015-0603

A NEURAL NETWORK MODEL FOR FORECASTING PRODUCTION TIME SERIES IN BRAZILIAN INDUSTRIES

FÁBIO AUGUSTO MOLLIK ZOUCAS

FEI University Center

Av. Humberto de Alencar Castelo Branco, 3972, Assunção, São Bernardo do Campo - SP, Brasil, Zip code: 09850-901

e-mail: fabiozoucas@hotmail.com

PATRÍCIA BELFIORE

FEI University Center

Av. Humberto de Alencar Castelo Branco, 3972, Assunção, São Bernardo do Campo - SP, Brasil, Zip code: 09850-901

e-mail: prspbelfiore@fei.edu.br

JOÃO CHANG JUNIOR

FEI University Center

Av. Humberto de Alencar Castelo Branco, 3972, Assunção, São Bernardo do Campo - SP, Brasil - Zip code: 09850-901

e-mail: chang.joao@gmail.com

POMS 21st Annual Conference

Vancouver, Canada

May 7 to May 10, 2010

Abstract

This paper aims to propose a neural network model for forecasting the production time series of eleven different industries in Brazil. These data were collected from IBGE (Brazilian Institute of Geography and Statistics). Firstly, we study different networks topologies that have been implemented in the literature in recent years, such as perceptron, linear networks, multi-layer perceptron (MLP), probabilistic network, Hopfield model, Kohonen model, time delay neural network (TDNN), Elman and Jordan Network, in addition to the backpropagation and Levenberg-Marquadt algorithms. Studying the behavior of these time series and the main characteristics of the each network topology, we conclude that the TDNN with multi-layer perceptron is the best to estimate the production time series of eleven industrial segments. The neural network was then applied considering two different strategies of structural model. We conclude that the neural network model proposed was effective for forecasting production time series in these industries.

1. Introduction

Sales forecasting plays an important role in business strategy. An effective sales forecasting can help the companies to calculate the production costs and to determine the sales price of the product (Kuo and Xue, 1998).

For Slack (1997), good forecasting systems are essential for capacity planning, but it is also crucial to comprehend the uncertainty of demand, because it allows to judge the risks of the service level.

According to Werner and Ribeiro (2003), the study of time series forecasting is an important activity, because it reveals market trends, contributes to strategic planning, helps to solve more immediate problems, allows a better understanding of the behavior of individual customers and constitutes additional source of information in supporting decisions about investments and team size.

For Stevenson (2001), the best forecasting system is not necessarily the most accurate, and not the cheapest one. Actually, it reflects the best combination of accuracy and cost.

Characteristics of the forecasting system, especially in industrial sectors, have a flexible behavior and are dependent of numerous variables, and often inherent to the process. By providing a set of points that can’t be relevant for conventional techniques, the neural network models used has been commonly used to predict time serious.

Artificial Neural Networks (ANN) are flexible structures that can be applied to a wide range of forecasting problems with a high degree of accuracy. However, the neural networks need a large amount of historical data to reach the highest level of accuracy of the results and have as the main advantage the ability of modeling nonlinear systems (Khashei et al., 2008).

The neural networks can also be defined as computer systems containing many simple nonlinear units or nodes interconnected by links. The scope of this work is to verify the effective application of this tool to production forecasting of the industrial sector. The construction of networks and existing models for different activities will be detail in this work.

Problem solving using Artificial Neural Networks (ANN) is quite attractive, whereas the way they are represented internally by the network and the natural parallelism inherent the ANN architecture creates the possibility of increasing the performance over conventional models (Braga et al., 2007).

According to the same author, the ability of a neural network of learning through a limited set of examples and then give coherent answers to unknown data is a demonstration that the ability of Neural Networks goes beyond mapping input and output relations.

2. Methodology

The purpose of this study is to apply a neural network model for forecasting production time serious of the industrial segment, where the variables are strongly influenced by external factors and not scheduled, in different areas of activities

For this analysis we collect real time series of eleven industry sectors, extracted from the data of IBGE (Brazilian Institute of Geography and Statistics), presenting different characteristics, cycles and behaviors, which are: Pharmaceuticals, Mineral Extraction, Manufacturing Industry, General Industry, Textile Industry, Capital Goods, Intermediate Goods, Consumer Goods, Durable Goods, Non-Durable Goods and Beverage Industry.

The initial data collected for each sector from the period of January 1975 to September 2008 correspond to 405 points. Firstly, we analyzed the data behavior of each sector and then different topologies of neural network were tested in order to find the best one.

Through the auto-correlation function of each segment, we determine the best network structure compatible with a generic model presented, and then implementing an algorithm for commercial use (consistent with the data used), we propose a neural network model to production forecasting using the Matlab software.

Using measures of forecast error in time series analysis, we will justify the effectiveness of the model proposed.

3. Neural Network Approach

According to Braga et. al. (2007), Artificial Neural Networks are distributed parallel systems consisting of simple processing units (artificial neurons) which calculate certain mathematical functions (usually nonlinear) and are arranged in one or more layers connected by a large number of connections, usually one-way and with two processing phases, including learning and use.

In a neural network, data are connected in the first layer and the processed results are declared on the last node of output layer, being discovered by the modification of the weights assigned to it.

According to Wray et. al. (1994), the analysis of neural networks has three primary advantages. The first one affirms that the development of the network does not require knowledge of their fundamental relationships between the input and output variables. Secondly, its associative ability makes the neural network more robust and able to work with situations of data loss or uncertainty. The third advantage points that the performance of neural networks is not affected by multicollinearity problems.

Neural network models can be distinguished by different aspects as: their architecture, propagation rule, learning rule, types of algorithms, number of layers and their activation function.

The combination of architecture and neurodynamics defines the paradigm of neural network, and theoretically a network with one hidden layer and enough hidden neurons satisfies any continuous function (Kaastra and Boyd, 1996).

The fundamental classification of neural structures considers the propagation rule of the information received where the distinction is related to the feedback topology (Feedforward and Recurrent) as presented in figure 1.

Figure 1 – Feedback topology of neural networks

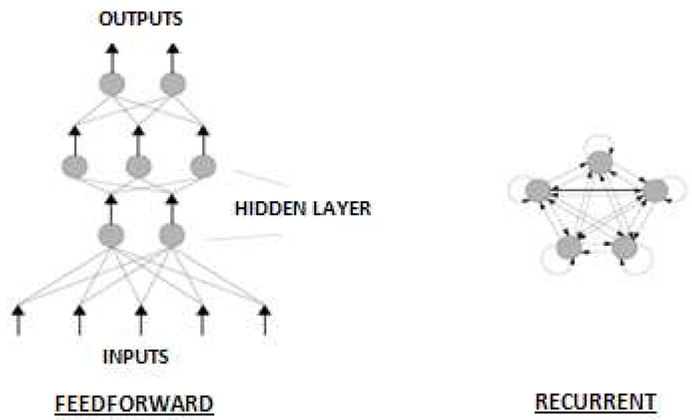

An important property of neural networks is the ability of learning and improving their performance by adjusting their parameters, called neural network training. The training can be classified as supervised or unsupervised.

The learning algorithms are a set of well-defined rules for solving a neural network problem. All knowledge is stored in the synapses, i.e. in the weights assigned to the connections between neurons.

The activation function of a neural network is responsible for generating the resulting output from the application of the weight vectors in the network input. Figure 2 shows the most common examples of activation functions.

Figure 2: Examples of activation functions

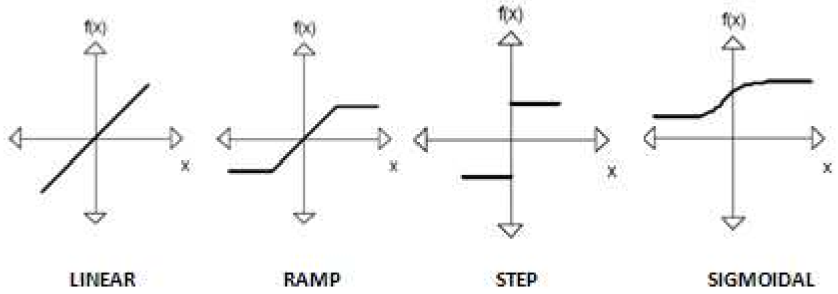

One of the most important tools for training neural networks to perform some task is the use of an algorithm, in order to reduce the error between the actual and desired output. There are two types of algorithm commonly used for time serious forecasting which include backpropagation and Levenberg-Marquadt.

Several topologies of neural networks models were developed from the original model proposed by MacCulloch and Pitts (1943), such as Perceptron, Adaline (Adaptive Linear Network), feedforward network as Multilayer Perceptron (MLP), Bi-directional Associative Memory (BAM), Adaptive Resonance Theory (ART), Hopfield model, Kohona model, and Time Delay Neural Network (TDNN). In this work we use a Time Delay Neural Network with Multilayer Percepton.

4. Model proposed

The historical data of each of the eleven segments studied correspond to a period of 405 months. To represent this model we used basically two types of neural networks: Time Delay Neural Network (TDNN) with Multilayer Percepton and Recurrent Neural Network (RNN). The software used to design, implement and simulate the neural network model proposed was Matlab.

