



Forecasting of traffic origin NO and NO₂ concentrations by Support Vector Machines and neural networks using Principal Component Analysis

István Juhos^a, László Makra^{b,*}, Balázs Tóth^a

^a Department of Computer Algorithms and Artificial Intelligence, University of Szeged, H-6701 Szeged, P.O. Box 652, Hungary

^b Department of Climatology and Landscape Ecology, University of Szeged, H-6701 Szeged, P.O. Box 653, Hungary

ARTICLE INFO

Article history:

Received 16 August 2007

Received in revised form 7 August 2008

Accepted 12 August 2008

Available online 20 August 2008

Keywords:

Artificial neural networks
Multi-Layer Perceptron
Support Vector Machines
Support vector regression
Principal Component Analysis
Dimension reduction
Forecast

ABSTRACT

The main aim of this paper is to predict NO and NO₂ concentrations four days in advance comparing two artificial intelligence learning methods, namely, Multi-Layer Perceptron and Support Vector Machines on two kinds of spatial embedding of the temporal time series. Hourly values of NO and NO₂ concentrations, as well as meteorological variables were recorded in a cross-road monitoring station with heavy traffic in Szeged in order to build a model for predicting NO and NO₂ concentrations several hours in advance. The prediction of NO and NO₂ concentrations was performed partly on the basis of their past values, and partly on the basis of temperature, humidity and wind speed data. Since NO can be predicted more accurately, its values were considered primarily when forecasting NO₂. Time series prediction can be interpreted in a way that is suitable for artificial intelligence learning. Two effective learning methods, namely, Multi-Layer Perceptron and Support Vector Regression are used to provide efficient non-linear models for NO and NO₂ times series predictions. Multi-Layer Perceptron is widely used to predict these time series, but Support Vector Regression has not yet been applied for predicting NO and NO₂ concentrations. Grid search is applied to select the best parameters for the learners. To get rid of the curse of dimensionality of the spatial embedding of the time series Principal Component Analysis is taken to reduce the dimension of the embedded data. Three commonly used linear algorithms were considered as references: one-day persistence, average of several-day persistence and linear regression. Based on the good results of the average of several-day persistence, a prediction scheme was introduced, which forms weighted averages instead of simple ones. The optimization of these weights was performed with linear regression in linear case and with the learning methods mentioned in non-linear case. Concerning the NO predictions, the non-linear learning methods give significantly better predictions than the reference linear methods. In the case of NO₂ the improvement of the prediction is considerable; however, it is less notable than for NO.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

Nitric oxide (NO), as one of the nitrogen oxides (NO_x), is a highly reactive gas. Human activity has drastically increased the production of nitric oxide by traffic. It is produced by the chemical union of O₂ and N₂ in the cylinders of internal combustion engines. (However, the catalytic converter in automobile exhaust systems reduces air pollution by oxidizing hydrocarbons to CO₂ and H₂O and, to a lesser extent, converting nitrogen oxides to N₂ and O₂.) Nitric oxide plays a major role in photochemical reactions, which lead, among other things, to the formation of nitrogen dioxide (NO₂) and photochemical smog. Since

* Corresponding author. Tel.: +36 62 544 856; fax: +36 62 544 624.

E-mail addresses: paper@juhos.info (I. Juhos), makra@geo.u-szeged.hu (L. Makra), toth.balazs@yahoo.co.uk (B. Tóth).

NO₂ absorbs in the visible wavelength region – creating brown cloud over megacities (e.g. Mexico City and Beijing) – and can be photolysed and yield oxygen atoms that can react with molecular oxygen to create ozone, NO₂ and the NO/NO₂ ratio is important in tropospheric chemistry. Nitrogen dioxide is formed primarily from burning fuel in motor vehicles, power plants, and other industrial, commercial, and residential sources that burn fossil fuels. Nitrogen oxides, reacting with other substances in the air, form acid rain, accelerate the corrosion of buildings and monuments, and reduce visibility.

Exposure to nitrogen dioxide can irritate the lungs and may lower resistance to respiratory infections. Sensitivity increases for people with asthma and bronchitis. NO₂ is also a major source of fine particulate pollution, which is a significant health concern.

Due to the harmful effects of these pollutants on human health, it is important to have reliable methods enabling the prediction of their concentrations several hours in advance, so that the public authorities could avoid the harmful consequences of severe air pollution episodes.

The more accurate prediction of future values of a time series will improve performance in each field of everyday life. The classical statistical procedures, as well as neural network have already been applied for short-term prediction of air pollutants by several authors.

Artificial neural networks appear as useful alternatives to traditional statistical modelling techniques in many scientific disciplines. They are composed of a large number of possible non-linear functions (neurons) each with several parameters that are fitted to data through a computationally intensive training process. Some statisticians and forecasters, who prefer the statistical approach to forecasting, may disregard the performance of neural networks because of their lack of rigorous statistical foundation. However, neural networks do fit comfortably with the heterogeneous background of alternative forecasting techniques.

Gardner and Dorling [1] present wide fields of recent applications of the Multi-Layer Perceptron as one type of artificial neural network in the atmospheric sciences. They applied MLP neural networks to model hourly NO_x and NO₂ concentrations in Central London from basic hourly meteorological data [2]. The results of the models perform well when compared to those received by using regression-based models. They also demonstrated that MLP neural networks offer several advantages over traditional multivariate linear regression models. Jorquera et al. [3] developed an accurate forecast of ozone episodic days for downtown Santiago, Chile. The simple model structure included a combination of persistence and daily maximum air temperature as input variables. The model was validated by comparing the outcome of three different modelling schemes: linear model, fuzzy models and neural networks. The three forecasts developed present significant improvement of successful forecasts compared with pure persistence. Predictions of PM_{2.5} [4], as well as NO and NO₂ [5], plus SO₂ concentrations [6] were compared, produced by three different methods: persistence, linear regression and Multi-Layer Perceptron neural networks. Furthermore, Perez and Reyes [7] improved PM_{2.5} predictions several hours in advance with a type of neural network which was equivalent to a linear regression. The effect of meteorological conditions was included by using real values of temperature (*T*), relative humidity (*H*) and wind speed (*W*) at the time of the intended prediction as inputs to the different models. It was revealed that a three-layer neural network gave the best results to predict concentrations of the pollutants in the atmosphere of downtown Santiago, Chile several hours in advance, when hourly concentrations of the previous day were used as input. A multivariate regression model is also used by [8] for comparison with the results obtained by using the neural network model. Their results indicate that the neural network is able to give better predictions with less residual mean square error than those given by multivariate regression models. Mechaqrane and Zouak [9] compared a linear model with MLP, when predicting indoor temperature of a building. Maqsood et al. [10] used data of temperature, wind speed and relative humidity to train and test seven different models for weather forecasting. With each model, they made 24-h ahead forecasts for all seasons. In comparison, they found the ensemble of neural networks to produce the most accurate forecasts.

Agirre-Basurko et al. [11] developed two Multi-Layer Perceptron-based models and one multiple linear regression-based model. The models utilised traffic variables, meteorological variables and O₃ and NO₂ hourly levels as input data. The performances of these three models were compared with persistence of levels and the observed values. The results indicated improved performance for the Multi-Layer Perceptron-based models over the multiple linear regression model, if they considered predictions for more than 3 hours in advance. Hansen et al. [12], by using neural network techniques, achieved impressive increases in forecasting accuracy. According to their results, genetic-algorithms-guided selections of neural network architectures displayed distinct improvements over statistical refinements and other heuristic architectures. Additionally, they concluded that neural networks could improve forecasting performance dramatically and found structure in data which remained hidden to other techniques. According to the investigations of Small and Tse [13], an artificial neural network is particularly well suited to modelling chaotic time series data. Castillo and Melin [14] also use neural networks for simulation and forecasting economic time series. The performance of the neural networks (MLPs) was compared with classical regression models. Time series prediction gave the best result when neural networks were used compared to that by the other regression models. Kukkonen et al. [15] evaluated five neural network models (MLPs), a linear statistical model and a deterministic modelling system for the prediction of urban NO₂ and PM₁₀ concentrations. They found that the non-linear neural network models performed slightly better in terms of the model performance values than the deterministic model. Furthermore, the results also showed improved performance for most of the neural network models, compared with the linear statistical model, both for predicting NO₂ and PM₁₀ concentrations. Ordieres et al. [16] compared three different topologies of neural networks to two classical models: a persistence model and a linear regression. The results clearly demonstrated that the neural approach not only outperformed the classical models but also showed fairly similar values among different topologies.

