Intelligent Data Analysis 5 (2001) 23–39
IOS Press

A neural network-based software sensor for coagulation control in a water treatment plant

Nicolas Valentin[a,b] and Thierry Denœux[b]
[a]*CIRSEE – Suez Lyonnaise des Eaux, Technopolis, ZAC de Merci ères, 14, rue du Fonds Pernant, F-60471 Compiègne, France*
E-mail: Nicolas.Valentin@lyonnaise-des-eaux.fr
[b]*Université de Technologie de Compiègne, UMR CNRS 6599 Heudiasyc, BP 20529, F-60205 Compiègne, France*
E-mail: Thierry.Denoeux@hds.utc.fr

Accepted 2 October 2000

Abstract. This paper reports on the application of Artificial Neural Network techniques to coagulation control in drinking water treatment plants. The coagulation process involves many complex physical and chemical phenomena which are difficult to model using traditional methods. The amount of coagulant ensuring optimal treatment efficiency has been shown experimentally to be non-linearly correlated to raw water characteristics such as turbidity, conductivity, pH, temperature, etc. The software sensor developed is a hybrid system including a self-organising map (SOM) for sensor data validation and missing data reconstruction, and a multi-layer perceptron (MLP) for modelling the coagulation process. A key feature of the system is its ability to take into account various sources of uncertainty, such as atypical input data, measurement errors and limited information content of the training set. Experimental results with real data are presented.

Keywords: Data validation, missing data reconstruction, outlier rejection, self-organising maps, multi-layer perceptrons, hybrid system, coagulation control, water treatment

1. Introduction

In water treatment, as in many other domains, process monitoring and control relies heavily on accurate and reliable sensor information. Whereas many process parameters can be measured continuously using relatively simple and cheap physical sensors, the determination of certain quantities of interest requires costly laboratory analyses which cannot be performed on-line. Such high level information may, however, sometimes be inferred from available measurements of observable quantities using a statistical model usually referred to, in this context, as a "software sensor" [1].

This paper addresses the problem of building a software sensor for on-line determination of optimal coagulant dosage from raw water characteristics such as turbidity, pH, conductivity, etc. Previous studies [2,3] have shown the potential effectiveness of such an approach based on artificial neural networks. The innovative aspect of the present work resides primarily in the integration of various techniques in a global system allowing for data validation and reconstruction, prediction of the quantity of interest, and analysis of uncertainties.

Given the high variability of the inputs and the low reliability of available sensors, an important requirement for such a system is robustness against erroneous sensor measurements or unusual water

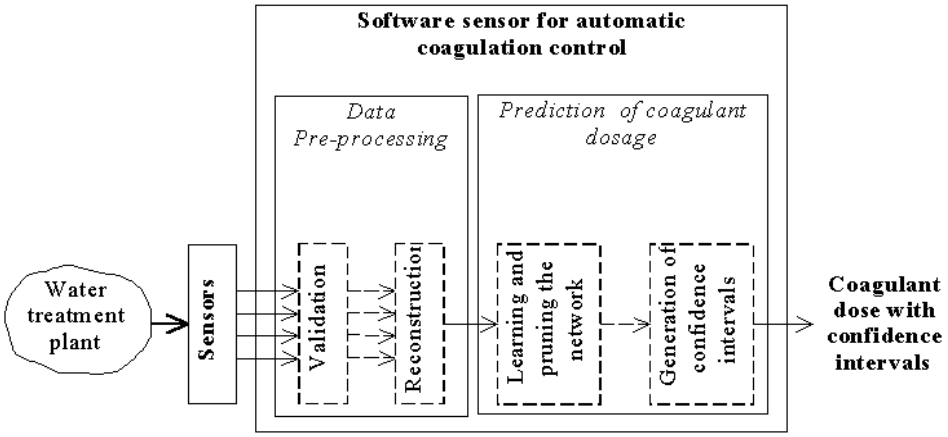

Fig. 1. Structure of the system for automatic coagulation control.

characteristics, due to accidental pollution for instance. In our system, such a robustness is achieved using a modular architecture composed of two levels: a pre-processing level responsible for outlier rejection and missing data reconstruction, and a prediction level involving the determination of the optimal coagulant amount from raw water characteristics (Fig. 1).

A second important requirement from the considered application is the possibility to install the system at low cost in various sites, which necessitates a methodology for designing and training the neural networks automatically from new data, including the phases of data validation and model choice. Our system uses pruning and re-sampling techniques for automatic determination of the network architecture and computation of confidence bounds for the predictions.

This paper is organised as follows. Section 2 presents a general introduction to the application domain and states the problem addressed in the rest of the paper. The two main modules of the system, based on self-organising maps and multi-layer perceptrons, are then described in Section 3 and 4, respectively. Finally, experimental results are presented and discussed in Section 5.

2. Problem description

2.1. Overview of water treatment operations

Water treatment involves physical, chemical and biological processes that transform raw water into drinking water. However, contrary to most industrial processes, for which the quality of the input raw material is under control, the quality of the given raw water source may fluctuate due to natural perturbation or occasional pollution.

Figure 2 depicts the main processes in a typical plant for surface water treatment (the Viry-Chatillon plant, near Paris). Raw water is abstracted from the resource (a river in this case) and pumped to the treatment works. A water treatment plant invariably consists in two main process units: clarification and filtration; other units may be required depending on the quality of the water source.

The coagulation process is brought about by adding a highly ionic salt (aluminium sulphate) to the water. A bulky precipitate is formed which electrochemically attracts solids and colloidal particles. The solid precipitate is removed by allowing it to settle to the bottom of the tank and then periodically

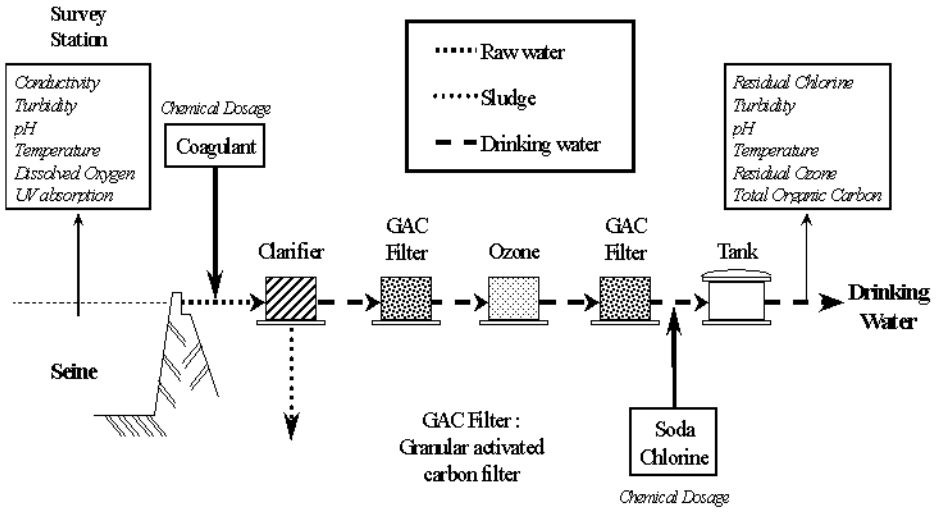

Fig. 2. Simplified synopsis of the Viry-Chatillon water treatment plant.

removing it as sludge. The coagulation process accounts for the removal of most of the undesirable substances from the raw water and hence tight monitoring and control of this process is essential. The next stage is filtration, where the particles passing through the previous stages are removed. Filtered water is also treated by ozonation to eliminate the last micro-pollutants. The final stages in the process are chlorination and pH adjustment. The water is then stored in a tank and ready to be transported through the water supply network.

