

Using Smoothing to Reconstruct the Holocene Temperature in Lapland

Lasse Holmström, Panu Erästö and Petri Koistinen *

Rolf Nevanlinna Institute
University of Helsinki, Helsinki, Finland

Jan Weckström and Atte Korhola †
Department of Ecology and Systematics
University of Helsinki, Helsinki, Finland

Abstract

Small arctic and subarctic lakes are known to be sensitive to climatic variation. Changes in external conditions are continuously recorded in their sediments in the form of aquatic organisms. The abundance of such organisms can therefore be used to reconstruct past environmental conditions. We use data collected from the Finnish Lapland to reconstruct post ice age temperatures. Non-parametric smoothing has not been used often in this context. We find smoothing a viable tool both in the actual reconstruction phase and in the subsequent time series smoothing, where the SiZer method is used.

1 Introduction

During the past 100 years the global temperature has risen about 0.6°C (1.0°F) [7]. Some computational models predict that if the atmospheric CO_2 concentration doubles during the next 100 years, a conservative estimate if the world proceeds on a business as usual path, then for example the United States could experience $3\text{--}6^{\circ}\text{C}$ ($5\text{--}10^{\circ}\text{F}$) increases in the average temperature [12]. How much of the recent rise in temperature can be explained by natural variation? How big temperature changes has the earth experienced over a longer period of time? What were the reasons for the changes? Some answers may be found by considering the global temperature variation during the Holocene, a period that spans about 11 000 years from the end of the last ice age to the present and that exhibits more climatic dynamics than previously thought.

The arctic lakes of northern Fennoscandia are ecolog-

ically very sensitive and hence excellent natural laboratories for the study of climatic changes [3]. Variation in the external conditions are continuously recorded in the lake sediments for example in the form of fossils of aquatic organisms. The abundance of such organisms can be strongly affected by environmental variables such as lake water pH, temperature, and alkalinity.

Perhaps the most commonly used organisms in environmental reconstruction studies are diatoms. They are microscopic unicellular algae whose cell walls are made of silica and whose fossils are typically well preserved in lake water. The size of these fossils is of the order of tens of micrometers and the different taxa exhibit rich morphological variation. A model for the dependence between the relative abundances of different diatom taxa and an environmental variable such as air temperature can be built based on surface sediment samples and environmental data from a set of lakes representing a wide range of conditions [13]. The model can then be applied to diatom fossil samples obtained from a sediment core to reconstruct past temperatures.

The focus of this paper is to show how nonparametric smoothing can be used in diatom-based temperature reconstruction as well as in helping to interpret the results. For a more thorough discussion on the ecological aspects of this research, details of data collection, and references to relevant paleoecological literature, see [8].

The rest of the paper is organized as follows. Section 2 describes the data used in our study. Section 3 discusses some popular classical approaches to temperature reconstruction and explains the particular smoothing technique we propose as an alternative. Temperature reconstructions are presented in Section 4. We use the SiZer methodology of Chaudhuri and Marron to evaluate the significance of the Holocene warming and cooling periods suggested by the reconstruction. In Section 5 the results are discussed further and a summary is presented in Section 6.

*Research partially supported by Academy of Finland Grant No. 37551.

†Research partially supported by the EC Environment Climate Research Program (contract: ENV4-CT97-0642) and the Academy of Finland Grant No. 1017383.

2 The Data

Our goal is to reconstruct the Holocene mean July temperature at a certain lake in Finnish Lapland going back about 10 000 years in time. The predictor vector X is the vector of relative abundances of different diatom taxa and the response Y is the mean July temperature. (Notice that in much of the environmental reconstruction literature the symbol X is used for the environmental variables and Y for the taxon abundances.) Our notation does not try to suggest that the taxon abundances influence the mean temperature; on the contrary, the taxon abundances should be thought to depend on the mean summer temperature, the values of other relevant environmental variables and random influences.

Our data were collected during summer 1998 in the northwest part of the Finnish Lapland. The training set $(X_1, Y_1), \dots, (X_n, Y_n)$ consists of relative diatom taxon abundances X_i for 226 taxa in the surface sediments of $n = 38$ lakes and the mean July temperatures Y_i for the same lakes. Thus, the dimensionality of each X -vector is 226.

A different lake, lake Tšuołbmajavri, was used for temperature reconstruction. A 291 cm long cylindrical sediment core was obtained from the bottom of this lake. The lake deglaciated about 10 300 years ago and the sedimentation rate has been nearly constant. (This and all the subsequent dates are calibrated radiocarbon dates [cal yr B.P.].) The core was sampled at 2 cm intervals corresponding to about 50–70 year time intervals. In this way the relative abundance X_j^* of the 226 taxa was determined at 142 different depths and this constitutes the prediction set X_1^*, \dots, X_m^* , $m = 142$. The aim is to estimate the corresponding (unobserved) mean July temperatures Y_1^*, \dots, Y_m^* . The lakes in the training set resemble geologically and environmentally the reconstruction lake so that we can transport information obtained from the training set to infer past environmental conditions in the reconstruction lake. Further details on the data set can be found in [8].

We consider each observation (X_i, Y_i) in the training set to form a random vector whose distribution does not depend on index i so that we can speak of the regression function

$$(1) \quad m(x) = E[Y_i | X_i = x].$$

The idea is to get an estimate \hat{m} for this function and then to use it to reconstruct the past temperatures by applying \hat{m} to the prediction set abundances,

$$\hat{Y}_j^* = \hat{m}(X_j^*).$$

Our motivation is the fact, that under the quadratic loss, $m(X_j)$ is the best predictor of Y_j among all predictors

depending only on X_j . If the regression function of Y_j^* given X_j^* is the same as m and provided that \hat{m} estimates m sufficiently well, then $\hat{m}(X_j^*)$ should be a good estimate of Y_j^* .

Figure 1 shows the relative abundances of some of the most common diatom taxa found at different depths in the sediment core. The relative abundances change in time and this is believed to reflect the the changing environmental conditions. Due to the heavy dominance of the samples by only a few taxa, the four lowermost samples were discarded in the temperature reconstruction.