Besides the good non-linear regression abilities of neural networks, they also have some drawbacks. During the neural network optimization process we have to move on a surface, having many local optima. Neural network's learning/optimizing algorithms cannot avoid from being stuck in a local optimum, which can lead to a sub-optimal solution. Another important deficiency is that the structure of a network is not given in advance; therefore, we have to optimize it as well. It means that we have to decide how many neurons and hidden layers would be necessary and what kind of activation function or functions would be appropriate and how to connect neurons with each other to form a network. Fortunately, an MLP with two hidden layers can approximate an arbitrary function [17–19], which negates some drawbacks mentioned above; however, the others remain unsolved.

Support Vector Regression addresses these limitations and gives promising results [20]. The basic idea, which is behind the SVR technique, is to start with linear regression, which has no parameters or only a few, is able to control the possible hypotheses (capacity) by considering the width of the margin of the regression plane. It would be useful to extend this technique in order to hold the almost parameterless property and the capacity control of the possible hypotheses, as well. An MLP, which is also an extension of the linear regression technique, chooses an explicit way for a non-linear description. It takes several linear regression methods and non-linearizes them by an activation function to get building blocks of the model. Then, it connects these blocks together to make a comprehensive non-linear model. This approach makes it hard to control the good properties of the regression mentioned above. Applying implicit mapping via a kernel function, an SVR redefines the dot product in the linear regression method, in order to get a linear regression-like method. Due to the implicit non-linear mapping, this regression becomes non-linear. Thus, the simple and good properties of the linear regression are inherited. The implicit mapping provides an implicit description instead of the neural network's explicit one, which expresses the model required with an explicitly defined composite function. While SVR solves several drawbacks of the MLP (optimizing many parameters, choosing topology; being stuck in a local optimum [21]), it brings a new problem; namely, choosing an appropriate kernel function [22]. Nevertheless, this problem can be solved easily by choosing a proper kernel function from a small set. We have to maintain some additional parameters; namely, capacity control and parameters of the chosen kernel function [23]. Several papers suggest that SVR performs well in many time series prediction problems [24–29].

Both MLPs and SVRs have several successful applications in the field of prediction (see above). Therefore, they are readily chosen for predicting NO and NO₂ series. When a Gaussian kernel is used in an SVR, it corresponds to a radial basis MLP with Gaussian radial basis functions ('rbf') and one hidden layer. While the size of the hidden layer is unknown in the MLP approaches, the SVR automatically sets it. The size, i.e. the number of hidden neurons, is obtained as a result of the SVR optimization procedure. Hidden neurons and support vectors correspond to each other. Thus, the centre problem of the radial basis network is solved, since the support vectors serve the centres of the basis functions. Considering this fact, radial basis network is not used in this paper. Instead of the radial basis function, a sigmoid activation function was used in our MLPs. The use of sigmoid-like functions is popular in the practice for both MLP and SVR [30,31]. The application of the sigmoid function is even more justified, since training algorithms of these neural networks do not require positive definite property of the activation function, as opposed to an SVR, which needs positive definiteness for its kernel function [21]. In case of SVR the parameters of the sigmoid function were set so that the function has positive definite property [32]. It allows us to compare the results of the application of the sigmoid-like function by two learning techniques.

In spite of the fact that neural networks and SVR are different learning techniques, the learnt hypotheses can be the same [33]. This is also a reason for comparing these methods.

The aim of this paper is to predict hourly averages of NO and NO₂ concentrations in a traffic junction in Szeged downtown, which has heavy traffic especially in rush hours. Furthermore, the methods mentioned above are compared with the reference ones. Reference methods do not have any parameters; nevertheless, our MLP and SVR methods have some. They were set by preliminary tests on the historical data by grid search in the parameter space of the algorithms considering the suggestions of the program libraries used. Generally, these suggestions lead to good enough results, which are proved by our experiments as well. Nevertheless, parameter tuning by grid search leads to further considerable improvements, but our applied dimension reduction can bring additional improvements.

Our aim is to combine these effective techniques with different spatial embedding schemes, which have not yet been discussed together in comprehensive form in the literature. Then, we can see which schemes and which regression techniques perform well on the subject and what kind of benefits can be achieved using the different combinations.

2. Geographical, topographical, climatological and air quality characteristics of Szeged

2.1. The geographical position and topographical characteristics of Szeged

Szeged, as the largest town in SE Hungary (20°06'E; 46°15'N), is located at the confluence of the Tisza and Maros Rivers characterized by a landscape of extensive flats and an elevation of 79 m a.s.l. (Fig. 1). The built-up area covers a region of about 46 km² with about 165,000 inhabitants.

Szeged and its surroundings are not only characterized by extensive lowlands but this city also has the lowest elevation value not only in Hungary but in the Carpathian Basin as well, rendering it a so-called "basin in the basin" or "double basin" situation. This special situation favours the development of stronger anticyclonic currents, enabling higher concentrations of pollutants in the air.

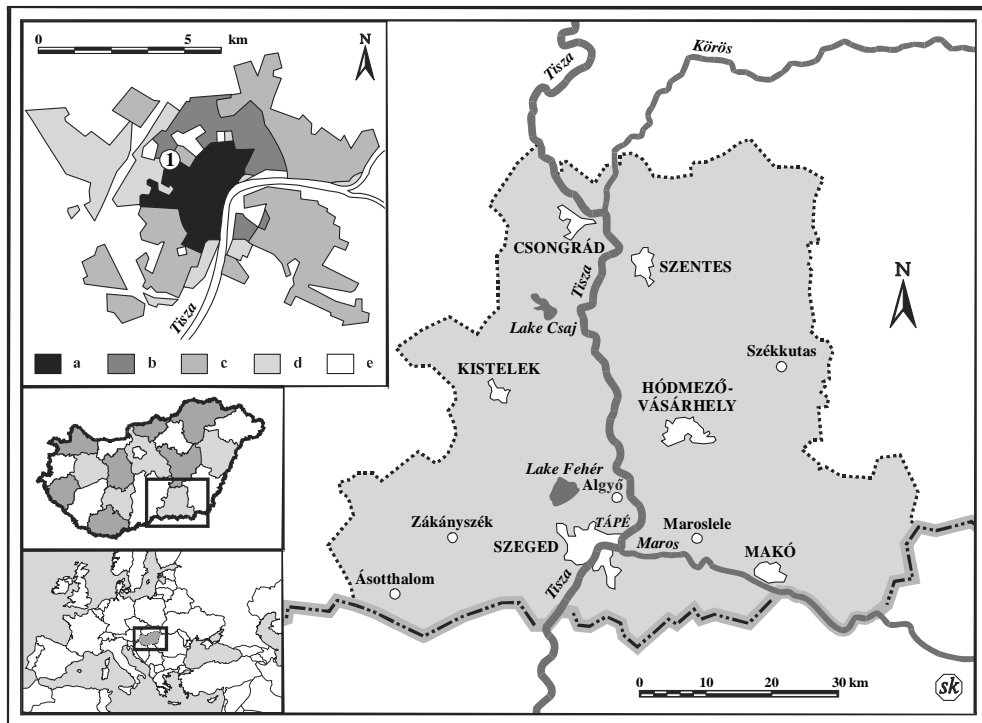


Fig. 1. The geographical position of Szeged, Hungary and built-up types of the city (left, up) [(a) city centre (2–4-storey buildings); (b) housing estates with prefabricated concrete slabs (5–10-storey buildings); (c) detached houses (1–2-storey buildings); (d) industrial areas; (e) green areas; (1) monitoring station].

2.2. The climatic conditions of Szeged

The climate of Szeged is characterised by hot summers and moderately cold winters. The distribution of rainfall is fairly uniform during the year: with a share of 29% and 19% for the summer (JJA) and the winter (DJF) seasons, respectively. Mean daily summer temperatures are around 22.4 °C, while the mean daily winter temperatures are 2.3 °C. The irradiance values also exhibit large-scale variances; with an average of 20.2 and 4.2 MJ m⁻² in summer and winter days, respectively. The most frequent winds blow along the NNW–SSE axis, with prevailing air currents arriving from NNW (42.3%) and SSW (24%) in the summer and from SSE (32.6%) and NNW (30.8%) during the winter. Due to its unique geographical position, Szeged is characterised by relatively low wind speeds with average daily summer and winter values of 2.8 and 3.5 m s⁻¹, respectively. The highest hourly wind speeds have been recorded during the spring with a rate of 5 m s⁻¹ [34].

