Monotonic Regression for the Detection of Temporal Trends in Environmental Quality Data

Mohamed Hussian[1], Anders Grimvall[1], Oleg Burdakov[1], and Oleg Sysoev[2]

[1]Department of Mathematics, Linköping University, SE-58183 Linköping, Sweden
E-mail: mohus@mai.liu.se, angri@mai.liu.se, olbur@mai.liu.se
Web site: www.mai.liu.se

[2]Faculty of Control and Applied Mathematics, Moscow Institute of Physics and Technology,
Institutskij per. 9, Dolgoprudnyj, Moscow region 141700, Russia
E-mail: osysoev@mail.ru

Abstract

Monotonic regression is a non-parametric method designed especially for applications in which the expected value of a response variable increases or decreases in relation to one or more explanatory variables. Our research group recently demonstrated how the pool-adjacent-violators (PAV) algorithm can be generalised from fully to partially ordered data, which has made it practical to handle multiple regression data that include one or more continuous variables. Here, we show that this generalisation facilitates the estimation of monotonic response surfaces and temporal trends in environmental quality data. In particular, we present new methods for simultaneous extraction of a monotonic trend and seasonal components and for normalisation of time series of environmental quality data that are influenced by random variation in weather conditions or other forms of natural variability. Theoretical descriptions of our methodology are

accompanied by examples of trend assessments of water quality data and normalisation of the mercury concentration in cod muscle in relation to the length of the analysed fish.

1. INTRODUCTION

Monotonic responses are widespread in all types of environmental systems. For example, the rates of chemical and microbial processes normally increase with increasing temperature. Also, the concentrations of many contaminants in living organisms increase with the age or size of the analysed individual, and fluxes of substances through terrestrial and aquatic systems can increase with the amount and intensity of precipitation. The simplest forms of monotonic relationships can easily be described by using an appropriate parametric model, and numerous algorithms have been developed to fit such models to observed data. However, for more complex relationships involving two or more explanatory variables, it can be necessary to perform non-parametric modelling especially if the response includes a threshold effect or is strongly non-linear in some other respect.

Monotonic regression, which comprises the special case of isotonic regression, is a non-parametric method designed for applications in which the expected value of a response variable (y) increases or decreases in relation to one or more explanatory variables (x_1, …, x_p). The most commonly used computational method for this type of regression is what is known as the pool-adjacent-violators (PAV) algorithm [1, 2, 3]. When $p = 1$, this algorithm is computationally efficient, and it provides solutions that are optimal in the sense that the mean square error is minimized. When $p > 1$, the PAV algorithm has proven useful for estimating monotonic responses to explanatory variables that are varied at only a few levels [4, 5, 6, 7]. However, it was not until Burdakov and colleagues [8, 9] recently generalised the PAV algorithm from fully to partially ordered data that it became feasible to handle typical regression data that include one or more continuous variables.

Here, we briefly describe a generalised pool-adjacent-violators (GPAV) algorithm that operates on the Hasse diagram [10] of the points in R^p that are determined by the observed values of the p explanatory variables. Furthermore, we show that this algorithm has important applications in several areas of environmental science and management, especially in the following contexts:

(i) Estimation of response surfaces that are required to be monotonic in two or more variables.

(ii) Simultaneous extraction of seasonal components and a monotonic trend from a univariate time series.

(iii) Normalisation of time series of environmental quality data.

The first of these tasks is also highly relevant in many areas other than environmental science; for example, monotonic regression is often appropriate for estimating dose-response curves in experimental studies [6, 7]. The second task entails time series decomposition, which is a classical undertaking in official statistics (e.g., [11]). The method we present takes into account that many seasonal patterns in the environment can be decomposed into one increasing and one decreasing phase. The third task, normalisation or adjustment, aims to clarify the human impact on the environment by removing weather-dependent fluctuations or other natural variability in the collected data [12].