2.2. Coagulation control

Coagulation is one of the most important stages in surface water treatment, allowing for the removal of colloidal particles. The main difficulty is to determine the optimum quantity of chemical reagent related to raw water characteristics. Poor control leads to wastage of expensive chemicals, failure to meet the water quality targets, and reduced efficiency of sedimentation and filtration processes. In contrast, good control can reduce manpower and chemical costs and improve compliance with treated water quality targets.

The traditional method of controlling coagulant dose, called the jar-test, relies heavily upon human intervention. It involves taking raw water samples and applying different quantities of coagulant to each sample [4,5]. After a short period of time each sample is assessed for water quality and the dosage that produces the optimal result is used as a set point. Operators change the dose and make a new jar test if the quality of treated water changes. Disadvantages associated with such a procedure are the necessity to rely on manual intervention, and lack of adaptation to abrupt changes of water characteristics.

More recently, an automatic device, called a streaming current detector (SCD) [6,7] has emerged. This device is based on the measurement of the net residual charge surrounding turbidity and colloidal particles in water. It requires a set point to be entered, assumed to represent an optimum water-quality standard. Streaming-current values above the set point reveal an excess of coagulant, while values below the set point indicate insufficient coagulant dosage for full flocculation to occur. A jar test needs to be carried out to determine the set point. Disadvantages associated with the SCD are its operation cost and limited efficiency for certain types of raw water quality.

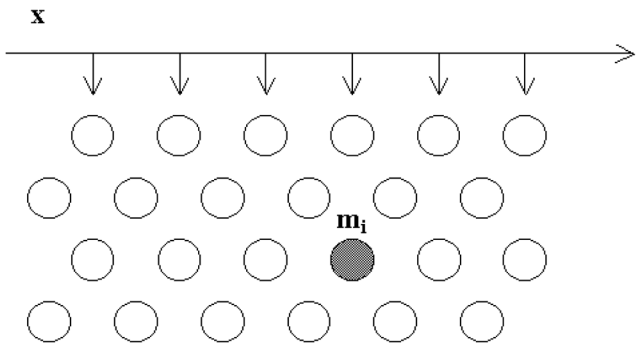

Fig. 3. Principle of the SOM.

The objective of this paper is to propose an alternative to the jar-test and SCD methods allowing for the automatic determination of optimal coagulant dose from raw water characteristics, using an artificial network approach. This approach requires the availability of on-line water quality measurements at an upstream survey station. An important objective of the present study is to automatically validate the sensor measurements performed at the survey station so as to provide reliable inputs to the automatic coagulation control system.

3. Data validation and reconstruction

3.1. General approach

Applications in the environmental domain such as the one considered in this paper generally rely on complex sensors located at remote sites. The processing of the corresponding measurements for generating higher level information (such as predictions of optimal coagulant dosage) must therefore account for possible sensor failures and incoherent input data.

In many cases, even domain experts are unable to categorise off-line data as “valid” or “invalid” with complete certainty, because of incompleteness of the available information or lack of knowledge of the underlying physical phenomena. For this reason, data validation should not be merely seen as the suppression of spurious data, but as the determination of a degree of confidence in the data, based on a comparison between observations and models incorporating background knowledge about the application domain. For example, the validity of a temperature measurement may be put in doubt because (1) the value is surprisingly high or low; (2) the rate of change between consecutive time steps is not within certain limits; (3) the temperature measured at successive time steps normally follows a certain pattern that is not present in the data; or (4) the observed value is not compatible with other measurements of the same quantity obtained by an independent device, or of other quantities whose values are normally related to temperature, etc.

In our system, data validation is performed at two levels, referred to as single parameter and multi-parameter validation. The first level is quite simple and is based on the comparison of each input parameter and its derivative to a distribution of historical values obtained in the absence of sensor faults. A confidence level is calculated, and raw data with confidence level below a user-defined threshold are declared as invalid data.

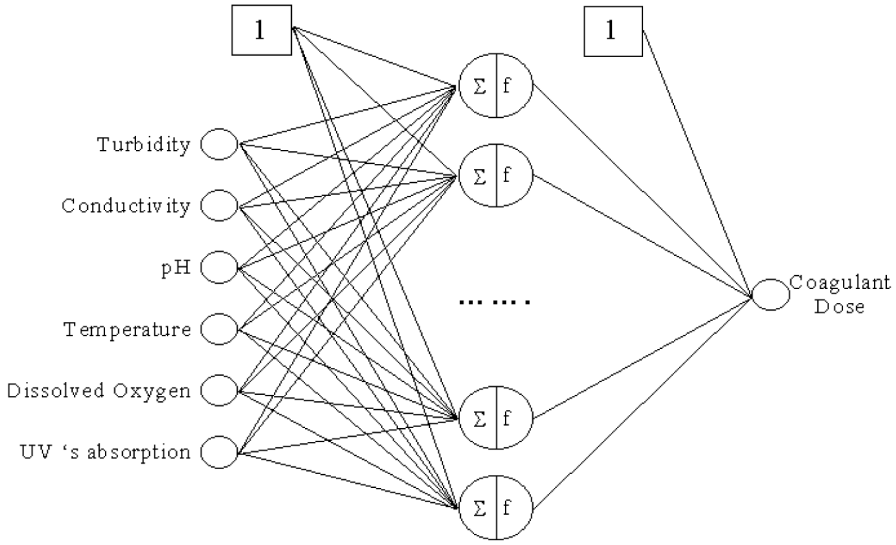

Fig. 4. Multi-layer perceptron network for prediction of coagulant dosage.

Although this simple approach allows in many cases to detect gross measurement errors, the detection of inconsistencies in data involving more than one parameter requires the use of more sophisticated multidimensional techniques such as the self-organising feature map model implemented in our system.

3.2. Multi-parameter data validation and reconstruction

3.2.1. Basic approaches

Neural network approaches to signal failure detection and reconstruction include autoassociative neural networks (AANN) and SOM's [8].

