A Two-Stage Multi-Agent System to Predict the Unsmoothed Monthly Sunspot Numbers

Mak Kaboudan

Abstract—A multi-agent system is developed here to predict monthly details of the upcoming peak of the 24th solar magnetic cycle. While studies typically predict the timing and magnitude of cycle peaks using annual data, this one utilizes the unsmoothed monthly sunspot number instead. Monthly numbers display more pronounced fluctuations during periods of strong solar magnetic activity than the annual sunspot numbers. Because strong magnetic activities may cause significant economic damages, predicting monthly variations should provide different and perhaps helpful information for decision-making purposes. The multi-agent system developed here operates in two stages. In the first, it produces twelve predictions of the monthly numbers. In the second, it uses those predictions to deliver a final forecast. Acting as expert agents, genetic programming and neural networks produce the twelve fits and forecasts as well as the final forecast. According to the results obtained, the next peak is predicted to be 156 and is expected to occur in October 2011- with an average of 136 for that year.

Keywords—Computational techniques, discrete wavelet transformations, solar cycle prediction, sunspot numbers.

I. INTRODUCTION

THE Zurich sunspot number is an index representing daily appearances of huge dark areas on the Sun's visible surface. A daily sunspot number is computed using information gathered from observatories located around the world. The daily number is $N_d = A (10 g + f)$, where A is an adjustment factor that accounts for observing conditions as well as differences among observatories and among observers, g is the number of groups of sunspots, and f is actual count of visible spots. The daily numbers are averaged monthly (N_m) and annually (N_a). Figure 1 has plots of N_a of the past 23 cycles. As the figure shows, the sunspot number has a clear cyclical pattern where the observed cycles differ in length and magnitude. Predicting the lengths and magnitudes of them remain a challenge. Perhaps this is why the annual sunspot number is among the most analyzed and forecasted time series in history. [1-8] are among the many who forecasted the series. Having an accurate forecast of the next cycle peak is critical to plan satellite orbits and space missions as well as manage power and communication systems on Earth [7].

Although a link between the "annual" sunspot number and disruptions to high-frequency radio communications, radars, and power systems on Earth does exist, it seems that such link should be even stronger when evaluating the impact of "monthly" solar magnetic activity. Figure 2 shows the more pronounced variations in the average monthly counts (N_m) relative to the annual counts (N_a) for the last few cycles. Clearly, the monthly average sunspot number naturally provides more detailed information than the annual number. Logically then, predicting short term (i.e., monthly rather than annual) variation in magnetic activities, particularly when they are strong, may help plans to ward off impacts and economic damages the strong magnetic activities may cause.


Fig. 1. Historical representation of the Zurich annual sunspot relative number 1749-2008.

Fig. 2. The dotted line is of the unsmoothed monthly sunspot numbers (N_m) while the continuous one depicts the annual (N_a) sunspot numbers for 1951-2008.

The focus in this paper is on predicting the "unsmoothed" monthly average sunspot number. The unsmoothed monthly number (N_m) is a simple summation of the observed daily numbers for a month divided by the number of days in that month. Investigating and forecasting unsmoothed monthly sunspot numbers is not the norm. Estimating and forecasting the smoothed monthly numbers (SN_m) is (see [9-13] for example). SN_m are calculated as an annual moving average using thirteen months. More specifically,

$$SN_{m7} = [(N_{m1}/2)+(N_{m2}+N_{m3}+...+N_{m12})+ (N_{m13}/2)]/12 \qquad (1)$$

where if N_{m1} = January of current year, N_{m7} = July and N_{m13} = January of next year, SN_{m7} = the smoothed average for July

Mak Kaboudan is professor of statistics with the School of Business, University of Redlands, Redlands, CA 92373 USA (phone: 909-748-8772; fax: 909-335-5125; e-mail: Mak_Kaboudan@Redlands.edu).

[14]. Other monthly averages are computed similarly.

In general, predictions of the sunspot number are obtained using either time series analyses (rooted in statistical methods as well as computational techniques) or precursor methods (rooted in solar physics and pioneered by [1]). Statistical methods capture the dynamics of the observed cyclical behavior of the numbers recorded. Precursor methods rely on hypothetically possible physical changes. Because neither time series analyses nor precursor methods are based on the true fundamentals underlying the solar magnetic cycles (because they are unknown), obtaining accurate predictions using computational powers is warranted. While there are already 45 predictions of solar cycle 24 tabulated in [15], this paper introduces yet another method to predict the magnitude and timing of cycle 24. One key assumption distinguishes this effort from others, specifically, what happens in a current time period follows what happened in a prior set of periods at a distant past, i.e.; a lag structure. An examination of the observed data suggests that it takes about four years from a low point defining the beginning of a cycle for a maximum or cycle peak to be reached, and from that peak, it takes six to seven years to decay reaching its end and the birth of a new one [13]. A model can then be constructed such that periods with a given trend are explained by the preceding cycle's similar trend. Thus, observed values along the range when the sunspot numbers are rising (or when a cycle is at its beginning) should follow a path similar to the prior period of time when they were also rising, and so on. Figure 3 is a demonstration of such a lag structure scheme. If all cycles were eleven years long, the lag structure would also be eleven years long and a logical model to capture such dynamics using monthly data would be $N_{m,t} = f(N_{m,t-132})$ where t = 1, …, M months and t-132 represents the eleven-year lag. Since cycles are not all eleven years long (the mode is ten), the lag length may then be structured such that it captures a combination of lags that may account for variations in cycle-lengths . The lag structure then should capture variations that represent cycles between six and twelve years in length and the functional model would be $N_{m,t} = f(N_{m,t-72}, …, N_{m,t-144})$.

