Forecasts Modeling in Industrial Applications Based on Artificial Intelligence Techniques

Agostino Bruzzone, Roberto Mosca, Alessandra Orsoni, Roberto Revetria
DIP University of Genoa
Via Opera Pia 15, 16145 Genova-Italy
Tel. +39 019 97398 – Fax +39 019 97 600
URL http://st.itim.unige.it – Email: agostino@itim.unige.it

Abstract

The management of industrial systems involves decision making with respect to complex processes that are often stochastic in nature. Simulation is frequently the only effective mean to model the complexity of such industrial processes. Simulation enables detailed scenario testing and, thus, is well suited for "what if" analysis. However, industrial users often need to solve inverse problems, such as optimization or decision analysis, which cannot be handled by simulation alone. This paper proposes the integrated use of simulation and Artificial Intelligence techniques in hybrid system architectures for advanced industrial problem solving. Hybrid Decision Support Systems (DSSs), combine the complementary strengths of different techniques for integrated forecasting, modeling, and optimization.

Keywords: Decision Support Systems, Artificial Intelligence, Neural Networks, Forecasting, Scenario Testing.

1 Introduction

Artificial Intelligence (AI) techniques are often difficult to implement due to a lack of formalized methodologies and systematic procedures for their application in specific contexts (Fahlman, 1988). In response to these issues, the authors propose an integrated approach for the application of AI techniques as anticipatory components of hybrid Decision Support Systems (DSSs). In this context, hybrid DSSs are intended as specialized tools which combine the complementary strengths of multiple techniques to address complex industrial problem solving. Specifically, the approach proposed in this paper focuses on the integration of simulation modeling techniques and AI techniques based on Artificial Neural Networks (ANNs) to combine the "what if" analysis capabilities offered by traditional simulation tools and the anticipatory capabilities offered by ANNs.

The paper illustrates these concepts through a set of four example applications, each one relevant to a specific industrial context, where the approach was successfully implemented. The specific problems addressed in these applications include: energy management in industrial buildings, Master Production Scheduling (MPS) for electronic component manufacturing, management of purchasing processes in a public transportation company, inventory and resource management in a large retail chain. This set of industrial cases provides a context for the application of AI techniques and for the assessment of the benefits that can be accrued from the synergy of complementary techniques (e.g. AI and conventional modeling techniques) when dealing with complex decision processes. For each application the paper systematically addresses the issues associated with the design and implementation of ANNs in relation to their designated anticipatory function within the particular DSS. Depending on the specific nature of the application, the ANN may be required to perform strong anticipatory functions (e.g. forecasting), or just to predict some trends (e.g. weak anticipatory functions). The authors'experience in the design, implementation and testing of ANNs as part of anticipatory DSSs, along with a significant core of experimental results, produce lessons and guidelines that are summarized in this paper for future use.

2 Structure of the Approach

In most industrial contexts demand for a product or a service is the main decision driver around which activities are planned, scheduled and optimized. As illustrated in further detail in the following, in many areas of business the variability of demand, and the difficulty to anticipate with sufficient accuracy its fluctuations over time, make it extremely difficult not only to formulate strategic plans, but also to make day-to-day management choices. In the proposed architecture ANNs are used as anticipatory components of the hybrid DSSs. The strength of ANNs as anticipatory systems lies in their ability to establish sensible experiential correlations among the set of variables that describe a given phenomenon or process. This ability makes ANNs especially effective when the complexity of the application excludes the possibility of developing suitable analytical models. This is typically the case when attempting to model the evolution of demand. The hybrid DSS architecture proposed in this paper bases its strength on the anticipatory capabilities of ANNs to produce accurate demand forecasts. As shown in the schematic of Figure 1, the ANN-based module is used to project the current situation into a future scenario, which sets the boundary conditions for the assessment of alternative management strategies with respect to the relevant measures of performance. In particular, current data from the actual system and historical trends are fed into the neural model which produces in return a demand forecast. The forecast is either fed into a simulation module for testing of alternative strategies or straight into a control system for automated correction of the current management policies.