2. ALGORITHMS FOR ISOTONIC OR MONOTONIC REGRESSION

Let $\boldsymbol{x}_i = (x_{1i}, \ldots, x_{pi})$, $i = 1, \ldots, n$, denote the observed values of p explanatory variables and let y_i, $i = 1, \ldots, n$, represent the associated responses. Then isotonic regression implies that we wish to find a set of fitted response values $\{z_i, i = 1, \ldots, n\}$ such that

$$S = \sum_{i=1}^{n} (y_i - z_i)^2$$

is minimized under the constraints

$$z_i \leq z_j, \text{ if } x_{ki} \leq x_{kj} \text{ for all } k = 1, ..., p.$$

The set of observed values $\mathbf{x}_1, \ldots, \mathbf{x}_n$ is partially ordered by the relation

$$\boldsymbol{x}_i \preceq \boldsymbol{x}_j \text{ if } x_{ki} \leq x_{kj} \text{ for all } k = 1, ..., p$$

Therefore, isotonic regression can also be regarded as a smoothing method that transfers this partial order to the fitted response values z_i, $i = 1, \ldots, n$. Monotonic regression,

which allows the response to increase in some of the explanatory variables and decrease in others, may appear to be more general. However, it can easily be reduced to isotonic regression by changing the sign of some of the explanatory variables.

2.1 The ordinary PAV algorithm

The ordinary PAV algorithm for isotonic regression in one explanatory variable [2, 3] assumes that the data set $M_n = \{(x_i, y_i), i = 1, \ldots, n\}$ is presorted so that $x_1, \ldots, x_n$ form a non-decreasing sequence. In that case, the optimal solution is a vector whose components z_i are nondecreasing in $i \in \{1, \ldots, n\}$, and we can form clusters of adjacent indices for which the fitted z-values are identical. The PAV algorithm identifies these clusters and the associated z-values in a recursive procedure. The (optimal) solution for M_n is constructed by starting from the solution $z_1 = y_1$ for M_1, subsequently modifying it into the solution for M_2, and so on. Moreover, the PAV algorithm has the following characteristics:

(i) If $z_1, \ldots, z_r$ denote the solution for M_r, then a preliminary solution for M_{r+1} is formed by setting $z_{r+1} = y_{r+1}$.

(ii) The optimal solution for M_{r+1} is derived by pooling adjacent preliminary z-values that violate the monotonicity constraints; to be more precise, the same value $(y_{r+1} + \ldots + y_{r+1-k})/(k+1)$ is assigned to each of $z_{r+1-k}, \ldots, z_{r+1}$, where k is the smallest integer such that $z_1, \ldots, z_{r-k}$ along with new values of $z_{r+1-k}, \ldots, z_{r+1}$ form a non-decreasing sequence.

To pave the way for our generalisation of the PAV algorithm, we introduce the notation $I = \{i_1, \ldots, i_m\}$ for a cluster of indices $i_1, \ldots, i_m$, the symbol $\omega(I) = m$ for the number of elements in I, and the symbol $z(I)$ for the common value of all z_i, $i \in I$. Furthermore, we note that when two adjacent clusters I_1 and I_2 are joined to form a new cluster $I_1 \cup I_2$, the associated z-value is given by the expression

$$z(I_1 \cup I_2) = (\omega(I_1) \cdot z(I_1) + \omega(I_2) \cdot z(I_2)) / (\omega(I_1) + \omega(I_2))$$

2.2 A generalised pool adjacent violators (GPAV) algorithm for two or more explanatory variables

The GPAV algorithm developed by Burdakov and co-workers [8, 9] aims to provide optimal or close to optimal solutions to the isotonic regression problem for any number of explanatory variables. The basic idea is the same as in the ordinary PAV algorithm. Observations enter the calculations one at a time, and monotonicity violators are removed by pooling adjacent clusters of indices. However, the obtained solution may vary slightly with the order in which observations enter.