The relevance of air temperature as an environmental factor is illustrated in Figure 2, where the training set is used to display the relative abundance of some taxa as a function of temperature. It appears that in some cases there is a taxon-specific temperature at which the taxon fares particularly well. Note also that there are lots of zeroes in the abundances. In fact, there were many extremely rare taxa and this means that there should be lots of room for data compression.

3 Reconstruction Methods

3.1 Standard Approaches

Several methods are available for environmental reconstruction from fossil data, see the articles of Birks [1, 2] for reviews. One relevant way to group the methods is to divide them into classical and inverse approaches.

In a classical approach one models how the environmental variables influence the taxon abundances, estimates the parameters of the model and then inverts the model (in a statistical sense) to obtain estimates of the past environment on the basis of observed taxon abundances in the core sample. Often the relationship between the abundance of a taxon and an environmental variable is thought to be unimodal; each taxon thrives best at a particular value of the environmental variable and cannot survive when the value is either too low or too high. The expected value of the abundance of a single taxon given the value of the environmental variable is then typically modeled by a Gaussian curve. E.g., the Gaussian response model and the multinomial logit model discussed in [1] are based on this idea.

As an example of a recent, successful reconstruction method using the classical approach we mention the work reported in [10, 11], where Bayesian inference is applied to unimodal species response models. Of particular interest is the fact that like the present one, these studies consider temperature reconstruction for lake Tšuołbmajavri, however, using chironomids instead of diatoms as the study organisms and a different set of

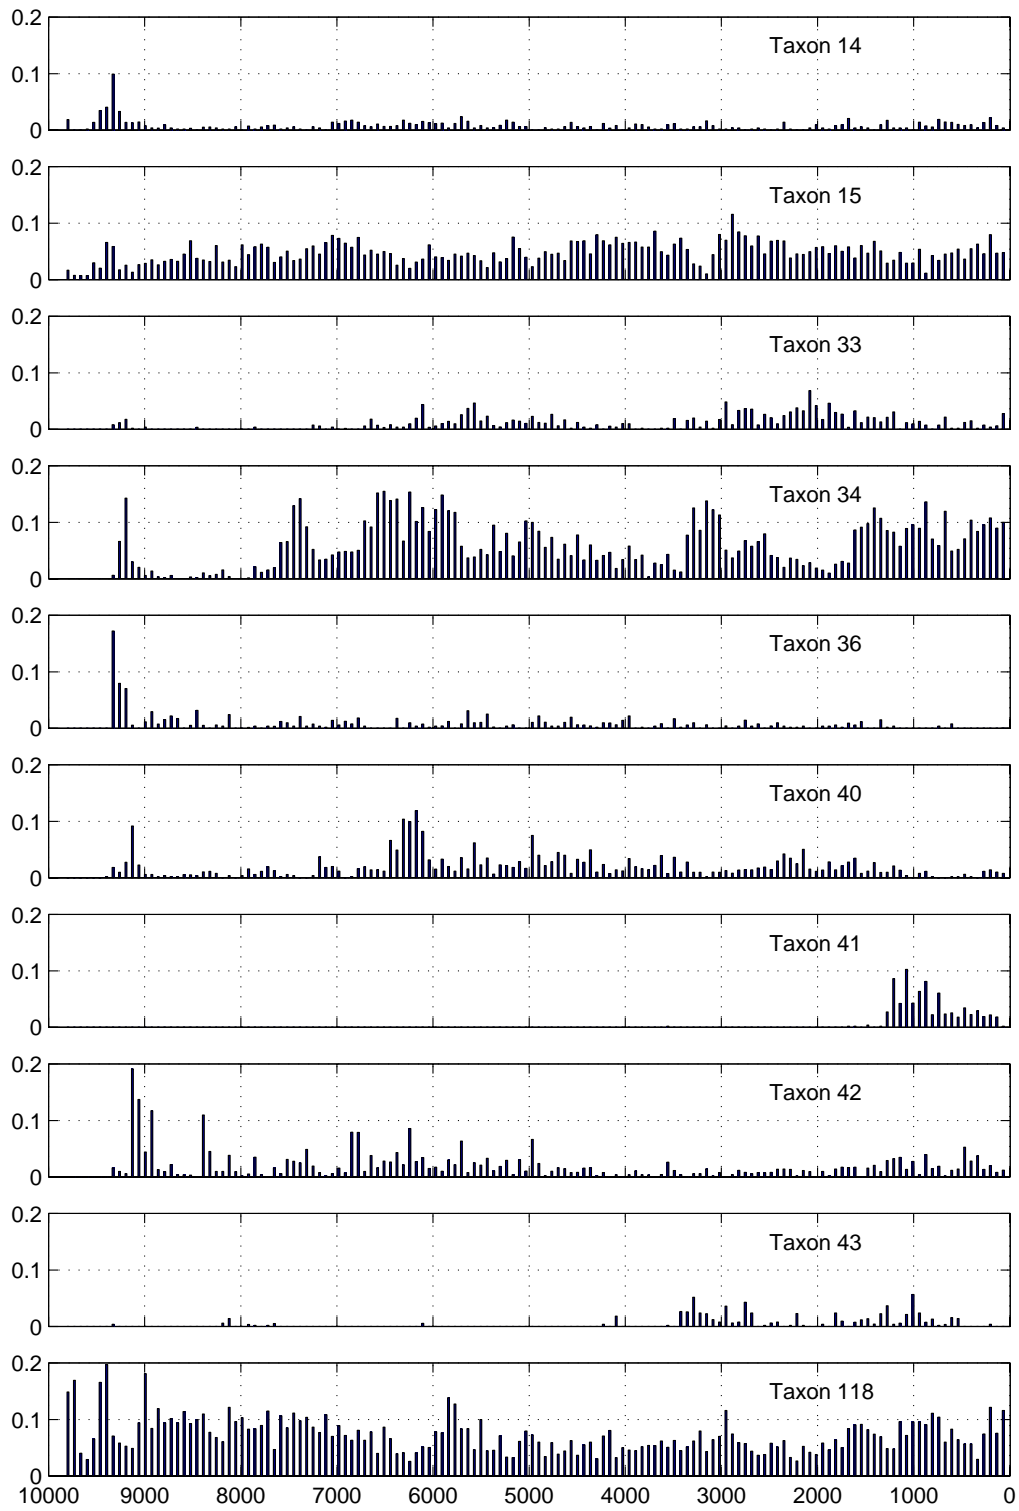


Figure 1: The relative abundances of some of the dominant diatom taxa recorded in the Tšuołbmajavri lake sediment. The horizontal axis is in years so that 0 is the present day and 10 000 corresponds to the end of the last ice age.