Prior to their implementation in hybrid DSSs, the ANNs were extensively tested as stand-alone. Systematic testing was conducted on multiple ANN designs to identify successful design criteria for specific performance requirements. The performance measures chosen for the assessment of the different ANNs and for the optimization of their architecture include precision, learning time, and generalization ability. For the purposes of this study, precision is measured by the RMS (normalized Root Mean Square) of the error with respect to the real system. The learning time is measured by the number of learning iterations (over a set of training data) to convergence. The generalization ability is a measure of the ability to anticipate the response of the real system for a combination of input data not previously included in the learning set. This ability is quantified by the difference in anticipatory precision measured on the learning set and on the testing data set, respectively. The following sections will illustrate the application of the methodology to four distinct industrial cases, focusing on the design of the DSS and on the experimental results obtained after implementation.

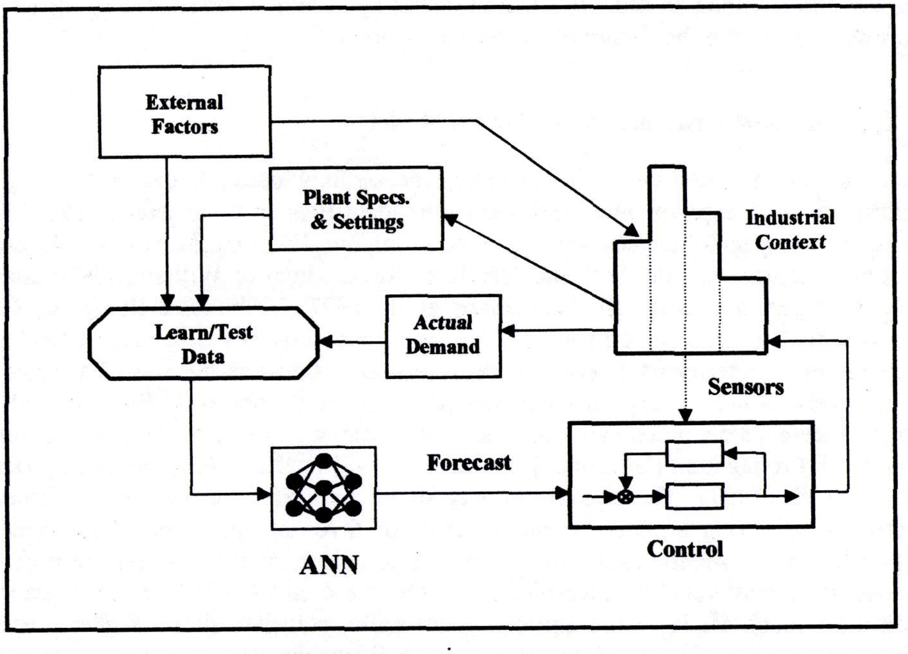

Fig. 1: Structure of the Approach

3 Energy Management in Industrial Buildings

The focus of this example application is on the development of an intelligent energy management system for a large telecommunication station located in Lagaccio, Genoa. This development effort fits in the more general context of re-engineering and management of building services with respect to their energy consumption. The main objective is to identify areas of potential savings and analyze their economical feasibility in terms of savings margins and corresponding investments. In this context, intelligent management systems provide dynamic decision support for complex decision problems that cannot be handled by traditional control systems alone. For the purposes of this application, the Authors developed a neural model based on Artificial Neural Networks (ANN) to work with the plant control system for dynamic decision support. The specific role of the ANN in this application is to provide forecasts of the plant energy requirements and of their evolution over time, taking into account the seasonal and daily fluctuation of demand. The choice of a hybrid control system with built-in anticipatory capabilities is shown to be most effective in this context where theoretical models and analytical tools would fail to capture the complexity of the interactions among the numerous demand drivers and their combined impacts on energy consumption. Forecasting accuracy and reliability of the tool are critical aspects of performance that highly depend on the particular network type, architecture, training data sets and methods. In order to ensure forecasting accuracy and reliability, significant effort has been put in the design of the neural model.

