Clim. Past Discuss., 8, 2477–2502, 2012
www.clim-past-discuss.net/8/2477/2012/
doi:10.5194/cpd-8-2477-2012


Climate of the Past Discussions



# Automated ice-core layer-counting with strong univariate signals

**J. J. Wheatley[1], P. G. Blackwell[1], N. J. Abram[2], J. R. McConnell[3], E. R. Thomas[4], and E. W. Wolff[4]**

[1]School of Mathematics and Statistics, The University of Sheffield, UK
[2]Research School of Earth Sciences, Australian National University, Australia
[3]Desert Research Institute, 2215 Raggio Parkway, Reno, Nevada 89512-1095, USA
[4]British Antarctic Survey, Cambridge, UK

Received: 7 June 2012 – Accepted: 12 June 2012 – Published: 3 July 2012

Correspondence to: J. J. Wheatley (j.wheatley@shef.ac.uk)

Published by Copernicus Publications on behalf of the European Geosciences Union.







## Abstract

We present an automated method for determining the annual layer chronology of an ice-core with a strong annual signal, utilising the hydrogen peroxide record from an Antarctic Peninsula ice-core as a test signal on which to count annual cycles. The signal is de-trended and normalised before being split into sections with a deterministic cycle count and those that need more attention. Possible reconstructions for the uncertain sections are determined and a method for assigning probability measures to each reconstruction is discussed. These methods could be adapted for use on much longer datasets, thereby reducing manual effort and providing a robust methodology.

## 1 Introduction

Ice-cores hold information about the Earth's past climate. This information is recorded by chemical and physical signals in the ice that reflect the environmental conditions when the ice was deposited. These signals are measured as a depth series, but need to be modelled against time. In some cases the chemical signals measured from the meltwater of ice-cores have annual cycles which can be counted to provide a chronology. Traditionally this has been undertaken by visual assignment of annual markers, with no robust method for ensuring consistency or assessing uncertainty. In some cases, ice-core annual layer chronologies have been counted manually for data spanning tens of thousands of years, see Andersen et al. (2006), requiring months to years of effort. This highlights a clear need to develop an automated, statistically-based methodology for application in this type of physical sciences research.

Previous attempts at developing automated cycle counting methods in the literature have generally given little consideration to uncertainty. Rasmussen et al. (2002) use Independent Component Analysis to combine chemistry data in conjunction with a low-pass filter; the resulting signal is dated by counting peaks that rise above a threshold value. Other methods range from iteratively picking the most likely annual peaks along





CPD
8, 2477–2502, 2012



a core with respect to their spacing and magnitude (Taylor et al., 2004; McGwire et al., 2011) to measuring periodicities in the signals (Rupf and Radons, 2004; Svensson et al., 2005).

Uncertainty has been addressed under the Bayesian framework, applying hidden Markov models (Winstrup, 2011) and Markov chain Monte Carlo methods (Wheatley et al., 2012) to univariate models based around sine waves. The method presented here has the advantage over these approaches that it is sufficiently straightforward computationally to use on full size cores, whilst still providing a measure of uncertainty.

Hydrogen peroxide ($H_2O_2$) is particularly suitable as a chronological marker in ice-cores. It is created in the atmosphere by a chemical reaction that requires ultraviolet light. Because of the very strong seasonality in Antarctica (from complete darkness in midwinter to 24 h daylight in midsummer), there is a strong and regular seasonality in ultraviolet light, and a theoretical basis for expecting a single quasi-sinusoidal variability in $H_2O_2$ concentration. Measurements of $H_2O_2$ dissolved in the ice therefore provide a good annual cycle.

In this study we utilise the $H_2O_2$ record from an Antarctic Peninsula ice-core as a test signal on which to count annual cycles. The Gomez core extends to 134 m (Thomas et al., 2008). We use 2 cm average $H_2O_2$ concentrations determined from continuous measurements along the core. This core has been manually dated to 153 yr by determining the depths of the lowest points ("nadirs") of the annual troughs. We present a simple method based on these data for counting cycles which is fully automated, requires no prior knowledge, and assigns probability measurements to the cycle count.

## 2 Data pre-processing

The Gomez $H_2O_2$ signal, indexed by its depth $i$, is transformed by taking logarithms to improve the symmetry in its annual periodicity, giving the log signal $\boldsymbol{x} = \{x_i\}$ which can be thought of as a sine wave on a non-linear time-scale with varying amplitude and mean.





CPD
8, 2477–2502, 2012



In order to simplify the process of cycle counting, $\boldsymbol{x}$ is de-trended and the annual cycle amplitudes are set to approximately 1. This is achieved by smoothing the signal with respect to an initial rough estimate of the average cycle length at each depth, see Sect. 2.1 for details. The precise shape of the normalised signal does not have much effect on the methods presented below. While hydrogen peroxide seems to (and is expected to) yield a sine wave, it may be necessary to consider other wave shapes for other chemical signals; Winstrup (2011) provides a discussion of some possibilities for visual stratigraphy data.

### 2.1 Smoothing

Bandpass or lowpass filters are frequently used to remove cycles with periods deemed to be larger or smaller than expected (see Rasmussen et al., 2002; Taylor et al., 2004; Rupf and Radons, 2004; McGwire et al., 2011). In Shimohara et al. (2003) the data are smoothed by taking a moving average over a range of half an estimated annual layer thickness. Here we use a point-wise standardisation: at each depth $i$

$$s_i = \frac{x_i - \mu_i}{\sqrt{2}\sigma_i}$$

where $\boldsymbol{\mu} = \{\mu_i\}$ is a moving average of $\boldsymbol{x}$ and $\boldsymbol{\sigma} = \{\sigma_i\}$ is a moving standard deviation of $\boldsymbol{x} - \boldsymbol{\mu}$, both over the range of one annual layer thickness based on a preliminary estimate – see Sect. 2.2 below.