2.3. The air quality conditions of Szeged

Urban air quality largely depends on the actual measured values of meteorological parameters. The recorded averages of these variables for the city of Szeged are the following: annual mean temperature: 11.2 °C; mean January and July temperatures: –1.2 °C and 22.4 °C, respectively; relative humidity: 71%; mean annual precipitation total: 573 mm; mean annual sunshine duration: 2102 h; and mean annual wind speed: 3.2 m s⁻¹.

The city structure is very simple, characterized by an intertwined network of boulevards, avenues and streets sectioned by the River Tisza (Fig. 1). However, this simplicity largely contributes to the concentration of traffic, as well as air pollution within the urban areas.

The industrial area is mainly restricted to the north-western part of the city. Thus, the prevailing westerly and northerly winds tend to carry the pollutants deriving from this area towards the centre of the city.

3. Data basis

3.1. Local meteorological and air pollution data basis

The data come from the monitoring station, which is located in Szeged downtown in a cross-road with heavy traffic (Fig. 1). The station is operated by the ATKÖFE (Environmental Protection Inspectorate of Lower-Tisza Region, Branch of

the Ministry of Environment). The database comprises the period between 12 March 1998 and 16 March 2001, including autumns, as well as winters of the years considered. Hourly average mass concentrations of NO and NO₂ (in µg m⁻³) as well as hourly means of temperature (°C), relative humidity (%) and wind speed (m s⁻¹) are considered for the period indicated. Sampling of the time series was half-an-hour, the hourly averages come from this sampled series where regularization of the series was performed, that is extreme and missing values were eliminated by simple averaging of their neighbours. On the other hand, prediction occurred three times using the threefold cross-validation scheme on the whole data set. Namely, two years were selected from the whole three years for training and the remaining year was used up for prediction. Result errors of each of the three predictions were averaged.

4. Time series prediction

In the time series prediction, the aim is to predict the value of a variable that varies in time using previous values and/or other variables. Typically, the variable is continuous, so time series prediction is usually a specialized form of regression. Several authors [35] transformed the temporal dimension of a time series into a spatial vector of the *l* dimension embedding space by taking a moving window over the last *l* elements of the series. We can define two kinds of forecasts: a) the first, when no other information is used apart from the time series being examined (that is, predicting without external variables); and b) the second, when other information is also available, (that is, predicting with external variables).

4.1. Prediction without external variables

Suppose we have a real valued time series {*y_t*}_{*t*=1}^{*n*}, i.e. the historical values of the series. When other time series, which can affect the *y* series are not given, the task is to predict *y_{n+k}* values with *k* > 0, based on the historical values. In other words:

$$y_{n+k} = h(y_n, y_{n-1}, \dots, y_{n-(l-1)}), \quad k, l > 0 \tag{1}$$

The usual way to make the prediction is to find an appropriate *l* and a function *h*, which describes the relation between *l* consecutive elements and the next element of the series. Here, *l* denotes the historical window size and *k* represents the horizon of the future.

Our first aim is to use the above prediction schemes to predict NO and NO₂ time series. Secondly, external influences (external variables) are considered to improve the results obtained, if possible.

4.2. Prediction with external variables

In some cases other time series are also known, which can influence the *y*-series under examination. They are called the external variables. If several factors are available, we can represent them as a vector series consisting of more than one scalar time series; namely, {*z_t*}_{*t*=1}^{*n*} = {*z_{t1}*, ..., *z_{tm}*}_{*t*=1}^{*n*}. We can include this information in Eq. (1):

$$y_{n+k} = h(\hat{z}_{n+k}, y_n, y_{n-1}, \dots, y_{n-(l-1)}), \quad k, l > 0 \tag{2}$$

Now suppose that the (*n* + *k*)th value of the *z*-series or its estimate is known from some source, which will be denoted by *z_{n+k}*. Unfortunately, the problem of obtaining information about *z_{n+k}* is similar to that for Eq. (1). Hence, the *z*-series prediction should be made before *y_n* prediction. When the aim is to predict only the future value of *y_{n+k}*, the prediction is called a one-step-ahead prediction. However, if we intend to estimate values beyond *y_{n+k}*, we have to use the previously predicted values in the *h* function and we call this a many-step-ahead prediction, especially with *N*-step-ahead predictions, where the intention is to forecast the next *N* values.

5. Inductive learning

The inductive learning of a concept requires recognizing a hypothesis for this concept after presenting training instances, which is supervised by a defined classification. The instances are generally given in the following format during the learning/training process:

$$\underbrace{\{x_{i1}, x_{i2}, \dots, x_{il}\}}_{\text{attributes}} \underbrace{\{y_i\}}_{\text{class}}, \quad i \in N \tag{3}$$

*i*th instance

In order to seek a relation (the hypothesis) between the attributes and their classes, a *h* function based on the training instances has to be approximated:

$$y_i = h(x_{i1}, x_{i2}, \dots, x_{il}) \tag{4}$$

The above formula is suitable for the prediction schemes in Chapter 4 if we make an appropriate replacement in the arguments of the *h* function. The *i*th instance means the *i*th data window, while *y_i* is the value to be predicted. From here on, this more general notation will be used. The Artificial Intelligence (AI) learners applied were Multi-Layer Perceptron and Support Vector Regression, both of them having good approximation characteristics [18,36].

5.1. Multi-Layer Perceptron (MLP)

The two-layered MLP is capable of approximating arbitrary finite sets of real numbers [18]. Hence, a maximum two hidden-layered MLP was used with a sigmoid activation function:

$$\sigma(x) = \frac{1}{1 + e^{-x}} \tag{5}$$

where the input and output layers have linear units. When the l attributes of the i th learning instance take the form x_{i1}, \dots, x_{il} , then the class of this instance produced by a perceptron is shown in Eq. (6), while the result of the one and two hidden-layered MLP is described in Eq. (7). Eq. (6) is a special one hidden-layered MLP called Perceptron, which is the building block of the Multi-Layer Perceptron

$$o_i^{1s} = \sigma \left(\sum_{t=1}^{l_1} w_t^{1s} x_{it} + w_{bias}^{1s} \right) \tag{6}$$

$$y_i = \sum_{r=1}^{l_2} w_r o_i^{1s}$$

$$y_i = \sum_{r=1}^{l_2} w_r \sigma \left(\sum_{s=1}^{l_1} w_s^{2r} o_i^{1s} + w_{bias}^{2r} \right) \tag{7}$$

o_i^{1s} is the output of the s th Perceptron for the i th instance; w_t^{1s} , w_t^{2r} , w_r are the weights in the first and second layers and output unit; w_{bias}^{1s} , w_{bias}^{2r} are the biases in the first and second layers; l_1 , l_2 are the numbers of Perceptrons in the first and second layers; and σ is the sigmoid activation function.

Changing the weights is the basis of the learning process. The well-known backpropagation method with momentum is used for adjusting the weights during the training process. An implementation was provided by the Weka software library [37,38].

5.2. Support vector regression (SVR)

There are two commonly used Support Vector Machines for regression; namely, the ε -SVR algorithm and its extension the ν -SVR algorithm [25]. We chose the ν -SVR because it has an advantage compared to ε -SVR. Namely, it is able to adjust automatically the width of the ε -tube around the function being approximated. An SVR maps the $\mathbf{x}_i = (x_{i1}, \dots, x_{il})$ attributes to a generally higher dimension space, called the feature space via a $\phi: R^l \rightarrow R^L$, $L \geq l$ map function. Then it makes a linear fit to certain accuracy by optimizing the weights $\mathbf{w} = (w_1, \dots, w_L)$ and w_{bias} :

$$y_i \approx \sum_{j=1}^L w_j \phi(x_{ij}) + w_{bias} \quad (= \langle \mathbf{w}, \phi(\mathbf{x}_i) \rangle + w_{bias}) \tag{8}$$

