AN IMPROVED ALGORITHM FOR CROSSDATING TREE-RING SERIES

MARTIN A. R. MUNRO

Palaeoecology Centre
The Queen's University of Belfast
Belfast, Northern Ireland

ABSTRACT

The CROS algorithm for crossdating tree-ring series has proved useful. Because it uses Student's t to test correlations which are not independent between autocorrelated tree-ring series, it does not give a good measure of the relative significance of high correlations. It can be improved by transforming the correlation coefficients to normally distributed values, and giving a conservative estimate of the significance of the highest of these by a revised t, derived from its studentized deviation from the mean value. The improved algorithm should help any dendrochronologist who routinely dates oak timbers.

Der CROS—Algorithmus zum wechselseitigen Vergleich von Jahrringfolgen hat sich bewährt. Er benutzt den t- Wert von Student zur Prüfung von Korrelationen. Da diese aber bei autokorrelierten Jahrringfolgen nicht voneinander unabhängig sind, stellt der t-Wert kein gutes Signifikanzmaß für hohe Korrelationen dar. Eine Verbesserung wird durch die Transformation der Korrelationskoeffizienten in normalverteilte Wert erreicht. Danach ist eine konservative Signifikanzschätzung für die größten Korrelationskoeffizienten möglich, und zwar mit Hilfe eines modifizierten t- Wertes, der sich von seiner nach Student bestimmten Abweichung vom Mittelwert ableitet. Dieser verbesserte Algorithmus ist hilfreich bei der routinemäßigen Datierung von Eichenhölzern.

L'utilisation de l'algorithme "CROS" pour interdater des séries dendrochronologiques a été prouvée. En utilisant le test "t" de Student pour estimer des corrélations qui ne sont pas indépendantes en raison de l'autocorrélation existant entre les séries de cernes, cet algorithme ne donne pas une bonne mesure de la signification relative des corrélations élevées. Il peut être amélioré par transformation des coefficients de corrélation en valeurs distribuées normalement et en donnant une première approximation de la signification de la plus élevée de ces valeurs, par l'utilisation d'un "t" revisé. Ce dernier est dérivé d'après sa déviation "studentized" de la valeur moyenne. L'algorithme alélioré devrait aider les dendrochronologues qui datent en routine des poutres de chêne.

INTRODUCTION

The CROS algorithm was first published over ten years ago (Baillie and Pilcher 1973), and has been widely used by dendrochronologists working with oak and other species that do not commonly have missing or multiple rings. The algorithm is now used in BASIC programs for crossdating on microcomputers as well as in several variants of the FORTRAN program which appeared in the original publication. Several people have pointed out defects in it (Barefoot and others 1978; Orton 1983; Steward 1983), but no one has yet produced a revised version for general use.

No current algorithm, whether encoded in a computer program or in some other form, can supplant the trained dendrochronologist in the crossdating of tree-ring series. This point has to be made clear at the outset, because an algorithm that could be applied by any unskilled person and produce the same results as a dendrochronologist would be very different from one that is merely a check on other methods of crossdating or a data-reduction tool.

Baillie (1982:85) has described how the results from a crossdating algorithm should be used in conjunction with other methods—the graphs of the two series should also be compared by eye, and further series should be crossdated with them to replicate the original results. Ideally there should be a range of different algorithms available to the dendrochronologist, ranging from ones that could be used to check large numbers of tree-ring series quickly to those that would take much longer to check (particularly important series) in a way that would give a more accurate result. Both the original version of CROS and the version described here fall into the former category of rapid data-screening algorithms. Neither of them will crossdate series much less than 100 years long, and neither will crossdate series from areas which are far apart, or separated from each other by natural boundaries such as mountain ranges.

