where n is the sunspot number. Monthly mean S numbers based on reliable sunspot observations before 1947 and on daily S-number measurements since then are shown in Figure 1.

Such observations reveal specific trends such as the well-known cycle that averages about 11.5 years in duration and was first suggested by Heinrich Schwabe in 1843. Also indicated is an apparent alternation in the amplitude of the cycle maximum and length from cycle to cycle that suggests a 22-year period and a period on the order of 100 years. Recently, evidence of a 155-day period has been suggested.[2] Although these trends provide useful insights into the behavior of the Sun, the period and amplitude of sunspots and solar flux vary significantly from cycle to cycle, making attempts to model solar behavior in a practical way more difficult. Moreover, other anomalies occurred, such as the so-called Maunder Minimum of the seventeenth century, during which few sunspots were observed and the cycles apparent in later epochs were missing or obscured.

On the basis of observed trends, attempts have been made to model solar activity empirically. One model widely used for predicting solar activity is based on Sargent's forecasting method[3] and attempts to fit the $F_{10.7}$ solar index for the last solar cycle (21) and to extrapolate the following cycle (22) as a series of cubic equations in time, as shown in Table 1. In Figure 2, predicted monthly means from Sargent's model are compared with observations obtained up to 1989.

Sargent's model itself produces noticeable discrepancies in the solar flux. As shown in Figure 2, the onset of the last solar maximum occurred about one year before the prediction. The model can be improved on, however, by applying linear regression techniques to the observation residuals based on the empirical model.[4] The linear
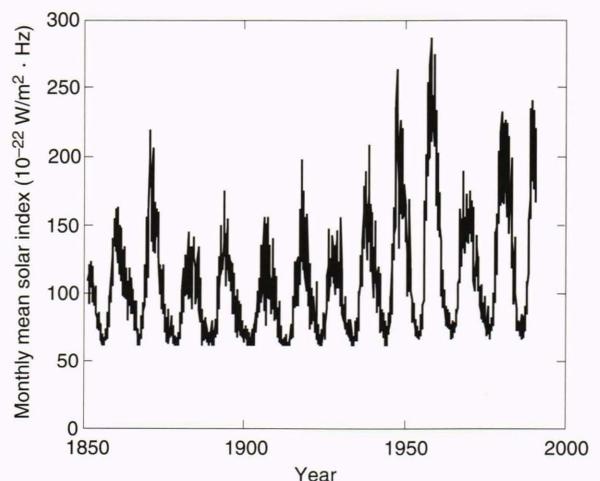
Figure 1. The history of solar activity as expressed by S number or $F_{10.7}$ solar flux for the period from 1850 to the present for which reliable observations of sunspots and direct measurements of solar flux are available.

Table 1. Cubic equations describing the best estimate for solar cycle 22 (Sargent's model).

Time frame	Time (t) (months)	Equation for 10.7-cm wavelength ($F_{10.7}$) solar flux (10^{-22}W/m^2·Hz)
Jan 82–Nov 86	1–59	$F_{10.7} = -3.3383 \times 10^{-5}t^3 + 4.0812 \times 10^{-2}t^2 - 4.444542t + 193.92$
Dec 86–Dec 87	60–72	$F_{10.7} = 66.33$
Jan 88–Apr 91	73–112	$F_{10.7} = -0.0023(t - 72)^3 + 0.1381(t - 72)^2 + 66.34$
May 91–Dec 92	113–132	$F_{10.7} = 140.1$
Jan 93–Apr 99	133–208	$F_{10.7} = -0.00034(t - 132)^3 + 0.0389(t - 132)^2 + 140.0$
May 99–Aug 02	209–248	$F_{10.7} = -0.00338(t - 208)^3 + 0.2025 (t - 208)^2 + 65.0$
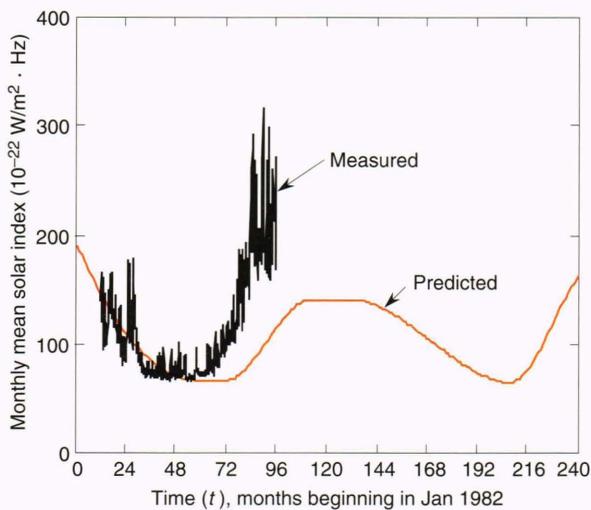
Figure 2. A comparison of the prediction of solar cycle 22 on the basis of Sargent's empirical model with monthly means derived from observations available through the end of 1989.

regression is performed over a period of at least one year for which measured S numbers are available. Predicted solar indices are then derived by extrapolating this linear function and adding it to the curve in Figure 2. The algorithm that performs the process is called a linear, or alpha–beta, tracker. This technique is the most widely used method for predicting solar activity.