Numerical experiments indicate that, if the elements $x_i \in R^p$ are grouped into levels as in a Hasse diagram of a partially ordered set [9] and then entered the calculations in accordance with these levels, the obtained solutions are close to optimal. This grouping implies that level 1 comprises all minimal elements. Furthermore, any element that dominates another element is assigned to a higher level.

The optimal solution for the level 1 data is obtained by setting $z_i = y_i$ for all minimal elements $x_i \in R^p$. The subsequent steps in the GPAV algorithm follow a recursive procedure that resembles the ordinary PAV algorithm. Observations (x_i, y_i) enter the calculations one by one in an upward movement through the different levels of the x-vectors, i.e., all elements of level k are entered prior to any element of order $k+1$ or higher. Moreover, monotonicity violators are eliminated in a downward movement by pooling adjacent violators. To be more precise, the GPAV algorithm has the following features:

(i)   Given that the clusters I_1, …, I_q and their associated values $z(I_1)$, …, $z(I_q)$ are a solution for M_r, a preliminary solution for M_{r+1} is formed by introducing the cluster I_{q+1} consisting of the integer $r+1$, and setting $z(I_{q+1}) = y_{r+1}$.

(ii)   A final solution for M_{r+1} is obtained by sequentially joining I_{q+1} with immediate predecessor clusters until the z-values violating the monotonicity constraints have been removed (a cluster I_j is called an immediate predecessor of I_l if there exists an $i \in I_j$ and a $k \in I_l$ such that x_i belongs to the lower cover of x_k).

In particular, we note that the solution $\{z_i, i = 1, …, n\}$ obtained for M_n consists of clusters of identical values.

3. ESTIMATING MONOTONIC RESPONSE SURFACES

Figures 1a and 1b illustrate how monotonic regression can be used to describe data on the concentration of mercury in Atlantic cod (*Gadu morhua*) in relation to sampling year and body length. The increase in mercury with increasing length of the fish is obvious in the two diagrams. In addition, the response surface in Figure 1b indicates a downward temporal trend.
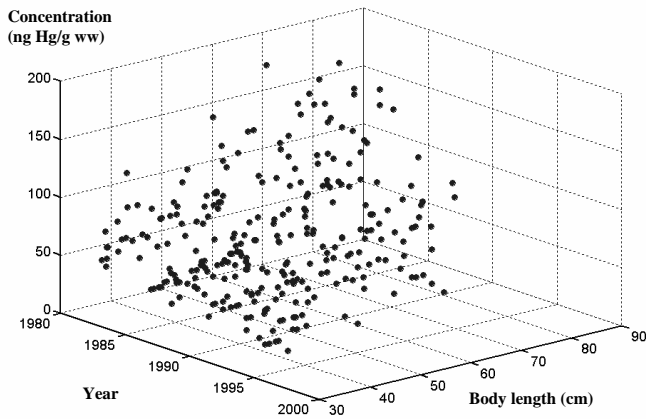

Figure 1a. Concentration of mercury in muscle tissue from Atlantic cod (*Gadu morhua*) caught in the North Sea (53° 10' N, 2° 5' E). The data represent observed concentrations (ng Hg/g ww) in relation to sampling year and body length of the analysed fish.
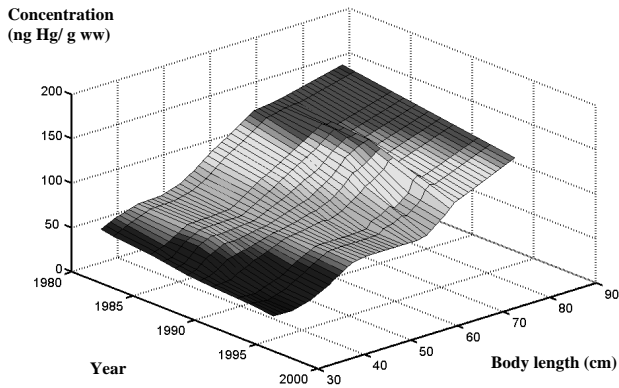
Figure 1b. Concentration of mercury in muscle tissue from Atlantic cod (*Gadu morhua*) caught in the North Sea (53° 10' N, 2° 5' E). The response surface was obtained by first using the GPAV algorithm for monotonic regression and then employing locally weighted scatter-plot smoothing to extrapolate the fitted regression values to a fine grid (see also section 4.3).