The basic idea behind the AANN approach is to train a multilayer feedforward network to approximate the identity function by using target values identical to the input values. One of the hidden layer typically has limited capacity and plays the role of a bottleneck, which forces the network to optimally encode the input vectors, thus performing information compression and dimensionality reduction. With a single hidden layer of linear units, this approach was shown by Bourlard and Kamp to be equivalent to principal component analysis (PCA) [9]. Consequently, more complex networks with non-linearities can be seen as implementing some form of "non linear PCA". Such an approach is proposed in Ref. [10], where a five-layer perceptron feedforward network is used for data validation. This network can be viewed as two independent three-layer neural networks connected in series. The first network mixes and compresses the n redundant measurements into a smaller number of characteristic variables which should ideally represents the essential characteristics of the process. The second network works in the opposite way and uses the compressed information to regenerate the original n redundant measurements. When trained on valid data, this network may be used to detect erroneous data, which are recognised from their higher reconstruction error [8]. However, the performance of such a system in the presence of incomplete input data is not fully predictable.

Another approach, which was adopted in this study, implies computing distances between input vectors and reference patterns, or prototypes. The determination of prototypes from data in an unsupervised way may be achieved using the Self-Organising Map (SOM) algorithm introduced by Kohonen [11]. The

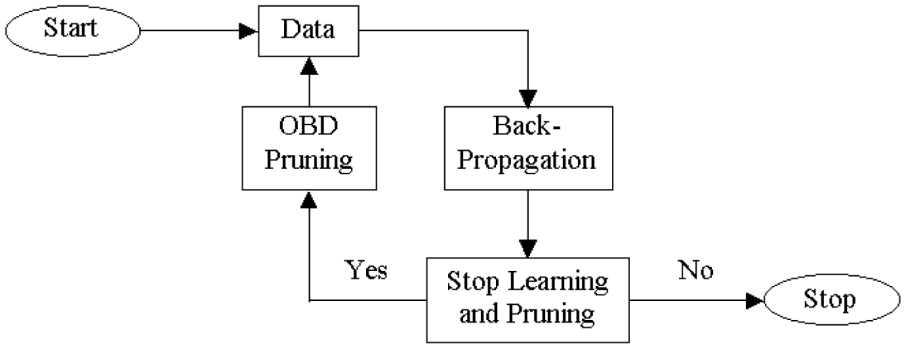

Fig. 5. Learning and pruning algorithm.

SOM model combines the goals of projection and clustering algorithms and may be seen as a method for automatically arranging high-dimensional data. It can be used at the same time to visualise the clusters in a data set, and to represent the data on a two dimensional map in a manner that preserves the non-linear relations of the data items, nearby items being mapped to neighbouring positions on the map. Applications of SOM's to water quality monitoring are described in [12,13].

3.2.2. Principles of SOM

The SOM defines a mapping from the input data space $\mathcal{R}^n$ (raw water quality parameters) onto a regular two-dimensional array of nodes (an hexagonal array in our case) as shown in Fig. 3. A weight vector (also called a reference vector or a prototype) $m_i \in \mathcal{R}^n$ is associated to every node i. Each input vector $x \in \mathcal{R}^n$ is compared to the m_i, and the best match m_c defines the winning prototype. The input is then mapped onto the corresponding location on the grid.

The process in which the SOM is formed is an unsupervised learning process. At each time step t, a data sample $x(t) \in \mathcal{R}^n$ is presented to the map. The node c that best represents the input is then searched for using, e.g., the Euclidean distance:

$$\|x - m_c\| = \min_i \{\|x - m_i\|\} \tag{1}$$

Next, the unit c as well as neighbouring units learn to represent the data sample more accurately. The weight vector of unit i is updated according to the following learning rule:

$$m_i(t+1) = m_i(t) + h_{ci}(t)\lfloor x(t) - m_i(t)\rfloor \tag{2}$$

where h_{ci} is a "smearing" or neighbourhood function expressing how much unit i is updated when unit c is the winner. The neighbourhood function typically is a symmetric, monotonically decreasing function of the distance between units i and c on the map grid. During repeated application of Eq. (2) with different inputs, the weight vectors of neighbouring units become gradually similar due to the neighbourhood function h_{ci}, eventually leading to global ordering of the model vectors. With time, the m_i then tend to become ordered along the array in a meaningful way.

3.2.3. Application to sensor failure detection and reconstruction

Self-organising maps allow not only to visualise the evolution of raw water quality in two dimensions, but also to detect atypical data or outliers by monitoring the distance between each input vector and its closest reference vector (which is a variant of the distance rejection concept introduced by Dubuisson

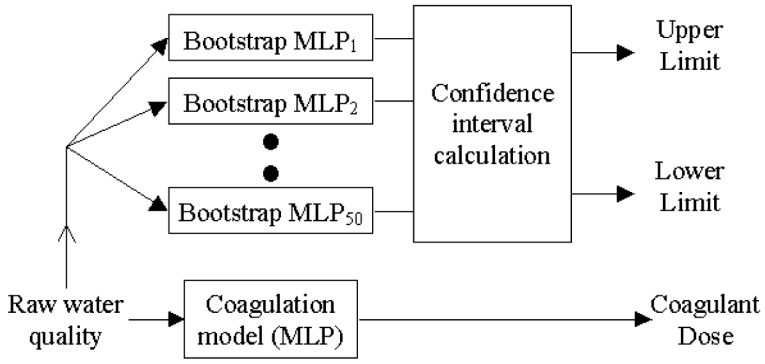

Fig. 6. Bootstrap sampling for the generation of prediction intervals.

and Masson [14]). More precisely, let us define the activation of unit i for input x using a Gaussian kernel as:

$$K(i) = \exp\left(\frac{-1}{2\sigma_i^2}\|x - m_i\|^2\right) \tag{3}$$

where σ_i^2 is a parameter defining the size of the influence region of unit i. σ_i^2 may be computed as the average empirical variance of the n input features, among the samples associated to unit i. If the activation $K(c)$ of the winning prototype is smaller than a specified threshold, the current sample is considered as invalid. The contributions of each of the components of vector x to the distance $\|x - m_c\|$ are then examined to determine more precisely which sensors should be declared as faulty. These sensor measurements are then disconnected to compute a new winning prototype with only valid parameters.