RATIONALE

My particular interest in modeling solar activity is centered on accurately predicting the orbital position of low-orbiting satellites (altitudes less than 1000 km) for periods of up to one year. Solar flux affects upper atmospheric density directly through heating and indirectly through geomagnetic activity. The density, in turn, affects the amount of drag exerted against satellites and other objects orbiting at altitudes of about 3000 km and lower. The Jacchia 1977 atmospheric density model[5] provides atmospheric density above 90 km as a function of altitude, solar activity, and other factors.

Accurate predictions of solar activity will be useful for planning and supporting future space missions. Precise lifetime estimates are important determinants in the design of spacecraft, driving possible requirements for station-keeping systems and fuel capacity. In addition to unmanned objects, manned space stations such as the proposed *Freedom* will require periodic refueling to maintain altitude. For *Freedom*, altitude must be kept between 240 and 400 km.[6] At the lower extreme, the station risks going out of control and reentering the Earth's atmosphere in less than 90 days; at the upper extreme, the level of radiation exposure to the crew begins to exceed safety limits. The space shuttle will make periodic flights to resupply the station with fuel for altitude maintenance. The planned frequency of the flights will depend heavily on the predicted level of solar activity over time. Because of the unexpected orbital decay and subsequent incineration of our first manned space station, Skylab, the accurate prediction of solar-induced drag forces is an issue of critical importance. The importance of solar flux forecasts on orbit predictions will be discussed later.

Prediction of solar activity is relevant for other applications. For example, RF communications become highly disrupted during periods of heightened solar activity. By knowing precisely when such periods are likely to occur, government and private organizations can plan to obtain access to alternative channels of communication or to vary the use of critical data links to minimize the impact of disruptions. It is well known that the Sun is the engine that generates climatic changes on Earth. The level of solar activity will be a critical parameter for both short-term weather forecasting and long-term atmospheric modeling.

NEURAL NETWORK APPROACH

By examining solar activity, as evidenced by sunspot number, as a nonlinear dynamical system, researchers at the University of Colorado demonstrated by means of the techniques of chaotic dynamics that the sunspot cycle is a low-dimensional system for which the interval of useful prediction can be extended theoretically from months to

years.[7] Neural networks provide a useful approach for estimating chaotic time series as noted in applications of neural networks to signal processing.[8] Given the experience at APL in applying neural networks to solving a variety of problems, this approach appears to be quite promising.

The prediction model includes a simulated multilayer feed-forward network (MLFN) with back propagation. The MLFN (see Fig. 3) consists of three layers of processing units. Signals are passed in one direction from the input through the hidden layer to the output layer. Each unit in the middle or hidden layer is connected to all units in both the input and output layers. The input pattern consists of a series of solar index measurements at regular intervals that have been averaged or smoothed over a given time period and normalized so that the measurements fall in the range from 0 to 1. The scale factors required to do this are retained and applied to the network output later to recover the predicted solar indices with proper scaling. The input signals are fed directly into the input layer. Each unit in the input layer transmits the input signal to every unit in the hidden layer. The connections between the input and hidden units are weighted, such that each signal received by a unit in the hidden layer is either diminished or augmented according to the weight of the connection. Such weights are adjusted during training in accordance with a process that will be described later. An additional input signal, which is held constant (i.e., a normalized value of 1), is also provided to each hidden unit to afford stability during training and an extra degree of freedom in the underlying function. This signal may be thought of as originating in a special input unit known as a bias unit.

Each hidden unit takes a sum of the weighted signals received from all the input units. The sum constitutes the activity of that unit. A transfer function is applied to the activity to determine the signal that will be sent from the hidden unit to each output unit. Signals from the hidden units are passed to the output layer over another set of weighted connections. The process of combining signals and applying a transfer function is repeated at the output layer. Signals from the output units form an output pattern that represents predicted solar indices over a time period subsequent to that represented by the input pattern. The transfer function used for the hidden and output layers is the sigmoid or "squashing" function $f(x) = 1/(1 + e^{-x})$ shown in Figure 3. The squashing function ensures that the output signals are normalized to the range 0 to 1. This characteristic of the MLFN provides the key to modeling nonlinear time series.