Adopting such a lag structure is a significant departure from existing time-series and physically-based methods. Time-series methods extract information from observed or measured historical dynamics to use in predicting future activity. Studies applying time-series methods to analyze and predict the sunspot numbers are plentiful [2-4, 16-18]. Physically-based or precursor-like methods forecast solar activity using estimation of the strength of the Sun's solar dynamo by monitoring geomagnetic precursors (geomagnetic fluctuations) near solar minimum [1].

The proposed adoption of a distant lag structure has an advantage. In $N_{m,t} = f(N_{m,t-72}, …, N_{m,t-144})$, values of a model's explanatory (right-hand-side variables) are known *a priori*. This means that values of the explanatory variables do not have to be forecasted when predicting a number of periods equal to the lowest lag length. Using actual values of explanatory variables reduces prediction errors. Using actual values is particularly beneficial when the system's dynamics are nonlinear as is the case with sunspot numbers.


Fig. 3. Determining lag structure according to the direction the numbers are moving. Increasing values assumingly follow increasing values while decreasing ones follow decreasing one.

The monthly data (N_m) employed start October, 1950 and end September 2008. The Estimation period starts February 1960 and ends May 2003. Model validation starts June, 2003, and ends September, 2008. The forecast period starts October, 2008 and ends February, 2014. The monthly data of unsmoothed sunspot numbers (N_m) were obtained from the National Geographic Data Center, NOAA [19].

Several agents are employed to model and predict the monthly sunspot number. The agents first compete then cooperate to produce what is postulated to be a more reliable forecast. All agents rely on detecting pattern-recognitions the historical monthly values may have using two computational techniques, genetic programming and neural networks. The methodology employed is presented in Section II. Estimation results are in Section III. Predictions of solar cycle 24 are in Section IV. The final section has the conclusion.

II. METHODOLOGY

This section introduces the structure of a two-stage multi-agent system to predict the monthly numbers. Acting as expert agents, genetic programming (GP) and neural networks (NN) produce twelve competing fits and forecasts of the monthly series first. In the second stage, the expert agents employ the twelve fits and forecasts as inputs to produce a final forecast. The development of the multi-agent system is based on two assumptions: (a) the dynamics of monthly sunspot numbers are characteristically nonlinear dynamics as Figure 2 above suggests, and (b) the observed values may be a combination of deterministic and random unknown functions. If these two assumptions are realistic, computational search techniques become rather attractive and well-suited to find models that should deliver reasonable forecasts. The twelve competing fits and forecasts are obtained by independent agents that utilize the modeling capabilities the two expert agents. The models they produce utilize the observed series and two different transformations of them. In the first transformation, observed series are decomposed using a discrete wavelet transformation (DWT). In the second, the observed series are decomposed using a discrete Fourier transformation (DFT). Prior applications of GP and NN to DWT are in [20-22]. DFT was never used to transform sunspot numbers. Both

transformations amount to adopting a divide-and-conquer strategy. The observed values are transformed into more than one component first. The components are then modeled and estimated separately. This strategy is attractive because reconstruction of the original observed series using the estimated values of the components can be obtained by simple inverse transformation processes. Using this strategy, twelve fits and forecasts of N_m are obtained independently. These twelve N_m outputs are then used as input to produce a reconciliatory final forecast of N_m. Figure 4 has a flow chart of the multi-agent system used to obtain the final forecast values of N_m.
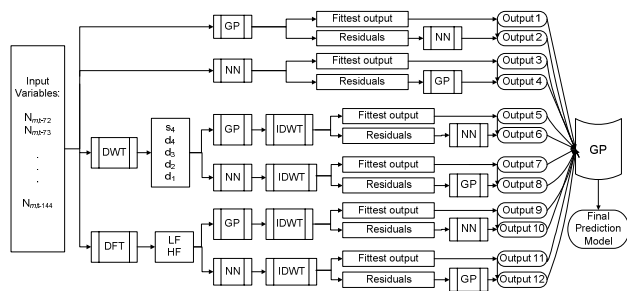

Fig. 4. A flow chart showing the twelve outputs the agents deliver to use in obtaining a final forecast of the monthly sunspot number.

Figure 4 start to the very left with a set of input variables. Output 1 (appearing near the end of the right side of the flow) is produced by an agent who applies GP to the raw data. Output 2 is produced by adding Output 1 to an output produced by an agent who uses NN to find a model that fits the residuals from Output 1. The residuals from Output 1 represent variations in the dependent variable that GP failed to explain. Using NN to fit those residuals (employing the exact same explanatory variables as GP) may succeed in capturing a portion of what the first agent missed. Output 2 is the summation of Output 1 values to the NN-fitted GP residuals. The assumption is that the new values should (but not necessarily) deliver results that are better than Output 1 would. The third agent uses NN to do exactly what the first agent did using GP. The fourth agent then fits the NN residual using GP to capture some of the variations missed in Output 3. The addition of the GP-fitted residuals to Output 3 gives Output 4. Outputs 5-8 are produced in a similar manner to that of Outputs 1-4 except that the input variables are wavelet-transformed first. Outputs 9-12 are produced in a similar manner to that of Outputs 5-8 except that the input variables are discrete-Fourier transforms (where LF is for the low-frequency and HF is for the high-frequency data to be described below).