According to Braga et al. (2007), the performance comparing TDNN and RNN is very similar. RNN tends to be a little better, but it is more difficult to implement and train. One strategy is to start using the TDNN, and only if its performance is not acceptable, the RNN must be applied.

So, we choose to use the TDNN with multi-layer perceptron or feedforward, whose inputs are samples of short time series represented by delay line, and the output is the forecast for the next month as described in Figure 3.

Figure 3 – Topology of a TDNN with Multilayer Perceptron

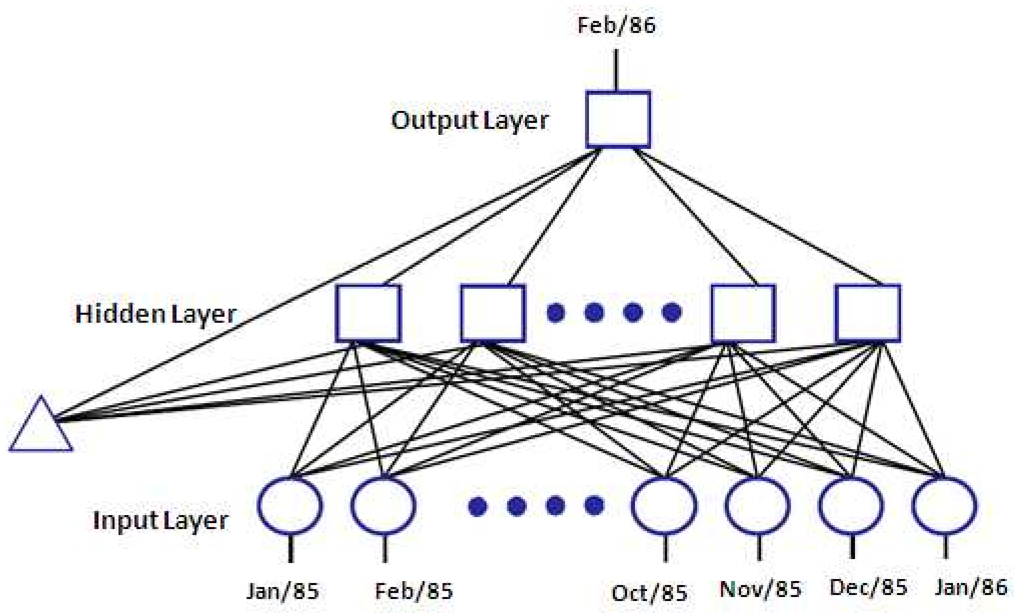

The topology of the TDNN model consists of 13 inputs that covers a 12-month period of seasonality and compares the last month of the previous period to the present period. The hidden layer consists of *n* neurons estimated from a residue analysis which are connected to a single output layer. All neurons have an input connected to a constant value called “bia”. The activation function of neurons in the hidden layer is sigmoid and in the output neuron is linear in order to generate continuous values.

The amount of thirteen inputs was obtained empirically. We realize that using less inputs, the network loses efficiency in the ability of modeling the variations in data. With more input, it becomes more difficult the training process, because results in a local minimum solution.

The model used presents a short-term memory that stores relevant past events, and an associator that uses the memory contents to classify or predict new data. A model with these two characteristics, according Makridakis et al. (1998), is suitable for processing patterns that vary over time.

5. Application and results analysis

5.1 Preliminary data analysis

Based on the calculation of the estimated auto-correlation function for time series of all eleven segments, it was found that all series studied are non-stationary (red line in figure 4). According to Pindyck and Rubinfeld (1991), a stationary time series has an autocorrelation function similar to a white noise signal (blue line in figure 4) where the correlation coefficient decreases rapidly over *k* periods (represented by months in figure 4).

As shown in figure 4, the autocorrelation functions of all segments have a different way, falling and fluctuating slowly over period *k*, as shown by red line in figure 4. In addition, most of segments present peaks every 12 months, proving the existence of seasonality annually.