We can reformulate Eq. (8). by expanding the weight vector as a linear combination of the instance vectors ($\mathbf{w} = \sum_{t=1}^n (\alpha_t^* - \alpha_t) \phi(\mathbf{x}_t)$, $\alpha_t^*, \alpha_t \geq 0$):

$$\sum_{t=1}^n (\alpha_t^* - \alpha_t) \langle \phi(\mathbf{x}_t), \phi(\mathbf{x}_i) \rangle + w_{bias} \quad \left(= \sum_{t=1}^n (\alpha_t^* - \alpha_t) \kappa(x_t, x_i) + w_{bias} \right) \tag{9}$$

where κ is a kernel function belonging to the ϕ mapping. To obtain α_t^* , α_t ; ν -SVR maximizes the following quadratic problem for $C, \nu > 0$:

$$-\frac{1}{2} \sum_{ij=1}^n (\alpha_i^* - \alpha_i) \kappa(x_i, x_j) (\alpha_j^* - \alpha_j) + \sum_{i=1}^n (\alpha_i^* - \alpha_i) y_i \tag{10}$$

subject to the constraints

$$\sum_{i=1}^n (\alpha_i^* - \alpha_i) = 0$$

$$\sum_{i=1}^n (\alpha_i^* + \alpha_i) \leq C \nu$$

$$0 \leq \alpha_i^*, \quad \alpha_i \leq C$$

Three kinds of well-known kernel functions were employed: (a) a radial basis function (rbf) $\kappa(\mathbf{x}_i, \mathbf{x}_j) = e^{-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2}$, (b) a polynomial one $\kappa(\mathbf{x}_i, \mathbf{x}_j) = -\gamma \langle \mathbf{x}_i, \mathbf{x}_j \rangle^d$ and (c) a sigmoid-like function, namely the hyperbolic tangent, where $\tanh(x) = (e^x - e^{-x}) / (e^x + e^{-x})$. A ν -SVR implementation was provided by the LibSVM software library [39,40].

5.3. Dimension reduction with Principal Component Analysis (PCA)

Handling the high dimensional embedding space with too few past instances can cause problem for the learning, namely one that is known as the curse of dimensionality [43]. In addition, there is usually noise in the data. A good way of dealing with it can be a reduction of the dimension of the data from l to Γ ($\Gamma < l$). From here on we will assume that the training attributes are centred $\sum_{i=1}^n x_i = 0$ as this is needed for the dimension reduction algorithm discussed below.

PCA is an optimal linear dimension reduction technique in the mean square sense [44]. After dimension reduction we have to deal with less attributes, hence the learning requires less computational effort and the data usually has reduced noise and is free of the curse of dimensionality. However, dimension reduction can bring several good properties, we can lose information during the process, which can have opposite influence, i.e. worse results than expected. In spite of this unpredictable outcome, PCA regularize our data and usually brings enough good results in the worse cases. The basic idea in PCA is to find the components $s_{i1}, s_{i2}, \dots, s_{i\Gamma}$, $i = 1, \dots, n$ so that they have the maximum amount of variance with Γ linearly transformed components. PCA can be defined in an intuitive way using a recursive formulation. Now let us define the direction of the first principal component by

$$v_1 = \arg \max_{\|v\|=1} E\{|\langle x_i, v \rangle|^2\}_{i=1}^n \quad (11)$$

where x_i -s are the population of the training attributes and E denotes the mean. Thus the first principal component is the projection in the direction in which the variance is maximized. Having determined the first $k - 1$ ($1 < k < l$) principal components, the k th principal component is found by calculating the principal component of the residual:

$$v_k = \arg \max_{\|v\|=1} E\left\{\left|\left\langle v, \left(x_i - \sum_{j=1}^{k-1} v_j \langle x_i, v_j \rangle\right)\right\rangle\right|^2\right\}_{i=1}^n \quad (12)$$

The principal components are then given by $s_{ij} = \langle x_i, v_j \rangle$. We have l components of the data, but many of them usually have a small variance. Depending on the variance, we keep the first Γ components which have a significant variance. In practice, the computation of v_i can be accomplished using the covariance matrix $C = E\{x_i x_i^T\}_{i=1}^n$ of the training attributes, assuming that x_i -s are column vectors. The v_i are the eigenvectors of C that correspond to the n largest eigenvalues of C .

5.4. Model selection by grid search (GS)

Reference model, i.e. linear regression and persistence, have no parameters to tune, therefore they are a good choice for making a basis of these experiments. However, MLP and ν -SVR are sophisticated techniques, they suffer from the parameter selection problem, which has large influence on the results. Both models have several parameters which have default values provided by the libraries used, but for the sake of good prediction these parameter values need modification according to the underlying distribution of the historical data. Among several parameter estimation techniques the grid search is the most reliable because it makes exhaustive search in the parameter space. Of course, only a subspace of the whole can be discovered due to the huge amount of computational efforts. LibSVM and Weka libraries have built in model selections using grid search technique, hence we relied on them. We compared prediction performance with and without model selection, i.e. using standard library settings of parameters and setting parameters by grid search on the “subspace”, more exactly on a grid of the possible parameter values.

In case of the MLP, for calculating the optimal number of neurons in the layers we used a grid with interval [1, 24] of integers for one coordinate and interval [0, 24] of integers for the other, respectively, to the number of neurons in the first and second layer of the MLP, where back propagation algorithm was used as training. The grid which was used in the optimization of the parameters of the ν -SVR was the following: ν parameter from 0.01 to 0.99 with initial steps of 0.05 and the γ parameter from 0.1 to 8.0 with initial steps of 1.0, where further steps were taken by the usual two factor exponential enlargements. In fact the grid search is extended by two additional parameters: the number of principal components kept and the possible kernels of the ν -SVR. The number of principal components which was kept in the dimension reduction were participated in the grid search as well where the grid is extended by the proportions of the variance (0.5–0.9 by 0.1 steps) of the retaining principal components. Predictions were done separately by rbf, poly2, poly3, sigmoid kernels and the best result of them was selected.

6. Experiments

In our experiments real meteorological data were used, which came from a monitoring station located in the city of Szeged in Hungary. The time series examined were 1-h averages of nitric oxide and nitrogen dioxide, which consisted of a 3-year data set in the period 12 March 1998–16 March 2001. Series were normalized according to get zero mean and unit variance before taking them for spatial embedding.

The diurnal cycles of NO and NO₂ have the shape of a double wave (Fig. 2), with bigger amplitudes for NO than for NO₂. Due to the traffic density, the concentration of NO is relatively higher on weekdays, than on weekends (Fig. 4). This effect can also be observed for the secondary substance NO₂ (Fig. 4). The average diurnal variations on weekdays are higher for NO than

for NO_2 , because NO_2 has a longer lifespan than the more reactive NO (Fig. 2). Generally, NO concentrations are higher in the morning, than in the evening (Fig. 2). This can be explained by the fact that in the morning the rush hour is shorter, and the atmosphere near the surface is more stable than in the evening. The low NO concentrations early in the afternoon result mainly from the reduction of O_3 by NO [41]. NO_2 concentration depends on that of NO ; hence, concentration of the latter pollutant is very useful to predict NO_2 levels (Fig. 2).

Our experiments showed that concentrations of NO are more precisely predictable than NO_2 . Therefore, the NO series was also used as an external variable to predict NO_2 [5,2]. Temperature (T), relative humidity (H) and wind velocity (W) as external variables were employed by several authors [11,2,15,5,42]. Thus, values of these meteorological variables might be included as inputs to an algorithm in order to improve the forecast of NO and NO_2 concentrations.

Prediction, as it was mentioned earlier, occurred three times using the threefold cross-validation scheme on the whole data sets. Namely, two years were selected from the whole three years for training and the remaining year was used up for prediction. Results of each (three) variations were averaged. One-step-ahead prediction was applied for each experiment (Fig. 2). Weekend data, because of the less traffic, were omitted from the database. Similar assumptions can be found in the following papers: [11,2,15,5,42]. The time series received have natural 24-hour periods, confirmed by their autocorrelation diagram (Fig. 3).

When external variables were applied in the experiments, their real values were used in order to avoid cumulative errors. Moreover, it is important to examine the exact relations between the time series considered and the external variables mentioned.

The performance of the mentioned AI methods was compared with three commonly used reference algorithms. The first reference algorithm is called Linear Regression (LR), where the past values of the data were weighted to result prediction. The next reference algorithm is the Persistence, which models the persistence of the values of the days. A prediction value of a future hour has the same value as that for the same hour of the previous day. This simple technique works well on this problem because the series has a 24-hour periodicity (Fig. 4). Average values of several past days can lead to a better persistence method; namely, to the Averages of Several-Day Persistence (Persistence Avg.). Consequently, the Persistence and the Persistence Avg. are not able to handle external variables. Agirre-Basurko et al. [11], Gardner and Dorling [2], Kukkonen




Fig. 2. Real values of the NO and NO_2 series at the forecasting (test) term.