$\boldsymbol{\mu}$ measures trend in the data, the annual seasonality is averaged out due to the interval length being a whole year. Subtracting $\boldsymbol{\mu}$ de-trends $\boldsymbol{x}$ and centres $\boldsymbol{s}$ on 0. $\sqrt{2}\boldsymbol{\sigma}$ is an estimate for the annual cycle amplitude at each depth. Dividing by $\sqrt{2}\boldsymbol{\sigma}$ sets the apices and nadirs of the annual cycles in $\boldsymbol{s}$ to a magnitude of approximately 1. In the case of missing values, the corresponding $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ values are linearly interpolated from the closest surrounding points where there are sufficient data.

The resulting $\boldsymbol{s}$ is essentially a sine wave on a non-linear time scale, with noise.





CPD
8, 2477–2502, 2012



## 2.2 Estimate of cycle length

An initial estimate for the average cycle length in any sufficiently large subsection of the signal can be found using the autocorrelation function (ACF); this is used for the smoothing, not directly for the actual layer-counting. $\boldsymbol{x}$ is split into $\beta$ non-overlapping subsections, each containing approximately the same number of cycles. This is achieved via an algorithm that perturbs the section boundaries with respect to the expected number of cycles in each section, estimated from the ACF. Each point is then assigned the average cycle length for its subsection to be used as the interval length for calculating $\mu_i$ and $\sigma_i$.

Other methods used to obtain this estimate as a precursor to layer counting include Fourier analysis McGwire et al. (2011) and ice flow modelling Shimohara et al. (2003).

Figure 1 shows the ACF for the second (circles) and fifth (stars) sections when $\beta = 6$. The second section has an estimated 51 points in an average annual cycle and is of length 1233 points. The fifth section has an estimated 32 points in an average annual cycle, and is of length 774 points. Both sections therefore contain an estimated 24.2 annual cycles, as do all 6 sections. This gives an initial estimate of approximately 142 for the number of cycles contained in $\boldsymbol{s}$, which is an underestimate due to the stretches of missing values.

Figure 2 illustrates the smoothing process for $\beta = 6$. The top plot shows the log signal $\boldsymbol{x}$, with its annual moving average $\boldsymbol{\mu}$ as a dotted line. The second plot shows the detrended signal, with its estimated amplitude ($\sqrt{2}\boldsymbol{\sigma}$) as a dotted line. The bottom plot is of $\boldsymbol{s}$, the standardised signal. There is still some variation in mean and amplitude visible due to the crude estimate of cycle length at each depth, however this is sufficient to serve as a starting point for our automated dating scheme.

The choice of $\beta$ affects the number of cycles available to estimate the average cycle length in each section; there is a trade-off between having sufficient data in each section and averaging over many cycles. In Sect. 5.2 we further examine the sensitivity of the method to the choice of $\beta$.





CPD
8, 2477–2502, 2012



## 3 "Certain" annual cycles

In large-scale manual layer counting exercises, annual cycles have been termed "certain" if they are judged through consensus of multiple counters to have probability at least 3/4, see Andersen et al. (2006). Here we present a method of automatically classifying quarter cycles, analogous to seasons, as being "certain" if they are well-defined in the smoothed signal.

### 3.1 Classification into runs

We aim to segment the points of $\boldsymbol{s}$ into non-overlapping subsections called *runs*, each representing either a "certain" quarter cycle or season, or an "issue" where manual intervention is required. Each run is a collection of consecutive points and has one of five labels:

P: peak / summer;

D: descending / autumn;

T: trough / winter;

A: ascending / spring;

$\chi$: issue.

This classification is a two-stage process. Firstly we find potential quarter cycles; these cannot contain missing values as we are unsure of their classification. For some sensible choice of cut-off parameter $\nu$: all runs of data points $s_i \geq \nu$ are labelled as $\mathrm{P}^*$, the star meaning potential; all runs of data points $s_i \leq -\nu$ are labelled as $\mathrm{T}^*$. Runs of data points for which $-\nu < s_i < \nu$ are labelled as potentially ascending and descending: $\mathrm{A}^*$ if they fall between a $\mathrm{T}^*$ and $\mathrm{P}^*$; $\mathrm{D}^*$ if they fall between a $\mathrm{P}^*$ and a $\mathrm{T}^*$; and $\chi$ otherwise. So the stretch of data in Fig. 3 with $\nu = 1/\sqrt{2}$ (top), has potential run label pattern:

$$\ldots, \mathrm{T}^*, \mathrm{A}^*, \mathrm{P}^*, \mathrm{D}^*, \mathrm{T}^*, \mathrm{A}^*, \mathrm{P}^*, \mathrm{D}^*, \mathrm{T}^*, \chi, \mathrm{T}^*, \mathrm{A}^*, \mathrm{P}^*, \chi, \mathrm{P}^*, \mathrm{D}^*, \mathrm{T}^*, \mathrm{A}^*, \mathrm{P}^*, \mathrm{D}^*, \mathrm{T}^*, \mathrm{A}^*, \mathrm{P}^*, \ldots$$