DATA TRANSFORMATION

Raw tree ring series do not lend themselves to crossdating, except in the case of people comparing graphs by eye. While several methods of transforming them can lead to better results, an inappropriate method can give dates that are not only wrong, but also seem to be supported at high levels of significance by several statistical tests. The best methods are undoubtably those that fit some kind of autoregressive (AR) or moving-average (MA) model to the series and then use the residuals from the model for crossdating. For example Laxton and Litton (1983) fit AR models, minimizing the value of Akaike's information criterion to determine the order parameter to use, and Steward (1983) fits ARIMA models, determining the parameters by a variety of criteria. These methods can be justified by the theory of the processes leading to variations in ring width (Fritts 1976:26), but unfortunately require a large amount of computation if they are to be sure of fitting the best model. On the other hand if the crossdating uses residuals from a running mean (such as the five year mean in the original CROS program) not much extra computation is needed. But it will be less justified on theoretical grounds, and may distort cyclic trends in the original series—for example a five year running mean will tend to enhance cycles with periods between 5 and 2.5 years. The best compromise seems to be to generalize the running mean into a symmetrical nonrecursive digital filter:

$$\acute{a}_i = \sum_{k=-m}^{+m} w_k a_{i+k}$$

where $a_0, a_1, \ldots a_{n-1}$ are the logarithms of the original series, the corresponding $\acute{a}$ values are the transformed indices, and $w_{-m}, w_{1-m}, \ldots w_0, \ldots w_m$ are the weights of the filter, with $w_j = w_{-j}$. Much has been written on filters of this kind (e.g. Hamming 1977), and they have already been applied to tree-ring series (LaMarche and Fritts 1972; Hughes et al. 1978; Guiot et al. 1982; Briffa et al. 1983), but it is important to bear in mind that filters constructed for other purposes (such as dendroclimatology) will not necessarily be the best ones to use for crossdating problems. I have tested several different transformation methods by using them to process groups of series whose relative ages are known, then crossdating every pair of series that is known to overlap in time to see which method produces the maximum number of correct datings. Such tests show filters can produce better results than either the method used in the original CROS program, or the differencing of logarithmically transformed series recommended by Hollstein (1980) and Steward (1983).

A good filter will obviously emulate other techniques (such as the fitting of splines or polynomials (e.g. Graybill (1982)) in removing long-term trends from the series (such as those due to the growth of individual trees and varying competition within stands). Experiments show that if the filters are more effective than the other techniques in removing low-frequency information from the series, the resulting datings will be improved, particularly in cases where the master chronologies used are based on simple averages of the raw ring-width series rather than indices. The autocorrelation functions of three such chronologies (Figure 1) have peaks and troughs at several different lags, and it is hardly surprising that two of them (Figure 1 a and b) crossdated correctly only when oscillations with periods of twelve years or longer were removed. Chronologies such as these are unfortunately in common use for archaeological dating applications in Britain and Ireland, so this is far from being an isolated bad example. Moreover, even in cases where the chronologies have used polynomial indexing methods it may be worthwhile discarding more of the low frequencies. One can demonstrate this using the seven index chronologies given by Baillie (1973), Pilcher