Remarks:

1. Note that this procedure for detecting outliers may be given a probabilistic interpretation, the input vectors x being assumed to be taken from a mixture of normal distributions (or classes) with means m_i and scalar covariance matrices $\sigma_i^2 I$ (I being the identity matrix). The activation of the winning unit i for input x may then be considered as an approximation to the i-th class-conditional probability density at x (up to a scaling factor). Our procedure for rejecting outliers then amounts to discarding feature vectors which were "implausible" under each of the class densities (see, e.g. [16, p. 25] on this topic). Other approaches to probability density estimation using, e.g., Gaussian mixture models and the EM algorithm [15], could also be used.
2. The above procedure allows for the rejection of atypical patterns and therefore implements some kind of "novelty detection". This constitutes a very conservative approach which prevents the prediction module of the system from blindly interpolating known relationships between water characteristics and coagulant dosage to previously unseen cases. It is clear however that the "atypicality" of rejected input vectors may have several causes, such as sensor failure, abrupt changes of water quality due to an accidental pollution, or merely lack of completeness of the training set. It is therefore necessary to store the rejected input patterns for subsequent interpretation by the user, and possible retraining of the system in case of undue rejection of "normal" patterns.

For reconstruction, each missing value of a given input variable is estimated by the value of the corresponding component of the winning prototype. In order to improve the reconstruction accuracy, a

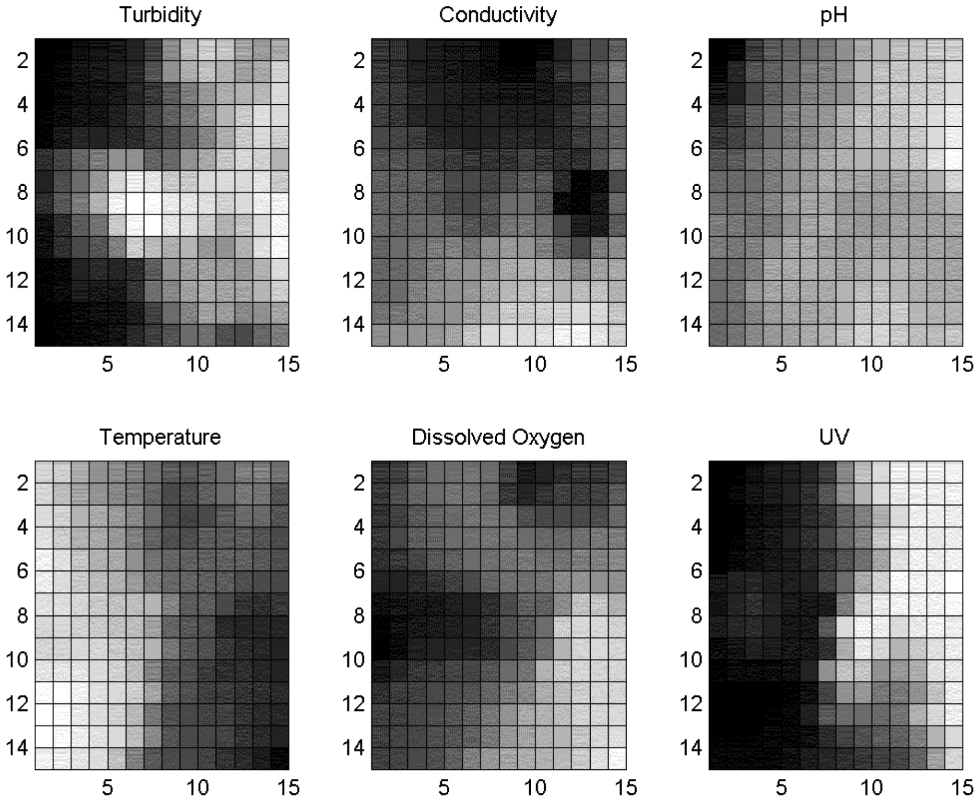

Fig. 7. Component planes of the SOM.

Table 1
Statistical summary of raw water conditions at the Viry treatment plant between Nov. 1997 and Nov. 1998

Property	Turbidity (NTU)	pH	Dissolved Oxygen (mg/l)	Conductivity (μS)	Temperature ($^\circ$C)	UV (do/m)	Coagulant Dose (ppm)
Maximum	166,5	8,37	14,3	560	25,3	19,2	7,8
Minimum	1,1	7,06	5,7	264	5,5	0,4	2
Mean	21,5	7,95	10,4	445	10,3	8,6	4,1
Standard deviation	28,4	0,2	1,8	48	5,5	4,9	1,7

combination of the k nearest nodes is used. Each missing or invalid value j is estimated by a combination of the corresponding component in the k nearest prototypes:

$$\hat{x}(j) = \frac{\sum_{i=1}^{k} K(i) m_i(j)}{\sum_{k=1}^{k} K(i)} \tag{4}$$

where $m_i(j)$ denotes component j of prototype i. Note that this method is similar in spirit to radial basis function networks, and to the fuzzy system approach described in [17].

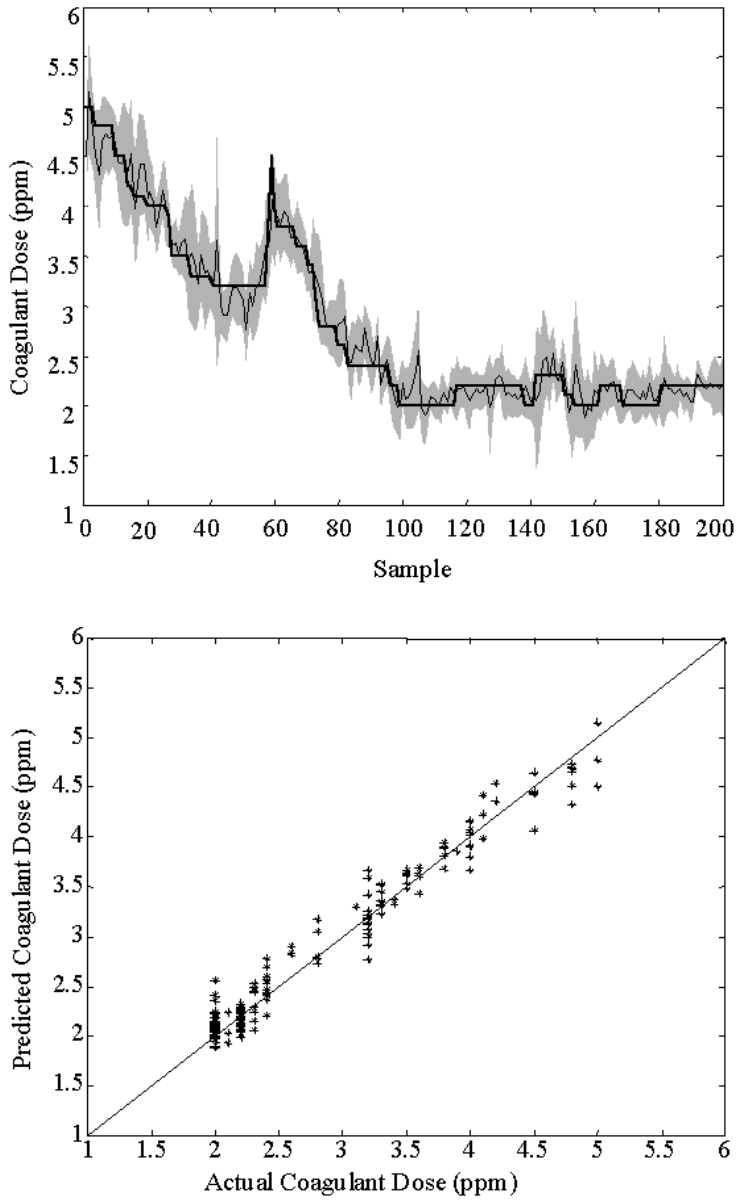

Fig. 8. Up: actual (thick line) versus predicted (thin line) coagulant dosage with ANN model on test data and confidence interval (shaded region). Down: predicted vs target coagulant dose.

