50

Physics of the Earth and Planetary Interiors, 27 (1981) 1–7
Elsevier Scientific Publishing Company, Amsterdam – Printed in The Netherlands

Fitting polar wander paths

R. Thompson [1] and R.M. Clark [2]

[1] *Research School of Earth Sciences, Australian National University, Canberra, A.C.T. (Australia)*

[2] *Department of Mathematics, Monash University, Clayton, Victoria (Australia)*

(Received January 16, 1981; revision accepted March 26, 1981)

Thompson, R. and Clark, R.M., 1981. Fitting polar wander paths. Phys. Earth Planet. Inter., 27: 1–7.

A formalized method of constructing a best-fitting smooth curve, with confidence limits, to sequential data on a sphere is proposed and a method of least-squares alignment of two time series on a sphere is discussed. The procedure, developed from Gould's regression technique for angular variates (corresponding to multiple regression analysis for normally distributed variates), involves fitting cubic splines using a weighted (robust) least-squares approximation. Paths of apparent polar wander (APW) are defined continuously, have continuous first and second derivatives and always lie on the sphere. The cross-validation mean square deviation is used to determine the degree of smoothing. North American and European APW paths, their 95% confidence bands and apparent polar accelerations for the last 500 Ma have been calculated.

1. Introduction

The palaeomagnetic record for any lithospheric plate may be considered as a noisy, unequally spaced time-series of points on a unit sphere. Because of this, workers have preferred to smooth the data by calculating disjoint Fisherian means or running Fisherian means and joining the means by a freehand curve. Such curve-fitting relies heavily on the subjective assessment of the individual, and so cannot be reliably assessed objectively. Instead, we propose a formalized method of constructing a best-fitting smooth curve, which involves a minimal amount of subjective assessment. The procedure involves least-squares fitting of cubic splines to the angular coordinates, the degree of smoothing being determined objectively from the data alone, using the cross-validation criterion. The reliability of the resulting curve is readily assessed by means of its corresponding confidence limits. As an illustration of these statistical methods we calculate:

(1) the North American and European (Baltic Shield) apparent polar wander (APW) paths and their 95% confidence bands for the last 500 Ma;
(2) their apparent polar accelerations.

Our procedure can also be used to solve the forward modelling problem of relative palaeolongitude, i.e. to obtain a unique continental reassembly from palaeomagnetic data. We plan to discuss this aspect of our procedure in a future paper along with the implications of the details of our APW paths, but to concentrate in this paper on the approach and assumptions of our method.

2. Mathematical approach

We treat the palaeomagnetic data as a time sequence of unit vectors defining a sequence of points on the unit sphere, to which we fit a smooth curve by fitting cubic splines by weighted least squares to the corresponding angular coordinates.

We assume that we have n independent observations $\{(\theta_i,\phi_i)\}$ of co-latitude and longitude, respectively, relative to some chosen coordinate system, such that the ith observation (θ_i,ϕ_i) may be regarded as a realization of a random variable having the Fisher (1953) distribution with mean direction (θ_{0i},ϕ_{0i}) and concentration parameter κ. The angles $\{(\theta_{0i},\phi_{0i})\}$ specify discrete points on the required APW path, but are unknown. We assume that these points lie on a smooth curve on the unit sphere, i.e.

$$\begin{matrix}\theta_{0i} = \mathrm{F}(t_i) \\ \phi_{0i} = \mathrm{G}(t_i)\end{matrix} \quad i = 1,2,\ldots,n$$

where F and G are smooth functions of 'time', and t_i denotes the age corresponding to the ith observation. Our problem is to estimate and place confidence limits on the smooth APW path defined jointly by F and G.

Our first step is to approximate the assumed Fisher distribution of (θ_i,ϕ_i). Our calculations show that, provided that κ is fairly large and θ_0 is not too close to either 0 or 180°, θ_i is approximately normally distributed with mean θ_{0i} and variance $1/\kappa$, while ϕ_i is approximately normally distributed with mean ϕ_{0i} and variance $1/\kappa \sin^2 \theta_{0i}$. Furthermore, θ_i and ϕ_i are approximately independent in probability distribution.