The Ottawa 10.7-cm S-number and Zurich sunspot number databases are used to form the required training and testing sets for the MLFN. Sunspot numbers are transformed as described above into S numbers to provide training data for periods before 1947. The daily S numbers are usually smoothed over periods of about one month. Also, it is possible to remove fluctuations in the data above a particular frequency with the aid of a low-pass filter that uses a fast Fourier transform. Smoothing and/or filtering facilitates the operation of the neural network model by reducing the required number of units

and connections. Moreover, short-period fluctuations in solar flux over intervals of less than a month have no utility in predicting long-term drag effects and are not of interest for our particular application.

Training is performed by back propagation from the output layer down to the input layer using the method of
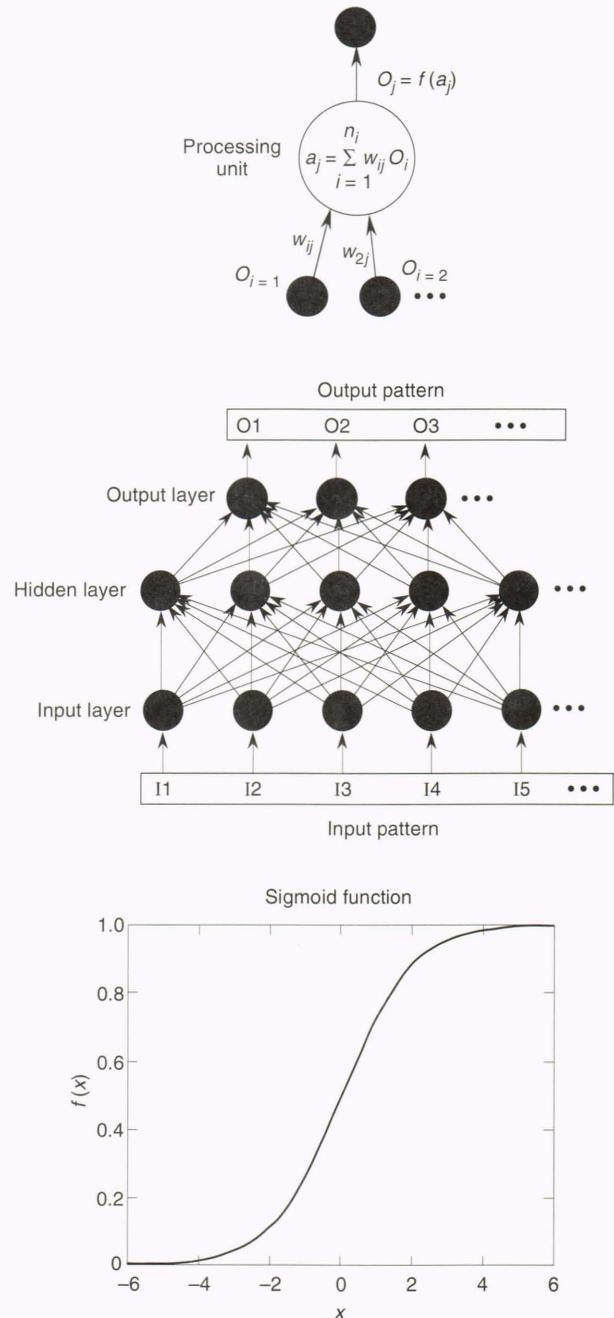
Figure 3. An overview of multilayer feed-forward network architecture. Units are interconnected as shown in the center of the figure. Processing in the hidden and output layers is performed in accordance with the equations shown at the top of the figure, where O_i represents a signal from a unit in a lower layer, w_{ij} is the connection weight between the units, a_j is the computed activity, and O_j is the output signal. The transfer function f, known as the sigmoid or "squashing" function and defined as $f(x) = 1/(1+e^{-x})$, is plotted at the bottom of the figure.

steepest descent.[9,10] Each training case consists of a set of input S numbers representing consecutive 32-day smoothed values and a corresponding set of output S numbers, known as the target pattern, for a subsequent time span of one or more consecutive 32-day periods. Both the input and target patterns are derived from measured Ottawa S numbers or Zurich sunspot numbers. Before training, random weights are assigned to all the network connections. During training, an output pattern is predicted for a given time period and compared with the target pattern. Adjustments are made to the connection weights to minimize the total squared error based on observation residuals across all the output units and given by

$$E_{\mathrm{T}} = \frac{1}{2} \sum_{p=1}^{N_p} \sum_{k=1}^{n_O} (t_k^p - O_k^p)^2 \,,$$

where t_k^p is the target value and O_k^p is the output signal for the kth output unit and pth training case, n_O is the number of units in the output layer, and N_p is the number of training cases. It is often convenient to normalize E_{T} with respect to all training or testing patterns to provide a consistent measure of effectiveness for training the network; the normalized total squared error is given by

$$\hat{E}_{\mathrm{T}} = \frac{\displaystyle\sum_{p=1}^{N_p} \sum_{k=1}^{n_O} (t_k^p - O_k^p)^2}{\displaystyle\sum_{p=1}^{N_p} \sum_{k=1}^{n_O} (t_k^p - \bar{t}_k)^2} \,,$$

where $\bar{t}_k$ represents the average target value for output unit k across all training or test patterns. The double sum in the denominator is sometimes called the "magic" number. When the unnormalized total squared error begins to fall below the magic number (i.e., $\hat{E}_{\mathrm{T}} = 1$), this indicates that the network has learned to estimate at least the mean of the data presented and that the training process is progressing normally.