Use of computational (or artificial intelligence) techniques as well as wavelet or Fourier data transformations to predict sunspot numbers is not new. GP was used in [22-26]. Use of NN in predicting sunspot numbers are [5,6, 21, 22 and 27-29] among many others. Wavelet transformations are in [8, 22 and 30-33]. However, using computational techniques as competing agents that ultimately produce a cooperative forecast is new. The balance of this methodology section contains a review of GP, NN, DWT, and DFT.

A. Genetic Programming

GP is a computationally intensive search technique designed to optimize a specified function. It is used here to obtain nonlinear-regression type models by minimizing their estimation mean square error (MSE). The GP employed here is a computer code written such that it evolves the best model using "survival of the fittest" Darwinian-like thought. The code is designed to evolve model specifications useful in forecasting [34]. A description of how GP is used in forecasting and its statistical properties are in [35]. The GP software used in this study is TSGP [36] written for Windows environment in C++. To execute TSGP, the user needs to provide two types of input: data input files and a configuration file. Data values of the dependent and each of the independent variables must be provided in separate files. The configuration file contains execution information including the name given to the dependent variable, the number of observations to train or fit, the number of observations to forecast, and other GP-specific parameters. TSGP produces two types of output files. One has a final model specification and the other contains actual and fitted values as well as performance statistics (R^2, historic MSE, and *ex post* prediction MSE).

The TSGP code is programmed to assemble a user defined fixed number of regression-like equations. Each equation is assembled by randomly selecting from the given explanatory variables and a set of operators. The operators used typically include $+$, $-$, $*$, $\div$, natural logarithm (Ln), exponential, sine, cosine, and square root. An equation is represented in the program by a parse tree. Figure 5 is an example of a hypothetical parse tree. The tree consists of nodes and arcs. The inner nodes take the operators while the end nodes (or terminals) take one of the explanatory variables or a constant. Constants are randomly generated numbers (between -128 and +127). Because of the many possible internal computations that take place during execution, standard protections are programmed. If in X/Z, $Z = 0$, then $X/Z = 0$. If in $X^{\frac{1}{2}}$, $X < 0$, then $X^{\frac{1}{2}} = - |X|^{\frac{1}{2}}$. If in $Ln(X)$, $X < 0$, then $Ln(X) = - Ln(|X|)$. These protections are designed to protect the computer from halting during execution.

When executed, TSGP starts by randomly assembling an initial population of (say) 1000 equations and solves each to obtain fitted and forecasted values. Fitted values (i.e., the solution using an assembled equation) are compared with historical values of the dependent variable and the residuals or errors are computed. TSGP evaluates the equations and ranks them according to their fitness (MSE). Equations with the lowest MSE (considered as members of the population with "good" genes) are kept in memory and are used to breed a new generation of equations that has the same population size. Typically the best 10% (user defined percentage) of the assembled equations are preserved. To breed a new generation, the program adds to the 10% held in memory new equations assembled using crossover and mutation. In crossover, two equations are randomly selected from the previous population and parts of those equations are swapped. This is similar to two parents having two offspring who inherit genes from each parent. If the offspring are fitter than the parents, they survive and the parents die. If not, the parents

survive since they are fitter. In mutation, part of a randomly selected equation is replaced by randomly-assembled operators and terminals. Breeding continues for a number of generations (defined by the user) before the program terminates and saves the best equation generated in the computer's memory. Figure 6 has a flow chart of the GP architecture used to obtain the fittest equation.


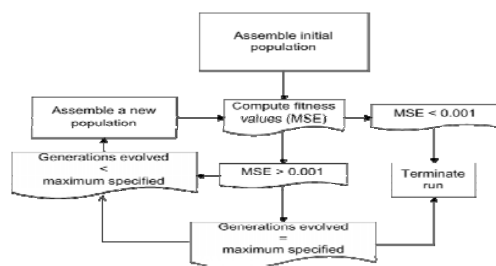Fig. 5. Example of a parse tree representation.

Fig. 6. The GP architecture to evolve regression-like equations. After assembling an initial equation, the program computes each equation's fitness (MSE). If an equation delivers an output with the specified threshold MSE (= 0.001 here), there is no reason for the search to continue. Else the search continues until a maximum number of generations (also prescribed by the user) is reached.

TSGP does not deliver a best-fit model every time it evolves an equation. It is therefore necessary to produce a large number of equations (typically 100) in a single run. Because many of the best-fit equations may or may not have predictive power, the following steps are used to find the best predicting model:

- Sort the 100 best-fit equations according to their MSE.

- Use the evolved models to predict out-of-sample values for which the outcome is already known (i.e., the *ex post* forecasts).

- Select the output among the lowest 10 (arbitrarily set) with the lowest *ex post* prediction MSE as the best model to use in forecasting the unknown future values of the dependent variable (i.e., the *ex ante* forecast).

B. Neural Networks

Neural networks (NN) are an information-processing paradigm based on the way the densely interconnected parallel structure of the human brain processes information. NN are known for their ability to detect structure in time-series. A neural network is a collection of mathematical models that emulates the nervous systems and draws on the analogies of adaptive learning.