Figure 4 – Autocorrelation function of all segments

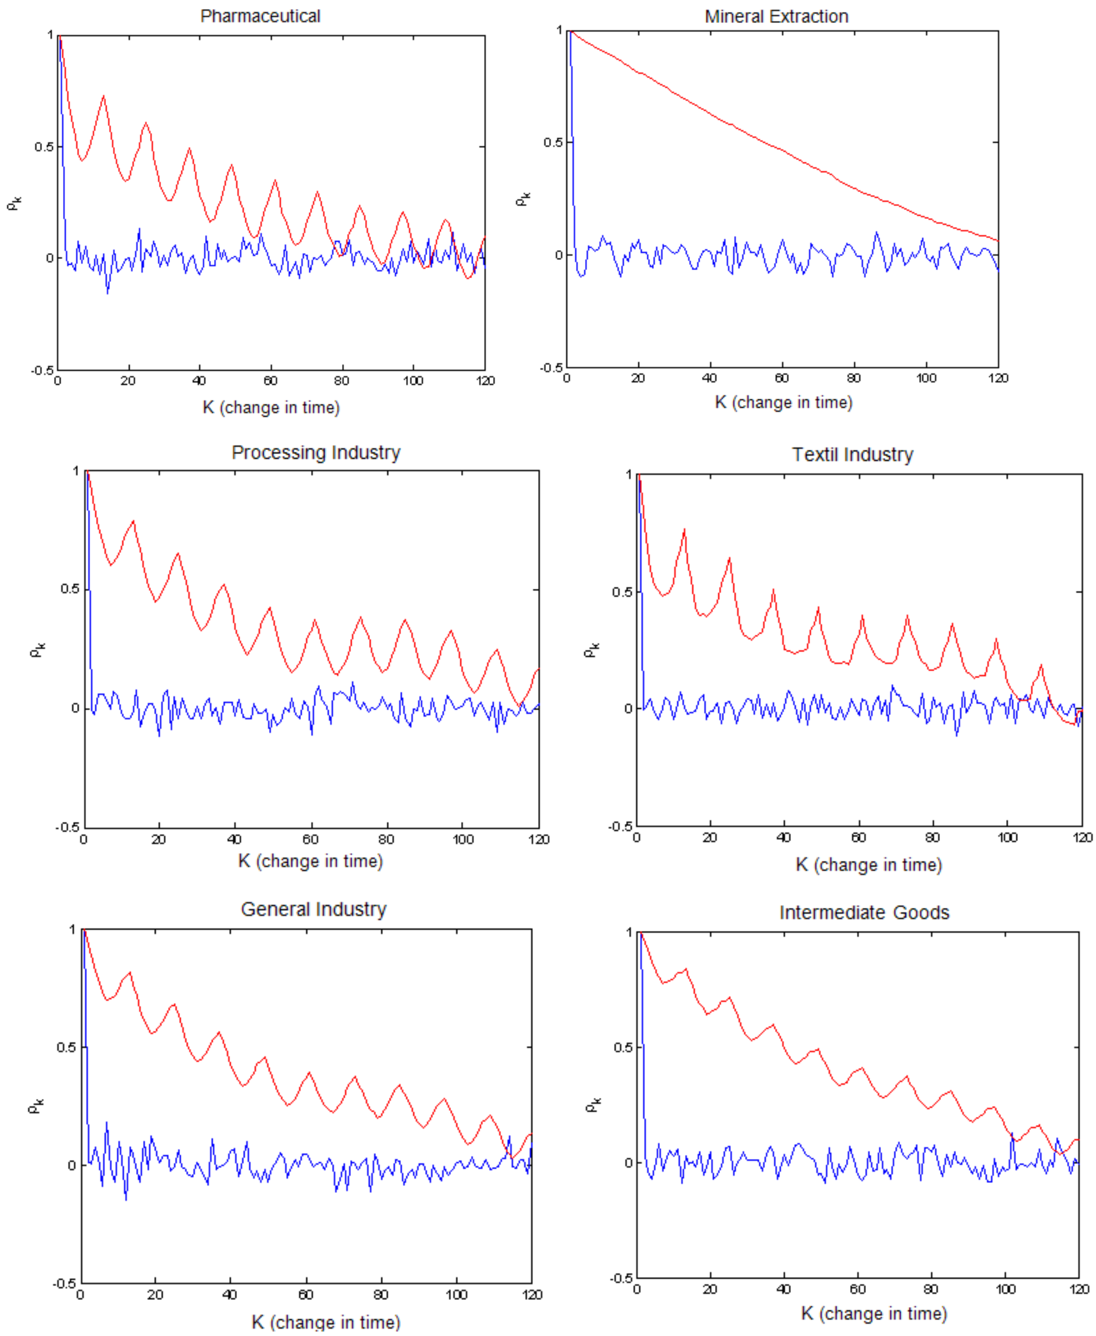

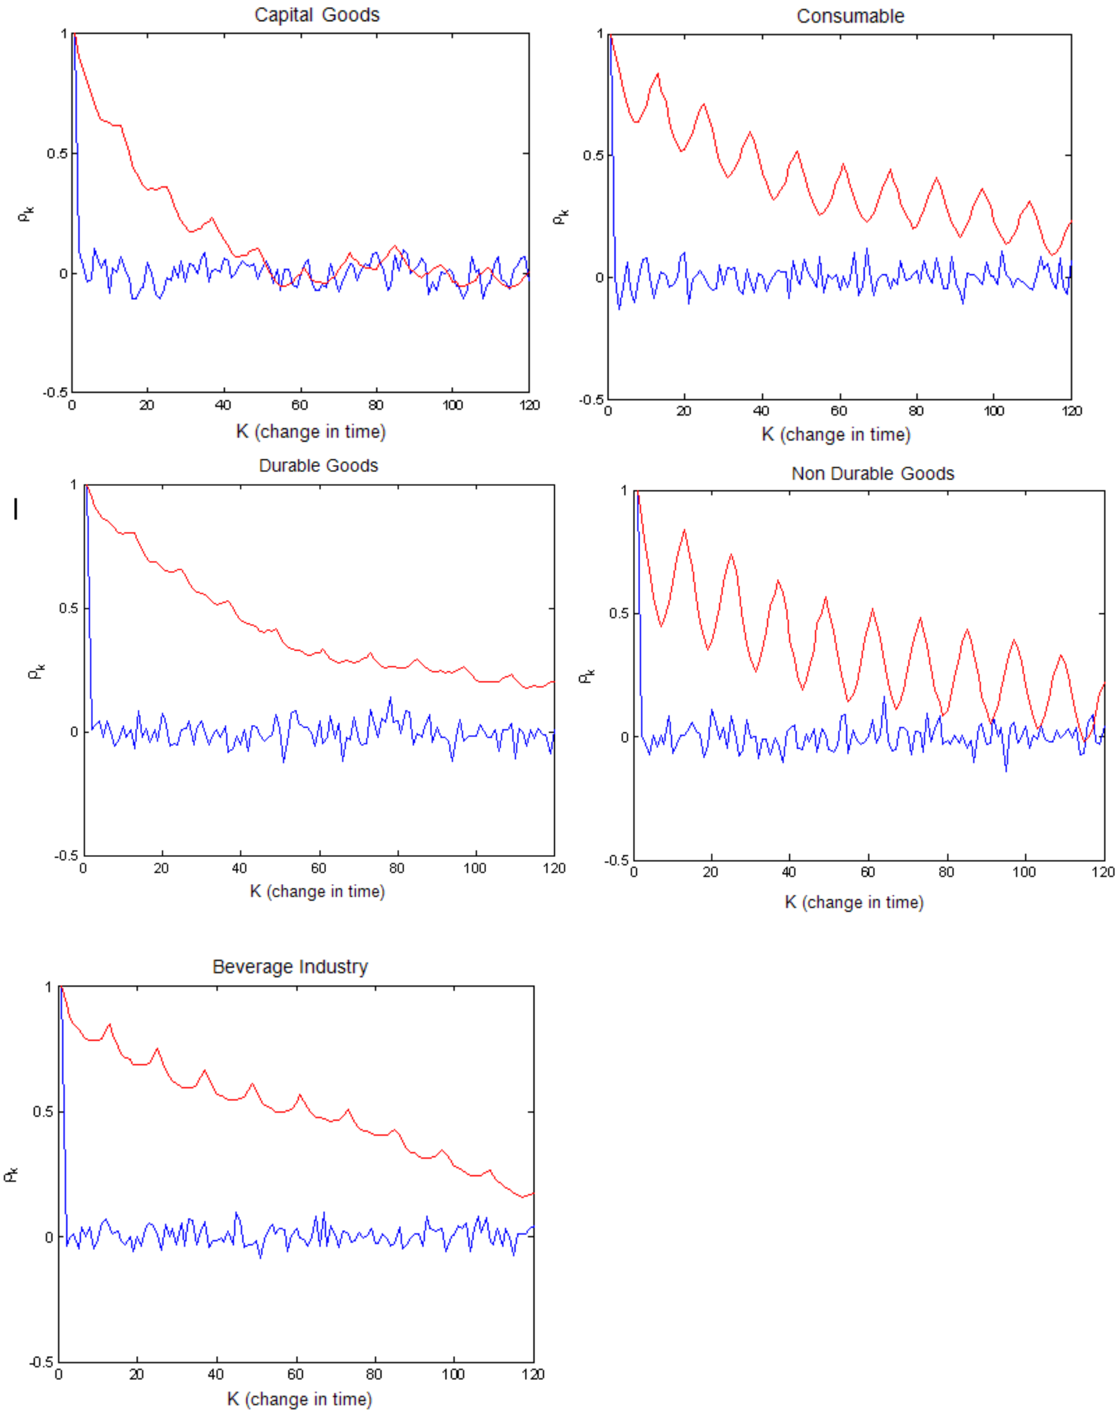

Legend: red line (non-stacionary time series)
blue line (stacionary time series)

Although the TDNN is able to work with non-stationary series and even with seasonality, we adjust the data calculating the variation of time series values in intervals of 12 months in order to eliminate seasonality and trends of increase and decrease, providing comparison between standardized time series as shown in Figure 5. Blue line is the autocorrelation function of a white noise signal and red line is the autocorrelation function of time series changed.

Figure 5 – Autocorrelation function of each segment calculated as the difference between values in intervals of 12 months