4. SIMULTANEOUS ESTIMATION OF A MONOTONIC TREND AND SEASONAL EFFECTS

When data are collected over several seasons, monotonic regression models may appear to be inadequate. However, many seasonal patterns can be decomposed into increasing and decreasing phases, which enables the use of various approaches based on monotonic regression. If we let $y_1, y_2, \ldots, y_n$ denote a time series of data collected over m seasons, and let $\hat{y}_i$ denote the sum of the trend and seasonal components at time i, it is possible to determine $\hat{y}_i$ by minimising

$$S = \sum_i (y_i - \hat{y}_i)^2$$

under a set of simple constraints, and we can also introduce these constraints by employing a monotonic regression model.

Let us, for the sake of clarity, consider a seasonal pattern of the type illustrated in Figure 2. Assume also that we would like to extract a non-increasing trend function

from the collected data. If we then perform a monotonic regression using sampling year and the variables x_1 and x_2 as explanatory variables, the fitted values $\hat{y}_i$ must be non-increasing for each season, i.e.,

$$\hat{y}_i \geq \hat{y}_{i+m}, \quad i = 1, ..., n - m$$

In addition, the fitted values representing different seasons in the same year must have non-increasing and non-decreasing phases with the same duration as in Figure 2.

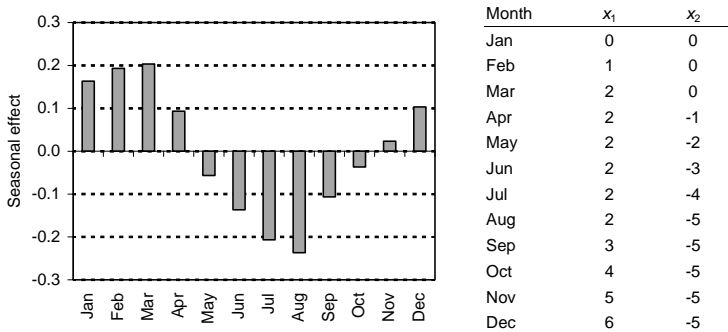| Month	x_1	x_2
Jan	0	0
Feb	1	0
Mar	2	0
Apr	2	-1
May	2	-2
Jun	2	-3
Jul	2	-4
Aug	2	-5
Sep	3	-5
Oct	4	-5
Nov	5	-5
Dec	6	-5

Figure 2. Seasonal pattern comprising increasing and decreasing phases, and a possible coding of these phases.