The training process is repeated over many iterations, typically numbering on the order of hundreds to thousands. During each iteration, weights are adjusted by a small amount to reduce the total squared error. Weight changes are described by

$$\Delta\mathbf{w}(t+1) = -\eta \nabla E_{\mathrm{T}} + \alpha \Delta\mathbf{w}(t) \,,$$

where $\Delta\mathbf{w}(t)$ is the vector of all weight changes from the tth iteration (initially zero), η is the learning rate (0 to 1, typically 0.01), ∇E_{T} is the gradient of the total squared error in the space whose dimensions represent individual connection weights (i.e., the solution space), and α is the momentum parameter (positive, 0.7 typically). The last term in the equation, often called the momentum term, is included to add stability to the training process by minimizing high-frequency fluctuations in the weights and by avoiding local minima in the solution space. A local minimum corresponds to a description of the underlying function for the solar flux that works well only for a specific case. The global minimum, or best solution for all cases in general, is the desired outcome of the training process. The concept of local and global minima is illustrated in Figure 4. Weight changes in accordance with the method of steepest descent are defined more precisely in the boxed insert.

Periodically, after a certain number of iterations, the performance of the network is tested independently with the test set. The training process continues until the total squared error across the outputs is minimized for the test set. At this point, typically after hundreds to thousands of iterations, the network has "learned" the functionality of the solar flux data as well as it can. Although further

BACKPROPAGATION EQUATIONS FOR THE METHOD OF STEEPEST DESCENT

The weights for the connections between the hidden and output layers are adjusted as follows:

$$w_{jk}(t+1) = w_{jk}(t) + \eta \sum_{p=1}^{N_p} \delta_k^p O_j^p + \alpha \Delta w_{jk}(t) \,,$$

where

$$\delta_k^p = (t_k^p - O_k^p) O_k^p (1 - O_k^p) \,,$$

$w_{jk}(t)$ is the connection weight between the jth hidden unit and kth output unit for the tth training iteration, η is the learning rate, N_p is the number of training cases, O_j^p is the output signal from the jth hidden unit for the pth training case, α is the momentum parameter, $\Delta w_{jk}(t)$ is the hidden-to-output weight change from the tth iteration (initially zero), and t_k^p and O_k^p are the target value and output signal, respectively, for the kth output unit and pth training case.

The weights for the connections between the input and hidden layers are adjusted in accordance with the following equation:

$$w_{ij}(t+1) = w_{ij}(t) + \eta \sum_{p=1}^{N_p} \delta_j^p O_i^p + \alpha \Delta w_{ij}(t) \,,$$

where

$$\delta_j^p = \left[\sum_{k=1}^{n_O} \delta_k^p w_{jk}(t) \right] O_j^p (1 - O^p) \,,$$

$w_{ij}(t)$ is the connection weight between the ith input unit and jth hidden unit for the tth training iteration, O_i^p is the output signal from the ith input unit for the pth training case, $\Delta w_{ij}(t)$ is the input-to-hidden weight change from the tth iteration (initially zero), and n_O is the number of units in the output layer. The terms α, η, N_p, O_j^p, $w_{jk}(t)$, and δ_k^p are defined as above.
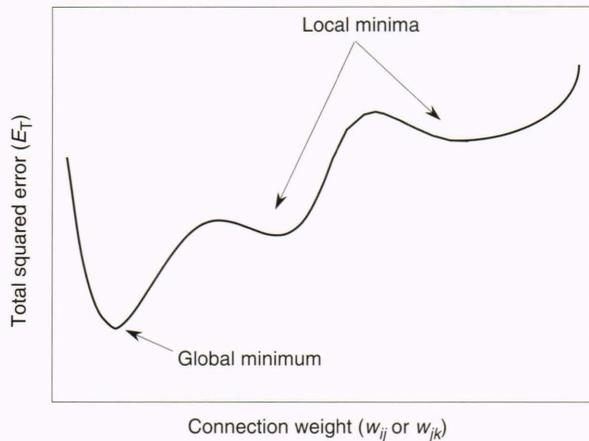
Figure 4. A two-dimensional representation of N-dimensional solution space, where E_T is a function of the N connection weights modeled.