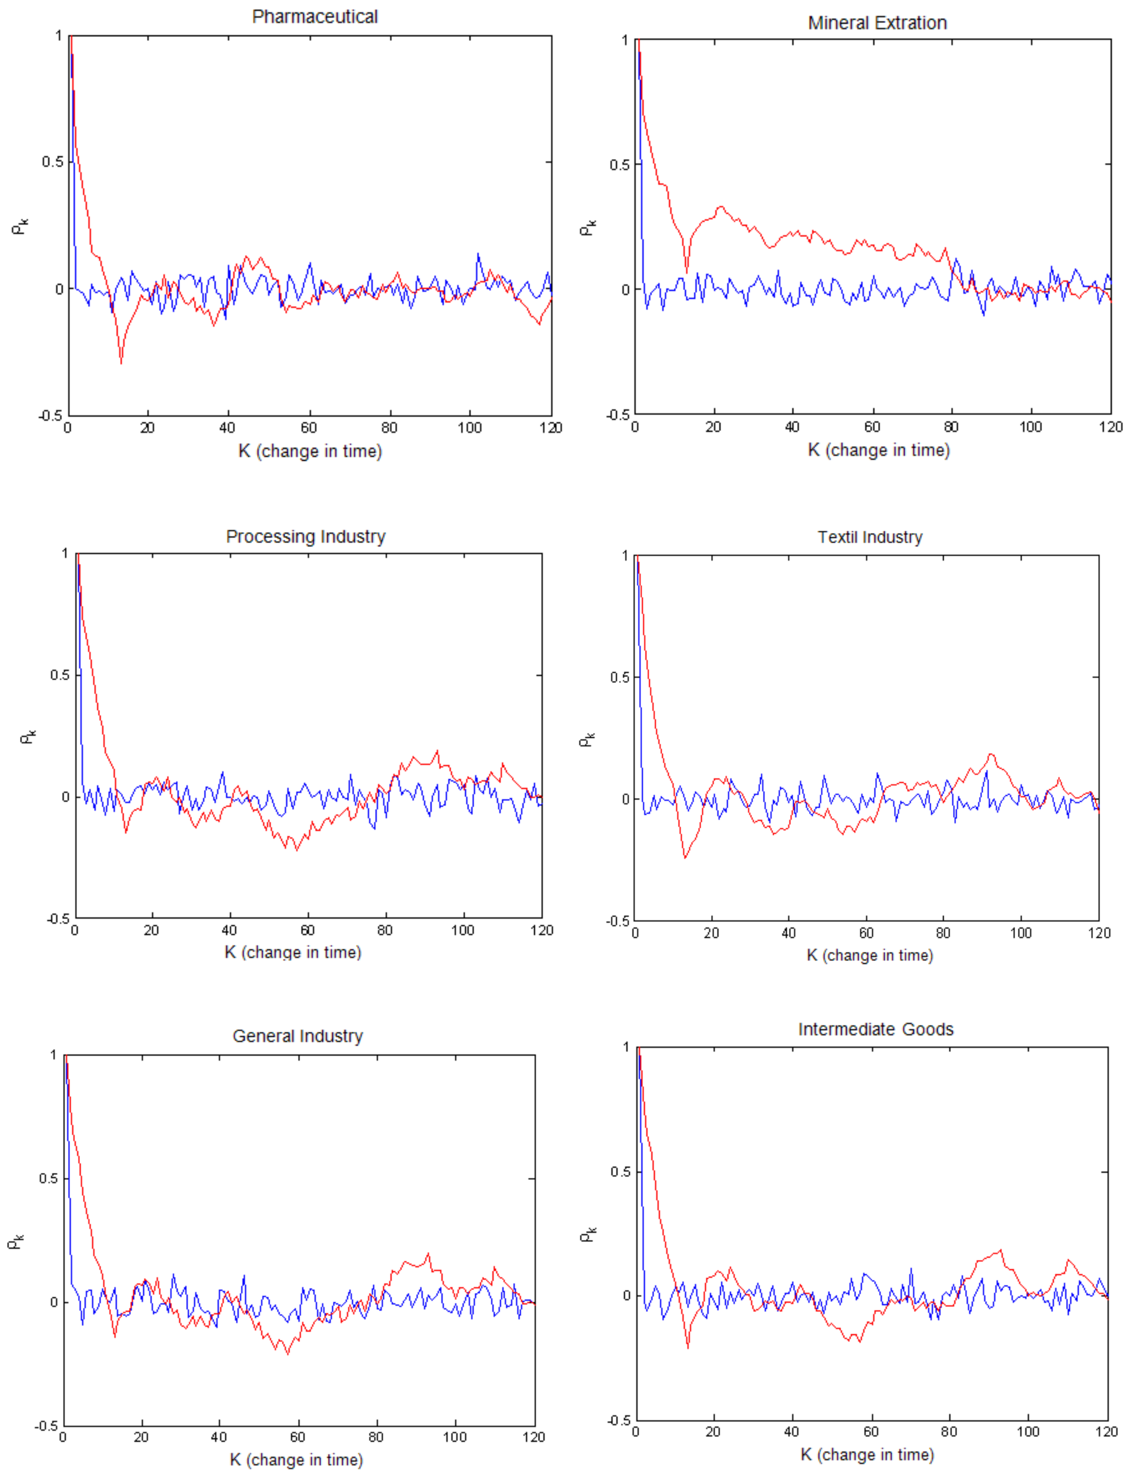

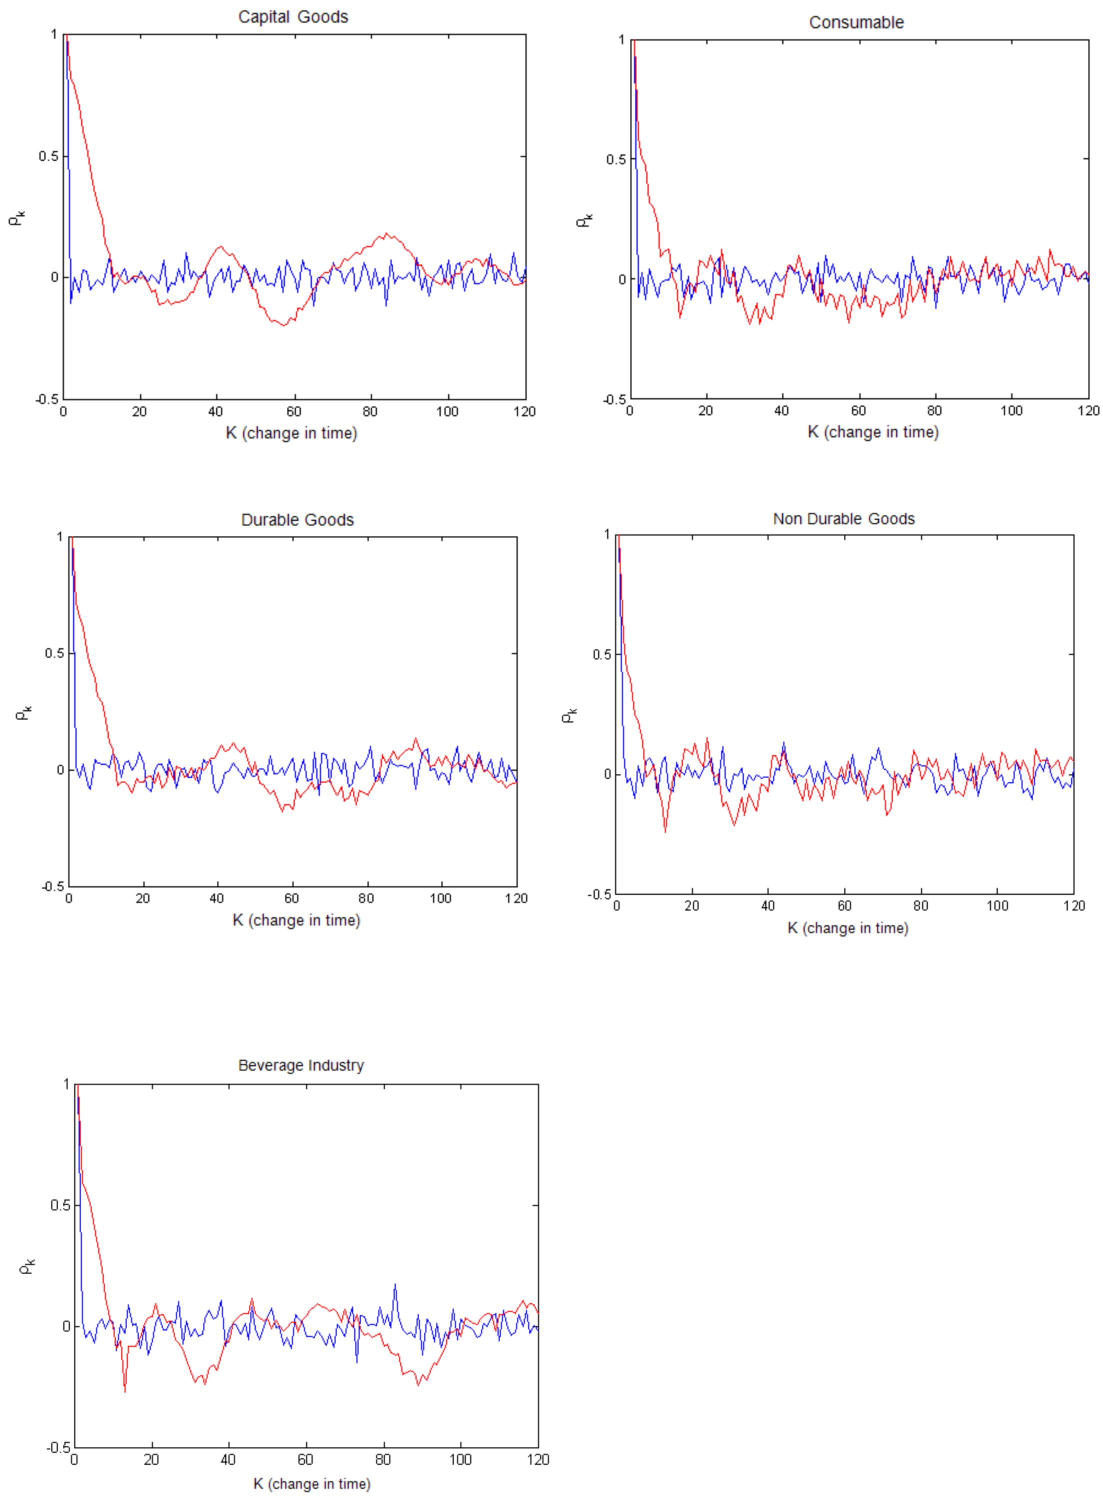

Legend: blue line (autocorrelation function of a white noise signal)

red line (autocorrelation function of time series change)

5.2 Strategies proposed for using the Time-Delay Neural Network

From the preliminary analysis, firstly we decided to apply the TDNN in the actual data of industrial production (1st strategy), and secondly the same topology on the data changed (2nd strategy). The type of training was supervised, divided into two parts: training and simulation, both with input and output data.

The algorithm used for training was the Levenberg Marquardt, requiring only 25 epochs (complete pass through all data) of training to achieve a good convergence. Firstly, the total error was approximately 4,000, being reduced to about 8 after 25 passages, as shown in figure 6.

Figure 6 – Learning curve of neural network considering 25 epochs

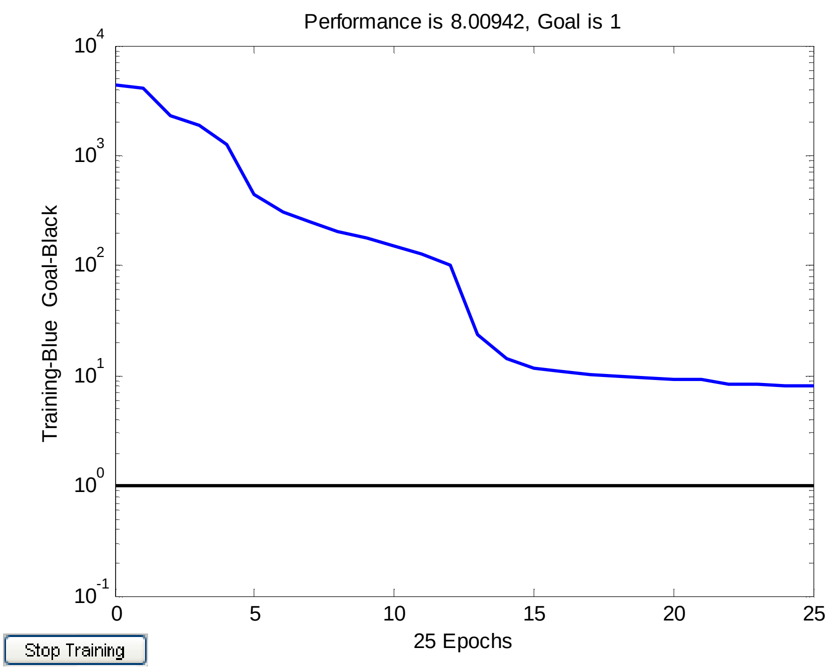

To define the number of neurons in the hidden layer of the network, we analyzed residues of simulations ranging its value from 1 to 25, through a statistical test of normality and an autocorrelation function, assuming that all processes are stochastic.