Fig. 3. Autocorrelation of the NO (left) and NO_2 series (right).

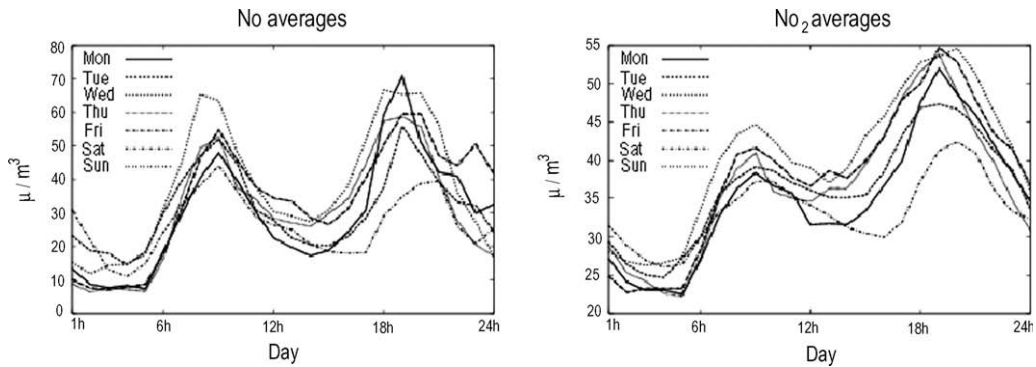


Fig. 4. Hourly average values of NO (left) and NO₂ (right) in the historical period.

et al. [15] and Perez and Trier [5] showed that all the reference methods mentioned work well on NO and NO₂ predictions. They also revealed that these reference methods can be as good as an MLP or, in certain cases, even better.

Our experiments showed that the Persistence Avg. performed well, namely, better than the Persistence itself, especially in case of prediction NO-series. Based on the good results of Persistence Avg., a prediction scheme was introduced, which forms not equally weighted averages instead of simple Persistence Avg. Optimization of these weights was performed by Linear Regression in linear case and by the learning methods mentioned in non-linear case. Another difference between the Persistence Avg. and the schemes introduced (Schemes 3 and 4 in Table 1 and Schemes 7 and 8 in Table 2) is that 2–10-day looking back periods were employed instead of the whole historical days averages. For example, in case of a 3-day looking back period it is $NO(t + k) = h(NO(t), NO(t - k), NO(t - 2k)) = h(NO[t, t - k, t - 2k])$, $k = 24$. These schemes can be applied to the NO and NO₂ series predictions. However, the weights were adjusted by optimization learning weights.

Different kinds of MLP learners were used according to the number of neurons in the hidden layers. They are denoted as MLP_{l_1, l_2} where l_1 and l_2 means the number of neurons in the first and second hidden layers, respectively. Beside the MLPs, different kinds of ν -SVRs were trained by the mentioned kernel types for each hour of a day. They provide learner-specific hypotheses, which will be denoted by MLP_{l_1, l_2} and ν -SVR_{rbf, poly2, poly3, sigmoid} (degree is shown in the subscripts in case of the polynomial kernel, e.g. degree is 2 when *poly2* appears). Linear Regression, Persistence and Persistence Avg. have no parameters but MLP and ν -SVR have some, which need to be set properly. Standard library settings (LS) were chosen for each in order to analyse their standard behaviour on the data. In the case of MLP, these settings were as follows: learning rate was 0.3, momentum was 0.2, and the number of training epochs was 500. Weka software library [37] suggests $MLP_{1,0}$, $MLP_{a,1}$, $MLP_{a,a/2}$ to use, where “*a*” is the number of attributes. A scheme was chosen for the best model of these three different models, which was denoted by MLP (LS). Due to the different library suggestions MLP (LS) can varied from scheme to scheme. This is the case for the ν -SVR (LS) as well, where the four kinds of kernels (rbf, poly2, poly3, sigmoid) gave alternatives; however, in fact, rbf and poly2 performed good enough on the problems. For other parameters of the ν -SVR learner the following settings were chosen: γ is 1/*a*, ν is 0.5, and ϵ is 0.1. When grid search tuned the parameters of the learners, then abbreviation (GS) was used to define the presence of automatic parameter selection process.

First, an integrated result is shown characterizing the performance. The Normalized Root Mean Squared Error (see NRMSE in Eq. (13)) gives rough but important integrated information about the performance. Normalization of the error is important to compare the predictions of the time series obtained in different places or times.

$$NRMSE = \frac{\sqrt{\sum_{i=1}^n \frac{(y_i - \hat{y}_i)^2}{n}}}{stdev(y_i)} \tag{13}$$

where \hat{y}_i is the estimation of the y_i -series and n is the length of the series, while *stdev* is the abbreviation of the standard deviation. Beside NRMSE, index-of-agreement (see IA in Eq. (14)) is also calculated, which verifies the results from another aspect, where not only the difference-prediction error, but the differences of the prediction and the real value from a certain basis are also considered.

$$IA = 1 - \frac{\sum_i |y_i - \hat{y}_i|^2}{\sum_i (|\hat{y}_i - \bar{y}| + |y_i - \bar{y}|)^2} \tag{14}$$

where \bar{y}_i is the mean of the series, that is 0 due to the normalization process. It has been recommended by [45] as a valuable parameter to describe the ‘agreement’ between the two time series.

In order to get a detailed comparison of the predictions, the daily averages of the hourly absolute errors ($|y_i - \hat{y}_i|$) were also analyzed in the forecast period. Thus, we can define those hours of the day, where a method performs well.

Since the series have 24-hour periods, it is expedient to choose the size of the embedding dimension, i.e., the window size for a prediction scheme 24 (e.g. $NO(t + k) = h(NO[t, t - 1, \dots, t - 23])$, $k = 1, \dots, 24$).

Table 1

NRMSE and IA averages of the one-step-ahead predictions of the NO series come from three-fold cross-validation on the 3-year dataset, according to different prediction schemes, using reference model (italic letters), and learning models with selected library settings (LS) and optimized settings by grid search (GS), moreover beside this (PCA) denotes its usage

NO	Scheme 1 (S1) NO($t+k$) = NO[$t, t-1, \dots, t-23$] $k = 1, \dots, 24$ each hour of the previous day NRMSE/IA	Scheme 2 (S2) NO($t+k$) = NO[$t, t-1, \dots, t-23$], $W(t+k)$, $k = 1, \dots, 24$ each hour of the previous day and a known factor for the prediction time NRMSE/IA	Scheme 3 (S3) NO($t+k$) = NO[$t, t-k$] $k = 24$ two- day step back for the same hour NRMSE/IA	Scheme 4 (S4) NO[$t, t-k$] $k = 24$ two-day step back for the same hour and a known factor for the prediction time NRMSE/IA
<i>Lin. regression</i>	0.149/0.558	0.172/0.529	0.134/0.523	0.133/0.554
<i>Persistence</i>	0.145/0.301	0.145/0.301	0.145/0.301	0.145/0.301
<i>Persistence Avg.</i>	0.162/0.598	0.162/0.598	0.162/0.598	0.162/0.598
MLP ^a (LS) ⁴	0.244/0.419	0.250/0.441	0.140/0.534	0.140/0.571
MLP ^b (GS)	0.134/0.597	0.144/0.615	0.125/0.563	0.114/0.641
MLP ^c (PCA, GS)	0.132/0.605	0.130/0.607	0.124/0.537	0.115/0.622
v-SVR ^d (LS)	0.134/0.541	0.132/0.558	0.135/0.495	0.134/0.512
v-SVR ^e (GS)	0.125/0.584	0.122/0.558	0.111/0.590	0.109/0.637
v-SVR ^f (PCA, GS)	0.126/0.586	0.121/0.563	0.113/0.591	0.110/0.642

The best is emphasized by bold letters in a scheme, where performance indicators NRMSE and IA were taken account in equal weight in the goodness decision. One-step-ahead predictions mean 24-hour predictions in case of Schemes 3 and 4, because the forecast horizons are constant 24-hour between the last known historical value and the predicted future value. Schemes 1 and 2 use dynamic range of forecast horizons, since these schemes use the values of the whole day to predict the following day values, hence the one-step-ahead prediction means 1-h if the first hour of the next day is predicted and 24-hour if the 24th hour of the next day is predicted. Meaning of the symbols: NO – Nitrogen oxide; W – Wind; H – Humidity; T – Temperature; t – current time; k – horizon.