Figure 5. An example of the convergence of the network during training. The normalized total squared error ($\hat{E}_T$) is plotted versus training iterations for the training and test sets. When the normalized total error for the test set reaches a minimum (indicated by the dashed line), the network has "learned" the rules governing the underlying behavior of the solar fluxes as well as it can. Training set = black line. Test set = red line.

training will continue to reduce the error in the training data, such training constitutes overtraining of the network, and the performance as measured by the test set will begin to deteriorate. This deterioration indicates that the network is beginning to "memorize" the noise in the training data instead of generalizing the underlying function controlling the behavior of the training and test data. Figure 5 shows an example of training, where the vertical broken line denotes the approximate optimal number of iterations at which to end training and to freeze the weights.

In each instance, the optimal set of weights is then saved and used with the network to predict solar indices over a span not used in either training or testing, as described above. The network is employed recursively so that the outputs obtained from one prediction form a portion of the input set for the next prediction, which results in predictions of solar activity for as long as several years.

PRELIMINARY RESULTS

Various combinations of parameters dealing with the network configuration and training method have been used to generate a meaningful prediction of solar activity, including numbers of units employed for each of the three layers of the MLFN, the cutoff frequency used for processing the S numbers before training, the size of the training and test sets, and the interval covered by each training and test case. The learning rate was fixed at 0.01, and the parameter in the momentum term was set to 0.7. All weights were randomized initially to values between –0.3 and 0.3. Solar index values of $F_{10.7}$ were scaled so that the lowest and highest possible indices would correspond to input values of 0.1 and 0.9, respectively, to avoid saturation of units in the hidden and output layers. In all cases, we used 32-day average solar indices with 120 units in the input layer, so that nearly a full solar cycle could be used by the network to ascertain subsequent solar indices. Also, on the basis of initial studies, six units were determined to be adequate for the hidden layer.

Beyond a required minimum value, the number of hidden units used has been found to have a negligible effect on all results obtained thus far. An additional bias unit was included in each hidden and output layer in each instance. Training was performed using solar data from 1900 to 1985; data from earlier epochs (1850 to 1899) formed the test set used for weight adjustments. Training and testing over shorter intervals worsened the predictive capabilities of the neural network considerably. Therefore, we found it beneficial to train on as much reliable data as possible. After completion of training, the prediction performance of the network was tested over recent periods since the end of 1985. Use of a low-pass filter to preprocess the solar data before training was limited to a cutoff frequency of about six cycles per year. Although lower cutoff frequencies have been suggested[7] to facilitate the neural network prediction, we have found that overfiltering data tends to impede the prediction of solar activity, particularly during periods of solar maxima.