As the pharmaceutical industry presented a complete set of common characteristics to other segments, as shown in figure 4, we selected it for normality tests due to offer the best conditions of adaptability for all segments.

Considering the 1^{st} strategy, the cases with the lowest variance of the average of the observations were found for 8 neurons (variance equal to 7.64) and 5 neurons (variance equal to 7.67). As the variance values were very similar, we chose the case with fewer neurons in the hidden layer, i.e. 5 neurons.

In the 2^{nd} strategy, almost all cases presented distribution close to a normal curve, although there were few residues approved in the normality test, but curiously normality was achieved with the amount of 3 neurons (variance of 8.53).

Therefore, the final network structure in the first strategy had 13 inputs, in order to observe the seasonal behavior of the model in periods of 12 months, 5 neurons in the hidden layer and one neuron in the output layer, identified by N^{13-5-1} as shown in figure 7.

Figure 7 – Structure of TDNN used in the first strategy

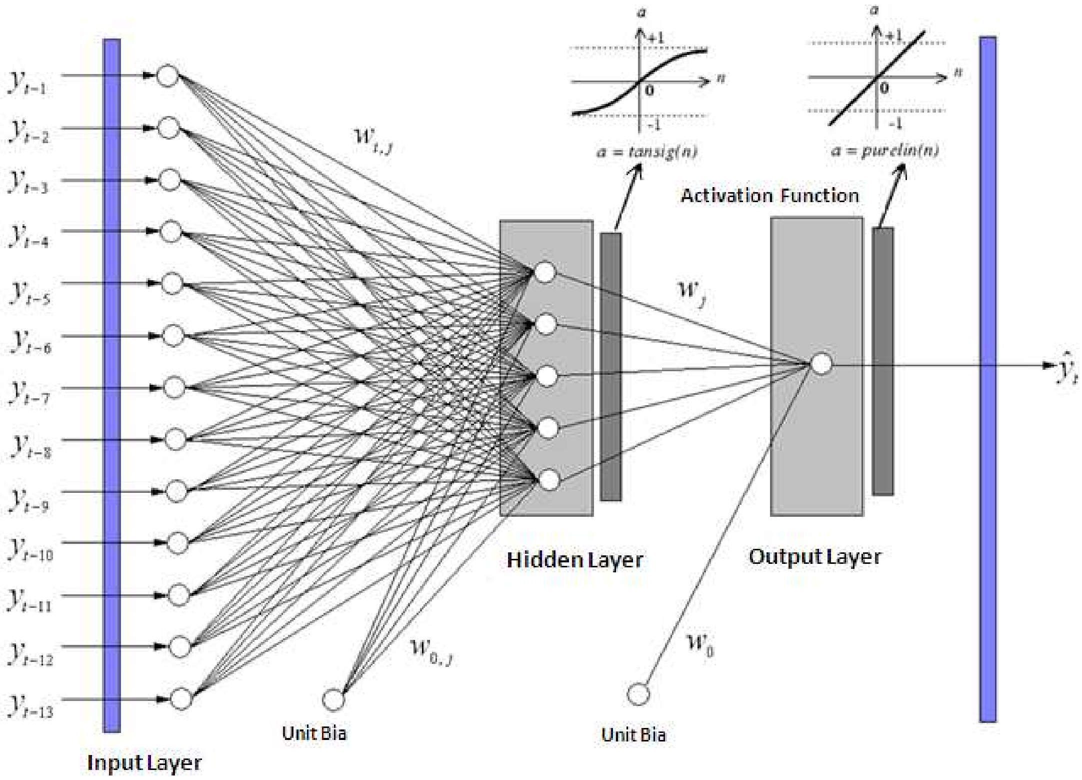

As in the 2nd strategy we eliminated the effect of seasonality, we did 5 training and simulations for each of the eleven segments to find the number of inputs that could vary from 13 to 3. In each processing, we analyzed the normality of the residuals and calculated the mean square error (MSE) and mean absolute error (MAE).

The results of the analysis for each segment were standardized and all the information was consolidated from the definition of a variable penalty related to the normality of the residuals and to errors found, according to figure 8. The best structure for the 2nd strategy with the lowest penalty factor has 4 inputs and it is identified by N^{4-3-1}. The structure of this network is presented in figure 9.

Figure 8 – Results of simulations to choose the number of inputs of the neural network for the second strategy

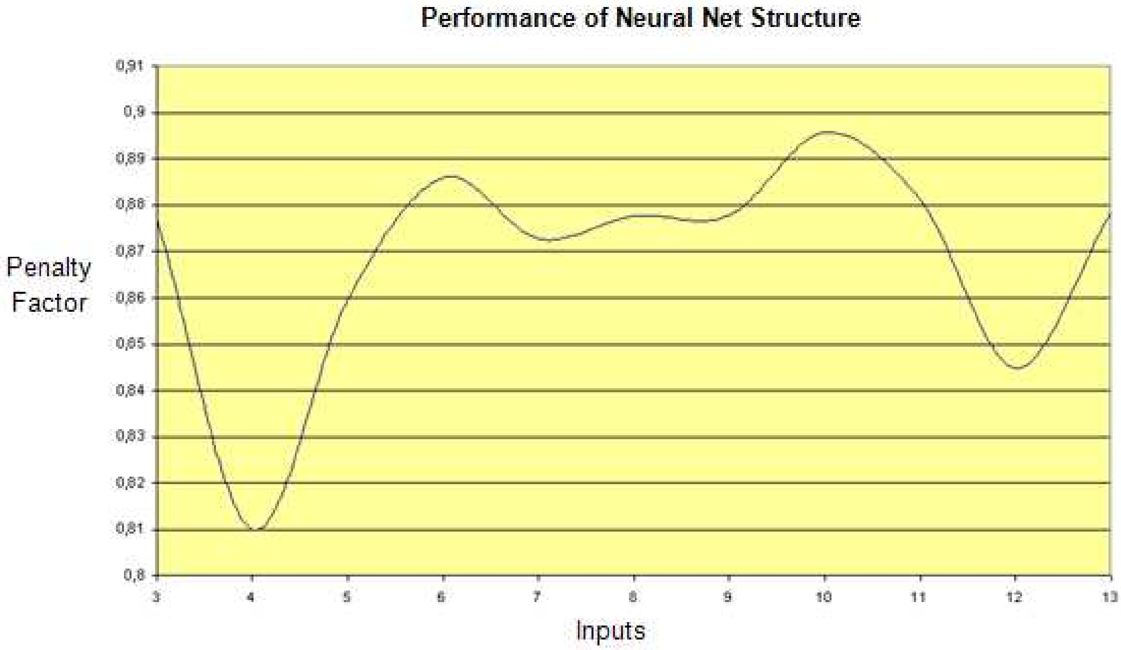

Figure 9 – Structure of TDNN used in the second strategy

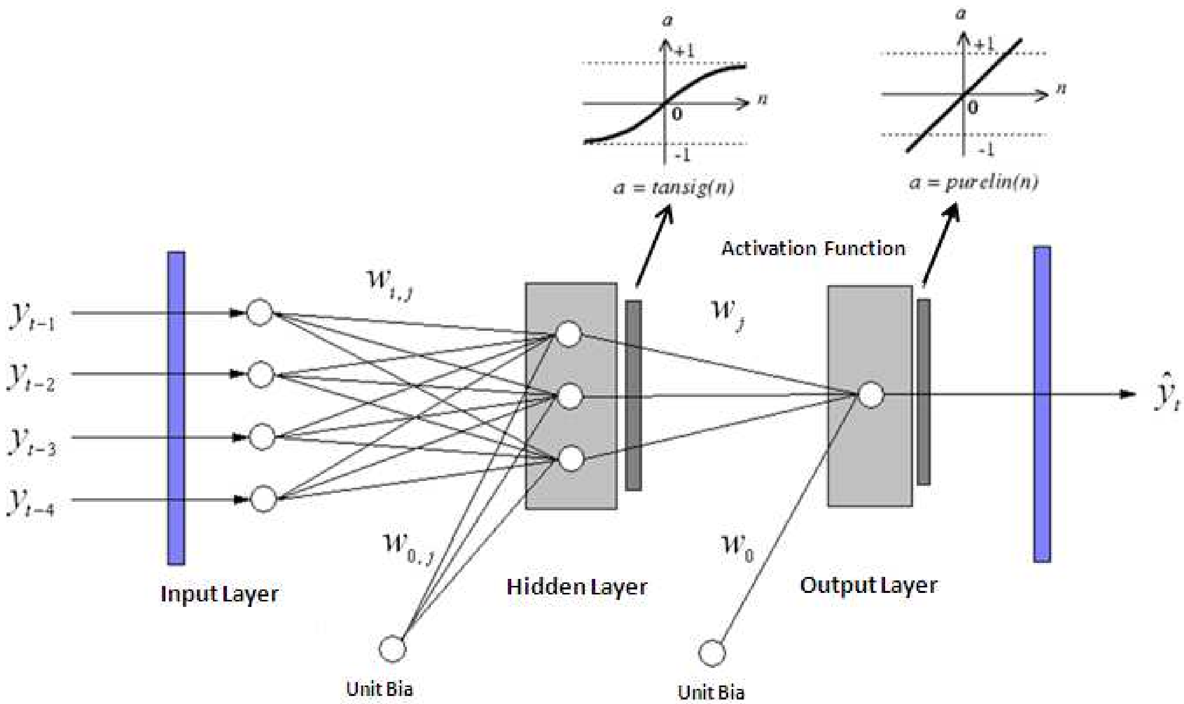

5.3 Analysis of two modeling strategies

Firstly, in both strategies we evaluate the segment of mineral extraction for presenting little seasonality and constant growth as well as the pharmaceutical industry for presenting an evident seasonality and slower growth. Furthermore, results of the beverage industry will also be presented. These three segments will give us a first idea of the performance of the neural network model proposed.