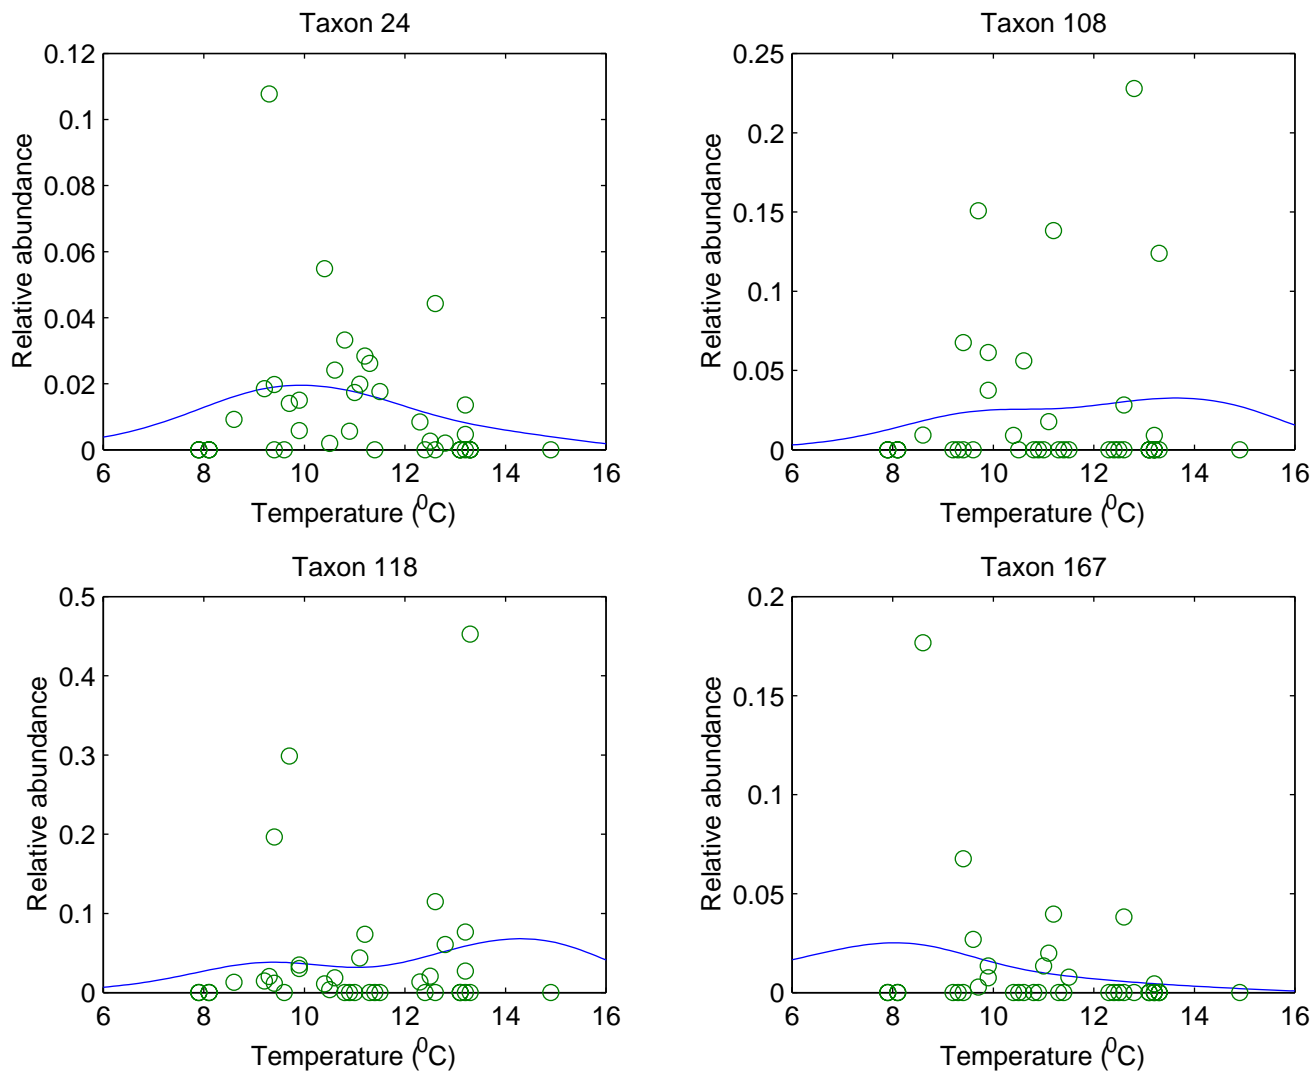


Figure 2: The relative abundance of some taxa as a function of temperature. For each taxon, the relative abundances in the training set lakes are shown as circles together with a regression smooth of the scatter plot.

training lakes.

The inverse approach bypasses the modeling of how the environment influences the taxon abundances, e.g., by estimating the regression function m defined in equation (1) and by applying the same function in reconstruction. Popular methods utilizing the inverse approach include (inverse) linear regression, principal components regression (PCR), partial least squares regression (PLS), k -nearest neighbor regression (often called modern analogue techniques, MAT in this context), and different versions of weighted averaging (WA). Reconstruction methods using the inverse approaches have generally performed better than those using the classical approach (see, e.g. [2]).

3.2 Our Approach

Our idea was to use local linear regression as a key component in an inverse approach to reconstruction. Consideration of sample sizes and the dimensionality of the predictor vector led to the conclusion that we need to reduce the dimension of the predictor drastically before this approach would work. The outline of our method is the following. In a preprocessing step, we first replace the relative abundances by their square roots. Then we reduce the dimension of our X -vectors by using components from one of two multivariate methods, principal component analysis (PCA) or partial least squares (PLS), and finally we apply local linear regression using one of two different weighting schemes. Both the final dimension in dimension reduction and the smoothing parameter in local linear regression are selected using cross-validation.