The results illustrated in Figures 6 and 7 are among the best obtained thus far. Figure 6 shows the prediction performance for a network that contains a single output unit representing a 32-day period. Predictions are propagated recursively starting from the years shown. The top graph shows the actual predictions compared with 32-day averages of the measured $F_{10.7}$ solar index over the period from the beginning of 1986 to the present. The bottom graph reveals the residuals or differences between the various predictions and the measured $F_{10.7}$ solar indices over the first year of each prediction. Note that the performance is much better for predictions beginning before the solar activity reached a maximum in early 1989. Unfortunately, the larger residuals during peak periods in the solar cycle have a much greater impact from the
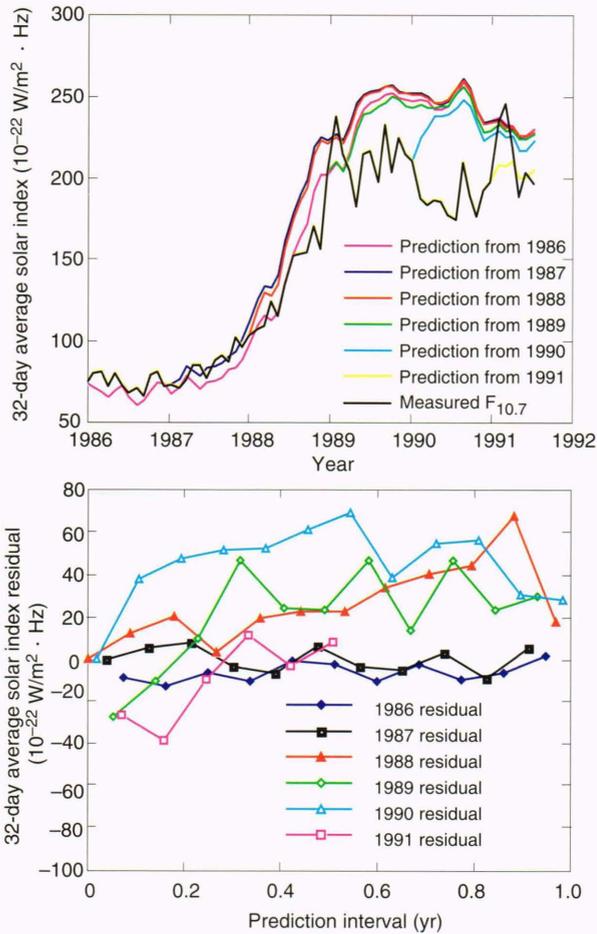
Figure 6. Performance of a neural network (120 input units plus a bias unit, 6 hidden units plus a bias unit, 1 output unit) in predicting 32-day average solar fluxes from 1986 to the present. About 900 iterations were sufficient for convergence.
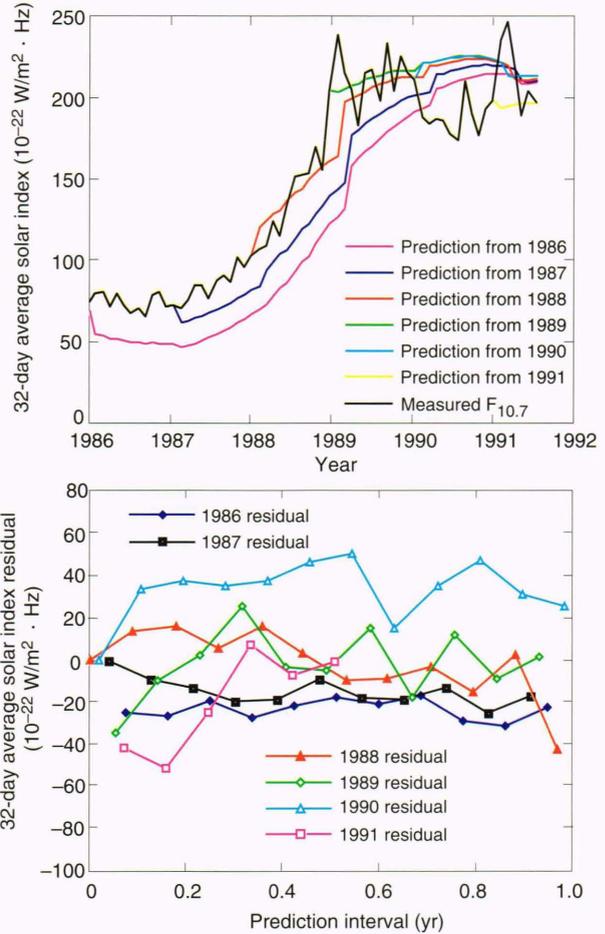
Figure 7. Performance of a second neural network (120 input units plus a bias unit, 6 hidden units plus a bias unit, 12 output units) in predicting 32-day average solar fluxes. About 250 iterations were needed for convergence.

standpoint of estimating drag on satellites. If the number of output units is increased to twelve so that an entire year of solar flux predictions is made during a single pass through the network, the results shown in Figure 7 are obtained. Again, the top graph shows the actual predictions compared with 32-day averages of the measured $F_{10.7}$ solar index over the period from the beginning of 1986 to the present; the bottom graph reveals the corresponding residuals for the first year of each prediction. The performance during peak periods from 1989 onward is noticeably improved with some worsening of performance for earlier periods. For satellite orbit prediction, however, the performance shown in Figure 7 is preferable because it results in smaller errors in the overall drag calculation. For both networks, however, the 1990 prediction was particularly poor. Perhaps the network was unable to anticipate the relative drop in solar activity during that year because the underlying trend was not embedded in the data available for training.

Although the problem of predicting solar activity during maxima has not been resolved satisfactorily, the neural network generally performs better than the linear regression method in predicting solar activity. Figure 8

shows predictions over a period of one year based on the linear regression tracker applied to Sargent's empirical model (top graph) compared with predictions provided by the neural network with twelve output units (bottom graph) as described earlier. Predictions shown are for all years since 1988, during which solar activity has been relatively high. The measured values of solar index shown are based on daily values that have been smoothed by applying a sliding average of approximately one month. Comparison of the two graphs reveals that the neural network generally is doing at least as well as, if not better than, the linear regression technique in following the trend of solar activity. The improvement in prediction capability is evident in Figure 9, where the differences in the residual magnitudes of Figure 8 are shown. The curves usually fall above zero on the vertical axis, suggesting an improvement in the prediction of solar indices overall.