Mineral Extraction Segment

According to figure 10, the first strategy applied to the mineral extraction segment is represented by the blue line (actual data), and the overlap represents the transformed data (pink line). From the 395 to 405 months we can see the green line representing the demand forecasting.

Figure 10 – Results of simulations of the mineral extraction segment

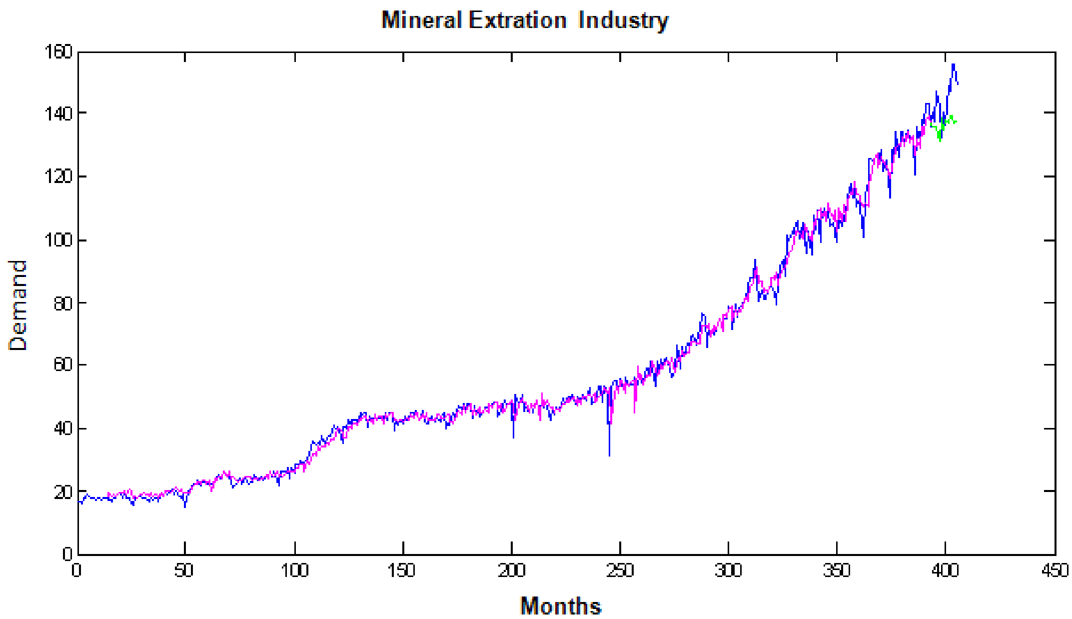

Legend: blue line – actual data of the 1st strategy

pink line – output the network from the data changed

green line – demand forecasting

Through the residues analysis (figure 11), we can see that the forecast is farther from the actual data than the simulation data, though the model has shown a normal distribution, with a large number of events around the average. The standard deviation (σ) of the residue is used to calculate uncertainties in the estimated demand.

For the first forecast, we use the same σ of the residue, and then to subsequent forecasts, we use $\sigma\sqrt{l}$ where l is the number of periods forward. This calculation is based on the fact that the variances σ^2 of each estimate add up, besides the residues are uncorrelated in time (Oja, 1982).

The demand forecasting of the mineral extraction segment comparing actual demand values of the first strategy to the values estimated by the neural network is also presented in figure 11.

Figure 11 – Probability distribution of the residue with μ = -0,0142 e σ =2,8765 (left) and demand forecasting of the mineral extraction segment using the 1st strategy. MSE=5,8403 e MAE=16,6178 (Right).

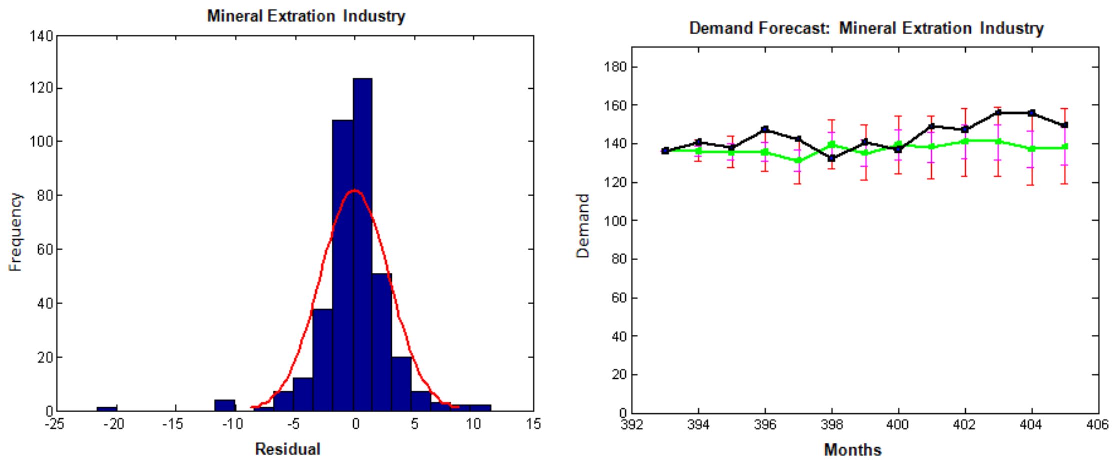

Legend of the right chart: black line – actual data of the 1st strategy

green line – estimated values by the neural network

Table 1 shows the weights used to calculate the TDNN, generating the demand forecasting for the mineral extraction segment.

Table 1 – Weights and Biases of TDNN ($N^{13\text{-}5\text{-}1}$) in the 1st strategy of the mineral extraction segment

WEIGHTS (inputs)					WEIGHT (hidden layer)	Biases	
$w_{i,1}$	$w_{i,2}$	$w_{i,3}$	$w_{i,4}$	$w_{i,5}$	w_j	$w_{0,j}$	w_0
-0,0017	-6,9300	-0,4944	18,5079	1,7358	126,5049	-0,6073	28,6583
0,0039	-6,3551	0,6830	51,0023	-3,5346	0,8638	-125,1319	
-0,0011	22,0324	0,1007	-10,2418	2,2065	18,0697	8,4731	
0,0006	-3,5472	-2,2646	18,6822	3,2288	-0,1158	-100,9507	
0,0000	-14,0444	0,2887	10,1820	-0,5146	-27,1206	-2,9353	
-0,0008	29,2165	-0,1688	7,2768	6,7214			
0,0014	3,5420	-0,3529	11,8190	-18,8722			
-0,0001	-19,3653	0,5940	-42,0550	8,8137			
-0,0009	18,8064	-0,3146	-25,7108	1,6403			
0,0008	-20,8188	0,8716	26,1314	-4,5609			
0,0014	-29,6160	1,3460	-60,9601	4,8853			
0,0020	3,8463	-1,0612	13,9632	-14,0026			
0,0031	23,9751	1,7118	-23,1733	-0,8097			

In conclusion, we can see in figure 10 that the forecasting (green line) is farther to the actual data (blue line) than the simulation data (pink line). This shows a degree of difficulty of the neural network in generalizing the behavior of the data using the first strategy. Although the neural network is able to represent the characteristics of seasonality, it can not follow the growing demand. Moreover, figure 10 shows that the variation of the actual data compared to the forecasting extrapolated the limits of uncertainty, concluding that the neural network was unable to represent this stochastic process.

Pharmaceutical Segment

Similarly to mineral extraction segment, we will present the same results for the pharmaceutical segment. Figure 12 shows the simulations results for the pharmaceutical sector. It can be easily seen the annual seasonal cycles.