Secondly, we assume that F and G can each be represented by a cubic spline with suitable knots. Cubic splines (De Boor, 1978) are piecewise cubic polynomials with smooth joins at the join points or knots, such that their first and second derivatives are continuous everywhere. Cubic splines are known to be extremely successful and adaptable functions for representing empirical relationships (Greville, 1969; Wold, 1974). In particular, they are to be preferred to polynomials or trigonometric functions, which have the property that their behaviour in any small region determines their behaviour everywhere. In contrast, empirical functions (such as F and G) may be expected to have just the opposite property, as do cubic splines, by their piecewise nature.

Under these two assumptions, the data can be expressed by two independent multiple regressions (one for θ, one for ϕ); in each case the regressors are the corresponding values of the B-splines (Wold, 1974; De Boor, 1978) used to define the cubic spline in question. We first fit the cubic spline $\hat{\mathrm{F}}$ to the data $\{(\theta_i,t_i)\}$ on co-latitudes, using the method of least squares (or, equivalently, maximum likelihood) to estimate the necessary coefficients. We then fit a separate cubic spline $\hat{\mathrm{G}}$, with possibly different knots, to the longitudes, using weighted least squares, i.e. by minimizing $\Sigma w_i\,(\phi_i - \hat{\mathrm{G}}(t_i))^2$, where $w_i = \sin^2 \hat{\mathrm{F}}(t_i)$.

Gould (1969), in considering a similar problem of regression with angular variates, used a different method involving a complex iterative procedure that does not provide confidence bands for the fitted curve. In contrast, our method is straightforward and yields valid confidence bands directly, as long as our normal approximation to the Fisher distribution is reasonable. When θ_0 is close to either 0 or 180° this approximation cannot hold, even for very large κ. However, theoretical calculations and numerical integration show that the approximation improves rapidly as θ_0 tends to 90° and as κ increases. For example, when θ_0 is 50°, the approximation is excellent for κ greater than 40, and is tolerable for κ as small as 15.

These difficulties associated with extreme latitudes can be avoided by positioning the unit vectors so that the latitudes are reasonably small. We did this for the palaeomagnetic data by first rotating the data set so that its mean and trend lay near the equator. After curve fitting and the computation of associated confidence limits, the corresponding fitted curves and confidence limits were rotated back to the original coordinate system.

In practice, the degree of smoothing, determined in this case by the number and location of knots for the two curves F and G, is more crucial than the actual method of smoothing. If smoothing is too heavy, fine details of the APW path may be lost, while too little smoothing may produce spurious kinks or loops. We prefer to assess the level of smoothness using the internal evidence of the data, rather than by assuming a priori a particular level. The cross-validation method (Stone, 1974) is a remarkably adaptable and successful technique for objectively determining the appropriate degree of smoothing and can be applied to our problem. The method is temporarily to set aside data (a validation sample) and construct a curve, with a certain degree of smoothing, to the remaining data. The performance of the curve is then assessed by the deviations, from the curve, of the data set aside. An oversmoothed curve will be a poor estimation of the validation data, as it will not have followed the main data set sufficiently closely. An undersmoothed curve will also be a poor estimation as it will have followed the random noise rather than approximating the true underlying function. In the full cross-validation method, which we employed, one data point is set aside at a time and its deviation from the fitted curve calculated. This process is repeated for every data point and a cross-validation mean square deviation is computed from all the deviations. The whole procedure is repeated for other levels of smoothness to give a cross-validation mean square deviation for each level. The minimum mean square deviation determines the appropriate level of smoothness, and the final best-fitting curve is computed, at this level of smoothness, using the full data set. The entire procedure is carried out with the co-latitudes and longitudes separately, since F and G do not necessarily have the same degree of smoothing or the same knots.

In our initial formulation of the problem, the concentration parameter κ was assumed implicitly to be constant and, in particular, independent of the age t_i of the ith observation. Inspection of the data indicates, however, that the older measurements are more variable than the younger. The initial assumption of constant κ is not crucial; our method can be easily modified in the usual way (Seber, 1977) if, for example, κ is assumed to be a simple, slowly varying function of time, of known form. An alternative and equivalent method of dealing with data of unequal precision is to perform weighted least-squares estimation, each observation being weighted separately and independently according to its perceived precision. The bi-square weighing method of Mosteller and Tukey (1977) is a robust and resistant way of doing this. It has the added advantage that extreme or outlying observations are automatically allowed for, by being given low weight.

We shall now illustrate our method by applying it to real data before discussing further (Section 5) the underlying assumptions and comparing alternative smoothing methods.