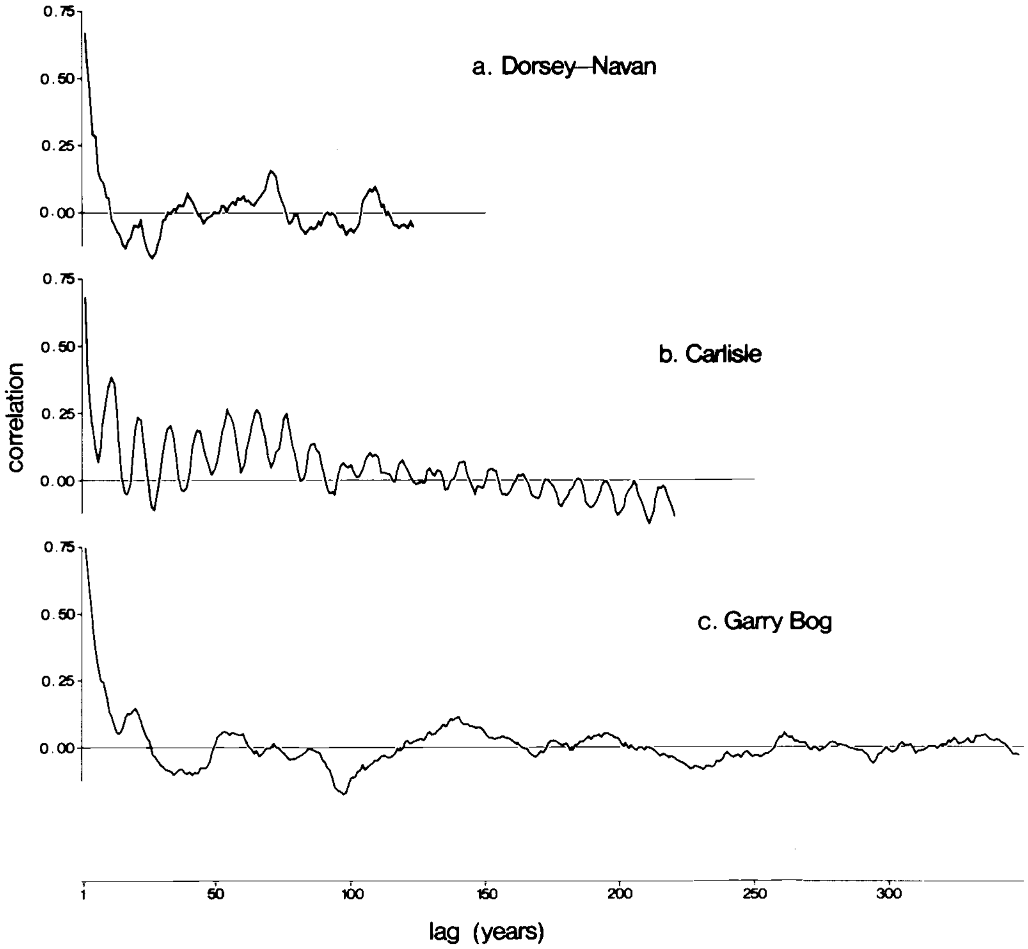

Figure 1. Autocorrelation functions of some master chronologies.

(1976), and Pilcher and Baillie (1980), which are now deposited in the International Tree Ring Data Bank. If one truncates these to the period 1850-1969 and crossdates every pair of chronologies, one gets a greater number of correct datings after discarding frequencies with a period greater than or equal to twelve years than with the original chronologies.

If the filter removes too much low frequency information the dating will again suffer, and since the removal of information is equivalent to a reduction in the degrees of freedom of the series, statistical tests may give a spurious significance to these incorrect datings. In a test of three filters, arranged so that they removed periods of over eight, nine and ten years respectively, on the dating of all 120 possible pairs of 16 samples from the beginning of the subfossil oak chronology from Swan Carr in the north of England (M.G.L. Baillie, unpublished data), the filter that removed periods of over ten years gave rise to the best datings, both in terms of the number of correct datings and the significance levels for these (see Tables 1 and 2). Figure 2d shows the transfer function for this filter: tests show that for crossdating purposes filters such as this, which have good spectral properties (in the sense of having narrow transition bands and no large ripples in their transfer functions), seem to be better than filters that have good statistical properties (such as the Gaussian tapered filters used by Briffa et al. (1983)).

Table 1. Weights for three filters.

Weight number	Cutoff period (years) 10	9	8
w_0	0.200000	0.222222	0.250000
w_1	0.183152	0.200291	0.220333
w_2	0.138929	0.143860	0.146079
w_3	0.083054	0.075628	0.061750
w_4	0.032927	0.019160	0.000000
w_5	0.00000	−0.012451	−0.025742
w_6	−0.013667	−0.020137	−0.023252
w_7	−0.013569	−0.014051	−0.010089
w_8	−0.007808	−0.005277	0.000000
w_9	−0.002515	0.000000	0.003025
w_{10}	0.000000	0.001213	0.001888
w_{11}	0.000340	0.000569	0.000409

Table 2. Percentages of correct datings produced by using the three filters on 16 ring-width series from sub-fossil oaks from Swan Carr, England.