Figure 12 – Simulations results for the pharmaceutical segment

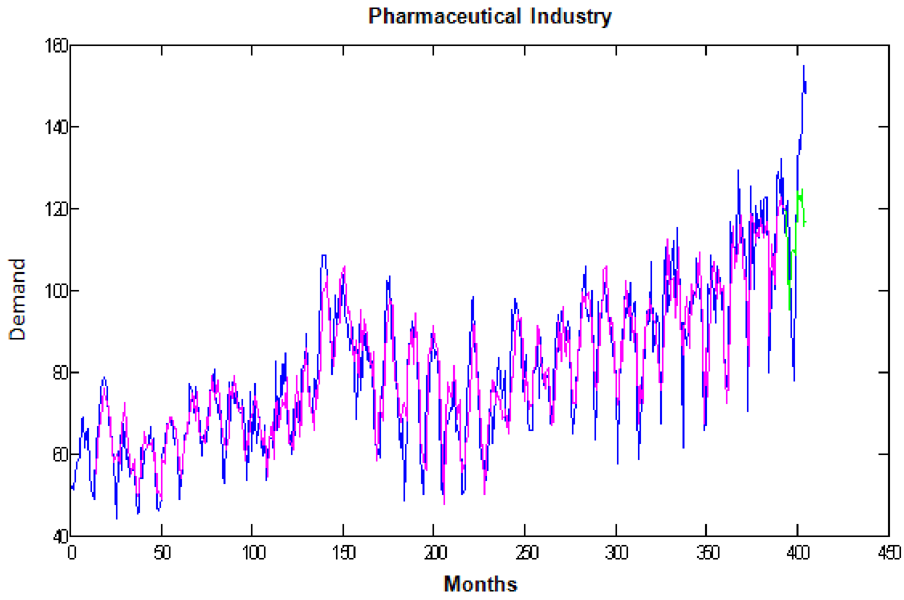

Legend: blue line – actual data of the 1st strategy

pink line – output the network from the data changed

green line – demand forecasting

Figure 13 shows the probability distribution of the residue and actual demand values of the first strategy compared to the values estimated by the neural network, both for the pharmaceutical segment.

Figure 13 – Probability distribution of residue with μ =0,0031 e σ =7,8158 (left) and Demand forecast of the pharmaceutical segment using the 1st strategy. MSE=5,8403 and MAE=16,6178.

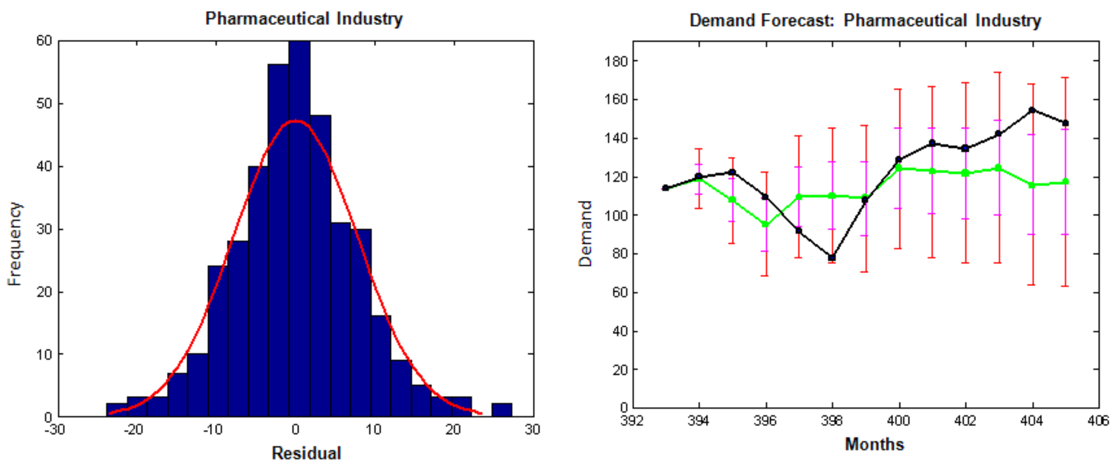

Table 2 shows the weights used to calculate the TDNN, generating the demand forecasting for the pharmaceutical segment.

Table 2 – Weights and Biases of TDNN ($N^{13\text{-}5\text{-}1}$) in the 1st strategy of the pharmaceutical segment

WEIGHTS (inputs)					WEIGHT (hidden layer)	Biases	
$w_{i,1}$	$w_{i,2}$	$w_{i,3}$	$w_{i,4}$	$w_{i,5}$	w_j	$w_{0,j}$	w_0
97,83995	-0,21902	-1,02725	-0,00474	4,89337	-0,08996	246,29189	85,51668
-98,78128	-1,99744	-1,44104	0,00750	-7,05858	-0,42920	285,89226	
91,35870	-1,56001	5,27727	0,00387	7,23419	0,97392	-19,46200	
-42,85439	1,31064	0,92957	-0,00191	-1,88510	62,15531	-1,48335	
106,12335	0,46509	-5,49845	0,00208	5,34686	-2,84345	68,98586	
-75,41247	0,10952	5,70046	-0,00247	-12,75571			
-86,43458	-0,73637	-0,75436	0,00159	7,05579			
-90,95014	-0,20782	0,04243	0,00063	-0,96218			
32,58869	0,76141	-1,47557	-0,00155	-1,31476			
334,34178	-1,95373	-4,25216	0,00069	2,13308			
-89,57207	-1,88741	3,44475	-0,00014	-0,02269			
-211,33687	-3,08808	1,73796	0,00243	-0,13788			
71,38521	1,02377	-0,98157	0,00888	-4,76130			

As the neural network is trained considering a data set whose standard in time is not repeated in the future, the second strategy that works with differences in time, instead of the original time series, presented a better result. Thus, the second strategy should present better results to other segments.

Beverage Industry

Applying the second strategy to the beverage industry, the seasonality effects were eliminated calculating the differences between demand values in intervals of 12 months. The new series appears in figure 14, showing a closer approximation of the transformed data (blue dashed line) and the neural network adjusted (red solid line).

Figure 14 – Transformed data of the beverage industry (blue dashed line) and neural network adjusted (red solid line).

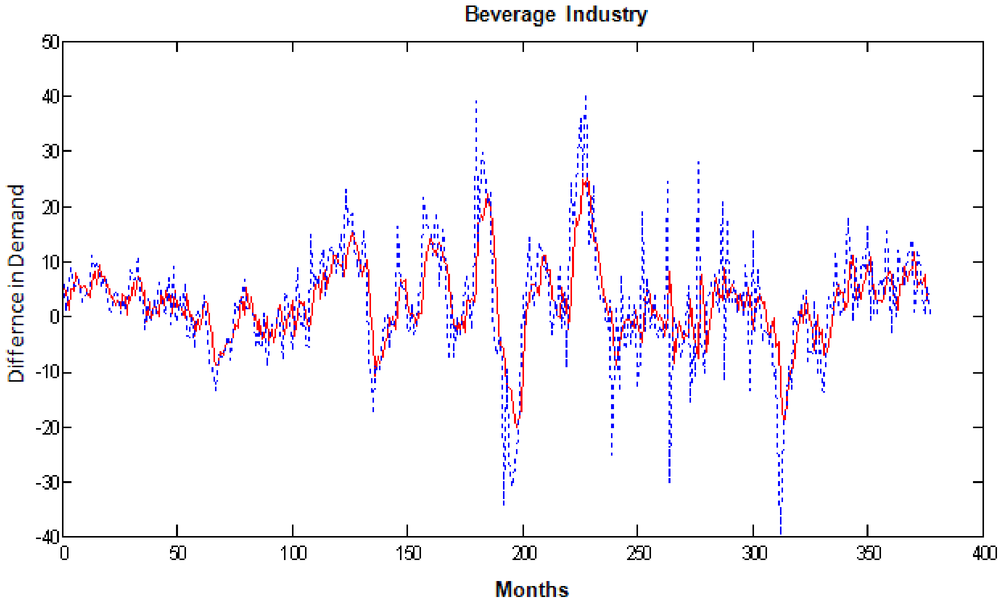

According to figure 14, the neural network model achieved a good approximation of the behavior of the time series differences, but smoothing the signal peaks generated from the actual data.

Making the signal recovery from the estimated differences in the simulation of the neural network, we obtained the results shown in Figure 15.

Figure 15 - Results of simulations of the beverage industry adjusted from time series differences.