4. Prediction of coagulant dosage

4.1. The model

The prediction of optimal coagulant dosage from water characteristics is a non linear regression problem which can be tackled using multilayer perceptrons (MLP). We used a conventional MLP with

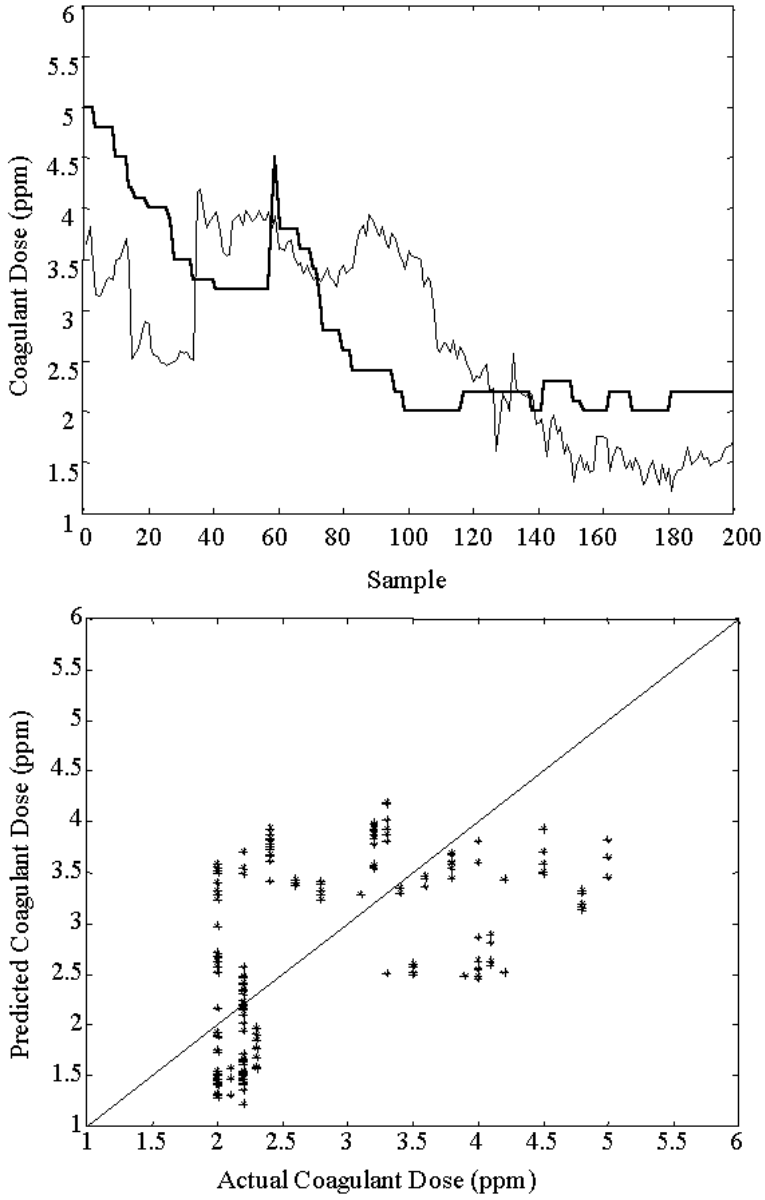

Fig. 9. Up: actual (thick line) versus predicted (thin line) coagulant dosage with linear regression model. Down: predicted vs target coagulant dose.

one hidden layer of sigmoidal units trained by minimisation of the mean squared error function (Fig. 4). This approach is known to provide estimates of the conditional average of the output variable (here, the optimal coagulant dosing rate), given the observed values of the input variables (the raw water quality parameters) [18].

For the determination of the architecture, a pruning approach was used, starting from a relatively large network and then removing connections in order to arrive at a suitable network size. Several approaches

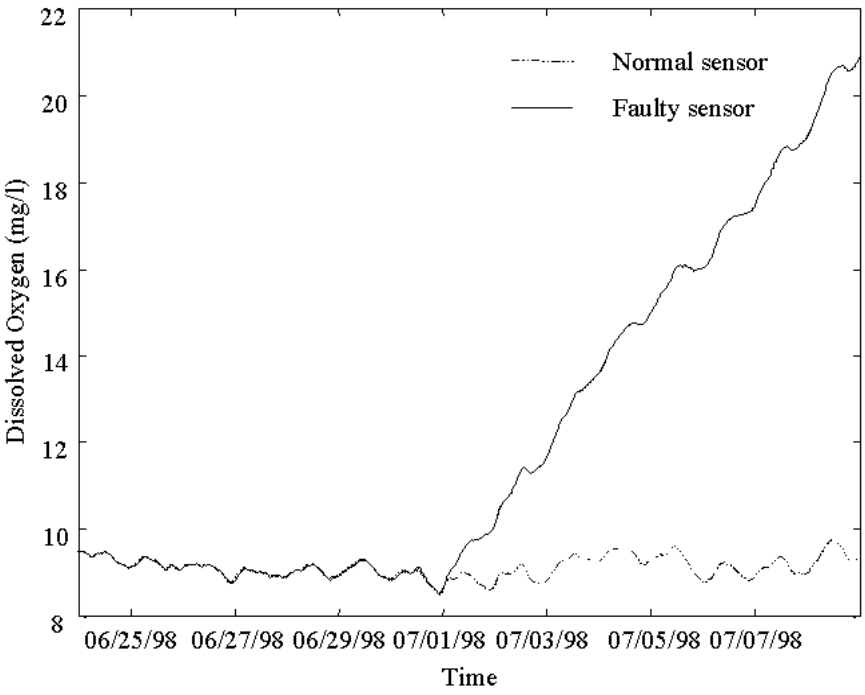

Fig. 10. Simulated dissolved oxygen sensor fault.

to network pruning are based on the following general procedure [18]. First a relatively large network is trained using one of the standard training algorithms. Then the network is examined to assess the relative importance of the weights, and the least important are deleted. Typically this is followed by some further training of the pruned network, and the procedure of pruning and training is repeated for several cycles. Clearly, there are various choices to be made concerning how much training is applied at each stage, which fraction of the weights are pruned, and so on. These choices are usually made on a heuristic basis. The most important consideration, however, is how to decide which weights should be removed. For that purpose, some measure of the relative importance, or saliency, of weights has to be defined. The Optimal Brain Damage (OBD) method [19] provides such a measure. This method is briefly recalled here.