Figure 3 (a and b) illustrates a set of monthly flow-weighted concentrations of total nitrogen in the Rhine River and the monotonic trend and seasonal components that could be extracted from these data. The goodness-of-fit to observed data reached a maximum when we let the seasonal effects have a maximum in March and a minimum in August. Furthermore, it can be noted that the downward trend was particularly strong during the colder months of the year.
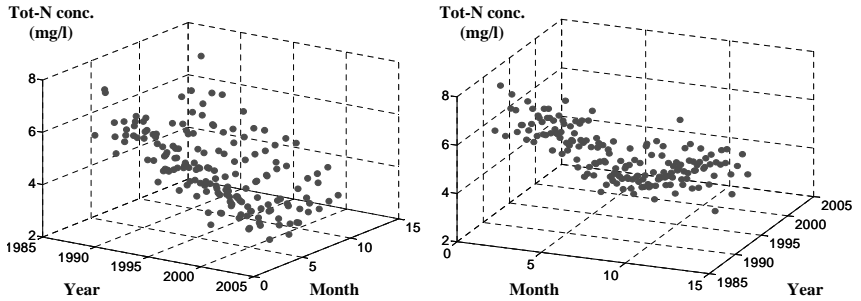
Figure 3a. Monthly mean concentrations of total nitrogen (Tot-N) measured in the Rhine River at Lobith on the border between Germany and The Netherlands.
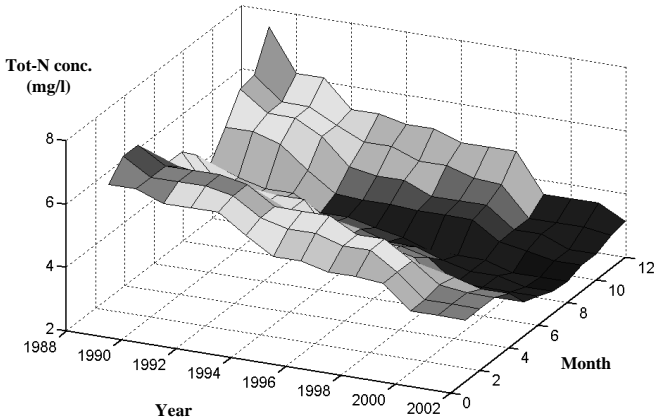

Figure 3b. Response surface obtained by applying monotonic regression to monthly mean concentrations of total nitrogen (Tot-N).

5. NORMALISATION OF TIME SERIES OF ENVIRONMENTAL QUALITY DATA

5.1 Normalisation formulae

The general aim of normalisation is to remove irrelevant variation in collected data. The basic idea is simple. If observations of meteorological or other naturally fluctuating

variables make us believe that the studied response variable takes a value c units higher than the mean response, then normalisation implies that we subtract this expected increase c from the observed response. A general probabilistic framework for normalisation was recently presented by Grimvall and co-workers [12]. Here, we discuss normalisation based on monotonic regression models.

Let us assume that the observed values of the response variable y have the general form

$$y_i = f(x_{1i}, \ldots, x_{pi}) + \varepsilon_i, \, i = 1, \ldots, n$$

where f is a deterministic function of p explanatory variables $x_1, \ldots, x_p$, and $\varepsilon_i, \, i = 1, \ldots, n$, depicts a sequence of independent, identically distributed random errors with mean zero. We can then normalise the observed responses with respect to $x_1, \ldots, x_q$ by forming

$$\widetilde{y}_i = y_i - \{\hat{f}(x_{1i}, \ldots, x_{qi}, x_{q+1,i}, \ldots, x_{pi}) - \hat{f}(\overline{x}_{1i}, \ldots, \overline{x}_{qi}, x_{q+1,i}, \ldots, x_{pi})\}$$

where $\hat{f}$ denotes an estimate of f, and $(\overline{x}_{1i}, \ldots, \overline{x}_{qi}), i = 1, \ldots, n$, is a sequence of given values of $x_1, \ldots, x_q$. Typically, these given values represent averages taken over the entire data set or subsets thereof. For example, if data are collected over several seasons, we can let $(\overline{x}_{1i}, \ldots, \overline{x}_{qi}), i = 1, \ldots, n$, represent seasonal means of $x_1, \ldots, x_q$.

Regardless of how the normalisation is carried out, we must be able to estimate the values of f at two sets of points: an estimation set

$$A = \left\{(x_{1i}, \ldots, x_{pi}), i = 1, \ldots, n\right\}$$

for which we have observed response values $\{y_i, i = 1, \ldots, n\}$, and an evaluation set

$$B = \left\{(\overline{x}_{1i}, \ldots, \overline{x}_{qi}, x_{q+1,i}, \ldots, x_{pi}), i = n+1, \ldots, n+m\right\}$$

for which no observations exist. The GPAV algorithm provides estimates of f for all points in the estimation set. It remains to extrapolate $\hat{f}$ to the evaluation set under the constraint that $\hat{f}$ is monotonic in each of the coordinates.