Another way to illustrate the potential of this improvement is in terms of its effect on orbit prediction. We can use the solar indices as input to a drag calculation in association with an analytic orbit integrator to predict orbit lifetime. Figure 10 demonstrates the effect on pre-
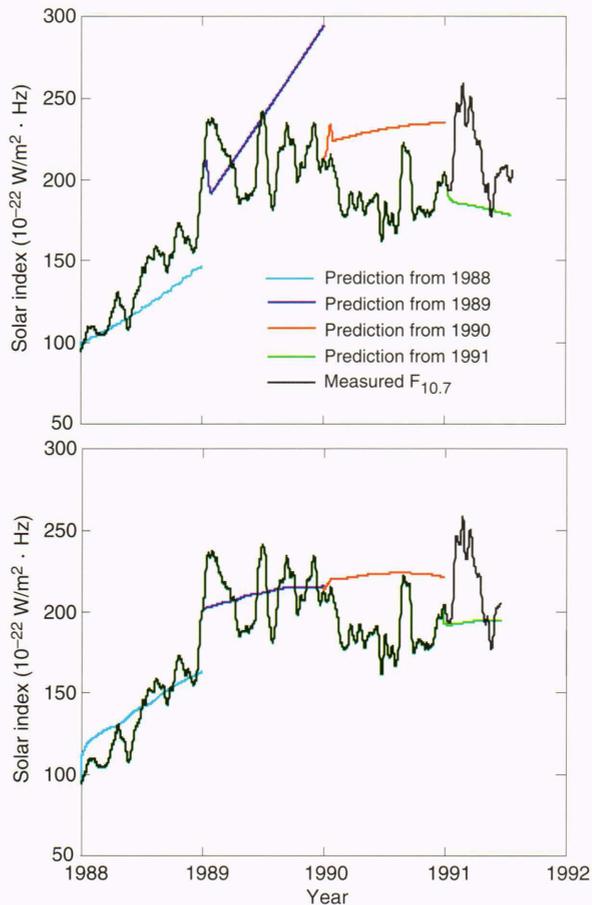
Figure 8. A comparison of one-year predictions from the linear tracker applied to Sargent's empirical model (top) and neural network (bottom). The neural network contains 120 input units plus a bias unit, 6 hidden units plus a bias unit, and 12 output units. The measured solar fluxes shown are based on the application of a 30-day sliding average to the daily values.
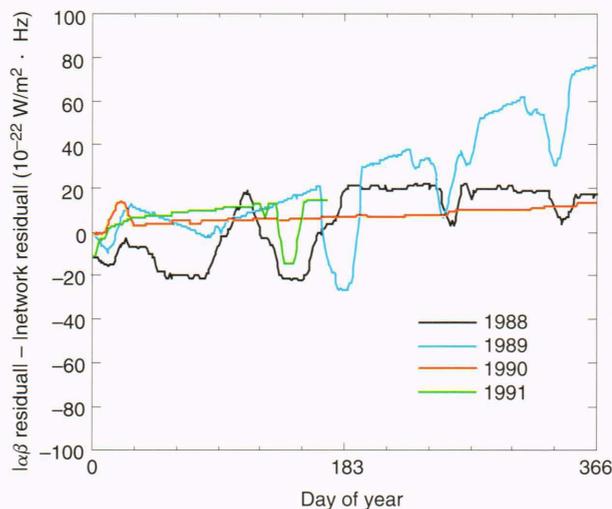
Figure 9. A comparison of residuals from solar flux predictions based on the neural network (120 input units plus a bias unit, 6 hidden units plus a bias unit, and 12 output units) and the linear tracker. A positive value indicates an improvement in the prediction capability afforded by the neural network.
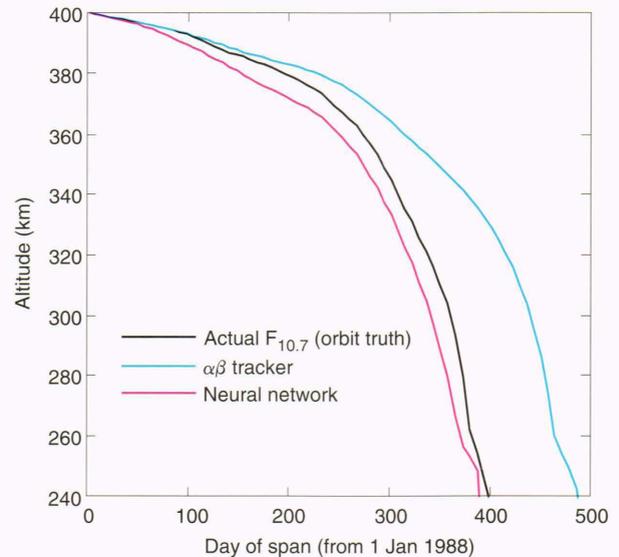
Figure 10. A comparison of orbit lifetimes based on different prediction methods with those based on actual solar data for an object similar to space station *Freedom* (area-to-mass ratio = 4.54 × 10^{-3} m^{-2}/kg^{-1}; coefficient of drag = 2.2).

dicted orbit lifetimes for an object similar in size and weight to the proposed space station *Freedom* as shown from 1988 onward, corresponding to the most recent solar maximum. The altitudes shown are based on orbit predictions generated with the aid of the analytic orbit integrator used in the personal computer orbit predictor.[11] As shown in Figure 10, the orbit altitude, based on perfect knowledge of average solar activity, decreases much faster than the altitude derived with activity predicted via linear regression of residuals from Sargent's model. The station will fall below the specified minimum altitude of 240 km on day 400 rather than around day 490. On the basis of the predicted fluxes provided by the linear tracker, the orbit lifetime would be overestimated by about 90 days. If the neural network solar flux prediction is used, however, the predicted lifetime falls at 390 days, just ten days short of the expected lifetime of the space station under such circumstances. The implication is that current solar flux prediction capabilities provide a very limited margin of error in planning for refueling of the space station in order to maintain orbit. Clearly, the neural network offers considerable improvement in the area of orbit predictions.