	Cutoff period (years) 10	9	8
	67.7%	63.3%	60.0%
(number)	(61)	(57)	(54)

The definition of the filtering method given above mentions a logarithmic transformation. Both Hollstein (1980:14-15) and the original version of CROS use this, which can be justified to a certain extent by the improvements it brings to the dating, and by the way it tends to bring the distribution of the ring widths closer to a normal distribution; but that is not to say that it produces normally distributed indices. For example if one tests the logarithms of the raw ring widths of the 16 samples from Swan Carr using Lilliefors' modification of the Kolmogorov-Smirnov test (Rohlf and Sokal 1981, Table 33), only 4 of the samples show no significant departures from normality ($P > 0.05$).

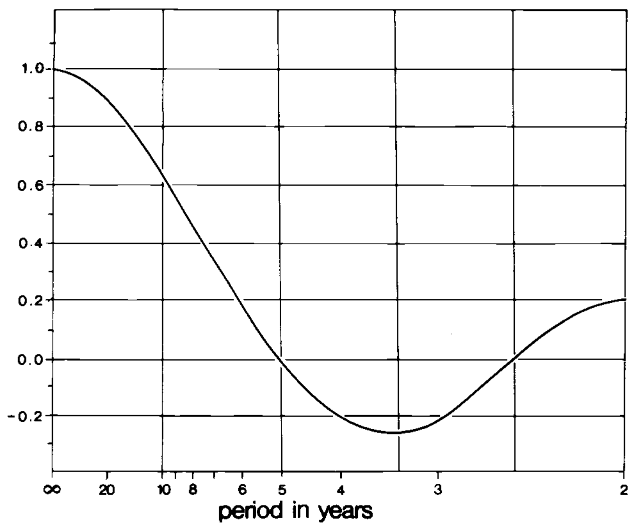

a. 5—year running mean

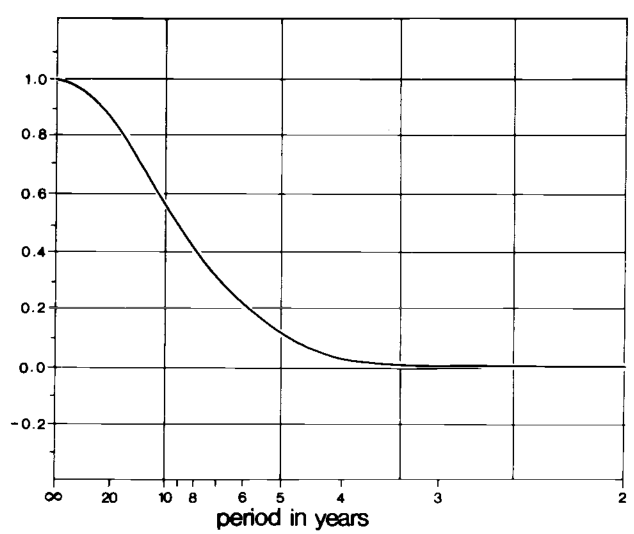

b. 13—term Gaussian tapered filter

Figure 2. Transfer functions for various methods of smoothing series, plotted against a linear frequency scale: the method used in the original CROS algorithm gives 2a, 2c is from the "low-A" filter from LaMarche and Fritts (1972:23), and the 10 year filter from Table 1 gives 2d.