Recall that both principal components regression (PCR) and partial least squares regression (PLS) operate, conceptually, by constructing a new set of predictor variables or components which are then used in ordinary linear regression to build a prediction model for the response variable. Our approach is a generalization of this idea where the final linear regression is substituted by a more flexible estimation method, local linear regression. We can still obtain PCR and PLS as special cases if we allow sufficiently large values for the smoothing parameter.

3.3 Preprocessing

We replaced each relative abundance by its square root as a preprocessing step. Such a square root transformation is standard practice in this area and in our case improved the results a little. Motivation for such a transformation is that the empirical distributions of the relative abundances have usually long tails and the square root

transformation shortens the tails and thus brings the distributions towards normal. From now on, X_i stands for the vector of square roots of relative abundances in the i 'th lake in the training set, and similarly, X_j^* stands for the vector of square roots of relative abundances measured from the j 'th core sample in the prediction set.

3.4 Dimension reduction

We next applied a dimension reduction mapping

$$\phi : \mathbb{R}^{226} \rightarrow \mathbb{R}^d,$$

where the final dimension $d \ll 226$ was selected using cross-validation as explained in Section 3.6. Dimension reduction was performed by projecting the square root transformed abundance vectors onto a subspace of the original 226-dimensional space. We denote our new lower dimensional variables by

$$\tilde{X}_i = \phi(X_i), \quad \tilde{X}_j^* = \phi(X_j^*).$$

Although the projection subspaces depend on the training data, this dependence is suppressed in the notation. Our dimension reduction mappings were of the form

$$\phi(x) = U^T x,$$

where the $226 \times d$ matrix U has orthonormal columns. The two dimension reduction methods tried were based on principal component analysis (PCA) or partial least squares (PLS).

Principal components were obtained by using the eigenvectors u_1, \dots, u_d corresponding to the d largest eigenvalues of the sample covariance matrix of the vectors X_i as the columns of the matrix U .

While the principal components are determined solely by the X_i -vectors, partial least squares (PLS) regression uses components which depend both on the X_i and the Y_i values in a manner which tries to capture that information in the X -vectors which is useful for predicting the Y -values. PLS algorithms produce orthonormal vectors (the w_i 's of [6]), the d first which form the columns of our matrix U .

3.5 Local linear regression

Suppose we are given a training set

$$(\tilde{X}_1, Y_1), \dots, (\tilde{X}_n, Y_n),$$

where each pair (\tilde{X}_i, Y_i) is distributed as the pair (\tilde{X}, Y) , and we want to estimate the regression function $m(x) = E(Y | \tilde{X} = x)$, where \tilde{X} is a random vector and Y a random variable. In local linear regression we approximate

m linearly in the vicinity of each evaluation point u by the linear function $x \mapsto a^T(x - u) + b$, where the vector a and scalar b are found from the weighted linear least squares problem

$$(2) \quad \min_{a,b} \sum_{i=1}^n \left(Y_i - [a^T(\tilde{X}_i - u) + b] \right)^2 w_{u,h,i}.$$

Here $w_{u,h,i}$ is the weight of the i 'th observation corresponding to the smoothing parameter h . The observations (\tilde{X}_i, Y_i) whose \tilde{X}_i lie closest to the evaluation point u receive the largest weights, while the value of the smoothing parameter h determines how fast the weight decreases as the distance $\|\tilde{X}_i - u\|$ increases. Small values of h correspond to local fits while large values of h emphasize all observations more equally, and in the limit when h goes to infinity, we obtain the ordinary linear least squares fit. For the fit at the evaluation point u we set $\hat{m}(u) = \hat{b}$, where \hat{b} is obtained from the solution of problem (2). We tried two different weighting schemes, one using Gaussian weights and one based on the k -nearest neighbor distance. The value of the smoothing parameter was selected by cross-validation, as explained in Section 3.6.

3.5.1 Gaussian weights

The Gaussian weights are given by

$$(3) \quad w_{u,h,i} = \exp\left(-\frac{1}{2} \frac{\|\tilde{X}_i - u\|^2}{h^2}\right), \quad i = 1, \dots, n.$$

3.5.2 Nearest neighbor weights

Here we follow the suggestions of Cleveland and Loader [5]. The number of nearest neighbors, k , depends on the smoothing parameter h as follows

$$k = \min(\lfloor hn \rfloor, n).$$

Denoting by d_k the distance of u to its k 'th nearest neighbor among the vectors $\tilde{X}_1, \dots, \tilde{X}_n$, we set

$$\Delta = d_k \max(1, h^{1/(d+1)}).$$

The nearest neighbor weights are given by

$$(4) \quad w_{u,h,i} = W\left(\frac{\|\tilde{X}_i - u\|}{\Delta}\right),$$

where W is the tricube function

$$W(s) = (1 - s^3)^3 1(0 \leq s \leq 1).$$

If $0 < h < 1$, then only the k nearest neighbors obtain nonzero weights, but if $h \geq 1$, then all the observations obtain positive weights with the closest observations still receiving larger weights than those lying farther away from the evaluation point.

3.6 Cross-validation

We used leave-one-out cross-validation as our parameter selection criterion. That is, we selected the parameters giving the lowest CV-score, i.e., we minimized the function

$$(5) \quad \text{CV}(d, h) = \sqrt{\frac{1}{n} \sum_{i=1}^n [Y_i - \hat{m}_h^{-i}(\tilde{X}_i^{-i})]^2}$$

simultaneously over d and h , where $\hat{m}_h^{-i}(\tilde{X}_i^{-i})$ is an estimate of Y_i corresponding to the dimension d and the smoothing parameter h , which is based on the $n - 1$ training points (X_k, Y_k) , $k \neq i$, with the i 'th pair left out. Explained in more detail, we determine for each i a dimension reduction mapping into dimension d based on the sample with the i 'th pair left out and denote the new d -dimensional vectors by \tilde{X}_k^{-i} , $k = 1, \dots, n$. Finally, $\hat{m}_h^{-i}(\tilde{X}_i^{-i})$ is the local linear fit at the evaluation point \tilde{X}_i^{-i} corresponding to the smoothing parameter h and the size $n - 1$ training data (\tilde{X}_k^{-i}, Y_k) , $k \neq i$.