CPD
8, 2477–2502, 2012



In the second stage, potential runs are labelled as runs only if they are central to 5 consecutive runs whose potential labels obey the pattern expected from a sinusoidal signal. In the example, this gives the following:

$$\ldots, \mathrm{T}, \mathrm{A}, \mathrm{P}, \mathrm{D}, \mathrm{T}, \mathrm{A}, \mathrm{P}, \mathrm{D}^*, \mathrm{T}^*, \chi^*, \mathrm{T}^*, \mathrm{A}^*, \mathrm{P}^*, \chi^*, \mathrm{P}^*, \mathrm{D}^*, \mathrm{T}, \mathrm{A}, \mathrm{P}, \mathrm{D}, \mathrm{T}, \mathrm{A}, \mathrm{P}, \ldots$$

Finally, consecutive data points which make up the unlabelled runs are collected together into runs labelled as $\chi$, and termed *issues*. This gives:

$$\ldots, \mathrm{T}, \mathrm{A}, \mathrm{P}, \mathrm{D}, \mathrm{T}, \mathrm{A}, \mathrm{P}, \chi, \mathrm{T}, \mathrm{A}, \mathrm{P}, \mathrm{D}, \mathrm{T}, \mathrm{A}, \mathrm{P}, \ldots$$

Each point of $\boldsymbol{s}$ is now in exactly one run, and each run that is not labelled as $\chi$ makes up a certain quarter annual cycle – these are termed certain runs. In Fig. 3 runs labelled P are coloured red, D runs are orange, T runs are blue, A runs are green, and issues are black.

The value of the threshold used here, $v = 1/\sqrt{2} \approx 0.707$, is natural in that it would lead to the same proportions of points in each of the four types of quarter cycles, in an ideal signal. Other choices have their merits; the value used is further discussed in Sects. 3.2 and 4.2, and robustness to the choice is considered in Sect. 5.1.

### 3.2 Issues

Runs labelled $\chi$, or *issues*, consist of consecutive data points that are missing values, that surround missing values, or that are in some way dubious in their periodicity. Sections of signal where there are no issues have a deterministic number of cycles: 1/4 of a cycle per certain run.

Figure 3 shows a stretch of $\boldsymbol{s}$ between 9.82 m and 18.58 m, $v = 1/\sqrt{2}$ (top) and $v = 1/2$. The issues here are caused by the run pattern alone and contain no missing values. At $v = 1/\sqrt{2}$ there is an issue between 13.12 m and 15.4 m: a peak dips below $1/\sqrt{2}$, and a trough rises above $-1/\sqrt{2}$. At $v = 1/2$ the points corresponding to this issue now make up 5 certain runs; however there is a new issue between 16.76 m and 18.16 m.





CPD
8, 2477–2502, 2012



Issues such as those seen in Fig. 3 occur only in the first 20 m of signal where, due to the very high resolution of the sampling with respect to the annual cycle length, small fluctuations are found in the annual cycles. Choosing different values of $\nu$ affects where data points are classified into certain runs and therefore the distribution of issues in this first part of the signal.

Figure 4 shows a stretch of $\boldsymbol{s}$ between 106.68 m and 110.2 m, $\nu = 1/\sqrt{2}$ (top) and $\nu = 1/2$. This issue is caused by a stretch of missing values, the non-missing data points that also form part of the issue are coloured black. The choice of $\nu$ affects only the issues length which decreases with $\nu$. Note how as $\nu$ decreases the length of extreme runs (P and T) increase and the length of central runs (A and D) decrease.

For $\beta = 6$: with $\nu = 1/\sqrt{2}$, $\boldsymbol{s}$ has 528 deterministic runs and 12 issues; with $\nu = 1/2$, $\boldsymbol{s}$ has 533 deterministic runs and 15 issues.

We can now split the signal into sections with a deterministic count and those that need more attention. At this stage, the "issues" could be presented to experts as in Figs. 3 and 4 so they can place certain and uncertain layer markers. However, the run pattern provides more information that could be exploited either to assist the expert analysis or to allow further automation.

## 4 Reconstructions and probabilities

### 4.1 Reconstructions

The labels of the certain runs that bound an issue provide insight into the label pattern of the certain runs that could replace it. It is convenient to refer to peak and trough runs collectively as *extreme* runs; similarly ascending and descending runs are termed *central*. Issues are always bounded by extreme runs, leading to four possible cases:

$$\ldots,\mathrm{P},\chi,\mathrm{P},\ldots \quad (1) \qquad \ldots,\mathrm{P},\chi,\mathrm{T},\ldots \quad (2)$$
$$\ldots,\mathrm{T},\chi,\mathrm{P},\ldots \quad (3) \qquad \ldots,\mathrm{T},\chi,\mathrm{T},\ldots \quad (4)$$





CPD
8, 2477–2502, 2012



The points of ***s*** that correspond to an issue can be replaced with a compatible section of sine wave, spread evenly over the points, referred to as a *reconstruction*. This could be used to fill in missing values or replace sections of data affected by a sampling problems.

The minimal reconstructions for cases (1) and (4) consist of 3 runs, as the minimal compatible sections of sine wave have run label patterns D, T, A and A, P, D, respectively. The minimal reconstructions for cases (2) and (3) consist of 1 run labelled as D and A, respectively. Further reconstructions can be found by adding in whole cycles.