3.1 Typology and Structure of the Neural Model

While no established rules exist to select the optimal network design for a given application, prior experimental work led to the identification of general design criteria that apply to specific classes and types of problems. In particular, for the Lagaccio telecommunication station, both the literature (Baba; Minai & William, 1993) and the Authors' experience in the field (Giribone et al., 1997; Giribone & Bruzzone, 1999) supported the choice of feed-forward, fully-connected networks with two hidden layers of neurons (i.e. intermediate layers of neurons placed between the input layer and the output layer), which ensure high generalization ability (Giribone & Bruzzone, 1997). Performance considerations narrowed the choice among possible learning algorithms to either Back Propagation (Giribone & Bruzzone, 1997) or Direct Random Search (DRS) (Mosca et al., 1997). The two algorithms sit on opposite sides of the performance spectrum, but strong arguments can be made in favor of either one. The former is simpler in nature, widely used, quickly trained and easily applied to real life problems because of its relatively short learning times (Mosca et al., 1997). The latter, although much more elaborate in nature and computationally intensive, shows higher precision and excellent generalization ability (Giribone & Bruzzone, 1999). Three major factors drive the architecture of the hidden layers. Specifically, the number of neurons required on the hidden layers is a function of the total number of data available, of the theoretical complexity of the problem, and of the range of analysis (as defined by the data ranges) (Mosca et al., 1997; Padgett & Roppel, 1992; Lippman, 1987). The results from an extensive experimental study conducted by the Authors on a large number of network architectures specifically developed for energy consumption forecasting (Giribone & Bruzzone, 1997) provided effective guidelines for the specification of the hidden layers. These include 9 neurons in the first hidden layer and 7 neurons in the second one.

3.2 Available Data and Learning Set

Data for the preliminary training and testing of the network was obtained performing systematic measurements at the Lagaccio station. Two sets of data were collected sampling the power consumption of the station every hour, on the hour, during two distinct periods: 7 days during the month of June and 4 days during the month of October. The two sets of data represent a limited portion of the overall behavior of the station and do not constitute an adequate training basis for the network with respect to its operational use in the station (for such purposes additional data points will be generated using a conventional simulation model of the plant.) However, the two sets are adequately representative of the monthly and daily fluctuations in energy demand and provide a challenging test-bed for the network. Not only should the network recognize the presence of two distinct periods, but is should also cope with different operating conditions due to the fact that the station underwent major power saving re-engineering activities in between periods. Thus, the network is in principle required to learn two distinct sets of correlations. The choice of input parameters was targeted to the enhancement of the self-adapting ability of the network to the evolution of the system. This objective is reflected by the choice of three input variables that trace the three most recent values of energy consumption and use them to generate a forecast of the consumption in the following one hour.

A total of 228 data points, measured over the two sampling periods, were available for training/testing purposes. Only 75% of the available data, namely 174 data points, were used during the network training phase, while all of them were used during the testing phase. This choice enabled the Authors to test not only the precision and learning time of the network, but also its generalization ability.

Testing of the neural models on the full set of available data points revealed an excellent overlapping between the behavior anticipated by the models and the data actually measured. The relatively small number of data points available for training, and their non-homogeneous nature, stress the importance of the network's generalization ability among other performance measures. In particular, the DRS-based model, known for its high generalization ability, maintains a very low error level across the whole set of test data. The error for this network is always lower than 7% with an average value of less than 2%.

3.3 Performance of the Neural Models

The results indicate that the forecasting abilities of the ANN-based model can effectively support the operations management and the strategic planning for the plant in the short/medium term. The results also highlight the potential benefits of future model developments for the design and management of industrial buildings.

One of the main concerns of this research is to develop general criteria to effectively design and customize ANNs for different application types and contexts.

ANNs are shown to perform very effectively in relation to the energy management problem. Their ability to self-adapt, based on simple re-training on extended/upgraded data sets, is a key strength of ANNs (Padgett & Roppel, 1992). This ease of upgrade makes them suitable to model a dynamic behavior, which is influenced by the evolution of the plant itself, and by the evolution of a large number of external factors. The thermal control of industrial buildings represents a very complex problem that is difficult to address in the most general case, due to the multiplicity of factors influencing energy consumption. However, this application shows that effective neural models can be developed, verified, and validated for a specified industrial facility.