The method is based on the computation of the change δE in the error function E due to small changes in the values of the weights [19]. If each weight w_i is changed to $w_i + \delta w_i$ the corresponding change in the error function E is given by

$$\delta E = \sum_i \frac{\partial E}{\partial w_i}\delta w_i + \frac{1}{2}\sum_i \sum_j H_{ij}\delta w_i \delta w_j + o(\delta w^3) \tag{5}$$

where the H_{ij} are the elements of the Hessian matrix

$$H_{ij} = \frac{\partial^2 E}{\partial w_i \partial w_j}. \tag{6}$$

If we assume that the training process has converged, then the first term in (5) vanishes. As proposed by Le Cun et al. [19], the Hessian matrix can be approximated by discarding the non-diagonal terms.

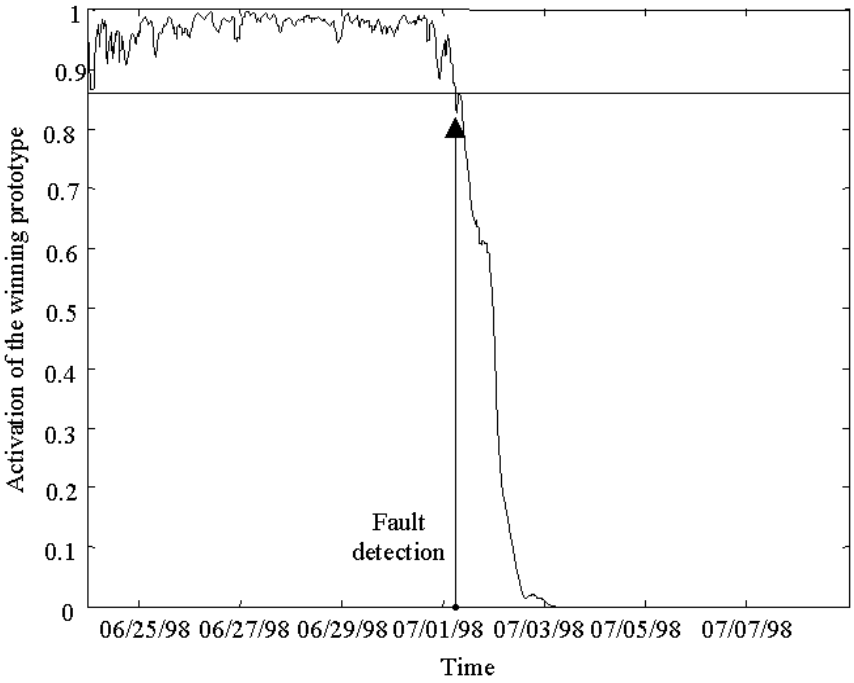

Fig. 11. Activation of the winning prototype.

Neglecting the higher-order terms in the expansion then reduces (5) to the form

$$\delta E = \frac{1}{2} \sum H_{ii} \delta w_i^2. \tag{7}$$

If a weight having an initial value w_i is set to zero, the increase in error is given approximately by (7) with $\delta w_i = -w_i$. The saliency values of the weights can therefore be defined as the quantities $H_{ii} w_i^2/2$.

The implementation of this technique consists in the following steps (Fig. 5):

1. Choose a relatively large initial network architecture.
2. Train the network using the back-propagation algorithm applied to the sum of squares error function.
3. Compute the second derivatives H_{ii} for each of weights, and hence evaluate the saliencies $H_{ii} w_i^2/2$.
4. Sort the weights by saliency and delete the q lowest saliency weights ($q = 1$ in our simulations).
5. Go to step 2 and repeat until the error measured using an independent validation set starts to increase.

4.2. Computation of confidence bounds

For practical use, a software system for the prediction the optimal coagulant amount should not only provide point estimates but also confidence intervals. Bootstrap sampling was used to generate confidence intervals for the system outputs [20]. The reader is referred to the book by Efron and Tibshirani [21] for a general presentation of the bootstrap.

Bootstrap is a statistical inference technique which uses training sets created by re-sampling with replacement from the original data (so that examples may occur more than once), and re-estimates

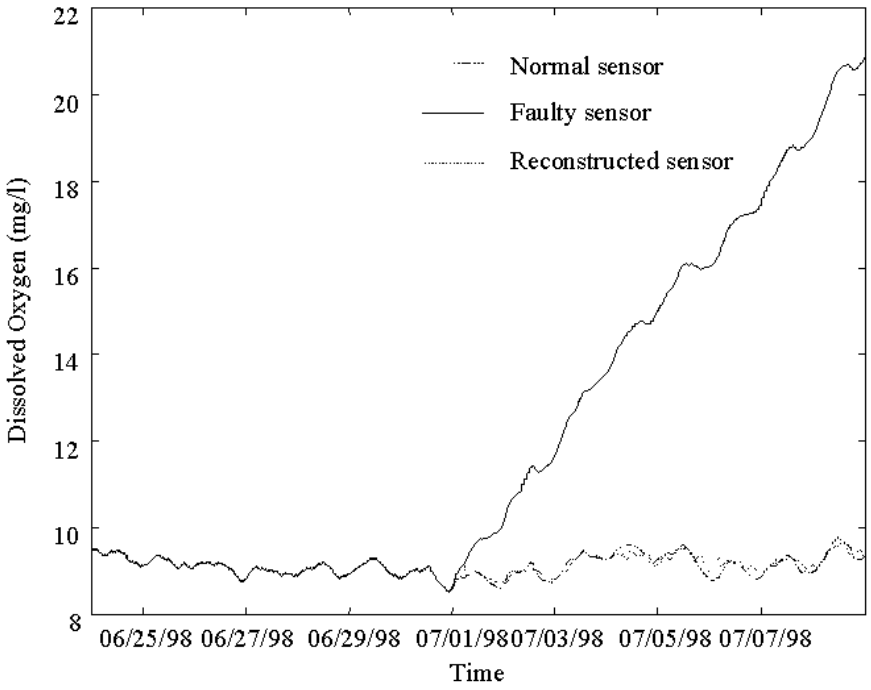

Fig. 12. Reconstruction of dissolved oxygen.

all the parameters on each bootstrap sample. The application of this technique to the computation of confidence bounds for the outputs of a neural network as proposed in Ref. [20] is illustrated in Fig. 6. In this approach, b bootstrap subsets of the initial training set are used to train b MLP models using the architecture and training procedure described previously. When a vector is fed into these networks, the b outputs provide an estimate of the distribution of the target variable for the current input. Lower and upper confidence limits for the prediction related to any given input vector may then be obtained by sorting these outputs and selecting, e.g., the 10% and 90% cumulative levels. As argued in [20], this approach allows to measure the variability due to the training algorithm and the limited training data. A theoretical discussion concerning the use of bootstrap percentile intervals may be found in [21, p. 170].