3. Palaeomagnetic data

North American and European palaeomagnetic data from rocks younger than 500 My were chosen for analysis. The data were taken from GJRAS compilations (McElhinny, 1972, 1973; McElhinny and Cowley, 1977, 1978, 1980), which include data published prior to 1979, and later additional European data from Briden and Duff (1981). The European palaeomagnetic data were taken from sampling localities west of the Ural mountains and north of the Alpine mountains. The minimum selection criteria of McElhinny (1973) and the minimum criteria (2) of Irving et al. (1976) for A-class poles were used in a sorting programme, written by B. Goleby, "to provide a first-stage filter, by which those results which can on common sense ground be considered of little use in tracing the past history of the field...can...be separated from the main body of the data" (Irving et al., 1976). Further data noted in McElhinny's catalogue as being remagnetized or of intermediate direction or from areas of recent tectonic activity (e.g. the Alpine belt, Spain, the Western Cordillera of North America) were also rejected. Ages in My were assigned to the palaeomagnetic pole positions by converting the geological age in McElhinny's catalogue according to the Geological Society of London time scale (Harland et al., 1964). Finally, the polarities of the magnetizations were determined from continuity considerations.

After data selection using our minimum criteria the North American data set consisted of 209 ordered unit vectors and the European set 326.

4. Apparent polar wander paths

The degree of smoothing of the data is specified, in our curve-fitting procedure, by the number of internal knots, or equivalently, the number of spline pieces. The initial cross-validation computations showed a clear optimum of four spline pieces for the North American data and six for the European data, for both co-latitude and longitude. We then applied the bi-square weighting technique to take account of both the unequal variability in the observations and the presence of outliers. This method is iterative, but it stabilised after only one

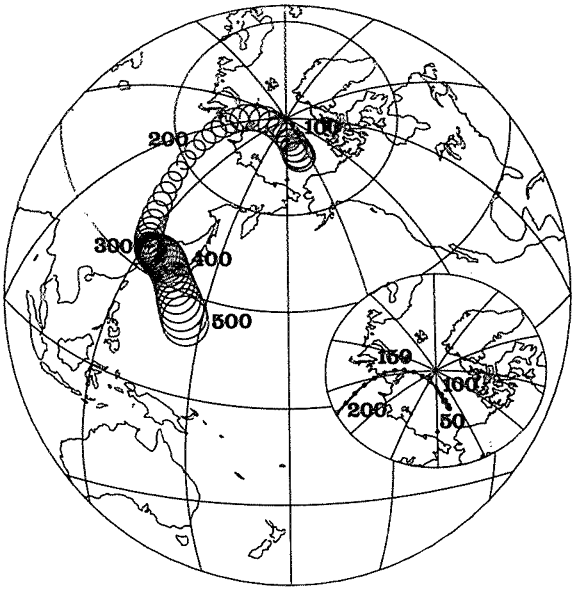

Fig. 1. North American apparent polar wander (APW) path for the last 500 Ma, with 95% confidence limits drawn at 10 Ma intervals. The tight grouping of confidence ellipses near 400 Ma corresponds to a period of slow APW. The inset shows the North American APW path from 200 to 10 Ma as a tight loop.

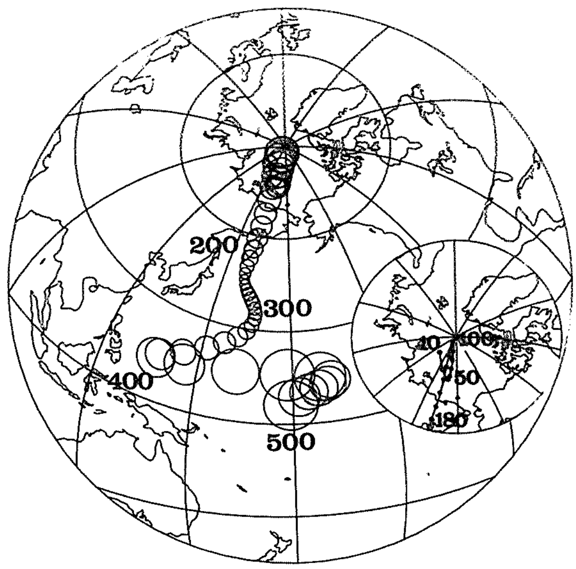

Fig. 2. European APW path for the last 500 Ma, with 95% confidence limits drawn at 10 Ma intervals. The inset shows the APW path from 180 to 10 Ma.

iteration on all four data sets. A thorough analysis of the residuals (Seber, 1977) showed that, with one exception, the subsequent fitted curves were a satisfactory fit to the data, taking account of the relative weights assigned to the observations. The exception was the curve for the European co-latitudes. Inspection and a runs test on the residuals showed it to be misfitted near 400 Ma at the Siluro-Devonian polar shift. We therefore used NEWNOT, De Boor's (1978) knot placement algorithm, adding one knot at a time to achieve a satisfactory fit to these data, using ultimately seven spline pieces (i.e. one additional internal knot).