NN have been around for more than thirty years now. Plenty of advancement took place and continues to take place as can be found in [37] among many others who describe their development and how NN are used in forecasting. NeuroSolutions [38] is the software used here in training the networks that forecast the monthly sunspot numbers.

Multilayer perceptions (MLP) and generalized feedforward (GFF) types of networks were found appropriate in training networks that would forecast the observed monthly numbers and their transformations. MLP is a layered feedforward network that learns nonlinear function mappings. It employs nonlinear activation functions. Networks are typically trained with static backpropagation and require differentiable, continuous nonlinear activation functions such as *hyperbolic tangent* or *sigmoid*. A network takes explanatory variables as input the activation function uses to produce fitted values of a dependent variable. Figure 7 has a network architecture with one-hidden layer. Although MLP trains slowly and requires a large number of observations, it seems to approximate well. GFF is a generalization of MLP with connections that jump over layers. GFF also trains with static backpropagation.


Fig. 7. Flow diagram of a NN architecture. In this flow, there are j = 1, …, k input variables, only one hidden layer with h processing elements and one output layer of predicted values.

C. Discrete Wavelet Transformations

Among the types of wavelets available to choose from, only the Haar wavelet is capable of reversing transformations of the observed data to obtain the exact original series [39]. This makes the Haar wavelet suitable for producing forecasts reconstructed from predictions of the transformed components. (More on this wavelet is in [40-42].)

A wavelet transform is a scaling function used to decompose a signal into father and mother wavelets. Father wavelets are representations of a signal's smooth or low-frequency component. Mother wavelets are representations of the details or high-frequency component in a signal. Haar wavelets are the simplest to use. They are a decimated process where half the number of observations disappears at successive levels of scaling. To generate Haar wavelet-transforms of a series Y_t, mid-point averages (s_1) and mid-point differences (d_1) of consecutive pairs of observations are obtained first. Averages preserve the main signal. Differences capture the series' detailed fluctuations. In turn, mid-point averages (s_1) are transformed into their mid-point averages (s_2) and their mid-point differences (d_2), and so on. The values of the mid-point averages and differences are known as "coefficients" in the literature on wavelets. Obtaining these coefficients is referred to as a *discrete wavelet transform* process (DWT). Alternatively, DWT maps a vector of Y_t values to a vector of wavelet coefficients w, or

$$w = \begin{pmatrix} s_J \\ d_J \\ d_{J-1} \\ \ldots \\ d_1 \end{pmatrix} \qquad (2)$$

where J is the number of scales, and

$$s_J = (s_{J,1}, s_{J,2}, \ldots, s_{J,T/2^J})'$$
$$d_J = (d_{J,1}, d_{J,2}, \ldots, d_{J,T/2^J})'$$
$$d_{J-1} = (d_{J,1,1}, d_{J-1,2}, \ldots, d_{J-1,T/2^J})' \qquad (3)$$
$$\ldots = \ldots$$
$$d_1 = (d_{1,1}, d_{1,2}, \ldots, d_{1,T/2^J})'.$$

To transform a series to a maximum of J components, that series must be of length $T = 2^J$. The $s_{J,t} = (s_{J-1,2t} + s_{J-1,2t-1})/2$ are averages and the $d_{J,t} = (s_{J-1,2t} − s_{J-1,2t-1})/2$ are the differences.

Given a series' wavelet transformed coefficients s_J and d_J, ..., d_1 in (3) above, and if a series has T = 512, setting J = 4 produces five series: s_4 and d_4 with 32 coefficients each, as well as d_3, d_2, and d_1 with 64, 128, and 256 coefficients in each, respectively. Figure 10 shows the DWT of 512 unsmoothed monthly sunspot numbers with scaling level J = 4 as well as their *inverse discrete wavelet transform* (IDWT). IDWT, which was constructed using s_4, d_4, d_3, d_2, and d_1, is identical to the original series before transformation.

In the figure, the transformed sequences (s_4, d_4, d_3, d_2, and d_1) portray different levels of complexity. Series s_4 represents the clearest or most predictable component (given that it is the smoothed point taken every 16 or 2^4 months). Series d_4 captures what may be viewed as systematic changes in the signal, and therefore may be fairly predictable as well. Series d_3 is more complex because it captures changes over a shorter period of time (2^3 or 8 months), and so on. Noise is captured mostly in series d_1. The search for models for each of the five data sets (s_4, d_4, d_3, d_2, and d_1) is possible if the number of observations in s_J is sufficiently large since the others (d_4, d_3, d_2, and d_1) must also contain a sufficient number of observations to fit models. Setting J = 4 is appropriate because if J = 5, s_5 and d_5 will have very few values – only 16 – to construct a model with.