5. Results

The water treatment plant in Viry-Chatillon was used as an application site for this study. This plant provides water to more than 300,000 inhabitants. It has a nominal capacity to process 120,000 m^3 of water per day and has been well instrumented for several years. Several raw water parameters are measured on-line (Fig. 2), and jar-tests are performed quite frequently, providing the necessary target values (the optimal dosing rates to be estimated from water quality data).

The raw database consisted of 100,000 measurements of 6 input variables (turbidity, conductivity, pH, temperature, dissolved oxygen and UV absorption) sampled every 5 minutes during a period of 12 months (from November 1997 to November 1998). Simple descriptive statistics of the data are given in Table 1. Note that this data set, which represented all the available data at the beginning of this study,

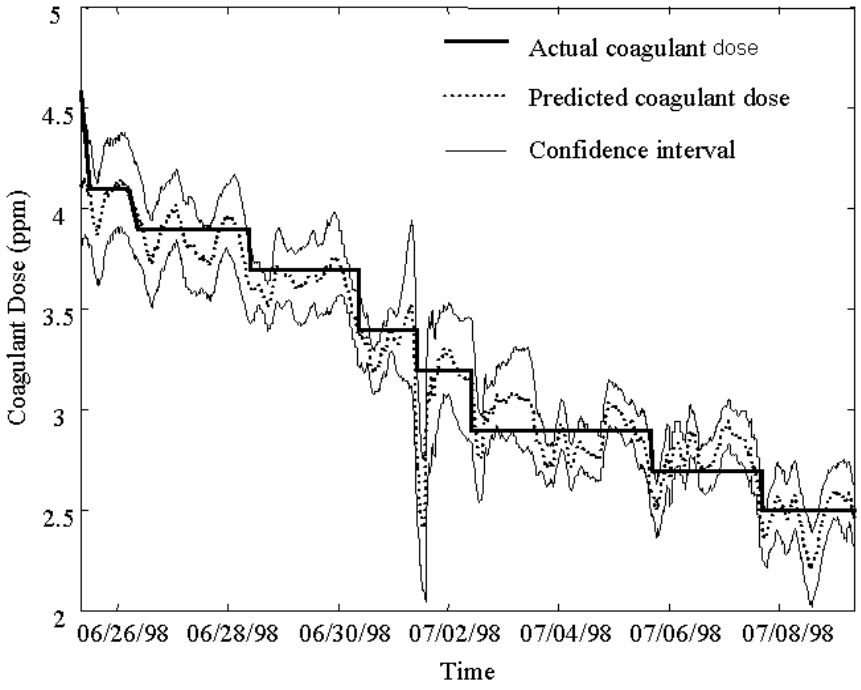

Fig. 13. Actual (thick line) versus predicted (doted line) coagulant dosage with ANN model and confidence bounds (thin lines).

covers a period of one year and so can be expected to account for seasonal variations of water quality. It is obvious, however, that periodic retraining of the system will be needed to ensure that the full range of water quality conditions (in the absence of dramatic pollution events) are accounted for, and to allow continuous adaptation of the system to any possible long-term evolution of water quality.

Component planes of a Kohonen map of size 15*15 trained on the whole data set are shown in Fig. 7. In this representation, individual components of reference vectors in the map are displayed as grey levels. Component planes are commonly used as a by-product of SOM's for "correlation hunting" in large data sets [22]. In our case, correlation between, e.g., turbidity and UV can be clearly seen on these displays. Such relationships between input variables are captured by the SOM and are exploited for the reconstruction of missing data.

For training the MLP's, a learning set of 1,600 complete measurement vectors was constructed by removing erroneous and incomplete data, and averaging the data over one hour time intervals. A total of 1,120 samples (about 70%) was exploited to build the model, the rest being used as an independent test set. Among the training data, approximately 30% was left out as a validation set for optimising the architecture. Training the model included: finding the best structure of the ANN and estimating the prediction accuracy by bootstrap. The prediction accuracy and confidence bounds computed on the validation set are shown in Fig. 8.

A linear regression model was also developed for comparison with the neural network model. Figure 9 shows the outputs of the linear model trained with the same data as the ANN. As expected, the prediction accuracy is clearly less than that of the ANN model, which confirms the non linearity of the relationship of interest. Performance results for the best model obtained from this data set with ANN model and linear regression are presented in Table 2.

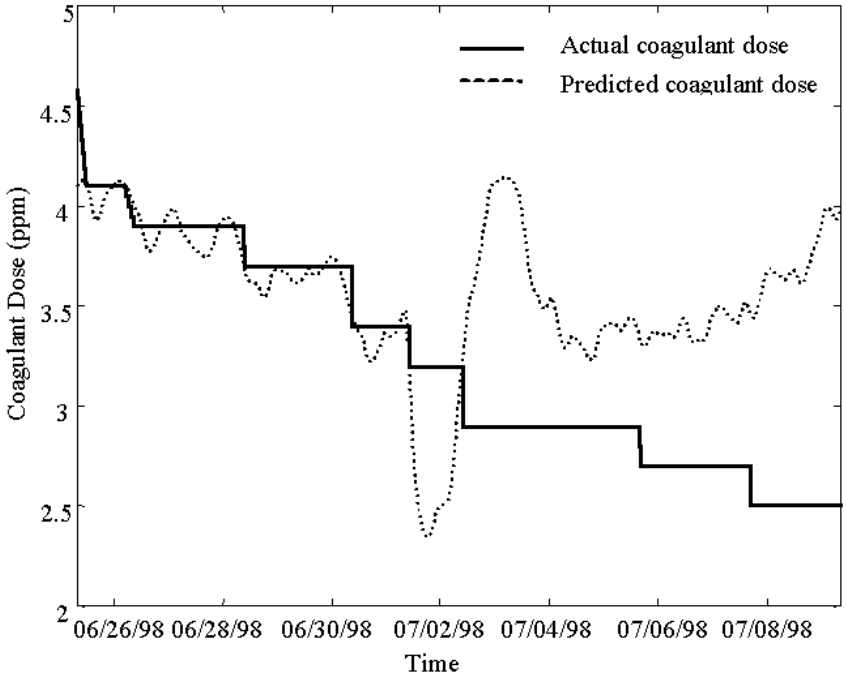

Fig. 14. Actual (thick line) versus predicted (doted line) coagulant dosage with ANN model without pre-processing.