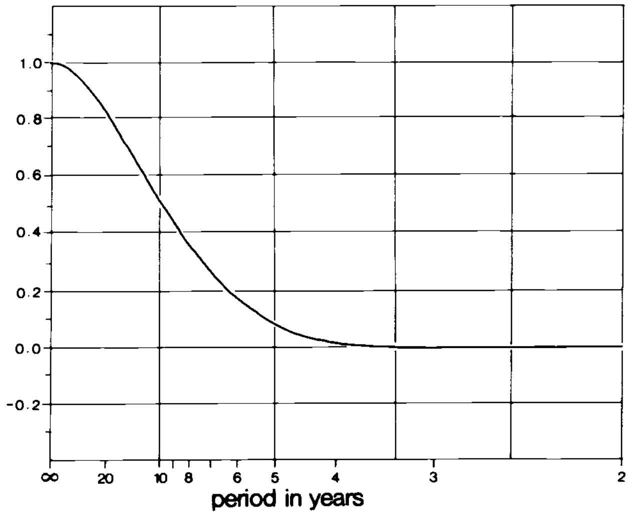

c. 13-term filter, cutoff at 8 years

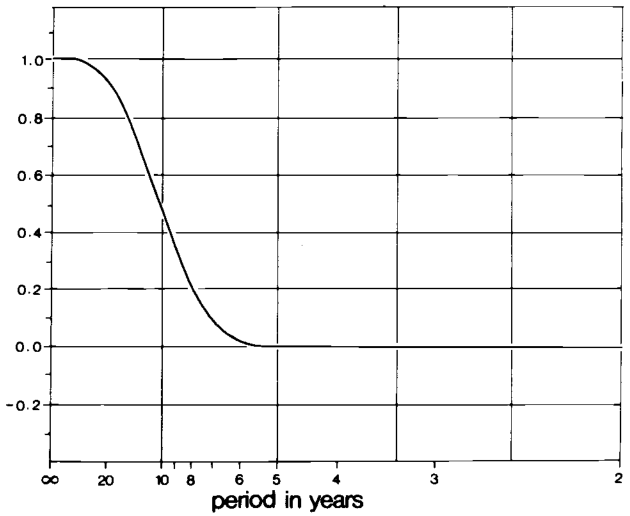

d. 23-term filter, cutoff at 10 years

Figure 2, *Continued.*

CALCULATING t VALUES

The procedure is to calculate the covariance of the two series for a particular overlap:

$$s_{ab} = \left(\sum_{\lambda=0}^{k-1} a_{i+\lambda} b_{j+\lambda} \right) / (k - 1)$$

where $a_0, a_1, \ldots a_{n-1}$ and $b_0, b_1, \ldots b_{m-1}$ are the two series of indices, with $n \geq m$; i and j indicate the start of the overlap interval in the a and b series respectively, with $0 \leq i < n$ and $0 \leq j < m$, and k is the length of the interval, where $\mu \leq k$ and ($i = 0$ and $j = m-k$) or ($j = 0$ and (($i = n-k$) or ($i \leq n-k$ and $k = m$))).

If we repeat this for all i, j, and k values that obey these relations we get a total of $1 + n + m - 2\mu$ covariance values, from which we can calculate t values:

$$t = (k - 3)^{1/2} (\tanh^{-1}(s_{ab} / s_a s_b))$$

There are several differences between this and the formula used in the original version of CROS. By assuming that the minimum number of years of overlap, μ, is not less than 60 we can use Fisher's z transformation to produce t values with ∞ degrees of freedom, i.e. values that should be distributed normally. In addition the s_a and s_b values used to standardize the covariance are the standard deviations of the a and b series over the same intervals used to calculate the covariance. The original version did not use these standard deviations, but values derived from the deviations from the means of the entire series rather than of the intervals over which they were calculated. This is the feature of the original algorithm which Baillie refers to (1982:84) when saying that it produces a "maximized t," but the two methods will produce similar results if $\mu \geq 60$ and the series have been through a high pass filter, and if the filter lets through some of the low frequencies the new version should produce more sensible results. Steward (1983) has also recommended these changes to the original method.