4 Temperature Reconstruction

The optimal dimension in dimension reduction and the optimal smoothing parameter in local linear regression needed for our temperature reconstruction method were obtained by finding the pair (\hat{d}, \hat{h}) that approximately minimizes the training set cross-validation score $\text{CV}(d, h)$ in (5). The dimension reduction mapping $\varphi : \mathbb{R}^{226} \rightarrow \mathbb{R}^{\hat{d}}$ was then formed based on all training data and the (square root transformed) abundance vectors both in the training set and the prediction set were mapped into new predictor variables $\tilde{X}_i = \varphi(X_i)$, $i = 1, \dots, n$ (training set) and $\tilde{X}_j^* = \varphi(X_j^*)$, $j = 1, \dots, m$ (prediction set). The post ice age temperatures were finally reconstructed as $\hat{Y}_j^* = \hat{m}(\tilde{X}_j^*)$, where \hat{m} was obtained from local linear regression based on \hat{h} and all training data (\tilde{X}_i, Y_i) .

The minimization of the cross-validation score $\text{CV}(d, h)$ was done by finding its smallest value on a two-dimensional (d, h) -grid. Our initial idea was to use a very large and dense grid and simply pick the point where the absolute minimum occurs. However, this resulted in an

unstable minimizer (\hat{d}, \hat{h}) close to a region where the CV-surface starts to rise rapidly (cf. Figure 3). Using such a minimizer sometimes even led to negative reconstructed temperature values. To alleviate this problem large d and small h values were simply cut off from the grid. This had a regularizing effect and resulted in a much smoother optimization problem (cf. Figure 4). The effect of a smaller grid size on the actual minimum value of the CV-score was negligible, less than 0.1°C .

Figure 5 shows temperature reconstructions corresponding to the two dimension reduction and the two kernel weighting schemes tested. The time scale is years ago so the present is on the right and the past on the left. The reconstructions all look quite similar although the principal components versions exhibit less variation. There appear to be several periods of cooling and warming and one can compare the timing of these features with findings in other temperature reconstructions made at different locations on the northern hemisphere; for references to relevant literature, see [8]. We also wanted to understand which of the temperature rises and falls really are there in a statistical sense and therefore decided to apply the SiZer method of Chaudhuri and Marron [4] to the first reconstruction in Figure 5 that uses principal components and Gaussian weights.

The upper panel of Figure 6 shows the reconstructed temperatures as dots together with several smooths based on different smoothing parameter sizes. Such a “family plot” can be thought to give information about the underlying curve at different levels of resolution. The smooth shown in red is based on an automatic smoothing parameter selection rule suggested by Ruppert, Sheather and Wand [9]. Based on confidence intervals for the derivatives of the smooths (computed using approximate simultaneous Gaussian quantiles, cf. [4]), the lower panel attempts to assess which of the features of the smooths are really there. The level of smoothing is on the vertical axis and the pattern of colors along a vertical line indicates significant increases and decreases in temperature. Red is for a rising temperature and blue for a falling temperature. Purple means that no conclusions about the slope can be made (zero within the confidence interval of the derivative) and grey means that the data are too sparse for drawing any conclusions about significance. The white horizontal line corresponds to the red smooth in the upper panel and the dash-dotted curves indicate the effective smoothing window size at each level of smoothing. The significance level used is 20%.

Our reconstructed present temperature is 11.0°C which is also the actual mean July temperature. Many of the cold and warm spells known from other studies can also be identified in our reconstruction. The so-called

Little Ice Age between 500 and 150 years ago together with the preceding Medieval Warm Period and the subsequent warming period between 1840 and 1950 are clearly visible. There is a known minimum at about 4200 years ago which is evident in our analysis, too, and the well-known warm period between 5000 and 7000 years ago is also there. There is a known cool period around 8200 years ago which roughly matches with the minimum between 8000 and 9000 years ago in our reconstruction. The minimum at 7200 years ago does not seem have support from other studies. Also, there is believed to be a broad downward trend in the Holocene temperature but at the chosen significance level we only get part of it in the form of the blue area in the upper part of the lower panel of Figure 6.

5 Discussion

We have shown how local linear regression can be used in organism-based environmental reconstruction. Compared, e.g., with the recent approach of [10, 11], the proposed method is simple and quick to use. The results obtained are also quite comparable with those obtained with standard algorithms, including the cross-validated prediction error, which is 1.1°C (cf. [8]). This figure was obtained with a double layer of cross-validation, where each observation in turn was left out, the parameters d and h were selected using an inner cross-validation loop and the prediction obtained for the left out observation was compared with the true observed value. The reported figure is the square root of the mean of the squared errors (often referred to as RMSEP in the environmental reconstruction literature). The use of the SiZer method in smoothing the reconstructed temperature time series is also novel.

Clearly, the proposed approach has best chance to succeed if the the dimension-reduced abundance-temperature distributions are equal in the training and the prediction sets, that is, if the pairs (\tilde{X}_i, Y_i) and (\tilde{X}_j^*, Y_j^*) are equally distributed for all i and j . At least one would hope to have equality of the regression functions

$$(6) \quad E(Y_i | \tilde{X}_i = x) = E(Y_j^* | \tilde{X}_j^* = x), \quad x \in \mathbb{R}^d.$$

One cannot check the validity of the equation (6) but one can explore visually the closeness of the distributions of \tilde{X}_i and \tilde{X}_j^* . Figure 7 shows the scatter plots of the first five principal components of the (square root transformed) abundance vectors in the training set and the prediction set. Red is for training and blue for prediction. The distributions look quite different for training and prediction vectors so one should not in fact as-




Figure 3: Example of the CV-score function when a grid with large d -values and small h -values is used. Here d is the number of principal components in dimension reduction and h is the smoothing parameter in local linear regression with Gaussian weights.




Figure 4: Example of the CV-score function with a reduced grid size. Here d and h are as in Figure 3.