4 Master Production Scheduling for Electronic Component Manufacturing

The electronic component market is characterized by large demand with high levels of variability and strong customization needs. In the telecommunication sector, for instance, the market presses component manufacturers with high volume demand characterized by strong variability and customization requirements. Manufacturers, making customer satisfaction their first priority, are forced to keep excess inventory in order to minimize the risk of stock-out on specific items. The high level of uncertainty associated with the demand makes inventory management a critical issue and calls for effective decision support tools capable of handling the different aspects of the problem in a more integrated way. The application presented in this section addresses the problems faced by a major Telecommunications Company, which manufactures electronic components for the telecommunication industry. This company produces circuit boards and assembles them in highly customized component configurations. Starting from a baseline product design, the different customers may specify the number of parts and modify their architecture to best suit their application needs, they may require open configurations, or simply order spare circuit boards. The large variety of possible product configurations, along with the stochastic nature of demand, make empirical forecasting methods, typically based on the definition of an "average product configuration", highly unreliable. The lack of accurate forecasts leads to poor inventory management and ineffective production planning. Typical consequences are excess inventory with high risk of stock-out on specific parts and components and, thus, average lead times stretching far beyond schedule. In response to the perceived needs, the Authors developed an integrated Master Production Scheduling (MPS) package specific to the production environment of the company. The package consists of a hybrid system combining an ANN-based demand forecasting tool and a simulation model of the production process. The system provides quantitative means to assess the impacts of alternative production management choices in terms of inventory costs and customer satisfaction. The package also includes a performance evaluation tool that systematically monitors the reliability of the demand forecasts and the effectiveness of the production plan in relation to customer satisfaction.

4.1 The Neural Model

The ability to anticipate product demand and to express it in terms of actual inventory items (i.e. components and part numbers) was identified as a first critical objective towards the effective management of the manufacturing process. The experience of the Authors in similar areas of application, (Mosca et al., 1993 and 1997; Giribone & Bruzzone, 1995; Bruzzone & Giribone, 1998) led to the development of a neural model capable of making demand forecasts for each inventory item. The model generates these detailed demand forecasts based on historical data and on the specific features of the different inventory items. Due to the dynamic nature of the telecommunication industry it was chosen not to entirely replace the existing forecasting methods with the new one. Rather, the outcomes of the two methods are combined using a weighing procedure. The sets of weights employed in this procedure are specific to the part type and number.

The neural model developed for this application consists of a feed-forward, fully connected, back propagation network for each inventory item. The ANN bases its predictions upon historical data relevant to the last 12 weeks of operation and upon the specific features of each item. Using this information the ANN generates demand forecasts for the following 6 weeks. A number of additional parameters are associated to each part number in order to fine-tune the neural forecast. These include the maximum acceptable error, a probability of occurrence (demand), and two ranges of variability, one relevant to the part quantity and the other one relevant to the date (week of occurrence.) Such ranges are expressed in terms of percentage variation (positive or negative) with respect to the neural forecast.

This way, each forecast is converted into a probability distribution which associates the highest likelihood of occurrence to the value originally anticipated by the network. These forecasts are then combined with the traditional marketing forecasts according to a weighing procedure that takes into account the specific part number. The model generates records of historical data and relevant forecasts in the form of data files that can be stored and used for future analysis of system performance.

The neural model was tested on a set of historical data provided by the company. The data is relevant to over one thousand inventory items and reflects an entire year of operation (November 1999 – November 2000.) As shown in Figure 2, the results of this preliminary test are highly encouraging, since the average error associated to the neural forecasts is less than 1% on over 90% of the part numbers considered. Future developments of the neural module include continuous upgrade and fine-tuning through systematic re-training on more recent data points. The ANN was integrated with a stochastic, discrete-event, dynamic simulation tool which models the production process as implemented in the production sites. The ANNs forcasts are fed into the simulation model for simulated scenario testing of alternative scheduling policies. Production data and records of met and unmet deadlines are stored in appropriate data files for performance evaluation purposes. Data from previous forecasts is systematically retrieved and compared to current production data to determine the level of confidence associated with the decision support tool. This analysis supports the identification of potential areas of improvement for future developments of the tool.