FUTURE PLANS

The foregoing represents only a cursory examination of the potential of the neural network prediction techniques. On the basis of preliminary findings, I believe that the prospects for improving solar predictions further via such techniques are promising indeed. I will continue to examine variations in the network configuration and the training technique to achieve an optimal prediction. Some variations under investigation include determining the numbers of units used in each layer, applying training and test sets covering different epochs, using various

cutoff frequencies in a low-pass filter applied to the solar data before training, sampling solar data as particular intervals rather than sequences of consecutive indices, and using additional hidden layers, multistage or recursive networks,[12] and higher-order neural networks. Another approach under study involves fitting the solar data to a differential equation that could then be integrated numerically, possibly as one component in a multistage solar index predictor.

REFERENCES

[1]Foukal, P. V., "The Variable Sun," *Sci. Am.* **262**(2), 34–41 (Feb 1990).
[2]Carbonell, M., and Ballester, J. L., "A Short-Term Periodicity Near 155 Day in Sunspot Areas," *Astron. Astrophys.* **238**, 377–381 (1990).
[3]Tobiska, W. K., Culp, R. D., and Berth, C. A., "Predicted Solar Cycle Twenty-Two 10.7 cm Flux and Satellite Orbit Decay," *J. Astronaut. Sci.* **35**(4), 419–433 (Oct–Dec 1987).
[4]McConahy, R. J., "Logic for S-Number Processing in PCOP," JHU/APL S1A-053-89 (20 Apr 1989).
[5]Jacchia, L. G., *Thermospheric Temperature, Density and Composition: New Models,* Special Report No. 375, Smithsonian Astrophysical Observatory, Cambridge, Mass. (15 March 1977).
[6]McDonald, B. M., and Teplitz, S. B., "Space Station Freedom Altitude Strategy," in *Proc. Flight Mechanics/Estimation Symp.*, NASA Goddard Space Flight Center, Greenbelt, Md., p. 165 (May 1990).
[7]Mundt, M. D., Maguire, B. II, and Chase, R. R. P., "Chaos in the Sunspot Cycle: Analysis and Prediction," *J. Geophys. Res.* **96**(A2), 1705–1716 (1 Feb 1991).
[8]Farmer, J. D., and Sidorowich, J. J., "Predicting Chaotic Time Series," *Phys. Rev. Lett.* **59**(8), 845–848 (24 Aug 1987).
[9]Rumelhart, D. E., Hinton, G. E., and Williams, R. J., "Learning Internal Representations by Error Propagation," Chap. 8 in *Parallel Distributed Processing*, Vol. 1, Rumelhart, D. E., and McClelland, J. L. (eds.), MIT Press, Cambridge, Mass., pp. 318–362 (1986).
[10]Sigillito, V. G., "Associative Memories and Feedforward Networks: A Synopsis of Neural-Network Research at the Milton S. Eisenhower Research Center," *Johns Hopkins APL Tech. Dig.* **10**(3), 254–261 (1989).
[11]Williams, K. E., "Personal Computer Orbit Predictor (PCOP) User's Guide," enclosure to JHU/APL S1A-061-90 (1 Aug 1990).
[12]Nguyen, D. H., and Widrow, B., "Neural Networks for Self-Learning Control Systems," *IEEE Control Sys. Mag.*, 18–23 (1990).

THE AUTHOR
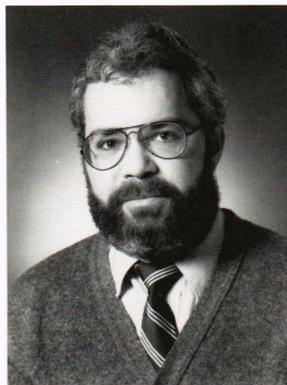
KENNETH E. WILLIAMS was educated at Indiana State University, where he received a B.S. in physics and mathematics in 1977 and an M.A. in physics in 1980. From 1981 to 1982, Mr. Williams was an instructor of physics at Eastern Illinois University in Charleston, Illinois. In 1982, he joined APL's Fleet Systems Department and in 1986 became a member of the Strike Warfare Group of the Naval Warfare Analysis Department. He joined the Computer Science and Technology Group in the Space Department in 1989, where he is involved in several projects, including the Personal Computer Orbit Prediction System, neural network applications, and the midcourse space experiment simulator.