^a Number of neurons in the hidden layers selected from the predefined LS-s.

^b Number of neurons in the hidden layers were optimized for each hour by GS.

^c Number of neurons in the hidden layers were optimized for each hour by GS.

^d Selected kernels from rbf, poly2, poly3, sigmoid by performance in accordance of the schemes: rbf, rbf, rbf, rbf.

^e Selected kernels from rbf, poly2, poly3, sigmoid by performance in accordance of the schemes: rbf, rbf, rbf, rbf.

^f Selected kernels from rbf, poly2, poly3, sigmoid by performance in accordance of the schemes: rbf, rbf, rbf, rbf.

Table 2

NRMSE and IA averages of the one-step-ahead predictions of the NO₂ series come from three-fold cross-validation on the 3-year dataset, according to different prediction schemes, using reference model (italic letters), and learning models with selected library settings (LS) and optimized settings by grid search (GS), moreover beside this (PCA) denotes its usage

NO ₂	Scheme 5 (S5) NO ₂ (t + k) = NO ₂ [t, t - 1, ..., t - 23] k = 1, ..., 24 each hour of the previous day	Scheme 6 (S6) NO ₂ (t + k) = NO ₂ [t, t - k] k = 24 two-day step back for the same hour	Scheme 7 (S7) NO ₂ (t + k) = NO ₂ [t, t - k], H(t + k), T(t + k), W(t + k) k = 24 two-day step back for the same hour and known factors for the prediction time	Scheme 8 (S8) NO ₂ (t + k) = NO ₂ [t, t - k], NO(t + k) H(t + k), T(t + k), W(t + k) k = 24 two-day step back for the same hour and known factors for the prediction time
	NRMSE/IA	NRMSE/IA	NRMSE/IA	NRMSE/IA
<i>Lin. regression</i>	0.178/0.667	0.181/0.607	0.175/0.665	0.147/0.834
<i>Persistence</i>	0.212/0.690	0.212/0.690	0.212/0.690	0.212/0.690
<i>Persistence Avg.</i>	0.210/0.233	0.210/0.233	0.210/0.233	0.210/0.233
MLP ^a (LS)	0.395/0.436	0.195/0.546	0.216/0.600	0.154/0.842
MLP ^b (GS)	0.260/0.563	0.165/0.592	0.157/0.684	0.126/0.873
MLP ^c (PCA, GS)	0.204/0.529	0.169/0.548	0.151/0.697	0.117/0.833
ν-SVR ^d (LS)	0.174/0.665	0.183/0.593	0.177/0.644	0.141/0.829
ν-SVR ^e (GS)	0.155/0.776	0.162/0.601	0.152/0.710	0.110/0.889
ν-SVR ^f (PCA, GS)	0.146/0.787	0.159/0.612	0.154/0.724	0.108/0.897

The best is emphasized by bold letters in a scheme, where performance indicators NRMSE and IA were taken account in equal weight in the goodness decision. One-step-ahead prediction is 24-h in case of Schemes 6–8, but dynamic 1–24-hour in Scheme 5. Detailed description is found in the caption of Table 1. Meaning of the symbols: NO₂ – Nitrogen dioxide; W – wind; H – humidity; T – temperature; t – current time; k – horizon.

^a Number of neurons in the hidden layers selected from the predefined LS-s.

^b Number of neurons in the hidden layers were optimized for each hour by GS.

^c Number of neurons in the hidden layers were optimized for each hour by GS.

^d Selected kernels from rbf, poly2, poly3, sigmoid by performance in accordance of the schemes: rbf,rbf,rbf,rbf.

^e Selected kernels from rbf, poly2, poly3, sigmoid by performance in accordance of the schemes: rbf,rbf,rbf,rbf.

^f Selected kernels from rbf, poly2, poly3, sigmoid by performance in accordance of the schemes: rbf,rbf,rbf,rbf.

6.1. NO series prediction

We found that the Schemes 3 type predictions gave the best results if we used only two-day looking back period in the past, where looking back trials were applied from 2 to 10 days. Based on Schemes 1 and 3 (factorless predictions), investigations were made by extending these schemes with external variables: wind velocity, humidity and temperature; first, one of them at a time, then all of them together. Results showed that only wind velocity could improve the accuracy of the predictions (Schemes 2 and 4 in Table 1). The other factors seem to have less significance to improve the results of Schemes 1 and 3. However, the factorless results are very impressive. The use of wind velocity brings general improvements in several cases. It is to be noted that if external variables are used, they also need to be predicted. Thus, the results received might be worse, due to the cumulative prediction error.

The SVR predictions with the appropriate kernel function performed significantly better than the others. SVR with Gaussian (rbf) kernel seems to be an efficient predicting tool with the schemes provided.

Schemes 3 and 4, which look back more than 1 day in the past, generally bring additional improvements without using external variables, and show better results in several cases when wind velocity was used. It is clear that this factor has some, probably non-linear, influences on the NO series (Table 1, Figs. 5 and 6).

The SVR error curves show smoother and more reliable predictions in average than the MLP, which produced several peaks (Figs. 5 and 6). However the external variable, the wind speed, makes improvements in the general results, but is not able to regularize the error curves of the MLP (see Fig. 5). We can remark that application of PCA method could bring significant improvements in the MLP results (see Table 1).

On one hand, reference methods gave good results; while, on the other, the applied SVR technique outperformed the results of the reference methods.

Furthermore, SVR and Linear Regression showed their best results using only the wind velocity factor. They could not improve their results using the three factors together with Scheme 1. While Linear Regression follows the persistence curve, the MLP and SVR reduced significantly the prediction error (Fig. 5). Moreover, Schemes 3 and 4 make the results of the MLP more unreliable (Fig. 6).




Fig. 5. Average prediction errors of the four prediction days of the NO series. The curves are related to the best results in Table 1. The best MLP and v-SVR results are compared to the best reference result by the Scheme 1 (left) and Scheme 2 (right) predictions.

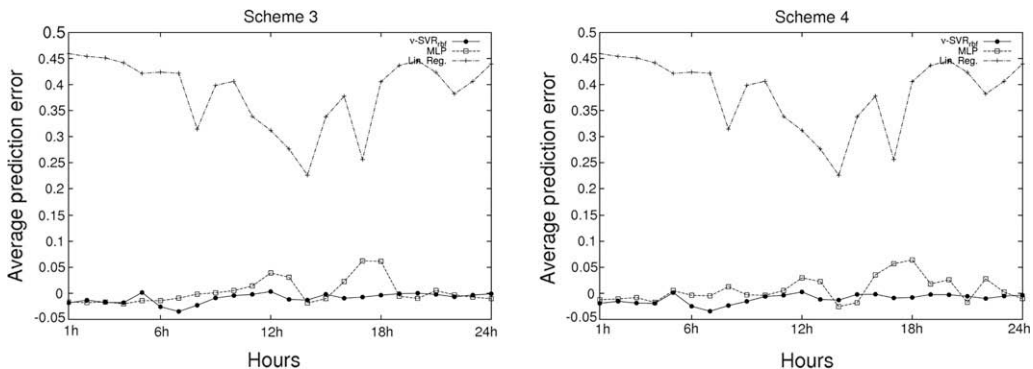


Fig. 6. Average prediction errors of the four prediction days of the NO series. The curves are related to the best results in Table 1. The best MLP and v-SVR results are compared to the best reference result by the Scheme 3 (left) and Scheme 4 (right) predictions.

6.2. NO₂ series prediction

Perez and Trier [5], furthermore Gardner and Dorling [2] suggested using NO series as an external variable for predicting NO₂ series. Relations of the two series are shown in Fig. 2. Perez and Trier [5] showed that the NO series can be predicted with more accuracy than the NO₂ series. Our results confirmed this statement, since there are worse results in Table 2 than in Table 1.

SVR with Gaussian (rbf) kernels gave again the best results in schemes (Table 2) as it has been displayed for the NO predictions (Table 1). Application of the PCA brings better SVR prediction for the NO₂ series using Scheme 5, as it could have been expected, because there are not so many attributes in the other schemes. PCA results are left behind our expectation in both the NO and NO₂ prediction; however, it can bring sometimes really significant improvements, but we could not produce general improvements with it; hence, its application is scheme and method dependent.