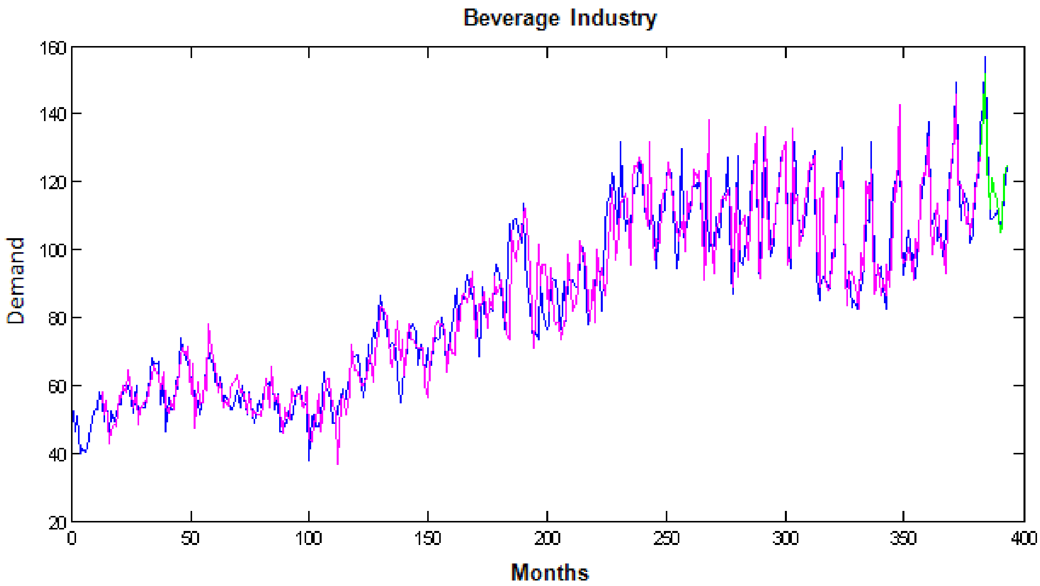

From figure 15, we can see that the model roughly followed and described the behavior of the real data. Moreover, the forecasting result (green) showed a better generalization of the seasonal behavior and growth of the real data.

Figure 16 shows the probability distribution of the residues and actual demand values of the second strategy compared to the values estimated by the neural network, both for the beverage industry. We concluded that the distribution of residues follows a normal curve with a 95% confidence level. The mean square error (MSE) was less than 2%, providing a good result for the model.

Figure 16 - Probability distribution of the residue with $\mu = 0.0158$ and $\sigma = 7.8261$ (left) and demand forecasting of the beverage segment using the 2nd strategy. MSE = 1.7104 and MAE = 4.8085 (right).

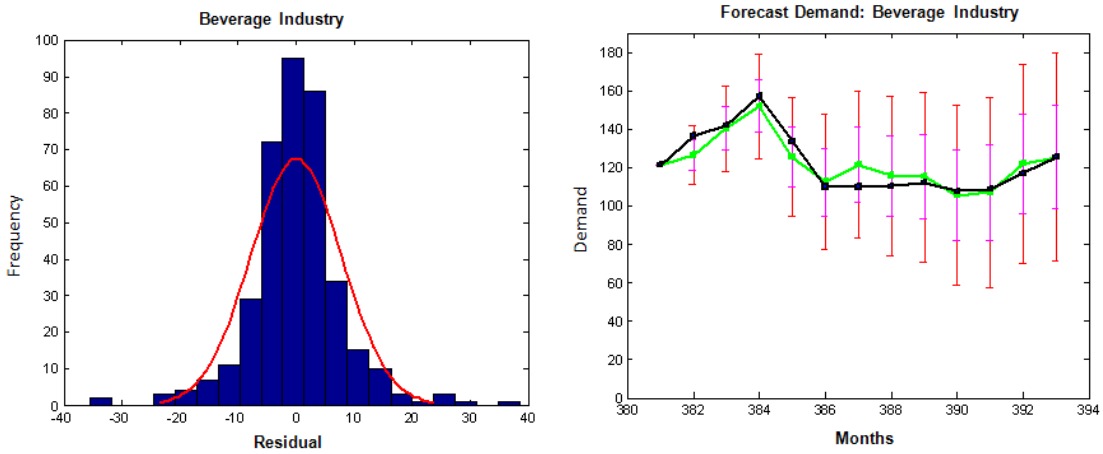

Table 3 shows the weights and biases of the TDNN in the 2nd strategy for the beverage segment.

Table 3 – Weights and biases of TDNN ($N^{13\text{-}5\text{-}1}$) in the 2nd strategy of the beverage segment

WEIGHTS (inputs)			WEIGHT (hidden layer)	Biases	
$w_{i,1}$	$w_{i,2}$	$w_{i,3}$	w_j	$w_{0,j}$	w_0
-0,5738	-0,0001	4,4431	1,1910	0,1802	8,1884
4,7981	-0,0013	-7,5446	-44,9871	0,1768	
-6,3155	-0,0090	4,6509	-0,7493	-2,8697	
4,7502	-0,0063	-1,2710			

Extended results

Similarly to the procedures applied to segments of mineral extractive, pharmaceutical and beverage, we applied them to other segments. We verified that the forecast model follows almost the same behavior of the actual data for the first strategy, but with some difficulty in following the increase or decrease of the data.

Applying the first strategy for all segments, it was found that the greatest differences were found in the segment of capital goods.

The smallest error between the forecasting model and the actual data occurred in the segment of intermediate goods and non-durable goods.

The application of the second strategy was effective in each industrial segment. Similarly to the first strategy, we applied tests of normality and auto-correlation of data. The results proved absence of autocorrelation in the distribution presented.

For all segments, we verify that the second strategy was more successful for demand forecasting, as shown by the values of MSE and MAE errors in Table 4.

Table 4 – Comparison of MSE e MAE between the TDNN (N^{13-5-1}) in the 1st strategy and the TDNN (N^{4-3-1}) in the 2nd strategy

SEGMENTS	1a. Strategy (N^{13-5-1})		2a. Strategy (N^{4-3-1})	
	MSE	MAE	MSE	MAE
Pharmaceutical	5,8403	16,6178	4,6268	12,2761
Mineral Extration	2,8893	8,7925	1,1954	3,9105
Processing Industry	2,6381	7,2603	1,5740	4,3921
Textil Industry	2,7127	6,8294	1,4719	4,2742
General Industry	2,5098	6,997	1,5276	4,2792
Intermediate Goods	1,7803	5,1607	1,2272	3,3421
Capital Goods	5,8575	17,906	4,9333	12,9632
Consumable	2,8293	7,3447	1,4335	4,0253
Durable Goods	5,3399	15,2202	3,4348	9,0343
Non Durable Goods	2,3334	6,708	1,1670	3,5839
Beverage Industry	4,1866	10,5392	1,7104	4,8085

6. Conclusion

This paper proposed a neural network model for forecasting the production time series of eleven different industries in Brazil. The time delay neural network (TDNN) with multi-layer perceptron was the best to estimate the production time series of the segments studied. The neural network model was then applied considering two different strategies varying according to the number of inputs and hidden layers.

In general, the model proposed was well adherent to the process of demand forecasting with better results in the 2nd strategy adopted transforming the original data and eliminating the effects of seasonality in intervals of 12 months.

The first strategy, even using a more complex neural network, was unable to achieve the performance of the 2nd strategy that used a simpler neural network processing the original data. The main reason for this result was the complexity of the original non-stationary time series, with a mean and variance dependent on time and seasonality.

Using a simpler network and with transformed and nearly stationary data, the 2nd strategy resulted in a greater potential for generalization, providing a more consistent learning within a maximum variation of 1 standard deviation, and confirming the greater ability of the neural network to forecast better demands.

Therefore, according to the results presented and errors obtained, we can prove and validate the effective application of the neural network structure proposed for both strategies, but with a better performance for the second strategy.

For future works, we suggest that new models and algorithms are tested in order to find a better efficiency in the results presented.

References

BRAGA, A. P.; CARVALHO, A. P. L.; LUDERMIR, T. B. **Redes Neurais Artificiais**: Teoria e Aplicações, 2nd edition, LTC, 2007.

KAASTRA, I.; BOYD. Designing a Neural Network for Forecasting Financial and Economic Time Series. **Neurocomputing,** Elsevier Science, p.215-236, 1996.

KHASHEI, M.; HEJAZI, S. R.; BIJARI, M. A New Hybrid Artificial Neural Networks and Fuzzy Regression Model for Time Series Forecasting. **Fuzzy Sets and Systems**, p. 769-786, 2008.

KUO, R. J.; XUE, K. C. A Decision Support for Sales Forecasting through Fuzzy Neural Networks with Asymmetric Fuzzy Weights. **Decision Support Systems**, v.24, p105-126, 1998.

MACCULLOSH, W. S.; PITTS, W. Alogical calculus of the ideas immanent in nervous activity. **Bulletin of Mathematical Biophysics**, v.5, p. 115-133, 1943.

MAKRIDAKIS, S.; WHEELWRIGHT, S.; HYNDMAN, R. J. **Forecasting Methods and Applications**. John Wiley & Sons, 3rd edition, New York, 1998.

OJA, E. Neural Networks: A Simplified Neuron Model as a Principal Components Analizer, **Journal of Math**, v. 15, p. 267-273, 1982.

PINDYCK, R. S; RUBINFELD, D. L. **Econometric Models and Economic Forecasts**, 3rd edition, McGraw-Hill, ISBN 0-07-050098-3, 1991.

SLACK, N. **Administração de Produção**. 1st Edition, Atlas, São Paulo, 1997.

STEVENSON, W. J. **Administração das Operações de Produção.** Rio de Janeiro, LTC, 2001.

WERNER, L.; RIBEIRO, J. L. D. Previsão de Demanda: Uma aplicação dos modelos Box-Jenkins. **Gestão & Produção**, v10, n.1, p.47-67, 2003.

WRAY, B.; PALMER, A.; BEJOU, D. Using Neural Network Analysis to Evaluate Buyer-Seller Relationships. **European Journal of Marketing**, v.28, n.10, p.32-48, 1994.