Each possible reconstruction is made up of an odd number of runs. If $m$ is the number of runs in its minimal reconstruction and $k$ cycles are added in, the issue would contain $d = m + 4k$ runs. Issues are always bounded by extreme runs and will therefore always contain $\frac{d-1}{2}$ extreme runs and $\frac{d+1}{2}$ central runs.

Figure 5 shows possible reconstructions to the issue from Fig. 4. The top plot shows the minimal reconstruction with one run, and the bottom shows the reconstruction with 3 additional cycles added in which consists of 13 runs. Plots like this could be used as visual aids to the manual assessment of issues. However, the available information can be used to assign probabilities to each reconstruction, which would provide an automated method of layer counting.

### 4.2 Assigning probabilities

The *length* of a run is used here to denote the number of points it contains, $\ell$ say. Provided that the time-depth relationship does not change too quickly, we would expect a run of a given type to have a similar length to other such runs near to it within the core. This concept is key to existing manual and semi-automated layer-counting approaches. Here we present a method of assigning probabilities to the possible reconstructions of each issue by comparing $\ell$ to its implied distribution. For simplicity we are assuming throughout that points are equally spaced in depth, as is the case in our example. Relaxing that assumption is straightforward in principle; it simply requires





CPD
8, 2477–2502, 2012



a slight extension of the notation, to work with lengths expressed in terms of differences of depth rather than just numbers of points.

We take $p(d|\ell) \propto p(\ell|d)$; that is, we take the probability of a reconstruction with $d$ runs, given that the issue contains $\ell$ data points, to be proportional to the probability of those $d$ runs having total length $\ell$. This is essentially a Bayesian statistical approach, with a flat prior distribution on $d$. Again, using an alternative prior distribution would be straightforward, but in practice, prior information is always likely to be dominated by the other information in the core, as represented by $p(\ell|d)$.

Groups of $d$ consecutive certain runs, of which $\frac{d-1}{2}$ are extreme, are analogous to issue reconstructions. Ideally, we would model the lengths of these directly to find the distribution of $\ell$ given $d$. However, issues are concentrated in the first 20 m due to the fluctuations discussed above, and the last 20 m where there are regular stretches of missing values. In either case, this results in the analogous sections being concentrated at the centre of $\boldsymbol{s}$, resulting in a poor fit to the sections with issues.

Instead, we need to make use of information on individual certain run lengths. Figure 6 is a plot of certain run lengths against depth for a range of $v$. Peak run lengths are equivalent in distribution to trough run lengths because of the symmetry in the seasonality of $\boldsymbol{s}$ and in the classification process. For this reason, both peak and trough runs – i.e. extreme runs – are denoted by stars. Similarly ascending and descending runs, or central runs, are both denoted by circles. When $v$ is equal to the 75th percentile of a sine wave ($1/\sqrt{2} \approx 0.7$) extreme and central run lengths are equivalent in distribution; for $v < 1/\sqrt{2}$ extreme run lengths are generally larger than central; and for $v > 1/\sqrt{2}$ central run lengths are generally larger than extreme. All these plots show the non-linearity in average length, and also the change in spread as a function of depth. This non-linearity in ice-core layer thickness is caused by vertical compaction of snow into ice and thinning of the ice layers caused by horizontal flow. Other proxy records where annual layer counting is used for establishing chronologies, such as tree rings, varves and corals, will not display this systematic reduction in layer thickness through the depth profile. However, after taking natural logs the individual certain run lengths





CPD
8, 2477–2502, 2012



show a linear trend. This is well described by a regression model which has the extreme or central run label as a factor, Fig. 7, meeting all the standard linear modelling assumptions. If we know the central depth of a run and its label, we can find its expected length from the model. This model is analogous to the results of Rasmussen et al. (2006) where annual layer thicknesses from the NGRIP core are shown to be log-normally distributed after a linear strain correction.

All of the $d$ runs that make up the reconstruction of an issue have implied central depths and labels. The expected value for the lengths of these runs, and therefore the total length of the issue implied by the reconstruction, $\hat{\ell}_d$, can be *interpolated* from the regression model. We assume that $\ln(\ell) \sim N(\ln(\hat{\ell}_d), \sigma_d^2)$, where $\sigma_d$ depends on $d$ and is estimated using groups of certain runs, analogous to that particular reconstruction, as discussed above. (Note that $\sigma_d$ can *not* be obtained directly from the model for individual runs, because of the dependence in lengths between consecutive runs.)

The issue from Fig. 5 contains $\ell = 64$ data points. The minimal reconstruction for this issue (top) is made up of one quarter cycle ($m = 1$); at this depth $\hat{\ell}_1 = 8$ and $p(\ell|d = 1) = 0$ to three decimal places after normalisation. The second reconstruction contains 5/4 cycles, $\hat{\ell}_5 = 38$, and $p(\ell|d = 5) = 0.005$. At $d = 9$: $\hat{\ell}_9 = 68$ and $p(\ell|d = 9) = 0.984$; and at $d = 13$: $\hat{\ell}_{13} = 95$ and $p(\ell|d = 13) = 0.011$. Continuing to add cycles in this way results in reconstructions with negligible probability.

## 5 Results

The probabilities obtained in Sect. 4.2 can be combined across the whole core, assuming separate issues to be independent given the certain parts of the classification.

Figure 8 shows the probability distribution for the number of annual cycle troughs in the example signal for $\beta = 6$, $\nu = 1/2$ (left) and $\nu = 1/\sqrt{2}$. This is found by combining the probabilities for each possible reconstruction of each issue. The manual count gave 153 yr exactly.