The extension of Scheme 5 with the variants of the factors of humidity, temperature, wind velocity and NO brings additional improvements.

The two-day looking back period scheme (Scheme 6) brought the best results among the 2–10-day looking back schemes, again. The same phenomenon was experienced for the NO series prediction. In addition, the above mentioned *H*, *W*, *T* and NO external variables brought significant improvements for Scheme 6 (Scheme 7 and 8 in Table 2), which is due to the presence of the NO series as external influence. PCA considerably helps the factorless prediction schemes; so that the factorless Scheme 5 can outperform some results in the factorhelped Scheme 8 in case of SVR. It should be noted that, in fact, the NO series and the other factors cannot support the prediction of its original values, since NO₂ prediction using NO data series comprises cumulative prediction errors. Hence, worse results are expected for Scheme 8.

Linear regression, Persistences and the MLP produce non-reliable predictions compared to the SVR smooth error graphs (Figs. 6 and 7). However, the non-linear MLP shows better results than the references, the SVR wins again all. Nevertheless, the MLP sometimes produces even worse results than the best reference method. MLP shows its best results in the early morning hours in the NO prediction (which is in agreement with those of Perez and Trier [5]). Nevertheless, this behaviour is probably data dependent. SVR produces smooth error curves with smallest errors during the whole day; while MLP gives less reliable estimations especially for the late afternoon hours. Both NO and NO₂ prediction show that, the average

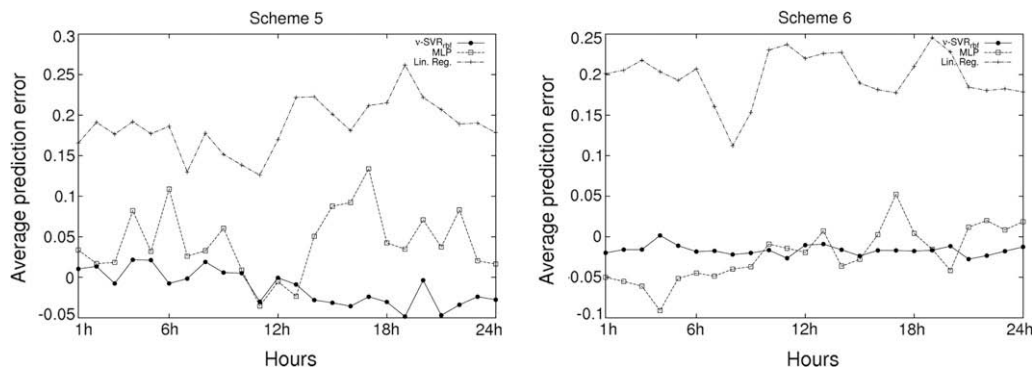


Fig. 7. Average prediction errors of the four prediction days of the NO₂ series. The curves are related to the best results in Table 1. The best and results are compared to the best reference result by the Scheme 5 (left) and Scheme 6 (right) predictions.

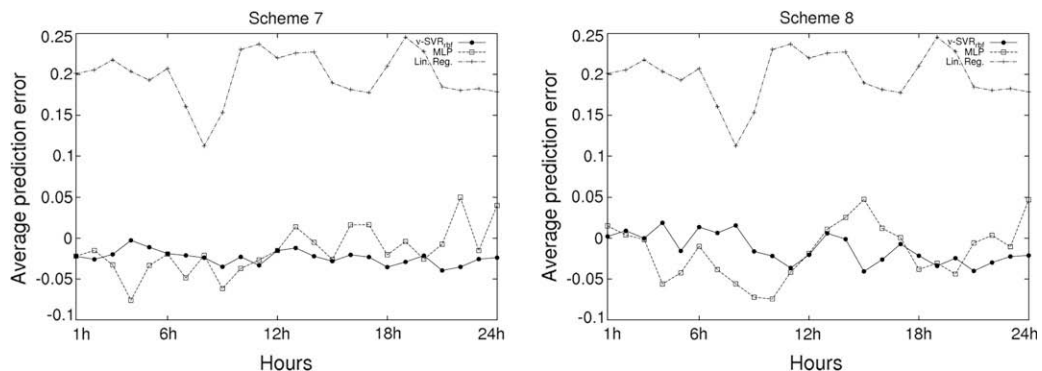


Fig. 8. Average prediction errors of the four prediction days of the NO₂ series. The curves are related to the best results in Table 1. The best MLP and v-SVR results are compared to the best reference result by the Scheme 7 (left) and Scheme 8 (right) predictions.

persistence follows enough good the real future curve according to its index-of-agreement values. However, SVR provides the most precise prediction according to NRMSE, but except of the precision MLP found better sometimes in the following of the curve indexes of the future values defined by the index-of-agreement for the NO₂. Therefore combination of them can lead to a hybrid method where not only the precision but the index of the curve can be considered also to help better a decision process according to a multi-objective forecast from several aspects (See Fig. 8).

7. Conclusions

The experiments clearly showed that the applied forecasting techniques could perform well on the prediction of NO and NO₂ concentrations. Forecasting these air pollutants is difficult because their concentrations fluctuate widely and depend on several factors. In many cases, the three reference algorithms proved to be successful at predicting the future values of the time series examined. These results are in accordance with those of Perez and Trier [5] and Kukkonen et al. [15]. Averages of several-day persistence performed well. In order to profit from this good performance, schemes based on this kind of persistence were introduced.

MLP and SVR using Grid Search can significantly improve the results of the best reference algorithms in both NRMSE and IA in the case of NO and NO₂ predictions, using each of Schemes 1–8. PCA brings solid improvements in the experiments, but its use is worse especially for NO₂ prediction. Both MLP and SVR produced equally good predictions; however, SVR gives more reliable forecast.

It was found that the applied $[t, t - 1, \dots, t - 23]$ scheme, where the previous day values were considered, was successful for both NO and NO₂ predictions. However, use of the several-day persistence motivated $[t, t - 24, t - 48, \dots]$ scheme (where several previous days were considered) provided better results.

External factors improved the factorless predictions significantly; however, their prediction can bring cumulative errors which were not considered in this study.

Undoubtedly, the application of machine learning techniques mentioned above can be relatively simple and is worth using.

There are several possibilities for future work. We can transform spatial embeddings of the historical values of NO and NO₂ into a lower dimensional space by non-linear methods in order to have possibility to exploit better the influence of the external factors and reduce the learning time, making more accurate and faster predictions. Furthermore, we can make a hybrid method using the applied prediction algorithms together. Pre-processing of data (e.g. other smoothing or de-noising techniques) can lead to more predictable structures. These methods are easily adaptable for forecasting other air pollutants.

Acknowledgements

The authors are indebted to Gábor Motika (Environmental Protection Inspectorate of Lower-Tisza Region, Szeged, Hungary) for handing monitoring data of the meteorological parameters and the air pollutants and Zoltán Sümegegy (Department of Climatology and Landscape Ecology, University of Szeged, Hungary) for digital mapping, as well as Rita Béczi (Department of Climatology and Landscape Ecology, University of Szeged, Hungary) for useful suggestions. This study was supported by the EU-6 Project “QUANTIFY” [No. 003893 (GOCE)] and the High Performance Computing Group of the University of Szeged.