Given that the five DWT series are independent, each should be modeled separately. The model specification assumed for s_4 is:

$$s_{4,t} = f(s_{4,t-5}, s_{4,t-6}, s_{4,t-7}, s_{4,t-8}). \qquad (4)$$

In (4), s_4 is assumed to be a function of four distant but consecutive lagged values. If s_4 has 32 observations, only 24 observations are used for training or fitting models after accounting for lags. With lags = (t-5) to (t-8), the estimated equation for this specification can produce a forecast for four periods without "predicted" values used as input. Similarly,

$$d_{4,t} = f(d_{4,t-5}, \ldots, d_{4,t-8}); \qquad (5)$$
$$d_{3,t} = f(d_{3,t-9}, \ldots, d_{3,t-16}); \qquad (6)$$
$$d_{2,t} = f(d_{2,t-17}, \ldots, d_{2,t-28}); \qquad (7)$$
$$d_{1,t} = f(d_{1,t-33}, \ldots, d_{1,t-56}). \qquad (8)$$

The number of observations used to fit and forecast d_4 is also 24, it doubles for d_3, doubles again for d_2, and then doubles again for d_1. The five resulting models are used to compute fitted and forecast values of the five series.


Fig. 8. The average monthly sunspot number Haar wavelet transformations and their inverse.

A forecast is obtained by shifting the starting point forward before inversion. In shifting, a number of older values are deleted and replaced by an equal number of forecast values added. Using this method, equations (4) and (5) provide four-step-ahead forecasts, equation (6) provides an eight-step-ahead forecast, equation (7) provides a 16-step-ahead forecast, and (8) provides a 32-step-ahead forecast. The final inverting step thus produces a 64-step-ahead forecast. The resulting inverted data set still contains 512 observations with the last 64 being forecast values.

D. Discrete Fourier Transformations

A DFT acts as a digital filter that divides the observed series into two components of low-frequency and high-frequency. The data observed values are signal in the time domain while DFT is the frequency domain. The process is a digital filter because some of the variations are filtered out to obtain the bigger picture. The low-frequency component is obtained first. Most of the amplitude is confined to the very smooth low-frequency component. The high-frequency component is the residual. It has nearly zero average and small amplitude but contains all the rapid variations

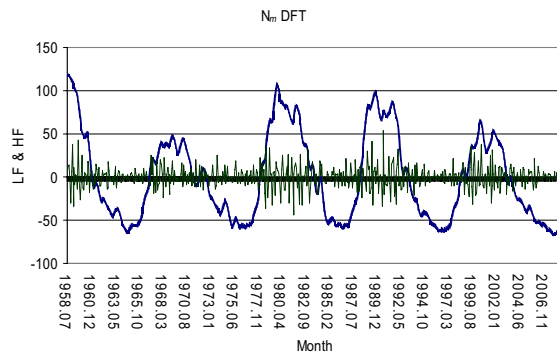Fig. 9. These are frequency transformations of the monthly numbers. The main signal is captured by the low-frequencies (with a clear cyclical pattern) while the high-frequency resembles what appears to be noise fluctuating closer to the zero line.

.Figure 9 shows the low- and high-frequency transformations of the monthly sunspot numbers. Full description of the method and software used to obtain the low-

and high-frequency series from the observed data are in [43]. GP and NN are then used to search for the best model that would fit the low- and high-frequency transformations obtained and forecast their future values. Fitted and forecast values of the observed series is a straight forward process once the estimates and forecasts from the best models of the transformed series are obtained. Inversion is completed by a simple addition of the low- and high-frequency fitted transformations.

III. ESTIMATION RESULTS

Only the final estimation results of the models obtained and their statistics are presented here. Their forecast results are presented in the next section. The input variables were identical when using GP and NN. Parameters to evolve GP models were set as follows: Population size = 1000, number of generations = 120, mutation rate = 60%, crossover rate = 20%, reproduction or cross self rate = 10%, selection method = tournament, maximum age = 2000, and the number of best-fit equations to evolve = 100. A trial-and-error NN search routine to find the best fit is implemented. In each run, one of parameters (the number of hidden layers, the momentum learning rule, or the maximum number of epochs) is varied. The number of hidden layers started at one and is increased by a second to evaluate if the increase made any difference. The momentum learning rule starts at 0.9 and is reduced by 0.1 using one and two hidden layers. With each combination, the maximum epochs starts at 1000 and are increased by 500 with each combination until the best output is obtained. The best results were found mostly using a single hidden layer with five input processing elements, a *TanhAxon* transfer function, used momentum learning rule (with step size = 1.0 and momentum = 0.7), and a number of epochs ranging between 1000 and 5000. For the most, a multilayered perceptron system with a layered feedforward network trained with static backpropagation was used.

Table 1 contains the estimation characteristics of the estimation results. In the first column, the agents participating in producing the forecasts are listed. The label 'GP' (or 'NN') alone indicates that only a GP (or NN) algorithm was used to obtain fitted values of the observed monthly sunspot numbers. 'GP-NN' indicates that the output combines the monthly series' fitted values using GP and GP's residuals fitted values using NN (i.e., cooperation between two agents). 'W' stands for wavelet transformations and 'F' stands for Fourier transformations. 'Mix' represents the final GP or NN equation that reconciliates between the above twelve outputs. The second column has the mean error of fitted values. Standard error and t-statistics follow. These statistics help test whether the mean error is significantly different from zero at the 5% level of significance or not. Only four of the twelve outputs have a mean error that is significantly different from zero (with t-stat > 1.96). Interestingly, the final output ('Mix') has a mean error that is not significantly different from zero at the 5% level of significance. MSE (in the fifth column) provides a comparison of how the different agents succeeded in fitting historical data. The results in the table suggest that the best

forecast should be obtained using GP to fit the wavelet-transformed data added to the NN-estimated residuals (where the residuals were obtained after inverting the GP-wavelet fitted values).