CPD
8, 2477–2502, 2012



The differences in these distributions are due to the fluctuations found in the first 20 m of signal as discussed in Sect. 3.2.

The most probable reconstruction ($p > 0.8$ in both cases) has 153 troughs and agrees closely to the manually counted solution. Figure 9 shows annual cycle lengths (nadir to nadir) as found by manual counting and from the models most probable reconstruction. Each nadir placed by the model corresponds uniquely to a manually determined one, and their placements agree closely except for two cases: at the top of the core where troughs are wide and determining their nadir is subjective; at the bottom of the core where a number of troughs have been marked in stretches of missing data.

The signal is standardised in sections with respect to a typical cycle length, which may not be sensitive to an abrupt change in frequency. Depending on the cut-off $\nu$ it is possible that a relatively short (and therefore uncertain) cycle could be either counted as certain or missed out altogether, without flagging an issue. Labelling a short cycle as certain would result in several consecutive low valued run lengths, whereas missing a short cycle out would result in one very high valued central run length. Outliers from the regression model have been assessed to test for this possibility and no examples were found.

### 5.1 Sensitivity to $\nu$

To test the sensitivity of this process to the choice of $\nu$ it was run for $0.2 \leq \nu \leq 0.8$. Figure 10 shows the resulting distributions of cycle counts, the area of the circles are proportional to the probability.

For $\nu < 0.3$ the resolution of the data is such that some of the A and D runs at the bottom of the core are of length 1 or missing. Similarly for $\nu > 0.8$ some of the P and T runs at the bottom of the core are of length 1 or missing. This adversely affects the regression model and causes the model assumptions to fail.





CPD
8, 2477–2502, 2012



### 5.2 Sensitivity to $\beta$

To test the sensitivity of this process to the choice of $\beta$ it was run for $2 \leq \beta \leq 20$, Fig. 11 shows the resulting distributions of cycle counts for $\nu = 1/2$ (left) and $\nu = 1/\sqrt{2}$.

For $\beta < 5$ the sections from which the annual cycle length is estimated are too long. The thinning of annual cycle length with depth causes the estimates to be unrepresentative of cycles at either end of the sections. This results in some of the mean trend and variation in amplitude still being present in $\boldsymbol{s}$. For $\nu = 1/\sqrt{2}$ a number of peaks and troughs are missed, causing an underestimate in the cycle count. However, $\nu = 1/2$ has many additional issues but still works well in this range.

For $\beta > 20$ there is not enough data in the first section with which to estimate the average cycle length.

## 6 Conclusions

With regards to developing an automated method for determining the annual layer chronology in an ice-core with a strong annual signal, we have presented a method to:

- split the signal into sections with a deterministic cycle count and those that need more attention
- display possible reconstructions for the uncertain sections
- assign probability measures to each reconstruction

which together provide a stable count with an uncertainty measure. Some of the work carried out for the example from Gomez would need to be repeated for each ice-core: for example, to determine the best value of $\beta$ for a particular case, or to assess the data for trends in layer thickness, which can arise for climatic as well as glaciological reasons.





CPD
8, 2477–2502, 2012



A further benefit of this method it that it provides an objective way to split the depth scale into seasons, allowing other variables to be discussed in seasonal terms. Assuming intra-cycle symmetry, the choice of $\nu = 1/\sqrt{2}$ splits each cycle into four equal parts. In the case of $H_2O_2$, where cycles are directly correlated to sunlight, these parts are analogous to seasons. We are currently developing a method that takes this idea further, providing a stochastic estimate for the time of year at each depth; however this granularity has a computational cost. We use a different model formulation than that described in Wheatley et al. (2012); as well as having added flexibility, our newer method generalises more readily to the case of phase shifted multivariate signals.

These methods could be adapted for use on much longer datasets, thereby reducing manual effort and providing a robust methodology. Future work to develop this method for broad application in physical science research, including but not limited to ice-core palaeoclimate research, may involve extending the methodology to take in information from multivariate datasets with more uncertain annual cyclicity and being able to provide solutions for optimally fitting annual chronologies between fixed points of known age.

*Acknowledgements.* This work was supported by the Natural Environment Research Council (grant reference: NE/I52797X/1).

## References

Andersen, K. K., Svensson, A., Johnsen, S. J., Rasmussen, S. O., Bigler, M., Röthlisberger, R., Ruth, U., Siggaard-Andersen, M.-L., Peder Steffensen, J., Dahl-Jensen, D., Vinther, B. M., and Clausen, H. B.: The Greenland ice core chronology 2005, 15–42 ka. Part 1: constructing the time scale, Quaternary Sci. Rev., 25, 3246–3257, 2006. 2478, 2482

McGwire, K., Taylor, K., Banta, J., and McConnell, J.: Identifying annual peaks in dielectric profiles with a selection curve, J. Glaciol., 57, 763–769, 2011. 2479, 2480, 2481

Rasmussen, S., Andersen, K., Siggaard-Andersen, M., and Clausen, H.: Extracting the annual signal from Greenland ice-core chemistry and isotopic records, Ann. Glaciol., 35, 131–135, 2002. 2478, 2480





CPD
8, 2477–2502, 2012



Rasmussen, S. O., Andersen, K. K., Svensson, A. M., Steffensen, J. P., Vinther, B. M., Clausen, H. B., Siggaard-Andersen, M.-L., Johnsen, S. J., Larsen, L. B., Dahl-Jensen, D., Bigler, M., Röthlisberger, R., Fischer, H., Goto-Azuma, K., Hansson, M., and Ruth, U.: A new Greenland ice core chronology for the last glacial termination, J. Geophys. Res., 111, D06102, doi:10.1029/2005JD006079, 2006. 2487