5.2 Extrapolation of monotonic responses to new points

Let $\boldsymbol{\chi} = (\chi_1, ..., \chi_p)$ be a point to which $\hat{f}$ shall be extrapolated from a given estimation set A. We can then define two subsets of A. The first subset

$$L_\chi = \left\{ (x_{1i}, ..., x_{pi}) \in A \; ; \; x_{ki} \leq \chi_k , k = 1, ..., p \right\}$$

contains all points in A that are dominated by $\boldsymbol{\chi}$. The second subset

$$U_\chi = \left\{ (x_{1i}, ..., x_{pi}) \in A \; ; \; x_{ki} \geq \chi_k , k = 1, ..., p \right\}$$

comprises all points in A that dominate $\boldsymbol{\chi}$.

Let us also for the moment assume that both L_χ and U_χ are nonempty. Then the expression

$$y_L = \max \left\{ \hat{f}(x_{1i}, ..., x_{pi}) \; ; \; (x_{1i}, ..., x_{pi}) \in L_\chi \right\}$$

provides a lower limit for the values of $\hat{f}(\boldsymbol{\chi})$ that can render $\hat{f}$ monotonic on the set $A \cup \{\boldsymbol{\chi}\}$. Furthermore, we can identify a point $\boldsymbol{\chi}_L \in L_\chi$ that minimises the distance to $\boldsymbol{\chi}$ under the constraint that $\hat{f}(\boldsymbol{\chi}_L) = y_L$. Likewise,

$$y_U = \min \left\{ \hat{f}(x_{1i}, ..., x_{pi}) \; ; \; (x_{1i}, ..., x_{pi}) \in U_\chi \right\}$$

defines an upper limit for the permissible values of $\hat{f}(\boldsymbol{\chi})$, and we can select a point $\boldsymbol{\chi}_U \in U_\chi$ that minimises the distance to $\boldsymbol{\chi}$ under the constraint that $\hat{f}(\boldsymbol{\chi}_U) = y_U$.

If $\boldsymbol{\chi}$ is on the straight line between $\boldsymbol{\chi}_L$ and $\boldsymbol{\chi}_U$, it would be natural to use linear interpolation to assign a value to $\hat{f}(\boldsymbol{\chi})$, in other words to set

$$\hat{f}(\boldsymbol{\chi}) = y_L + \frac{\|\boldsymbol{\chi} - \boldsymbol{\chi}_L\|}{\|\boldsymbol{\chi}_U - \boldsymbol{\chi}_L\|}(y_U - y_L)$$

where $\|\boldsymbol{u}\|$ denotes the length of the vector $\boldsymbol{u}$. Regardless of the location of $\boldsymbol{\chi}$, we can set

$$\hat{f}(\boldsymbol{\chi}) = y_U + \frac{(\boldsymbol{\chi} - \boldsymbol{\chi}_L, \boldsymbol{\chi}_U - \boldsymbol{\chi}_L)}{(\boldsymbol{\chi}_U - \boldsymbol{\chi}_L, \boldsymbol{\chi}_U - \boldsymbol{\chi}_L)}(y_U - y_L)$$

where $(\boldsymbol{u}, \boldsymbol{v})$ denotes the scalar product of the vectors $\boldsymbol{u}$ and $\boldsymbol{v}$.

If L_x or U_x is empty, we assign values to $\hat{f}(\chi)$ as follows:

$$\hat{f}(\chi) = \begin{cases} y_U, & \text{if } L_\chi \text{ is empty and } U_\chi \text{ is nonempty} \\ y_L, & \text{if } U_\chi \text{ is empty and } L_\chi \text{ is nonempty} \\ \overline{y}, & \text{if both } L_\chi \text{ and } U_\chi \text{ are empty} \end{cases}$$

where $\overline{y}$ is the mean response for the elements in the entire estimation set or a subset of elements within a fixed distance to χ.