Our best-fitting APW paths are shown in Figs. 1 and 2, with their 95% confidence limits drawn at 10 Ma intervals. These confidence ellipses were computed using the fact that, under our assumptions in Section 2, the estimated co-latitude $\hat{F}(t_0)$ and longitude $\hat{G}(t_0)$ at any arbitrary time t_0 are independently, normally distributed with variances computed by the usual formulae (Seber, 1977) from multiple regression. Our interpretation of these limits is that we can be 95% sure that the actual angular coordinates $(F(t_0), G(t_0))$ of the apparent palaeomagnetic pole at time t_0 lie somewhere within the corresponding confidence ellipse for that time. As expected, the sizes of the confidence regions vary along the path. In general, they are larger where the data are sparse and/or variable, and smaller where the data are more abundant and/or precise.

The European APW path (Fig. 2) is more complicated than the North American path (Fig. 1). Its longitudinal sweeps are the result of a clockwise rotation of Europe during the Ordovician and Silurian and a subsequent anticlockwise rotation. The central part of the path is dominated by the Permo-Triassic drift of Europe northward. Sharp changes in the direction of the European APW path occur in the mid-Cretaceous and mid-Tertiary (Fig. 2 inset). The North American APW path is markedly different in shape from the European path before 400 Ma, because of the relative movements between Europe and North America during the formation of the Iapetus Ocean. A trend corresponding to northerly drift similar to that of the European path is visible in the central section. For the last 150 Ma the North American path again has a different shape from that of the European path (compare insets of Figs. 1 and 2), owing to the recent phase of continental drift between North America and Europe.

The rate and acceleration of APW are also defined continuously for each path. The typical (median) absolute rates are 0.31 and 0.22° Ma^{-1} for North America and Europe, respectively. A maximum acceleration of 0.12° Ma^{-2} is found at 400 Ma on the European APW path. It is physically more interesting to estimate continental velocities than APW path velocities. Gordon et al. (1979) presented a technique for estimating minimum absolute continental velocities from APW paths and have applied their technique to Irving's (1977) APW paths. They noted the restriction imposed on their technique by the high degree of smoothing in the running Fisherian mean analysis, and the consequent loss of information, particularly when the direction of plate motion is changing rapidly. This restriction is minimized with our approach. The technique of Gordon et al. (1979) could also be used to estimate minimum continental accelerations from our APW paths.

5. Discussion

We shall now examine further the assumptions and properties of our method, and compare it with two alternative methods for smoothing a time sequence of noisy unit vectors.

5.1. Assumptions

Our method starts with the assumption that our observations follow the Fisher distribution with the mean direction (and possibly the concentration parameter) varying with time, and that this distribution can be adequately approximated by a particular normal distribution in the angular coordinates (θ,ϕ). However, it is likely (Onstott, 1980) that the original observations would have been only approximately Fisher-distributed in the first place. Nevertheless, we can expect our final curve and confidence limits to be robust to departures from the assumed normality of the data, by an extension of the results of Cochran (1947). In particular, the usually serious effects of outliers would be nullified by the use of bi-square weighting.

An implicit but important second assumption of our method is that the ages t_i assigned to the observations are known to within 10 Ma. Clearly, this is not so with palaeomagnetic data. To examine the effect of age uncertainties, we conducted several numerical experiments in which we modified the assumed ages by adding pseudorandom age errors with standard deviations of up to 15% of the assumed ages. These random perturbations had a negligible effect on the resulting APW path and its confidence limits.