Rupf, I. and Radons, G.: New approaches for automated data processing of annually laminated sediments, Nonlinear Proc. Geoph., 11, 599–607, 2004. 2479, 2480

Shimohara, K., Miyamoto, A., Hyakutake, K., Shoji, H., Takata, M., and Kipfstuhl, S.: Cloudy Band Observations for Annual Layer Counting on the GRIP and NGRIP, Greenland, Deep Ice Core Samples (Scientific Note), Mem. Natl. Inst. Polar Res., Spec. Issue, 57, 161–167, 2003. 2480, 2481

Svensson, A., Nielsen, S., Kipfstuhl, S., Johnsen, S., Steffensen, J., Bigler, M., Ruth, U., and Röthlisberger, R.: Visual stratigraphy of the North Greenland Ice Core Project (NorthGRIP) ice core during the last glacial period, J. Geophys. Res., 110, D02108, doi:10.1029/2004JD005134, 2005. 2479

Taylor, K., Alley, R., Meese, D., Spencer, M., Brook, E., Dunbar, N., Finkel, R., Gow, A., Kurbatov, A., Lamorey, G., Mayewski, P., Meyerson, E., Nishiizumi, K., and Zielinski, A.: Dating the Siple Dome (Antarctica) ice core by manual and computer interpretation of annual layering, J. Glaciol., 50, 453–461, 2004. 2479, 2480

Thomas, E., Marshall, G., and McConnell, J.: A doubling in snow accumulation in the Western Antarctic Peninsula since 1850, Geophys. Res. Lett., 35, L01706, doi:10.1029/2007GL032529, 2008. 2479

Wheatley, J. J., Blackwell, P. G., Buck, C., Lee, L., Röthlisberger, R., and Wolff, E. W.: Statistical Layer-Counting in Ice Cores, Environmental and Ecological Statistics, submitted, 2012. 2479, 2490

Winstrup, M.: An Automated Method for Annual Layer Counting in Ice Cores – and an Application to Visual Stratigraphy Data from the NGRIP Ice Core, Ph.D. thesis, Faculty of Science – University of Copenhagen, available at: http://www.nbi.ku.dk/english/research/phd_theses/phd_theses_2011/mai_winstrup/, 2011. 2479, 2480





CPD
8, 2477–2502, 2012



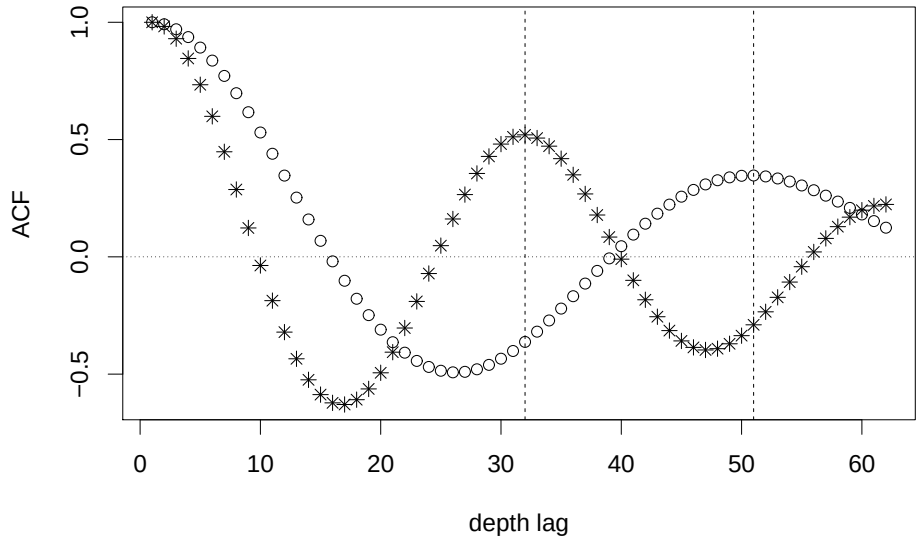


**Fig. 1.** The autocorrelation functions (ACFs) of the locally standardised signal (log concentration of $H_2O_2$) from the second (circles) and fifth (stars) of the $\beta = 6$ sub-sections of the Gomez ice-core, plotted against lag measured as a number of observations. The vertical dashed lines indicate the local maxima of the ACFs, used to obtain an initial estimate of cycle length.