SIGNIFICANCE PROBABILITIES

The t values produced in this way will be very similar to those from the original version of CROS, and will suffer from many of the same defects, the most severe of which is that they cannot be compared with tables of the quantiles of t_∞ to give significance probabilities. Orton (1983) has pointed out that if one of the t values is greater than a value of $t_{\infty\alpha}$ from the table, the significance probability should not be α but $1 - (1 - \alpha)^N$, where $N = 1 + n + m - 2\mu$ (i.e. the total number of t values calculated). He also suggests that if the t values are ranked in ascending order $t_1, t_2 \ldots t_N$ then the difference between the highest two values, $t_N - t_{N-1}$, can be used to test the significance of the highest, T_N. Unfortunately both these methods will work only if the t values from non-matching positions are distributed as t_∞, but empirical tests show that they are not. For example, of the 16 series from Swan Carr that were used to test the data transformation methods, 30 pairs did not overlap in time, and so the distributions of t values from crossdatings of these pairs were examples of the distributions from non-matching positions. The standard deviations of these distributions, s_t, were greater than 1.00 for 28 of the pairs, and significantly greater for 23 of these ($P < 0.05$ for a X^2 test on $(N - 1)\, s_t^2$ for the pair).

One way round this difficulty is to follow Laxton and Litton (1982) or Steward (1983) and use simulations of the series. If one uses random numbers to generate many different series that share autocorrelation or spectral properties with the series being dated, then the frequency with which t values as high as the one observed in the crossdating of the real series turn up in the crossdatings of the simulated series gives a measure of the significance of the observed value. The disadvantage of this method is that, as in the case of autoregressive model fitting, it needs large amounts of computation to be sure of producing a good result.

However there is another approach, which does not need so much computation, but still produces good results. If we consider the distribution of t values from non-matching positions to be normally distributed, but with an unknown mean and standard deviation, and the t value from the matching position (if it exists) to be from a different normal distribution with a higher mean, then we are dealing with a slippage model which forms the basis of many outlier tests. The best of these, in the sense of being the maximum likelihood ratio test statistic under this model (Barnett and Lewis 1978; Hawkins 1980), uses the studentized deviation of the maximum t value from the mean of all the values:

$$T_N = (t_N - \overline{t}\,) / s_t$$

It has a long history of use as an outlier test (Pearson and Chandra Sekar 1936), and there are tables of the fractiles of its distribution (Grubbs and Beck 1972). For values of $N > 147$ there are no tables available, but if we evaluate

$$\acute{t} = ((N(N - 2)\ T_N^2) / ((N - 1)^2 - NT_N^2))^{1/2}$$

we can use $\acute{t}$ to produce an reasonable approximation to the true significance probability for T_N. If we call this probability α, then it is given by the following inequality:

$$\alpha \leq N\ \Pr(t_{N-2} > \acute{t})$$

(Barnett and Lewis 1979:106), where t_{N-2} is the Student's t distribution with N-2 degrees of freedom. So we are again using the t distribution to approximate the significance of the best matching position of the two series, but the $\acute{t}$ value should not be confused with the t values calculated directly from the covariances, and in nearly all cases it will be less than the highest t value. Moses (1978) has published charts that give significance probabilities for $2 \leq \acute{t} \leq 7$, intended for use with the Bonferroni t statistic in several related applications. Although significance probabilities produced by way of $\acute{t}$ only give an upper bound on the true values, the cases where one can compare them with those given by Grubbs and Beck (1972) show that this is not an excessive upper bound.

The disadvantage of this approach is that it relies on the normality of the two hypothetical distributions of t values, and small departures from normality can have an adverse effect on statistics of this kind (Hawkins 1980:40). For example if two series did match at a particular position, autocorrelation effects might distort the distribution of t values from the other (non-matching) positions. However in practice it seems to be reasonably robust, and the significance probabilities it produces are only slightly conservative.