A desirable property of any procedure for smoothing a sequence of unit vectors is invariance under rotation, i.e. the operations of smoothing and rotation should be commutative. Equivalently, the resulting smooth curve on the unit sphere should be independent of the choice of coordinate axes used in defining both the input data and the resulting curve. By smoothing the angular coordinates relative to an arbitrary coordinate system, our method unfortunately is not invariant under rotation as required. We chose to sacrifice this property so that we could derive confidence limits for our fitted curve, by formulating the problem as one of multiple regression with approximately normally distributed errors. However, with these particular data, the departure from invariance is negligible in practical terms. Numerical experiments showed that, because of the relatively high variability of our data and the flexibility of cubic splines, there was little practical difference between smoothing first and then rotating, and rotating followed by smoothing.

5.2. Alternative methods

When few palaeomagnetic data points were available an appropriate and extremely successful approach was freehand construction of APW paths (Creer et al., 1954). With the production of further data the most common smoothing methods have been to compute disjoint Fisherian mean vectors (Irving, 1964; McElhinny, 1973) and more recently, with the proliferation of palaeomagnetic data, to compute running Fisherian mean vectors (Irving, 1977) using the observations in successive, generally overlapping time intervals, and then to draw a freehand curve through these successive means after plotting them on some suitable graph. The successive Fisherian means are of course invariant under rotation, but the final freehand smoothing may well depend on the coordinate system used to plot the means, as well as on the subjective judgement involved in freehand smoothing. The final curve, like ours, always lies on the unit sphere, and would be only marginally affected by minor perturbations in the assumed ages.

However, these methods too have disadvantages. Firstly, the degree of smoothing is determined primarily by the time span over which the successive mean vectors are computed. In the past, the choice of time span has been arbitrary and subjective, although in principle the method of cross-validation could be employed here also. If this time span is too large, the data are oversmoothed, and the concentration parameter κ can be considerably underestimated, since the systematic trend in the data would not have been completely removed. Consequently, some fine details of the APW path could be lost, and confidence cones for successive points on the path would be generally too wide. Secondly, all observations within any given time span are given equal weight when the corresponding Fisherian mean direction is computed. However, since the aim is to estimate the APW path at the midpoint of each time span, those observations farther from the midpoint should be given less weight. In addition, observations with less intrinsic precision should receive less weight, independent of their age relative to the midpoint age. This method is thus particularly sensitive to outlying observations. In contrast, both types of weights are handled automatically in our method.

Parker and Denham (1979) have outlined an alternative mathematical method for smoothing that is invariant under rotation. The method is an extension to spherical data of the smoothing cubic spline of Reinsch (1967, 1971) and Schoenberg (1964) which has the property of providing the best compromise between smoothness and goodness of fit, in a certain mathematical sense. Their procedure has a number of disadvantages, however. First, their cubic spline requires a knot at every point t_i, and so could be sensitive to errors in the assumed ages of the observations. Secondly, since this curve fitting does not correspond to either least-squares or maximum-likelihood estimation, it does not seem possible to perform other statistical analyses and, in particular, to construct confidence intervals. Thirdly, it has the theoretical drawback that the resulting curve is not constrained to lie on the sphere. In practical terms, however, we would expect their method to lead to smooth curves similar to those given by our method.

5.3. Prospects

A useful development, which in principle is straightforward, might be to combine more detailed data selection techniques with our least-squares fitting procedure by using a weighting technique. Paleomagnetic poles could be graded, as for example in the schemes of McElhinny and Embleton (1976) and of Briden and Duff (1981), and then weights associated with each grade for use in the regression calculation.

Our procedure can be adapted to help obtain a quantitative and mathematically reproducible continental reassembly using only palaeomagnetic data. The problem involves aligning two time series on a sphere. We have previously described a method of aligning two time series (Clark and Thompson, 1979) by examining the residual mean square error about a single curve fitted to the pooled data set of the two series. A search technique is used to find the minimum residual mean square error and hence the best alignment. In order to align two time series on a sphere we have to modify our method to deal with unit vectors rather than scalars, so we fit curves to the pooled palaeomagnetic data set for various reconstructions with the procedure outlined in this paper. Then the best alignment is found by minimizing the sum of the squares of the solid angle residuals about the APW path, instead of minimizing the R_1^2 of eq. 13 of Clark and Thompson (1979).

5.4. Summary

There is as yet no ideal method with all the necessary desirable properties for smoothing on a sphere. Clearly some compromise is necessary. We believe that there is little point in constructing a smooth curve unless one can make valid and objective statements, using confidence limits, about its reliability or precision. Accordingly, we have sacrificed invariance under rotation for simplicity and the ability to compute confidence limits, by reformulating the problem as a standard one of multiple regression. For the North American and European palaeomagnetic data, the lack of invariance turned out to be negligible in practical terms. By using a formal mathematical method, we have kept subjective assessments to a minimum, and it has been easy to conduct numerical experiments to examine the sensitivity of our results to our initial assumptions.