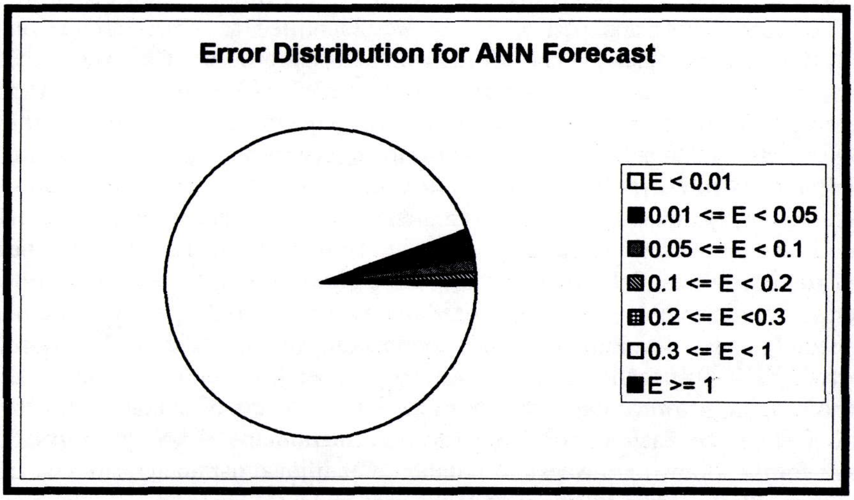

Fig. 2: Distribution of ANN Error on Test Data Set

5 Maintenance Management for Public Transportation Company

The application presented in this section consists of a decision support package for maintenance planning in a public transportation company. Effective maintenance planning in the transportation industry highly depends on the ability to anticipate parts and materials consumption. It is based on this information that a smart purchasing process can be defined to meet maintenance needs at the lowest possible costs. The decision support package was developed to address the specific planning needs of a mass transportation company. The company, which operates the entire ground transportation service of a large metropolitan area, counts approximately 3,000 employees and 1,200 vehicles. The most critical aspect of maintenance planning in this type of environment is the enormous number of different items and spare parts that need to be tracked in the definition of a purchase plan. A total of 28,000 different part numbers are currently required for maintenance operations purposes. Despite this complex maintenance part structure, until the most recent past purchasing decisions were entirely based on human experience. This research values experiential knowledge about the process and does not intend to replace human judgement with fully automated procedures. Rather, the research proposes an interactive DSS for the assessment of alternative purchasing strategies. The structure of the DSS is modular: the system integrates a simulation model, an ANN-based forecasting module, and a statistical analysis module, each one serving a specific purpose in the decision support process.

Specifically, the ANN-based module provides consumption forecasts relevant to the next 12 month period, the simulation model tests the impacts of different inventory management policies based on demand forecasts, and the statistical module analyzes inventory data to identify current consumption trends. The different modules were developed in C for consistency of data representation and effective integration within the existing company information system. A modular architecture was preferred over other modeling approaches to ensure high flexibility for future developments of the package in terms of incorporating new features (e.g. modules) or simply improving the existing ones.

5.1 Forecasting Item Consumption

Estimating the expected consumption for all the possible part types involved in this industrial application is a very complex forecasting task. The design of the neural model for these purposes is based upon simple feed-forward fully-connected neural networks with a single hidden layer of neurons (Hillis, 1989; Giribone & Bruzzone, 1997). This type of structure is known to provide a good level of approximation on a single unknown function. Given the large variety of inventory items and the highly stochastic nature of the respective consumption laws, easy and fast network training were the main objectives driving the choice of the learning algorithm. The choice of back propagation as learning algorithm entirely reflects this objective (Jones & Hoskins, 1987). A set of experiments was conducted to assess the precision of the chosen architecture as a function of the number of neurons in the hidden layer and of the numbers of learning iterations (learning time) required to convergence. The company database provided the data to train and test the ANN model. This data simply consists of the cost of each item and of a set of historical data on the monthly consumption. Automated procedures were developed to generate, train and test the large number of networks employed by the DSS. Three distinct networks are defined for each inventory item, which leads to a total of 60,000 networks for an average of 20,000 inventory items in stock during a given year. The forecasting procedures, in fact, are divided into three steps. The first step, performed by a single (1-25-1) ANN, consists of a preliminary consumption forecast for the next 12 months. This value is then fed into two (3-6-3) ANNs to generate forecasts for the next 4 months and 6 months, respectively. These forecasts are used to build an inventory scenario to test, through simulation, the effectiveness of different re-order policies.The rapidly evolving nature of item consumption does not allow to rely exclusively on forecasts based on monthly historical data. Therefore, a security algorithm was developed that keeps a set of ANNs under continuous training, to detect new consumption trends from the onset. This way, consistent increases in item consumption may be corrected ahead of time through the release of an emergency order, and the risk of stock-out is highly reduced. The main objective of this application is to provide accurate consumption forecasts in conjunction with an experimental set up where the effectiveness of alternative purchasing strategies can be assessed and compared. A prototype of the system has been implemented in the company and is currently being tested. In the meantime, a new company database is being developed, capable of performing multiple functions, such as statistical index management, data collection and upgrade to support the acquisition process. The system is expected to be fully operational as soon as the new database is finalized. The current version of the system runs both in Windows and in UNIX environments, and is suitable for multitasking and multi-user sessions. This way, highly computationally intensive procedures, such as the ANNs learning cycles, can be performed remotely on powerful workstations, while multiple users, working on their PCs, can access data and supervise the process through ftp and telnet sessions, respectively. This architecture is particularly efficient because it combines the benefits of a powerful UNIX platform while maintaining a user-friendly Windows interface.