TABLE I ESTIMATION RESULTS OF BEST OBTAINED FITS

	N_m Outputs				
	Mean Error	Std. Err.	t-stat	MSE	R^2
GP	8.92	1.80	4.97	1139.98	0.61
GP-NN	1.97	1.09	1.81	394.67	0.87
NN	1.49	1.22	1.23	488.58	0.83
NN-GP	2.13	1.12	1.90	420.24	0.86
GP-W	2.95	1.10	2.67	407.60	0.86
GP-W-NN	0.06	0.55	0.11	101.03	0.97
NN-W	0.88	1.05	0.84	362.75	0.88
NN-W-GP	-0.30	1.08	-0.28	386.76	0.87
GP-F	4.50	1.67	2.69	938.70	0.68
GP-F-NN	-1.74	1.33	-1.31	585.49	0.80
NN-F	3.75	1.86	2.02	1147.67	0.61
NN-F-GP	1.77	1.28	1.38	542.66	0.82
Mix	-0.12	0.98	-0.12	317.88	0.89

IV. PREDICTION RESULTS

Statistics on prediction results are summarized in Table 2. The table reports the MSE and Theil's U statistic. Theil's U is a measure of forecast performance and is defined as:

$$U = \frac{\sqrt{\frac{1}{F}\sum_{f=T+1}^{T+F}(Y_f - \hat{Y}_f)^2}}{\sqrt{\frac{1}{F}\sum_{f=T+1}^{T+F}\hat{Y}_f^2} + \sqrt{\frac{1}{F}\sum_{f=T+1}^{T+F}Y_f^2}} \quad (9)$$

where f = 1, 2, …, F, F = number of observations forecasted *ex post* (or with known outcomes), and $\hat{Y}_f$ are forecast values of Y_f. This statistic will always fall between zero and one where zero indicates a perfect fit [44, p. 387].

The results in the table suggest that GP-W has the lowest prediction MSE and lowest U-statistic and therefore produced the best *ex post* forecast among the competing twelve agents. However, 'Mix' produces the best forecast. Figure 10 shows the monthly fits of the last three cycle peaks and the forecast of the next solar peak of the unsmoothed monthly sunspot numbers 'Mix' delivers. Their annualized predictions (2000-2013 only) are in Figure 11.


Fig. 10. Fitted and forecasted values of the unsmoothed monthly sunspot numbers.

TABLE II THE *EX POST* N_M PREDICTION RESULTS

	MSE	Theil's U
GP	112.58	0.168
GP-NN	68.18	0.135
NN	165.75	0.189
NN-GP	125.19	0.199
GP-W	57.74	0.122
GP-W-NN	204.76	0.21
NN-W	210.82	0.218
NN-W-GP	171.21	0.198
GP-F	394.472	0.292
GP-F-NN	312.298	0.267
NN-F	591.888	0.297
NN-F-GP	153.828	0.195
Mix	29.07	0.088

TABLE IV COMPARISON OF FORECASTS OF CYCLE 24

Author	Date	Predicted	+/-	Expected
Clilverd *et al.*	2006	42	35	?
Javariah	2007	74	10	?
Baranovski	2006	80	21	2012
Schatten	2005	80	30	2012
Pesnell	2006	101	20	2012
Hirmath	2007	110	11	2012
Pesnell	2006	115	40	2011
Nevanlinna	2007	124	30	2010
Tritakis *et al.*	2006	133		2009
Kennewell & Patterson	2006	134	50	2011
Gholipour *et al.*	2005	145		2011/12
Hathaway & Wilson	2006	160	25	?
Mean		108		
Median	2006	113	28	2012

Table 3 presents the unsmoothed monthly number forecasts for 2011 and 2012. The unsmoothed monthly as well as the annual averages forecasts suggest that the next peak is 156 and it is expected to occur around the month of October in 2011. Taken annually, 2011 is expected to be year when the peak of cycle 24 occurs.

The forecast in Table 3 is significantly different from more than 45 forecasts of sunspot numbers already available. Table 4 contains a small sample of available forecasts. Information in the table was obtained from a summary provided in [15]. The selection covers the most recent forecasts only. Although not a comprehensive comparison by any measure, the table shows that most believe that the peak will occur in 2012 and at a level substantially lower than expectations forecasted by the best model here.
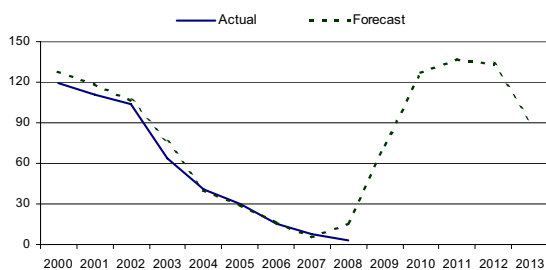

Fig. 11. Fitted and forecasted values of the annual average of the unsmoothed monthly sunspot numbers.