References

- [1] M.W. Gardner, S.R. Dorling, Artificial neural networks (the multi-layer perceptron) – a review of applications in the atmospheric sciences, *Atmos. Environ.* 32 (1998) 2627–2636.
- [2] M.W. Gardner, S.R. Dorling, Neural network modelling and prediction of hourly NO_x and NO₂ concentrations in urban air in London, *Atmos. Environ.* 33 (1999) 709–719.
- [3] H. Jorquera, R. Pérez, A. Cipriano, A. Espejo, M.V. Letelier, G. Acuña, Forecasting ozone daily maximum levels at Santiago, Chile, *Atmos. Environ.* 32 (1998) 3415–3424.
- [4] P. Perez, A. Trier, J. Reyes, Prediction of PM_{2.5} concentrations several hours in advance using neural networks in Santiago, Chile, *Atmos. Environ.* 34 (2000) 1189–1196.
- [5] P. Perez, A. Trier, Prediction of NO and NO₂ concentrations near a street with heavy traffic in Santiago, Chile, *Atmos. Environ.* 35 (2001) 1783–1789.
- [6] P. Perez, Prediction of sulfur dioxide concentrations at a site near downtown Santiago, Chile, *Atmos. Environ.* 35 (2001) 4929–4935.
- [7] P. Perez, J. Reyes, Prediction of particulate air pollution using neural techniques, *Neural Comput. Appl.* 10 (2) (2001) 165–171.
- [8] A.B. Chelani, R.C.V. Chalapati, K.M. Phadke, M.Z. Hasan, Prediction of sulphur dioxide concentration using artificial neural networks, *Environ. Modell. Softw.* 17 (2) (2002) 159–166.
- [9] A. Mechaqrane, M. Zouak, A comparison of linear and neural network ARX models applied to a prediction of the indoor temperature of a building, *Neural Comput. Appl.* 13 (1) (2004) 32–37.
- [10] I. Maqsood, M. Riaz Khan, A. Ajith Abraham, An ensemble of neural networks for weather forecasting, *Neural Comput. Appl.* 13 (2) (2004) 112–122.
- [11] E. Agirre-Basurko, G. Ibarra-Berastegib, I. Madariaga, Regression and multilayer perceptron-based models to forecast hourly O₃ and NO₂ levels in the Bilbao area, *Environ. Modell. Softw.* 21 (4) (2006) 430–446.
- [12] J.V. Hansen, J.B. McDonald, R.D. Nelson, Time series prediction with genetic-algorithm designed neural networks: an empirical comparison with modern statistical models, *Comput. Intell.* 15 (1999) 171–184.
- [13] M. Small, C.K. Tse, Minimum description length neural networks for time series prediction, *Phys. Rev. E* 66 (2002) 066701-1–066701-12.
- [14] O. Castillo, P. Melin, Hybrid intelligent systems for time series prediction using neural networks. Fuzzy logic and fractal theory, *IEEE Trans. Neural Networks* 13 (2002) 1395–1408.
- [15] J. Kukkonen, L. Partanen, A. Karpainen, J. Ruuskanen, H. Junninen, M. Kolehmainen, H. Niska, S. Dorling, T. Chatterton, R. Foxall, G. Cawley, Extensive evaluation of neural network models for the prediction of NO₂ and PM₁₀ concentrations, compared with a deterministic modelling system and measurements in central Helsinki, *Atmos. Environ.* 37 (2003) 4539–4550.

- [16] J.B. Ordieres, E.P. Vergara, R.S. Capuz, R.E. Salazar, Neural network prediction model for fine particulate matter (PM_{2.5}) on the US–Mexico border in El Paso (Texas) and Ciudad Juárez (Chihuahua), *Environ. Modell. Softw.* 20 (5) (2005) 547–559.
- [17] E. Blum, L. Li, Approximation theory and feedforward networks, *Neural Networks* 4 (1991) 511–515.
- [18] D. Chester, Why two hidden layers are better than one, in: L. Erlbaum (Ed.), *International Joint Conference on Neural Networks, Proceedings*, vol. 1, Washington, DC, 1990, pp. 265–268.
- [19] K. Hornik, Approximation capabilities of multilayer feedforward networks, *Neural Networks* 4 (2) (1991) 251–257.
- [20] H. Drucker, C.J.C. Burges, L. Kaufman, A. Smola, V. Vapnik, Support vector regression machines, in: M. Mozer et al. (Eds.), *Advances in Neural Information Processing Systems*, vol. 9, The MIT Press, Cambridge, MA, 1997, pp. 155–161.
- [21] J.A.K. Suykens, J. De Brabanter, L. Lukas, J. Vandewalle, Weighted least squares support vector machines: robustness and sparse approximation, *Neurocomputing* 48 (2002) 85–105.
- [22] C.J.C. Burges, A tutorial on support vector machines for pattern recognition, *Data Min. Knowl. Disc.* 2 (2) (1998) 121–167.
- [23] J. Shawe-Taylor, N. Cristianini, *Kernel Methods for Pattern Analysis*, Cambridge University Press, 2004.
- [24] V. Vapnik, S. Golowich, A. Smola, Support vector method for function approximation, regression estimation, and signal processing, in: M. Mozer et al. (Eds.), *Advances in Neural Information Processing Systems*, vol. 9, The MIT Press, Cambridge, MA, 1997, pp. 281–287.
- [25] B. Schölkopf, P. Bartlett, A. Smola, R. Williamson, Support vector regression with automatic accuracy control, in: L. Niklasson et al. (Eds.), *Proceedings of the International Conference on Artificial Neural Networks, Perspectives in Neural Computing*, Springer Verlag, Berlin, 1998, pp. 111–116.
- [26] N. Ancona, Properties of Support Vector Machines for Regression, Technical Report, Instituto Elaborazione Segnali ed Immagini, Bari, Italy, pp. 1–99, 1999.
- [27] K.R. Müller, A. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, V. Vapnik, Using support vector machines for time series prediction, in: B. Schölkopf (Ed.), *Advances in Kernel Methods – Support Vector Learning. Proceedings of the NIPS, Workshop on Support Vectors*, The MIT Press, Cambridge, MA, 1999, pp. 1–12.
- [28] B. Schölkopf, C.J.C. Burges, A.J. Smola (Eds.), *Advances in Kernel Methods – Support Vector Learning. Proceedings of the NIPS Workshop on Support Vectors*, The MIT Press, Cambridge, MA, 1999.
- [29] T. Van Gestel, J. Suykens, D. Baestaens, A. Lambrechts, G. Lanckriet, B. Vandaele, B. De Moor, J. Vandewalle, Financial time series prediction using least squares support vector machines within the evidence framework, *IEEE Trans. Neural Networks* 12 (4) (2001) 809–821 (Special Issue on Neural Networks in Financial Engineering).
- [30] R.D. Reed, R.J. Marks, *Neural Smoothing: Supervised Learning in Feedforward Artificial Neural Networks*, MIT Press, Cambridge, MA, 1999.
- [31] B. Schölkopf, *Support Vector Learning*, Ph.D. Thesis, Oldenbourg R. Verlag, Munich, 1997.
- [32] H. Lin, C. Lin, A Study on Sigmoid Kernels for SVM and the Training of Non-PSD Kernels by SMO-type Methods, Technical Report, Department of Computer Science and Information Engineering, National Taiwan University, 2003.
- [33] V. Kecman, *Learning and Soft Computing: Support Vector Machines, Neural Networks, and Fuzzy Logic Models*, The MIT Press, Cambridge, MA, 2001.
- [34] G. Péczely, *Climatology*, Tankönyvkiadó, Budapest, 1979, pp. 258–284 (in Hungarian).
- [35] A.S. Weigend, N.A. Gershenfeld (Eds.), *Time Series Prediction: Forecasting the Future and Understanding the Past*, Addison-Wesley, Reading, MA, 1994.
- [36] B. Hammer, K. Gersmann, A note on the universal approximation capability of support vector machines, *Neural Process. Lett.* 17 (2003) 43–53.
- [37] I.H. Witten, E. Frank, *Data Mining: Practical Machine Learning Tools with Java Implementations*, Morgan Kaufmann, San Francisco, 2000.
- [38] Weka Data Mining Software Library (accessed: February, 2005) <www.cs.waikato.ac.nz/ml/weka>.
- [39] C.C. Chang, C.J. Lin, Training nu-support vector regression: theory and algorithms, *Neural Comput.* 14 (8) (2002) 1959–1977.
- [40] LibSVM Support Vector Software Library (accessed: February, 2005) <<http://www.csie.ntu.edu.tw/~cjlin/libsvm>>.
- [41] L. Makra, The role of traffic in modifying air quality in a medium-sized city, Szeged, Hungary, *Epidemiology* 16 (5) (2005) S62.
- [42] P. Shi, R.M. Harrison, Regression modelling of hourly NO_x and NO₂ concentrations in urban air in London, *Atmos. Environ.* 31 (24) (1997) 4081–4094.
- [43] J.H. Friedman, An overview of predictive learning and function approximation, in: V. Cherkassky, J.H. Friedman (Eds.), *From Statistics to Neural Networks, Theory and Pattern Recognition Applications*, NATO ASI Series F, vol. 136, Springer, 1994, pp. 1–61.
- [44] I.T. Jolliffe, *Principal Component Analysis*, Springer-Verlag, New York, 1986.
- [45] C.J. Willmott, Some comments on the evaluation of model performance, *Bull. Am. Meteorol. Soc.* 63 (11) (1982) 1309–1313.