During operation the ANNs employed within the DSS are systematically trained and tested. Their outputs are subject to continuous statistical validation: should the ANNs fail the validation test, an additional learning session is proposed. In addition, the ANNs are periodically (typically every six months) re-trained on new data in order to keep them up to date with possible developments of new consumption trends. At the same time the simulation modules perform a systematic analysis on hypothetical scenarios to test the influence of the stochastic components and the effectiveness of the purchasing policies. This analysis explores the effects of possible deviations from the expected consumptions predicted by the ANNs. The simulation results are then passed on to the statistical analysis module which performs a set of statistical tests aimed at assessing the impacts of the different policies in terms of probability, risk, and robustness. Based on these statistics a multi-target optimization process built into the smart selection module (fuzzy logic) provides the final recommendations. While standard policies may result suitable for most categories of items, in some cases it may be recommended to proceed manually. In summary the co-operative use of A.I.-based and traditional techniques in a distributed network architecture showed great potential in the realization of an integrated package for maintenance management and planning.

6 Inventory and Resource Management in Large Retail Chain

Customer satisfaction has become a daily challenge in the competition among large retail businesses. Supermarket chains are striving to improve their logistics in order to meet higher standards of customer service. In this perspective there is a perceived need for new procedures supporting a better integration among the supply chain stakeholders. Central warehouses are being created as intermediate structures to establish closer links between the individual stores and the different suppliers. At the retail end of the chain, a large number of resources are continuously attending customers and re-shelving goods but several factors affect the efficiency of these operations. These include late deliveries, product damage/deterioration, wrong forecasts, and manpower shortage. This situation calls for innovative support tools capable of modeling the external logistic, monitoring the internal workload, forecasting sales and resource requirements. In response to these needs, the Authors developed an ANN-based model for sales forecasting and inventory management. The neural model was designed to anticipate future sales and stock-out events in relation to a set of critical parameters such as store location, sales history, day of the week, month, hour. In particular, the number of stock-outs, measuring the ability of management to react to changes in market behavior and boundary conditions, is a very good indicator of store efficiency and customer satisfaction. The model was implemented in a leading Italian supermarket chain. Extensive data collection performed in a number of reference stores enabled a detailed feasibility study including testing of multiple network architectures, multiple choices of training and testing data sets, multiple data pre-processing and post-processing techniques. The study aimed at assessing the full potential of ANNs in the specific retail sector but at the same time led to the consolidation of valuable design criteria for future applications.