COMPUTATION

We can avoid re-calculating the standard deviations of the two series for each new overlap between them by using a stable updating method to derive new standard deviations from ones at a previous position of overlap: Welford (1962) gives a method of updating the sum of squares; Hanson (1975) and Cotton (1975) give methods based on manipulating the standard deviations directly. Alternatively we can keep a separate record of Σa and Σa^2 for each series, which get updated at each new position and from which the standard deviation is calculated. These sums are best represented as fixed point numbers, with no rounding errors, but the summation formula of Kahan (1965, 1972) can be used for normalized floating point numbers if this is impossible (and also simplifies calculating s_t^2 from the t values). Updating will involve at most adding one value and subtracting another (if the new overlap position does not change k, the length of the overlap interval) and only a subtraction will be necessary in many cases (if k decreases), provided the starts of the overlap interval in the two series (i and j) are never advanced more than a year at a time.

We can completely eliminate the use of tables and charts by calculating the final probability directly from $\acute{t}$. All that is needed is a function that will give the probability integral for Student's t, such as the amended version of the algorithm devised by Hill (1970, 1981).

The best method of calculating the covariances of the series does not involve computing them separately using the sums of products for each overlap position, but produces a vector containing all the covariances indirectly using the discrete Fourier transforms of the series (Brigham, 1974): these can also simplify the initial filtering operation. At Belfast we still do not have an efficient implementation of the Fourier transform method, but an experimental program that implements the improved algorithm, calculates the standard deviations by a stable updating method, and which calculates the covariances directly from sums of products runs in less than a quarter of the time taken by a comparable program using the original CROS algorithm, when they are compared on an Apple IIe. Copies of this program, which is called Cross84, are available from the Palaeoecology Centre, Queen's University of Belfast.

SUMMARY OF PROCEDURES

(1) Take the logarithms of the original series.

(2) Use a high pass filter with good spectral properties to remove the low frequencies. A filter that removes frequencies with periods greater than ten years seems to give the best results: try experimenting with others if this does not work.

(3) Calculate the covariances of the transformed series for every overlap longer than ca. 60 years.

(4) Produce correlation coefficients by standardizing the covariances, using the standard deviations of the two series over the overlap interval. Use a stable updating method to avoid recalculating the means and standard deviations for each overlap.

(5) Use Fisher's z transformation on the correlation coefficients to produce t values.

(6) Find the mean and standard deviation of the t values.

(7) Divide the maximum deviation from the mean by the standard deviation to give an outlier statistic.

(8) Estimate the significance probability of this by using tables or a Bonferroni t statistic.

ACKNOWLEDGEMENTS

I would like to thank J. R. Pilcher for persuading me to publish this, and M. G. L. Baillie, D. M. Brown and E. Francis for their helpful comments, and for providing the raw series and chronologies used in the tests of various data transformation methods. C. R. Orton and N. R. J. Fieller provided useful comments about the statistical basis of crossdating, and the CROS2 program to be described in a forthcoming publication by T. M. L. Wigley, P. D. Jones, K. R. Briffa and D. A. Norton provided some interesting points of comparison with the methods described here.