TABLE III PREDICTED 2011 AND 2012 MONTHLY NUMBERS

Month	N_m	Month	N_m
2011.01	142	2012.01	136
2011.02	128	2012.02	142
2011.03	135	2012.03	134
2011.04	126	2012.04	133
2011.05	132	2012.05	126
2011.06	130	2012.06	137
2011.07	132	2012.07	137
2011.08	130	2012.08	141
2011.09	131	2012.09	136
2011.10	156	2012.1	126
2011.11	149	2012.11	106
2011.12	139	2012.12	108
Mean	136	Mean	130

V. CONCLUSION

A new system that employs multiple modeling agents and utilizes two computational modeling techniques was introduced to forecast the unsmoothed monthly sunspot numbers. The forecasts are of monthly data to help capture the dynamics of solar activity hidden when forecasting annual data. The models introduced in this paper were designed with the assumption that a repeated pattern of magnetic activities exists. Twelve modeling agents produced competing forecasts of the unsmoothed monthly sunspot numbers first. Four of the twelve series used the observed data. The other eight used transformations of the observed series. Transformations were obtained in two ways. First, the monthly series was transformed using the discrete Haar wavelet transform (DWT). It produced five series to predict. Using a history of 512 observations, the five series obtained had 256, 128, 64, 32, and 32 observations. The two computational modeling techniques – genetic programming and neural networks – then produced two competing fits and forecasts of the five series. Employing the Haar wavelet made it possible to inverse the process and to obtain the fitted and forecasted values of 512 observations. By shifting the starting point forward by 64 observations, it was possible to produce a 64-month ahead forecast. The series were then decomposed using a discrete Fourier transformation. Here only two series were produced: one captured what may be considered the signal while the other (the residuals) represented noise. The resulting series were then fitted and forecasted using the two computational techniques. The twelve competing forecasts produced were then used as input and fitted using genetic programming and NN. GP produced the final best forecast.

The final best forecast obtained suggested that the next annual peak of sunspot numbers will be at about 136 and will occur in 2011. The monthly peak is predicted to be at 156 and is expected to occur in October 2011. On average, a few recent predictions identified in the literature forecast a lower peak to occur in 2012.

REFERENCES

[1] K. Schatten, P. Scherrer, L. Svalgaard, and J. Wilcox, "Using the dynamo theory to predict the sunspot number during solar cycle 21," *Geophysics Research Letters*, vol. 5, 1978, pp. 411-414.

[2] M. Gabr, and S. Rao, "The estimation and prediction of subset bilinear time series with applications," *Journal of Time Series Analysis*, vol. 2, 1981, pp. 155-171.

[3] H. Tong, *Nonlinear Time Series Analysis: A Dynamical System Approach*. Oxford University Press: Oxford, 1990.

[4] H. Tong, S. Lim, "Threshold autoregressive, limit cycles and cyclical data," *Journal of the Royal Statistical Society B*, vol. 42, 1990, pp. 245-292.

[5] A. Weigend, B. Huberman, and D. Rumelhart, "Predicting the future: a connectionist approach," *International Journal of Neural Systems*, vol. 1, 1990, pp. 193-209.

[6] A. Dmitriev, Y. Minaeva, Y. Orlov, M. Riazantseva, and I. Veselovsky, "Solar activity forecasting on 1999-2000 by means of artificial neural networks," Reported on EGS XXIV General Assembly, 22 April 1999, The Hague, The Netherlands, http://dbserv.sinp.msu.ru/ RSWI/sun.pdf.

[7] L. Svalgaard, E. Cliver, and Y. Kamide, "Sunspot cycle 24: Smallest cycle in 100 years?" *Geophysical Research Letters*, vol. 32, L01104, 2005, pp. 1-4.

[8] H. Lundstedt, L. Liszka, R. Lundin, and R. Muscheler, "Long-term solar activity explored with wavelet methods," *Annales Geophysicae*, vol. 24, 2006, pp. 769-778.

[9] R. Wilson, D. Hathaway, and E. Reichmann, "On the correlation between maximum amplitude and smoothed monthly mean sunspot number during the rise of the cycle (from t=0-48 months past sunspot minimum)," *National Aeronautics and Space Administration, Technical Publications NASA/TP-1998-208591*.

[10] S. Sell, "Time series forecasting: A multivariate stochastic approach," 1999. Available: http://arxiv.org/PS_cache/physics/pdf/9901/9901050 v2.pdf.

[11] D. Hathaway, R. Wilson, and E. Reichmann, "A synthesis of solar cycle prediction techniques," *Journal of Geophysical Research*, 104, 1999, pp. 22375-22388.

[12] D. Hathaway, R. Wilson, and E. Reichmann, "Group sunspot numbers: Sunspot cycle characteristics," *Solar Physics*, vol. 151, 2002, p. 177.

[13] R. Wilson, and D. Hathaway, "Gauging the nearness and size of cycle maximum," *National Aeronautics and Space Administration, Technical Publications NASA/TP-2003-212927*.

[14] NOAA. Available: ftp://ftp.ngdc.noaa.gov/stp/solar_data/sunspot_ numbers/smoothed.

[15] W. Pesnell, "Predictions of solar cycle 24," NASA, Goddard Space Flight Center, Greenbelt, Maryland. http://www.swpc.noaa.gov/ SolarCycle/SC24/May_24_2007_table.pdf.

[16] D. Nordemann, "Sunspot number time series: Exponential fitting and solar behavior," *Solar Physics*, vol. 141, 1992, pp. 199-202.

[17] T. Lin, and M. Pourahmadi, "Nonparametric and non-linear models and data mining in time series: A case-study on the Canadian lynx data," *Applied Statistics*, vol. 47, Part 2, 1998, pp. 187-201.