6.1 The Neural Model

Sales forecasting in large supermarket chains is a very complex problem involving a large number of highly stochastic variables. For these types of industrial problems Neural networks present significant advantages over other modeling techniques. Traditional simulation models, for instance, require development times that are incompatible with the rapid evolution of the business and, in addition, they are relatively difficult to upgrade. Neural networks, instead, require minimal modeling effort due to their ability to learn and self-adapt to the representation of complex behaviors directly from a set of example data points (Padgett & Roppel, 1992; Hassoun, 1995; Hillis, 1989). A neural model can be upgraded, and virtually re-built, through appropriate re-training on new data sets. However, as the complexity of the modeled behavior increases, the training process becomes prohibitively time consuming and computationally intensive. In order to minimize training time for this complex application, the Authors developed an innovative ANN design based on digital logic. Neural networks, in fact, are designed to work with numerical values in the range between 0 and 1, and reach their best approximation performance when operating within that range (Anderson & Rosenfeled, 1988). Based on these considerations the Authors decided to conduct a series of experiments on several networks trained using binary data sets. Specifically, the Authors performed a statistical analysis on the corresponding neuron (processing element, also indicated as PE) activities to identify possible means of reducing training time. Feed-forward, fully-connected, back propagation networks with a single hidden layer of neurons were developed both in C++ and in Java. The networks were tested on a set of mathematical functions including linear, quadratic and hyperbolic functions. For each function two distinct sub-sets of data were selected, one for training and one for testing purposes. A first set of experiments was conducted on a 7-bit architecture which showed a precision of 2 significant digits. The remarkable overlapping between the network's predictions and the target values (indicating a high generalization ability) observed for this prototype network encouraged the authors to implement a more sophisticated 10-bit architecture which showed a precision of 3 significant digits. The networks developed in Java required longer training times compared to the ones developed in C++, however, for the particular application in exam, it was chosen to proceed with Java-based networks. The main reasons behind this choice were the fact that Java was already a well-established standard for Web-based applications and the fact that this would greatly facilitate the integration of the tool with other applications already under developed within the company (e.g. E-Commerce, Internet Services, Wide Area Network, Intranet Architecture.)

6.2 Experimental Results

Two ANNs were utilized for the purposes of correlating the number of stock-outs to the time of the day and to the day of the week.

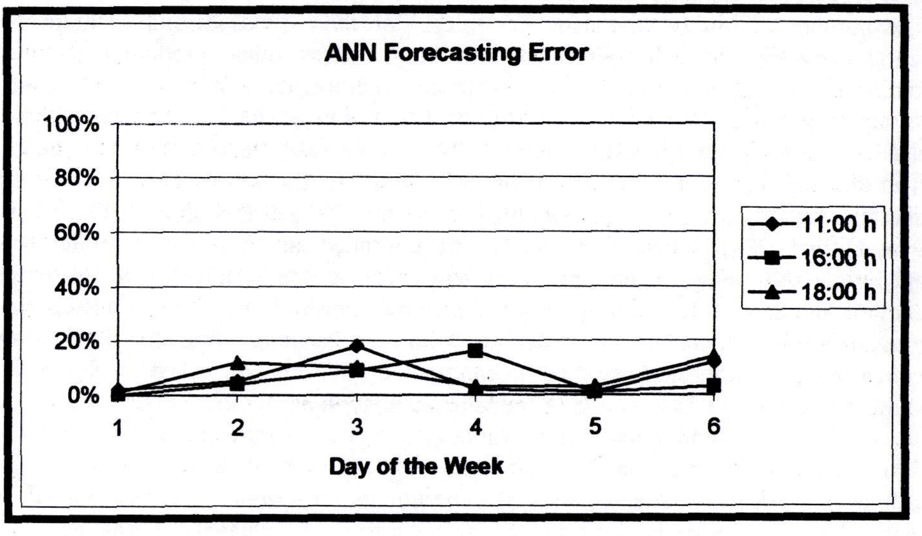

Fig. 3: ANN Error in the Approximation of Stock-Out Events

It is important to keep in mind that, as much as trends in sales can be expected, due to well-established consumer habits related to their life-style (in Italy, for instance, it is quite a diffused habit to do the main grocery shopping on Saturday afternoons and buy only a few items, as needed, during the rest of the week.) Stock-out events on specific products may occur suddenly and unpredictably due to a large variety of external factors (media advertising, especially TV commercials, play a key role in determining these kinds of phenomena.) Both the learning and the testing data sets are then affected by a noise element that makes the model intrinsically inaccurate. The experiments conducted by the Authors show that the use of binary architectures and binary data sets highly improves the precision of the network's forecasts, thus suggesting that binary coding may actually filter out some of the noise. Figure 3 shows the approximation error measured on the stock-out forecasts for two different times of the day over the period of one working week (Monday through Saturday.) The level of precision achieved is considered satisfactory, given that the purpose of this application is not to obtain exact sales/stock-out values, but rather the idea is to identify possible strategies to improve store logistics.

7 Conclusions

The set of industrial applications presented in this paper provided a context for the assessment of artificial intelligence techniques as forecasting tools in complex decision management. Extensive experimental analysis conducted on the specific models developed for these applications demonstrated that artificial intelligence techniques based on neural networks provide powerful decision support in real industrial contexts. However, when dealing with real problems, it is critical to keep in mind that these techniques do not represent the cyberage version of a universal crystal ball. Rather, they are able to provide effective support in specific contexts and with respect to specific problems.