The procedure described above can be repeated for an arbitrary set of points. However, it is important to note that the estimation set A must be updated from A to $A \cup \{\chi\}$ each time $\hat{f}$ has been extrapolated to a new point χ. Otherwise, there may be pairs of extrapolated values for which the monotonicity is violated.

If the evaluation set is large, the above-mentioned procedure can be computationally cumbersome. Hence there is also a need for extrapolation procedures that can provide a response surface that is nearly monotonic over a large set of points. For example, it can be convenient to use kernel smoothing or locally weighted scatter-plot smoothing [13] to extrapolate the fitted responses in a monotonic regression to a fine grid of values of the explanatory variables (see Figure 1b).

5.3   Normalisation of contaminants in fish

Simple time series plots of observed concentrations of mercury in Atlantic cod caught in the middle of the North Sea (53^o 10' N, 2^o 5' E) indicate a downward trend (Figure 4a). However, this may, at least in part, be a spurious trend caused by changes in the length of the analysed fish over time. Accordingly, it is of great interest to normalise the observed mercury concentrations with respect to body length. Figure 4b illustrates the results obtained by using the normalisation procedure described in sections 4.1 and 4.2. Apparently, the mercury trend after normalisation is considerably smaller than it is in the raw data.
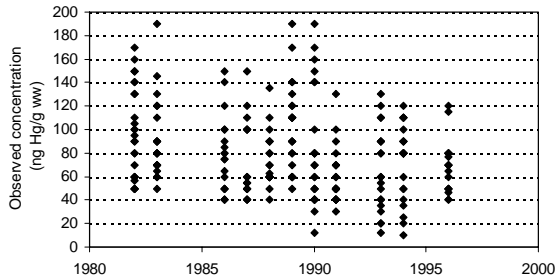
Figure 4a. Observed concentrations of mercury in muscle tissue from Atlantic cod (*Gadu morhua*) caught in the North Sea (53^{o} 10' N, 2^{o} 5' E).
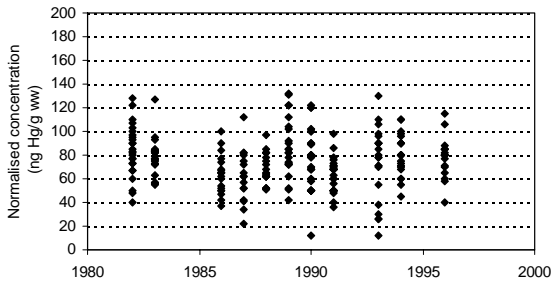

Figure 4b. Concentration of mercury in Atlantic cod (*Gadu morhua*) normalised to a body length of 49.6 cm.

6. DISCUSSION

The results presented in this article demonstrate that the recently developed GPAV algorithm for monotonic regression has great potential for applications in environmental science and management. The assumption of monotonic response often has a solid foundation in process-based modelling or strong empirical evidence, whereas in other contexts, such as trend assessment of environmental quality data, it can be more correct to regard monotonicity as a reasonable simplification. In either case, monotonic regression is convenient to use because issues concerning model selection are reduced to a minimum. Conventional parametric modelling of non-linear responses to two or more variables is usually more demanding, and the fit to observed data can be unsatisfactory even if a large number of models are tested. Kernel smoothing, locally

weighted scatter-plot smoothing, and other non-parametric regression techniques are often more viable alternatives to monotonic regression. However, the latter methods may produce very odd results if the set of points for which responses have been observed is very unevenly distributed. In addition, it may be unsatisfactory to obtain a non-monotonic response surface when there is theoretical support for monotonicity.

The calculations that were undertaken to simultaneously extract a monotonic trend and seasonal components from a time series of water quality data illustrate that monotonic regression usually produces a relatively smooth response surface. In this respect, our method is similar to the currently used procedures for time series decomposition (e.g., [11]). Nevertheless, the estimated seasonal effects can be different, because, in contrast to other techniques, our method includes constraints on how such effects can vary from one season to the next.