[18] T. Xu, J. Wu, Z. Wu, and Q. Li, "Long-term sunspot number prediction based on EMD analysis and AR model" *Chinese Journal of Astronomy and Astrophysics*, vol. 8, 2008, pp.337-342.

[19] National Geographic Data Center. Available: ftp://ftp.ngdc.noaa.gov/ stp/solar_data/ solar_radio/sunspot_numbers/monthly.

[20] A. Aussem, and F. Murtagh, "Combining neural network forecasts on wavelet-transformed time series," *Connection Science*, vol. 9, 1997, pp. 113-121.

[21] M. Kaboudan, "Extended daily exchange rates forecasts using wavelet temporal resolutions," *New Mathematics and Natural Computing*, vol. 1, 2005, pp. 79-107.

[22] M. Kaboudan, "Computational forecasting of wavelet-converted monthly sunspot numbers," *Journal of Applied Statistics*, vol. 33, 2006, pp. 925-941.

[23] M. Kaboudan, "Forecasting solar cycles with GP," in *Proceedings of Neural, Parallel & Scientific Computations, 2002, ISDA*, A. Abraham, N. Baikunth, M. Sambandham, and P. Saratchandran, Atlanta GA.

[24] A. Orfila, J. Ballester, R. Oliver, A. Alvarez, and J. Tintoré, "Forecasting the solar cycle with genetic algorithms," *Astronomy and Astrophysics*, vol. 386, 2002, pp. 313-318.

[25] H. Kwasnicka, and E. Szpunar-Huk, "Genetic programming in data modeling," in *Genetic Systems Programming: Theory and Experience*, N. Nedjah, and A. Abraham, Springer Berlin / Heidelberg, 2006, pp. 105-130.

[26] R. Jagielski, "Genetic programming prediction of solar activity, intelligent data engineering and automated learning — IDEAL 2000," *Data Mining, Financial Engineering, and Intelligent Agents, Lecture Notes in Computer Science*, 2008, pp. 191-210.

[27] S. Eklund, "Time series forecasting using massively parallel genetic programming," *Proceedings, Parallel and Distributed Processing Symposium 2003*, pp. 22-26.

[28] X. Li, C. Cheng, W. Wang, and F. Yang, "A study on sunspot number time series prediction using quantum neural networks," *Proceedings of the 2008 Second International Conference on Genetic and Evolutionary Computing*, 2008, pp. 480-483.

[29] J. Kyngas, "Forecasting sunspot numbers with neural networks," *Report A-1995-1, University of Joensuu, Department of Computer Science, Report Series A. Available:* ftp://ftp.cs.joensuu.fi/pub /Reports/A-1995-1.ps.

[30] P. Frick, D. Galyagin, D. Hoyt, E. Nesme-Ribes, K. Schatten, D. Sokoloff, and V. Zakharov, "Wavelet analysis of solar activity recorded by sunspot groups," *Astronomy and Astrophysics*, vol. 328, 1997, pp. 670-681.

[31] P. Kumar, and E. Foufoula-Georgiou, "Wavelet analysis for geophysical applications," *Reviews of Geophysics*, vol. 35, 1997, pp. 385-412.

[32] F. Boberg, H. Lundstedt, J. Hoeksema, P. Scherrer, and W. Lui, "Solar mean magnetic field variability: A wavelet approach to WSO and SOHO/MDI observations," *Journal of Geophysical Research*, vol. 107, 2002, pp. 15-1 – 15-7.

[33] H. Lundstedt, L. Liszka, and R. Lundin, "Solar Activity explored with new wavelet methods," *Annales Geophysicae*, vol. 23, 2005, pp. 1505-1511.

[34] J. Koza, *Genetic Programming*. Cambridge, MA: The MIT Press, 1992.

[35] M. Kaboudan, "Statistical properties of fitted residuals from genetically evolved models," *Journal of Economic Dynamics and Control*, vol. 25, 2001, pp. 1719-1749.

[36] M. Kaboudan, "TSGP: A Time Series Genetic Programming Software." Available: http://Bulldog2.Redlands. edu/fac/mak_kaboudan.

[37] J. Principe, N. Euliano, and C. Lefebvre, *Neural and Adaptive Systems: Fundamentals Through Simulations*. New York: John Wiley & Sons, 2003.

[38] NeuroSolutions[TM]. *The Neural Network Simulation Environment.* Version 4, NeuroDimensions, Inc.: Gainesville, FL, 2002.

[39] S. Mallat, "A theory for multiresolution signal decomposition: The wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, 1989, pp. 674-1989.

[40] A. Bruce, H. Gao, *Applied Wavelet Analysis with S-Plus*. New York: Springer, 1996.

[41] J. Walker, *A Primer on Wavelets and Their Scientific Applications.* Boca Raton: Chapman & Hall/CR, 1999.

[42] R. Gençay, F. Selçuk, and B. Whitcher, *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. San Diego, CA: Academic Press, 2002.

[43] T. Masters, *Neural, Novel and Hybrid Algorithms for Time Series Prediction*. New York: John Wiley & Sons, 1995.

[44] R. Pindyck R, and D. Rubinfeld, *Econometric Models and Economic Forecasting.* Boston: Irwin McGraw-Hill, 1998.