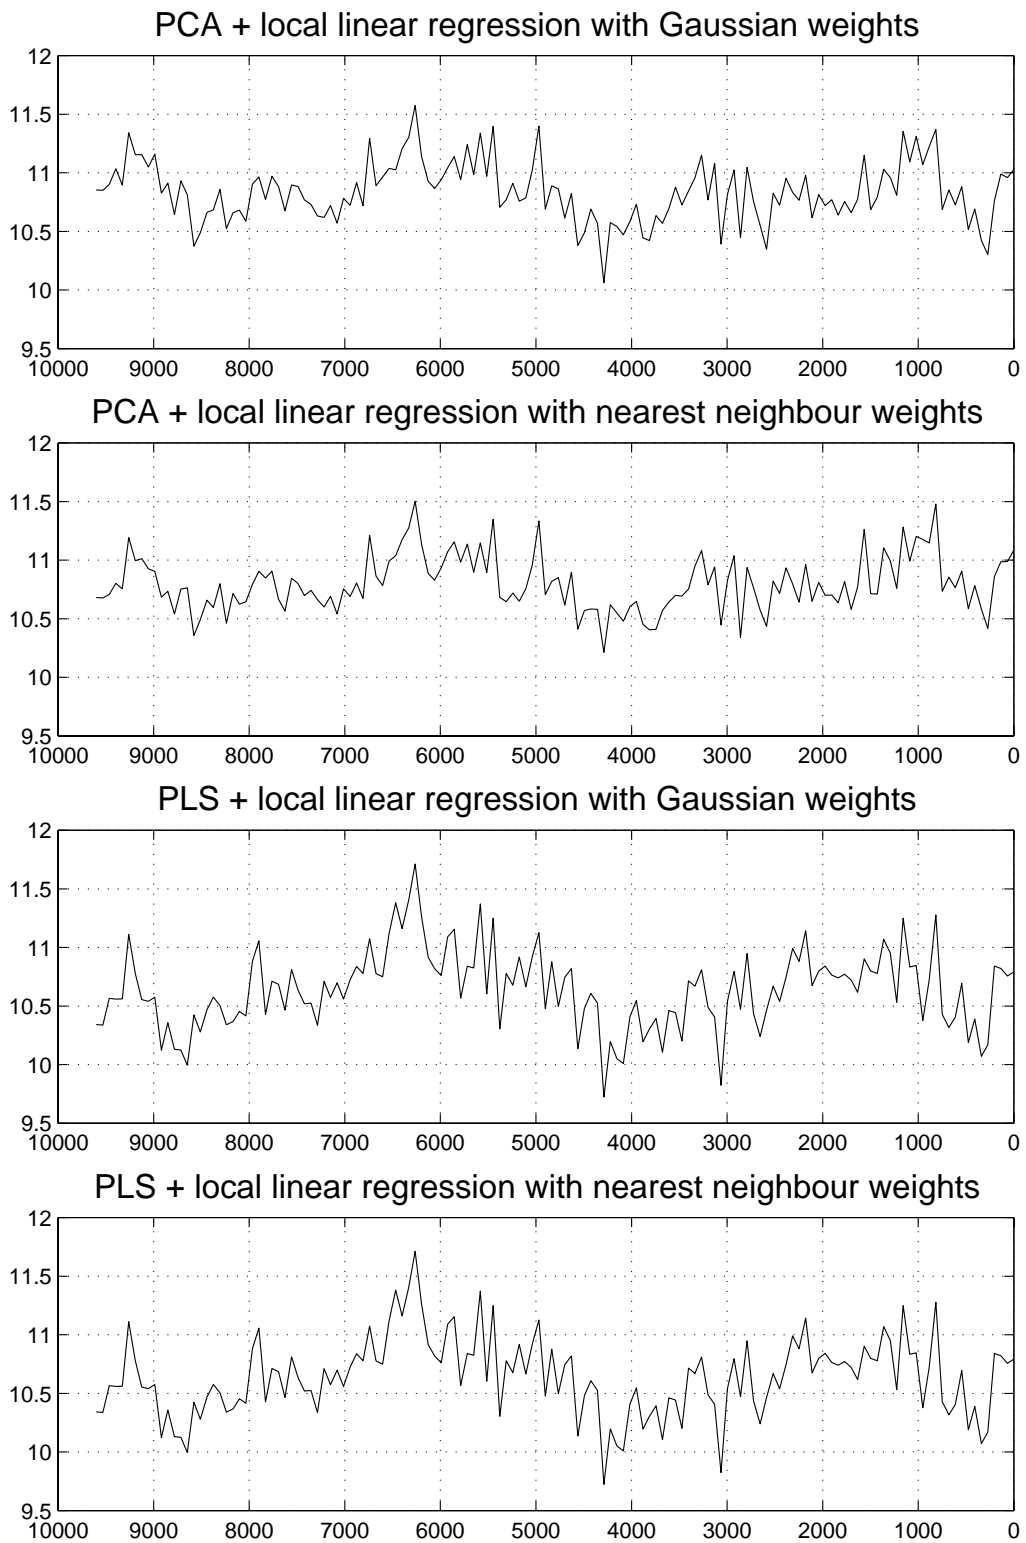


Figure 5: Reconstructed temperatures using two different dimension reduction and two different kernel weighting schemes.