Acknowledgements

We thank B.A. Duff, B.J.J. Embleton and M.W. McElhinny for helpful discussions about the palaeomagnetic data.

References

Briden, J.C. and Duff, B.A., 1981. Pre-Cenozoic palaeomagnetism of stable Europe and the British Isles. Rep. Int. Dynam. Commiss., 10 (in press).

Clark, R.M. and Thompson, R., 1979. A new statistical approach to the alignment of time series. Geophys. J.R. Astron. Soc., 58: 593–607.

Cochran, W.G., 1947. Some consequences when the assumptions for the analysis of variance are not satisfied. Biometrics, 3: 22–38.

Creer, K.M., Irving, E. and Runcorn, S.K., 1954. The direction of the geomagnetic field in remote epochs in Great Britain. J. Geomagn. Geoelectr., 6: 163–168.

De Boor, C., 1978. A Practical Guide to Splines. Springer, New York, NY.

Fisher, R.A., 1953. Dispersion on a sphere. Proc. R. Soc., Ser. A, 217: 295–305.

Gordon, R.G., McWilliams, M.O. and Cox, A.J., 1979. Pre-Tertiary velocities of the continents: A lower bound from palaeomagnetic data. Geophys. Res., 84: 5480–5486.

Gould, A.L., 1969. A regression technique for angular variates. Biometrics, 25: 683–700.

Greville, T.N.E., 1969. Introduction to Spline Functions. Academic Press, New York, NY. pp. 1–36.

Harland, W.B., Smith, A.G. and Wilcock, B., 1964. The phanerozoic time-scale. Q. J. Geol. Soc., London, 1205: 1–458.

Irving, E., 1964. Palaeomagnetism. Wiley, New York, NY.

Irving, E., 1977. Drift of the major continental blocks since the Devonian. Nature (London), 270: 304–309.

Irving, E., Tanczyk, E. and Hastie, J., 1976. Catalogue of paleomagnetic directions and poles. Geomagnetic Service of Canada, Ottawa Geomagn. Ser. 6; 70 pp.

McElhinny, M.W., 1972. Palaeomagnetic directions and pole positions XIII. Geophys. J.R. Astron. Soc., 30: 281–293.

McElhinny, M.W., 1973. Palaeomagnetism and Plate Tectonics. Cambridge University Press, London.

McElhinny, M.W. and Cowley, J.A., 1977. Palaeomagnetic directions and pole positions. XIV. Geophys. J.R. Astron. Soc., 49: 313–356.

McElhinny, M.W. and Cowley, J.A., 1978. Palaeomagnetic directions and pole positions. XV. Geophys. J.R. Astron. Soc., 52: 259–276.

McElhinny, M.W. and Cowley, J.A., 1980. Palaeomagnetic directions and pole positions. XVI. Geophys. J.R. Astron. Soc., 61: 549–571.

McElhinny, M.W. and Embleton, B.J.J., 1976. Precambrian and Early Palaeozoic palaeomagnetism in Australia. Philos. Trans. R. Soc. London, Ser. A, 280: 417–431.

Mosteller, F. and Tukey, J.W., 1977. Data Analysis and Regression, Addison–Wesley, Reading.

Onstott, T.C., 1980. Application of the Bingham distribution function in paleomagnetic studies. J. Geophys. Res., 85(B3): 1500–1510.

Parker, R.L. and Denham, C.R., 1979. Interpolation of unit vectors. Geophys. J.R. Astron. Soc., 58: 685–687.

Reinsch, C.H., 1967. Smoothing by spline functions. Numer. Math., 10: 177–183.

Reinsch, C.H., 1971. Smoothing by spline functions. II. Numer. Math., 16: 451–454.

Schoenburg, I.J., 1964. Spline functions and the problem of graduation. Proc. Natl. Acad. Sci., 52: 947–950.

Seber, G.A.F., 1977. Linear Regression Analysis. Wiley, New York, NY.

Stone, M., 1974. Cross-validatory choice and assessment of statistical predictions. J.R. Stat. Soc. B., 36: 111–147.

Wold, S., 1974. Spline functions in data analysis. Technometrics, 16: 1–11.