REFERENCES

Baillie, M. G. L.
1973 A recently developed Irish tree-ring chronology. *Tree-Ring Bulletin* 33: 15-28.
Baillie, M. G. L.
1982 *Tree-Ring Dating and Archaeology.* Croom Helm, London.
Baillie, M. G. L. and J. R. Pilcher
1973 A simple crossdating program for tree-ring research. *Tree-Ring Bulletin* 33: 7-14.
Barefoot, A. C., W. L. Hafley and J. F. Hughes
1978 Dendrochronology and the Winchester excavation. In *Dendrochronology in Europe* (British Archaeological Reports, International Series 51), edited by J. M. Fletcher, pp. 162-171. B. A. R., Oxford.
Barnett, V. and T. Lewis
1978 *Outliers in Statistical Data.* John Wiley, Chichester.
Briffa, K. R., P. D. Jones, T. M. L. Wigley, J. R. Pilcher and M. G. L. Baillie
1983 Climatic reconstruction from tree rings: part 1, basic methodology and preliminary results for England. *Journal of Climatology* 3: 233-242.
Brigham, E. O.
1974 *The Fast Fourier Transform.* Prentice-Hall, Englewood Cliffs.
Cotton, I. W.
1975 Remark on stably updating mean and standard deviation of data. *Communications of the ACM* 18 (8): 458, with corrigendum 18 (10) 591.
Fritts, H. C.
1976 *Tree Rings and Climate.* Academic Press, New York.
Graybill, D. A.
1982 Chronology development and analysis. In *Climate from Tree Rings,* edited by M. K. Hughes, P. M. Kelly, J. R. Pilcher, and V. C. La Marche, Jr., pp. 21-30.
Grubbs, F. E. and G. Beck
1972 Extension of sample sizes and percentage points for significance tests of outlying observations. *Technometrics* 14 (4): 847-854.
Guiot, J., A. L. Berger, A. V. Munaut and Cl. Till
1982 Some new mathematical procedures in dendroclimatology, with examples from Switzerland and Morocco. *Tree-Ring Bulletin* 42: 33-48.
Hamming, R. W.
1977 *Digital filters.* Prentice Hall, Englewood Cliffs.
Hanson, R. J.
1975 Stably updating mean and standard deviation of data *Communications of the ACM* 18 (1): 57-58.
Hawkins, D. M.
1980 *Identification of Outliers.* Chapman and Hall, London.
Hill, G. W.
1970 Algorithm 395. Student's *t*-distribution. *Communications of the ACM* 13 (10): 617-619.
1981 Remark on Algorithm 395. *ACM Transactions on Mathematical Software* 7 (2): 247-249.
Hollstein, E.
1980 *Mitteleuropäische Eichenchronologie* (Trierer Grabungen und Forschungen XI). Philip von Zabern, Mainz.
Hughes, M. K., P. Leggett, S. J. Milsom, and F. A. Hibbert
1978 Dendrochronology of oak in North Wales. *Tree-Ring Bulletin* 38: 15-24.
Kahan, W.
1965 Further remarks on reducing truncation errors. *Communications of the ACM* 8 (1): 40.
1972 A survey of error analysis. *Proceedings of the International Federation for Information Processing Congress 1971,* Invited Papers: 1214-1239.

LaMarche, V. C. Jr. and H. C. Fritts
1972 Tree-rings and sunspot numbers. *Tree-Ring Bulletin* 32: 19-33.
Laxton, R. R. and C. D. Litton
1982 Information theory and dendrochronology. *Science and Archaeology* 24: 9-24.
1983 Information theory and dendrochronology: the effect of pre-whitening. University of Nottingham Statistics Group Research Report, also published in *Computer Applications in Archaeology 1983.*
Moses, L. E.
1978 Charts for finding upper percentage points of Student's *t* in the range .01 to .00001. *Communications in Statistics* B7 (5): 479-490.
Orton, C. R.
1983 The use of Student's *t*-test for matching tree-ring patterns. *Bulletin of the University of London Institute of Archaeology* 20: 101-105.
Pearson, E. S. and C. Chandra Sekar
1936 The efficiency of statistical tools and a criterion for the rejection of outlying observations. *Biometrika* 28: 308-320.
Pilcher, J. R.
1976 A statistical oak chronology from the North of Ireland. *Tree-Ring Bulletin* 36: 21-27.
Pilcher, J. R. and M. G. L. Baillie
1980 Six modern oak chronologies from Ireland. *Tree-Ring Bulletin* 40: 45-58.
Rohlf, F. J. and R. R. Sokal
1981 *Statistical tables* (2nd. ed.). Freeman, San Francisco.
Steward, S.
1983 Statistical dendrochronology. M. Sc. thesis, University of Sheffield.
Welford, B. P.
1962 Note on a method for calculating corrected sums of squares and products *Technometrics* 4 (3): 419-420.