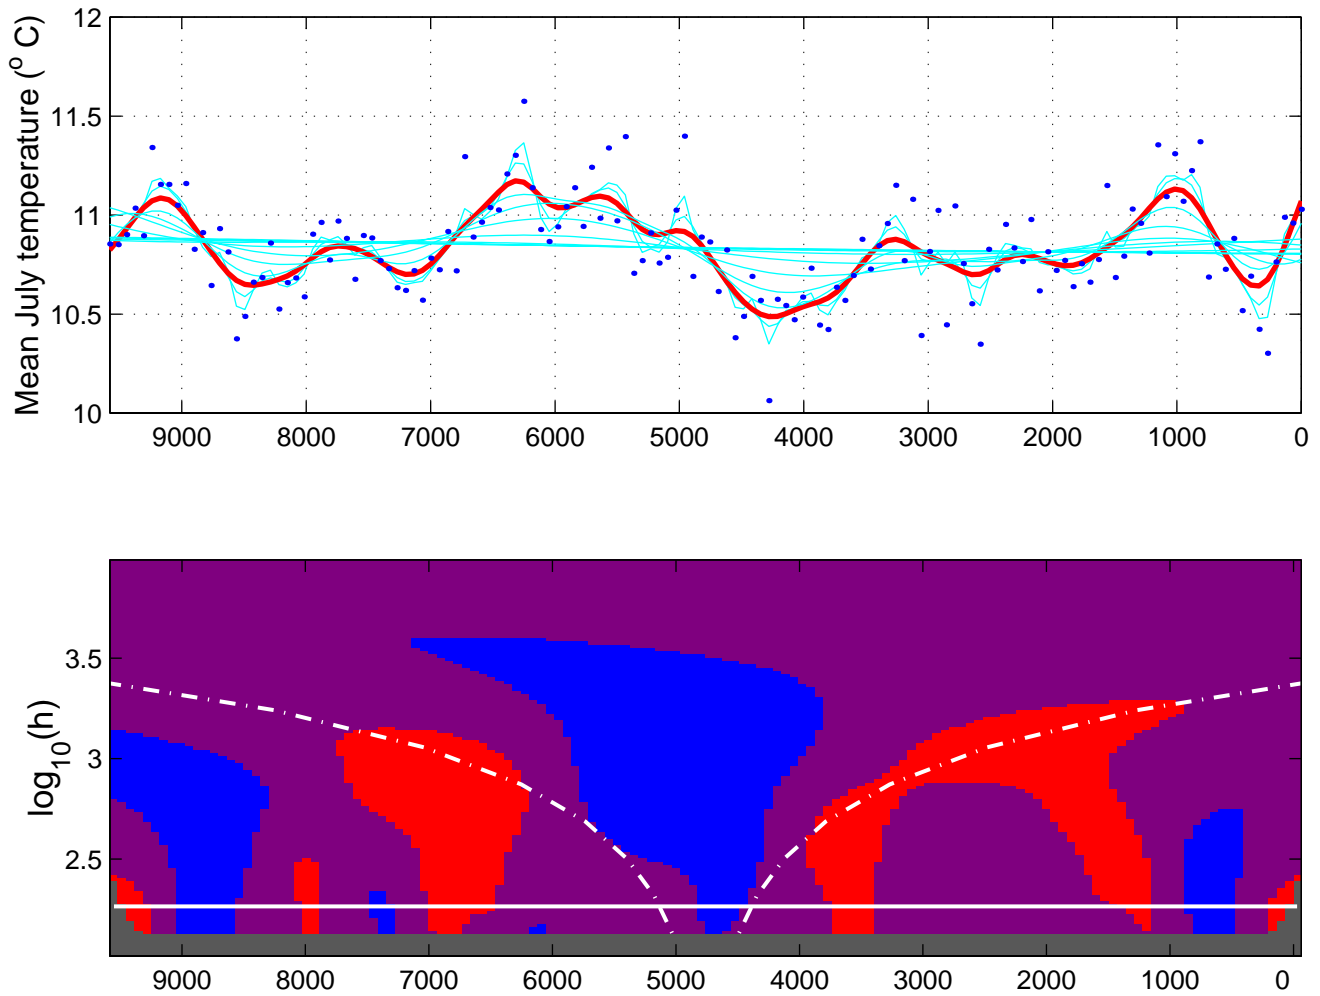


Figure 6: In the upper panel are the reconstructed temperatures based on principal components dimension reduction and local linear regression with Gaussian weights. The individual reconstructions are shown as dots together with several smooths. The horizontal time scale is years ago. In the lower panel the significance of the features of the reconstruction curve are assessed using the SiZer method. The white horizontal line corresponds to the smooth shown in red and the horizontal space between the two dash-dotted curves shows the effective smoothing window size. For more details, see the text.