CPD
8, 2477–2502, 2012



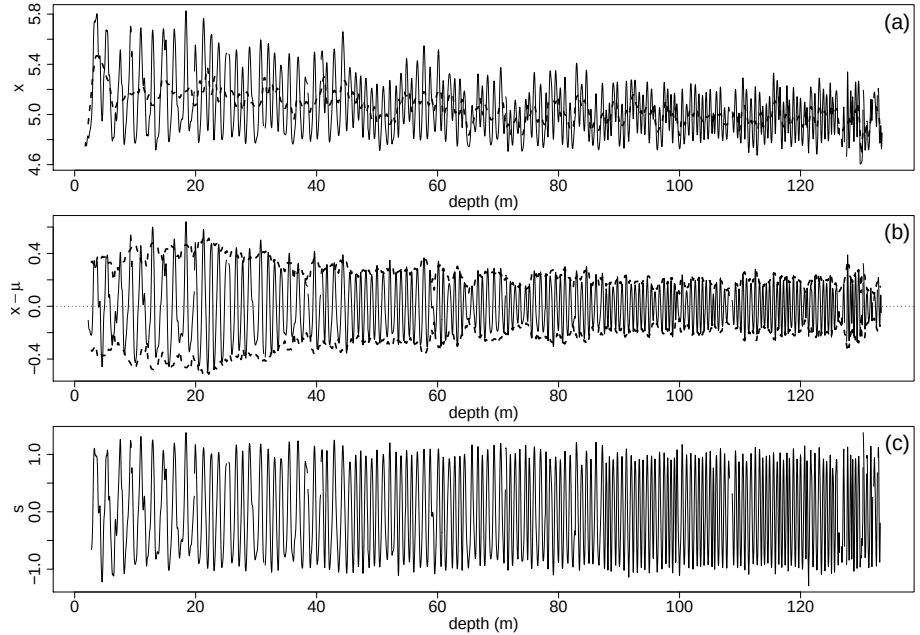


**Fig. 2.** The process of smoothing the $H_2O_2$ signal from the Gomez core, using an initial estimate of annual cycle length based on $\beta = 6$ sub-sections. **(a)** The log signal $\boldsymbol{x}$, with its annual moving average $\boldsymbol{\mu}$ as a dotted line. **(b)** The de-trended signal, with its estimated amplitude ($\sqrt{2}\boldsymbol{\sigma}$) as a dotted line. **(c)** The standardised signal $\boldsymbol{s}$.





CPD
8, 2477–2502, 2012



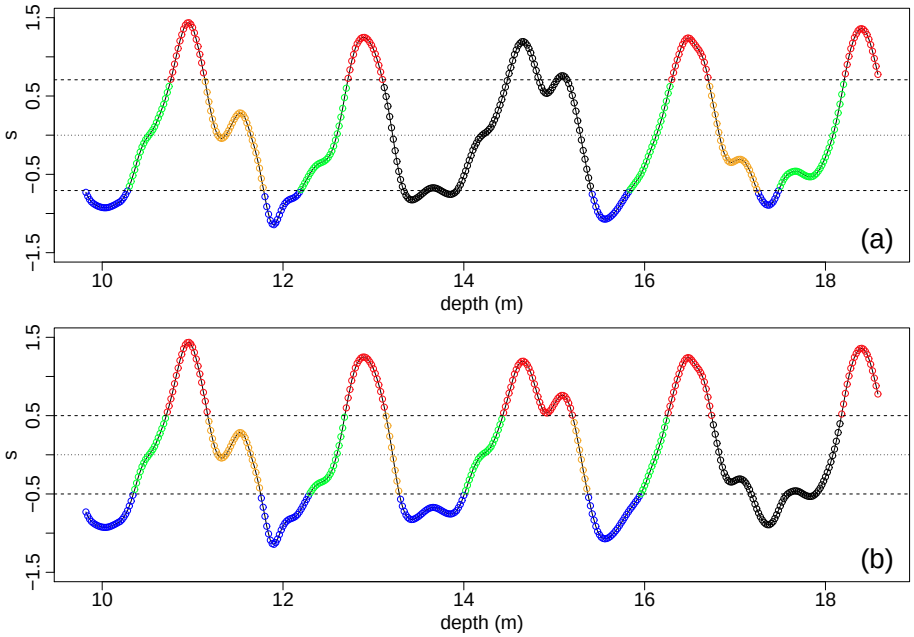


**Fig. 3.** Part of the standardised $H_2O_2$ signal, for depths between 9.82 m and 18.58 m, showing the "issues" (regions of dating uncertainty) caused by fluctuations in the signal and identified using two different values of the threshold $\nu$ for classifying individual points. Points within "issues" are black; points within peaks (labelled P in main text) are coloured red, descending points (D) are orange, troughs (T) are blue, and ascending points (A) are green. **(a)** Threshold $\nu = 1/\sqrt{2} \approx 0.707$. **(b)** Threshold $\nu = 0.5$.