Design, implementation and testing of these techniques for real-life applications highlighted some of the critical issues associated with the use of artificial intelligence techniques based on neural networks. Customization, for instance, along with proper experimental analysis, validation, and verification are critical aspects of the development process that need to be tailored to the particular problem at hand. In other words, effectiveness, accuracy, and reliability are properties that need to be built into the model through an accurate selection of the appropriate architecture, learning algorithms, and relevant data sets. In addition, due to the dynamic nature of most industrial contexts, such properties need to be constantly monitored through systematic data analysis and statistical verification/validation in order to identify the development of new trends that may require a model upgrade. Periodic upgrades of the model, which can easily be obtained through re-training sessions on more recent sets of data, are highly recommended in fast-changing industrial contexts. Most importantly, the precision and thus the effectiveness of the model strongly depend on its ability to reflect the specific characteristics of the problem. In this respect, data pre-processing and significant output testing are true corner stones in the successful application of these techniques.

Considering these complex implementation issues, it is highly encouraging to observe that the integrated use of ANNs, simulation, and other Artificial Intelligence techniques provides significant results on challenging management problems such as the ones outlined in the paper.

References

1. Anderson J., Rosenfeled E. "*Neurocomputing - Foundation of Research*", MIT Press, Cambridge MA, 1988
2. Baba N. "A New Approach for Finding the Global Minimum of Error Function of Neural Network" *Neural Networks*, Volume 2.
3. Bruzzone A.G., Giribone P. (1998) "Artificial Neural Networks as a Meta-Modelling Techniques for Car Component Industrial production", *Proceedings of Neurap 98*, Marseilles, 11-13 March
4. Fahlmann, Scott E (1988). "*An Empirical Study* to *of Learning Speed in Back-Propagation Networks*", CMU Technical Report, CMU-CS88, June 1988
5. Giribone P. & A.G.Bruzzone (1995) "Artificial Neural Networks as a System for Production Planning in an Alimentary Plant", *Proc. of Simulators International XII*, Phoenix, April 9-13
6. Giribone P., Bruzzone A.G. (1997) "Design of a Study to use Neural Networks and Fuzzy Logic in Maintenance Planning", *Proceedings of Simulators International XIV*, SMC'97, Atlanta, Georgia, April 6-10
7. Giribone P, Bruzzone A.G. M.Antonetti, G.Siciliano (1997) "Innovative Energy Management Techniques In Telecommunication Stations Through The Application Of Neural Models", *Proceedings of 1st World Congress on Systems Simulation*, Singapore, September 1-4
8. Giribone P., A.G. Bruzzone (1999) "Artificial Neural Networks as Adaptive Support for the Thermal Control of Industrial Buildings", International Journal of Power and Energy Systems, Vol.19, No.1, 75-78 (*Proceedings of MIC'97*, Innsbruck, Austria, February 17-19, 1997)
9. Hassoun M.H. (1995) "*Fundamentals of Artificial Neural Networks*", MIT Press, Cambridge MA
10. Hillis D.W. (1989) "*The Connection Machine*", MIT Press, Cambridge MA
11. Lippman R.P. (1987) "An Introduction to Computing with Neural Networks" *IEEE ASSP*, April 1987
12. Jones W.P. & Hoskins J. (1987) "Back-Propagation, A Generalized Delta Learning Rule" *Byte Magazine*, October
13. Minai A.A. & William R.D. "Acceleration of Back-Propagation Through Learning Rate and Momentum Adaptation", *International Joint Conference on NN*, vol I, January 1993
14. Mosca R., P.Giribone & A.G.Bruzzone (1993) "Application of Neural Networks to On-Line Simulations of Plant Processes", *Proceedings WNN93*, San Francisco, November 7-10
15. Mosca R., P.Giribone, A.G.Bruzzone, A.Orsoni, S.Sadowski (1997) "Evaluation and Analysis by Simulation of a Production Line Model Built with Back-Propagation Neural Networks", *International Journal of Modelling and Simulation*, Vol.17, no.2, pp.72-77
16. Padgett, Mary Lou & Roppel, T.A. (1992) "Neural Networks and Simulation: Modeling for Applications", *Simulation*, vol.58, no.5