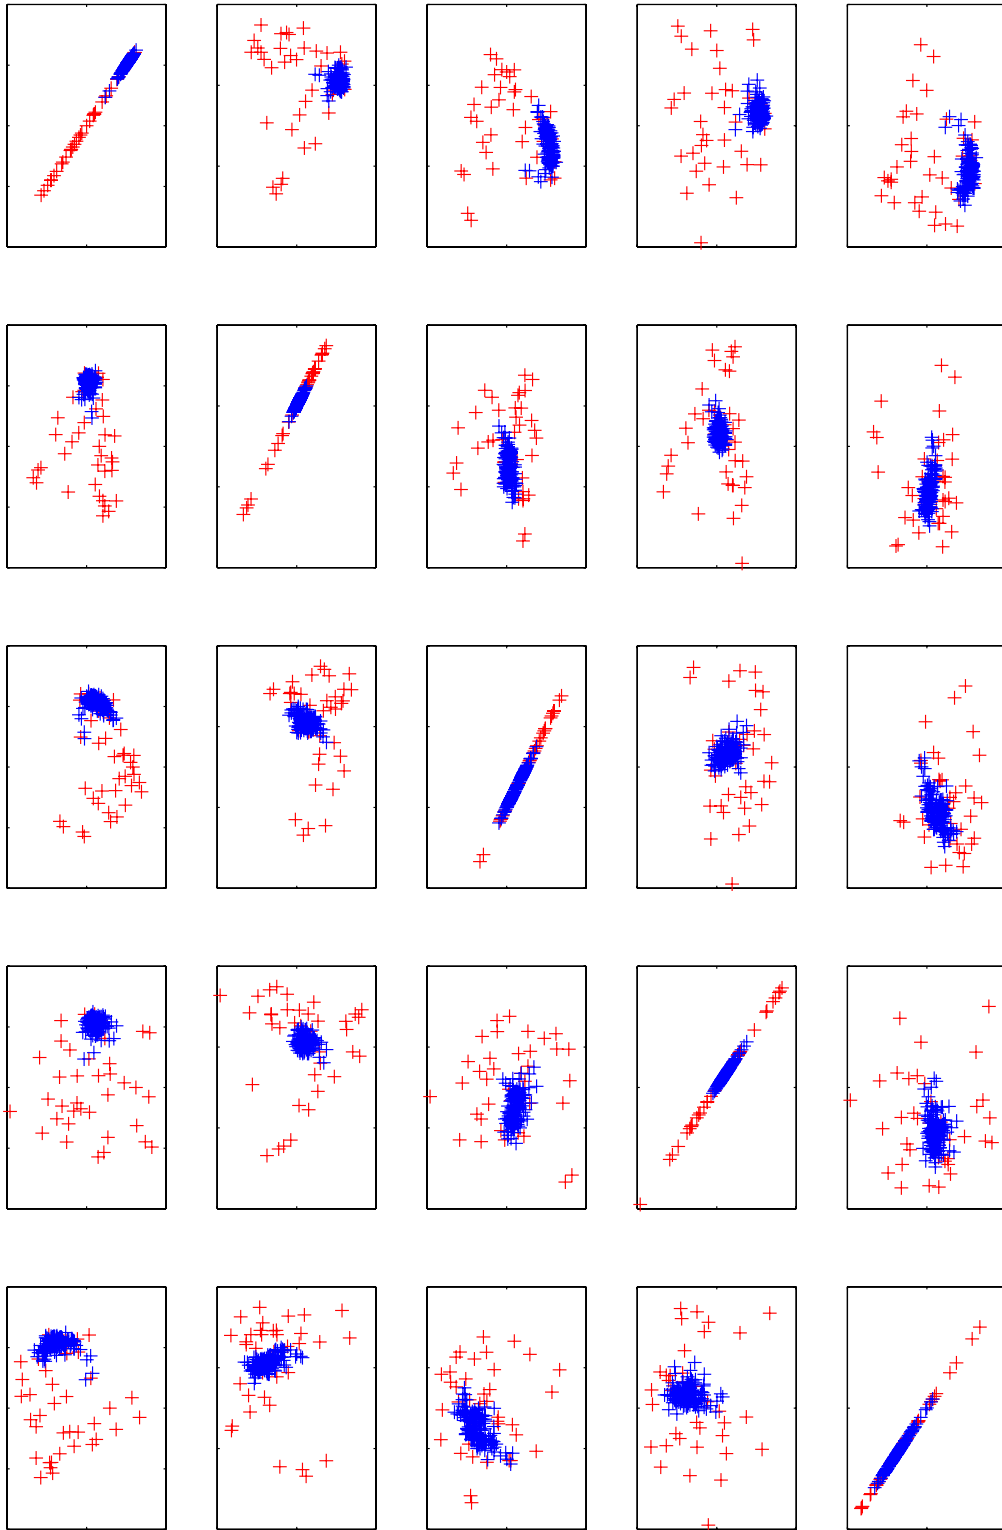


Figure 7: Scatter plots of the first five principal components of the abundance vectors in the training set and the prediction set. Red is training and blue is prediction.

sume that that distributions of (\tilde{X}_i, Y_i) and (\tilde{X}_j^*, Y_j^*) are equal.

The difference in distributions is not surprising if one bears in mind that 38 different lakes were used for the training set and that the prediction set is based on one lake only. An obvious solution would be to try more extensive sampling but here the problem is the cost of data collection. The analysis of sediment samples is very time consuming and identification of diatom taxa has to be done by a single analyst in order to maintain consistency. The preparation of the present data set took about one year. In the future more prediction data from other lakes will become available and the effectiveness of the proposed smoothing approach can then be re-evaluated. Until then one may try to improve regression function estimation in the prediction set by emphasizing the training points in its vicinity. Also, simulation studies can be carried out to estimate the potential of the proposed method. The development of these ideas is left for future work.

6 Summary

In this paper nonparametric smoothing was applied to the reconstruction of the Holocene July mean temperature in the Finnish Lapland. Smoothing was employed both in regressing the temperature on the diatom taxa abundances and in assessing the statistical significance of the observed cooling and warming periods. The reconstructed temperature variation matches quite well the climate changes observed in other studies. The proposed method therefore offers a viable alternative to the standard methods used in the field.

7 Acknowledgement

The graphics of Figure 6 were produced using software provided by Steve Marron. We are grateful to him for useful discussions on the use of SiZer methodology.

References

- [1] H. J. B. Birks. Quantitative palaeoenvironmental reconstructions. In D. Maddy and J. S. Brew, editors, *Statistical Modelling of Quaternary Science Data, Technical Guide 5*, pages 161–254. Quaternary Research Association, Cambridge, 1995.
- [2] H. J. B. Birks. Numerical tools in palaeolimnology—progress, potentialities, and problems. *Journal of Paleolimnology*, 20:307–332, 1998.
- [3] T. Blom, A. Korhola, J. Weckström, T. Laing, J. Snyder, G. M. MacDonald, and J. P. Smol. Physical and chemical characterisation of small subarctic headwater lakes in Finnish Lapland and the Kola Peninsula. *Verheissungen der Internationalen Vereinigung der gesamten Limnologie*, 27, 2000. In press.
- [4] P. Chaudhuri and J. S. Marron. SiZer for exploration of structure in curves. *Journal of the American Statistical Association*, 94(447):807–823, 1999.
- [5] W. S. Cleveland and C. Loader. Smoothing by local regression: Principles and methods. In W. Härdle and M. G. Schimek, editors, *Statistical Theory and Computational Aspects of Smoothing*. Springer, 1996.
- [6] Mike C. Denham. Implementing partial least squares. *Statistics and Computing*, 5, 1995.
- [7] J. T. Houghton, G. J. Jenkins, and J. J. Ephraums, editors. *Climate Change, the IPCC Scientific Assessment*. Cambridge University Press, 1990.
- [8] A. Korhola, J. Weckström, L. Holmström, and P. Erästö. A quantitative holocene climatic record from diatoms in northern Fennoscandia. *Quaternary Research*, 2000. To appear.
- [9] D. Ruppert, S. J. Sheather, and M. P. Wand. An effective bandwidth selector for local least squares regression. *Journal of the American Statistical Association*, 90(432):1257–1270, 1995.
- [10] H. T. T. Toivonen, H. Mannila, A. Korhola, and H. Olander. Applying Bayesian statistics to organism-based environmental reconstruction. *Ecological Applications*, 2000. In press.
- [11] K. Vasko, H. T. T. Toivonen, and A. Korhola. A Bayesian multinomial Gaussian response model for organism-based environmental reconstruction. *Journal of Paleolimnology*, 2000. In press.
- [12] Climate change: State of knowledge. An on-line White House document. <http://www.whitehouse.gov/Initiatives/Climate/content.html>.
- [13] J. Weckström, A. Korhola, and T. Blom. Diatoms as quantitative indicators of pH and water temperature in subarctic Fennoscandian lakes. *Hydrobiologia*, 347:171–184, 1997.