The normalisation in section 5 represents a type of statistical analysis that has attracted increasing attention in programs for environmental monitoring. Most of the methods presently used for that purpose are based on linear regression models for the removal of irrelevant variation in collected data [14], and model selection studies have shown that such models can perform well even if some of the underlying processes are non-linear [15]. However, it has also been demonstrated that non-linear features of a normalisation model can improve the performance if there is a strong, non-linear trend in the analysed data [16]. The monotonic regression described in this article provides yet another tool for normalisation that can be particularly useful if the natural fluctuations in collected data include non-linear responses to one or more covariates.

ACKNOWLEDGEMENTS

REFERENCES

[1] M. Ayer, H.D. Brunk, G.M. Ewing, W.T. Reid, and E. Silverman. An empirical distribution function for sampling with incomplete information. *The Annals of Mathematical Statistics*, **26**, 641-647, 1955.

[2] R.E. Barlow, D.J. Bartholomew, J.M. Bremner, and H.D. Brunk. *Statistical inference under order restrictions*. New York: Wiley, 1972.

[3] D.L. Hanson, G. Pledger, and F.T. Wright. On consistency in monotonic regression. *The Annals of Statistics*, **1**, 401-421, 1973.

[4] R. Dykstra and T. Robertson. An algorithm for isotonic regression for two or more independent variables. *The Annals of Statistics*, **10**, 708-716, 1982.

[5] G. Bril, R. Dykstra, C. Pillers, and T. Robertson. Algorithm AS 206, isotonic regression in two independent variables. *Applied Statistics,* **33**, 352-357, 1984.

[6] M.J. Schell and B. Singh. The reduced monotonic regression method. *Journal of the American Statistical Association*, **92**, 128-135, 1997.

[7] G. Salanti and K. Ulm. The multidimensional isotonic regression. *Proceedings of the International Society of Clinical Biostatistics*, 19-23 Aug 2001, Stockholm, p. 162, 2001.

[8] O. Burdakov, A. Grimvall, and M. Hussian. A generalised PAV algorithm for monotonic regression in several variables. *Proceedings of COMPSTAT*, 23-27 August, 2004, Prague, 2004a.

[9] O. Burdakov, O. Sysoev, A. Grimvall, and M. Hussian. *An algorithm for isotonic regression problems. Proceedings of the European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS)*, 24-28 July 2004, Jyväskylä, Finland.

[10] B.A. Davey and H.A. Priestley. *Introduction to lattices and order*. Cambridge: Cambridge University Press, 1999.

[11] A. Maravall and V. Gómez. Estimation, prediction and interpolation for nonstationary series with the Kalman filter. *Journal of the American Statistical Association*, **89**, 611-624, 1994.

[12] A. Grimvall, H. Wackernagel, and C. Lajaunie. Normalisation of environmental quality data. In: L.M. Hilty and P.W. Gilgen (Eds.) *Sustainability in the Information Society*. Marburg: Metropolis-Verlag, pp. 581-590, 2001.

[13] W. Härdle. *Applied non-parametric regression*. Cambridge: Cambridge University Press, 1997.

[14] M.L. Thompson, J. Reynolds, L.H. Cok, P. Guttorp, and P.D. Sampson. A review of statistical methods for the meteorological adjustment of ozone. *Atmospheric Environment*, **35**, 617-630, 2001.

[15] C. Libiseller and A. Grimvall. Model selection for local and regional meteorological normalisation of background concentrations of tropospheric ozone. *Atmospheric Environment*, **37**, 3923-3931, 2003.

[16] M. Hussian, A. Grimvall, and W. Petersen. Estimation of the human impact on nutrient loads carried by the Elbe River. *Environmental Monitoring and Assessment*, **96**, 15-33, 2004.