Table 2
Results for ANN and linear regression models

Models	ANN model	Linear Regression model
R^2 for training data	0.99	0.78
R^2 for validation data	0.96	
R^2 for test data	0.96	0.58
Root-mean-square-error for test data	0.19	1.5
Mean absolute error for test data	0.14	1.2
Maximum absolute error for test data	0.92	4.4
Correlation coefficient r for test data	0.98	0.76

In order to assess the robustness of the system, an off-line simulation study was conducted by artificially introducing faults in the original data at certain time steps. The simulation was based on two weeks of real data sampled every 5 minutes from 24th June 1998 to 9th July 1998. The dissolved oxygen was simulated to be degraded with a rising ramp of 0.005 mg/l per samples (every 5 minutes). The faults occurs on the 1st July at 8:00 at sample 2017 as shown in Fig. 10. Using the SOM, the fault was detected 72 samples (6 hours) later at 1st July 14:00 (Fig. 11), and the dissolved oxygen variable was correctly identified as being the faulty parameter.

Figure 12 shows the reconstruction of dissolved oxygen using the SOM approach. The prediction accuracy and confidence interval of the ANN are shown in Fig. 13 for the pre-processed data. This is to be compared with the prediction results without pre-processing as shown in Fig. 14. These results clearly demonstrate the robustness induced by the pre-processing module in our system.

6. Conclusion

An integrated coagulation control system based on unsupervised and supervised neural network models has been described. It has been demonstrated that process data can be used to develop and train a feed-forward controller in the form of a neural network to accurately predict a suitable coagulant dosing rate. Experimental results using real data have demonstrated the efficiency and soundness of this approach. Field testing is currently under way to fully validate the system before its widespread dissemination to other sites. Expected benefits are treated water of a more consistently high quality, together with improved security of service, as the system will respond reliably and effectively over long periods. Significant savings in coagulant usage can be obtained in certain cases.

The performance of the network is obviously dependent on the quality and completeness of the data available for training the system. Consequently, continuous updating of training data during operational use is expected to improve the performance of the system. This model, however, is only based on the previous behaviour of operators and jar-test results. Further work is needed to develop a model taking into account the dynamics of the process, and allowing to predict treated water parameters (mainly, turbidity) at the output of the clarification process. The concepts demonstrated in this paper will also be applied in the near future to the modelling of other water treatment processes such as filtration and chlorination.

Acknowledgements

The authors thank the management of CIRSEE, and particularly Farrokh Fotoohi for his constant support, as well as the two anonymous referees for their helpful and constructive comments.

References

[1] M.H. Masson, S. Canu, Y. Grandvalet and A. Lyngaard-Jensen, Software sensor design based on empirical data, *Ecological Modelling* **120** (1999), 131–139.
[2] A. Mirsepassi, B. Cathers and H.B. Dharmappa, Application of Artificial Neural Networks to the Real Time Operation of Water Treatment Plants, in: *Proceedings of the International Conference on Neural Networks*, (Vol. 1), Perth, Australia, 1995, pp. 516–521.
[3] J. Evans, C. Enoch, M. Johnson and P. Williams, Intelligent based auto-coagulation control applied to a water treatment works, in: *Proceedings of International Conference on Control*, 1998, pp. 141–145.
[4] C. Lind, Coagulation Control and Optimization: Part One, *Public Works* (October 1994), 56–57.
[5] C. Lind, Coagulation Control and Optimization: Part Two, *Public Works* (November 1994), 32–33.
[6] F. Bernazeau, P. Pierrone and J.P. Duguet, Interest in using a streamline current detector for automatic coagulant dose control, *Water Supply* **10**(4) (1992), 87–96.
[7] K.S. Dentel, Use of streaming current detector in coagulation monitoring and control, *Journal of Water Sciences Research and Technologies – Aqua.* **44** (1995), 70–79.
[8] T.J. Böhme, C.S. Cox, N. Valentin and T. Denœux, Comparison of autoassociative neural networks and Kohonen maps for signal failure detection and reconstruction, in: *Intelligent Engineering Systems through Artificial Neural Networks 9*, C.H. Dagli et al., eds, ASME Press, New-York, 1999, pp. 637–644.
[9] H. Bourlard and Y. Kamp, Auto-association by multilayer perceptrons and singular value decomposition, *Biological Cybernetics* **59** (1988), 291–294.
[10] M.A. Kramer, Nonlinear Principal Component Analysis using Autoassociative Neural Networks, *AIChE Journal* **37**(2) (1991), 233–243.
[11] T. Kohonen, *Self-Organizing Maps*, Springer Verlag, Heidelberg, 1995.
[12] T. Trautmann and T. Denœux, A constructive algorithm for S.O.M. applied to water quality monitoring, in: *Intelligent Engineering Systems through artificial neural networks 4*, Dagli, Fernandez, Ghosh and Kumara, eds, ASME Press, New-York, 1994, pp. 17–22.

[13] T. Trautmann and T. Denœux, Comparison of dynamic feature map models for environmental monitoring, in: *Proceedings of ICNN'95,* (Vol. 1), Perth, Australia, 1995, pp. 73–78.
[14] B. Dubuisson and M. Masson, A statistical decision rule with incomplete knowledge about classes, *Pattern Recognition* **26**(1) (1993), 155–165.
[15] G.J. McLachlan and K.E. Basford, *Mixture Models: Inference and Applications to Clustering,* Marcel Dekker, New-York, 1988.
[16] B.D. Ripley, *Pattern Recognition and Neural Networks,* Cambridge University Press, Cambridge, 1996.
[17] S. Petit-Renaud and T. Denœux, A neuro-fuzzy system for missing data reconstruction, *Proc. IEEE Workshop on Emerging Technologies, Intelligent Measurement and Virtual systems for instrumentation and measurement,* Saint-Paul, USA, May 1998, pp. 78–84.
[18] C.M. Bishop, *Neural Network for Pattern Recognition,* Oxford University Press, Oxford, 1995.
[19] Y. Le Cun, J.S. Denker and S.A. Solla, Optimal Brain Damage, in: *Advances in Neural Information Processing Systems 2,* D.S. Touretzky, ed., Morgan Kaufman, San Mateo, CA, 1990, pp. 598–605.
[20] R.P. Lippmann, L. Kukolich and D. Shahian, Predicting of Complications in Coronary Artery Bypass Operations using Neural Networks, in: *Advances in Neural Information Processing System 7,* G. Tesauro et al., eds, MIT Press, Menlo Park, CA, 1995, pp. 1055–1062.
[21] B. Efron and R.J. Tibshirani, *An Introduction to the Bootstrap,* Chapmann & Hall, New York, 1993.
[22] J. Vesanto, SOM-based data visualization methods, *Intelligent Data Analysis* **3**(2) (1999), 111–126.