CPD
8, 2477–2502, 2012



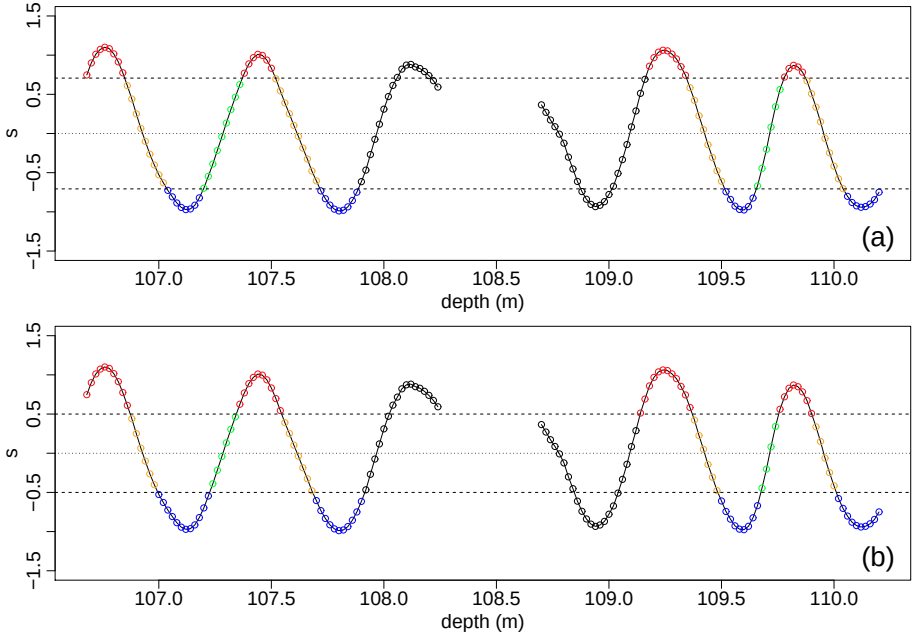


**Fig. 4.** A short stretch of the standardised $H_2O_2$ signal, for depths between 106.68 m and 110.2 m, showing an "issue" (region of dating uncertainty) caused by missing values in the signal and defined using two different values of the threshold $\nu$ for classifying individual points. Points within the "issue" whose classification is uncertain because of their proximity to the missing values, are black; points within peaks (labelled P in main text) are coloured red, descending points (D) are orange, troughs (T) are blue, and ascending points (A) are green. **(a)** Threshold $\nu = 1/\sqrt{2} \approx 0.707$. **(b)** Threshold $\nu = 0.5$.





CPD
8, 2477–2502, 2012



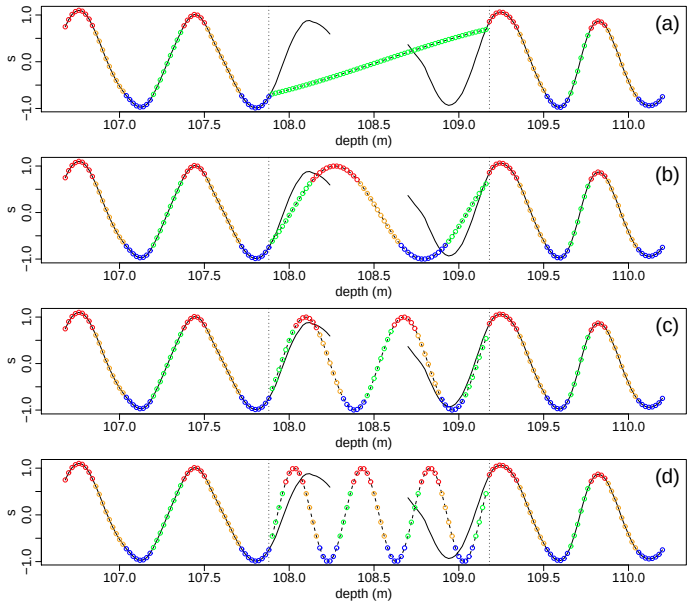


**Fig. 5.** Four possible reconstructions of an "issue" (region of dating uncertainty, bounded by the dotted lines) caused by missing values in the standardised $H_2O_2$ signal (black curve) at a depth of around 108.5 m. Coloured points within the issue represent imputed values for missing observations, or conjectured "corrected" values for nearby observations; coloured points outside the issue indicate the classification of actual observations. In each case, points classified as being within peaks (labelled P in main text) are coloured red, descending points (D) are orange, troughs (T) are blue, and ascending points (A) are green. The four different reconstructions are indexed by $d$, the number of runs (sets of consecutive points classified in the same way) used to reconstruct the issue, including the adjacent runs; adding an extra year to the reconstructed chronology increases $d$ by 4. **(a)** The minimal reconstruction, with no additional annual cycles, involving $d = 1$ runs. **(b)** The reconstruction with one additional annual cycle and $d = 5$. **(c)** The reconstruction with two additional annual cycles and $d = 9$. **(d)** The reconstruction with three additional annual cycles and $d = 13$.





CPD
8, 2477–2502, 2012



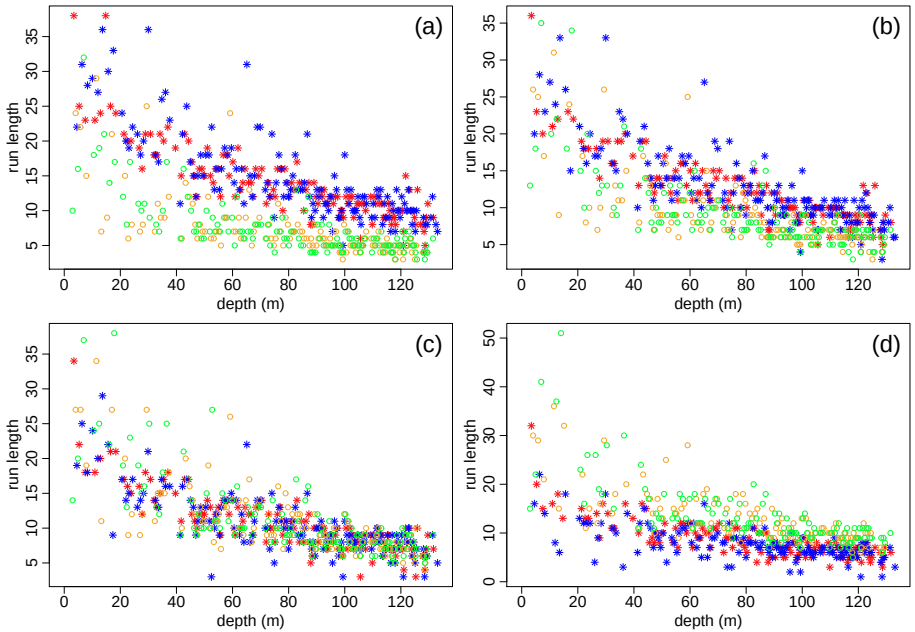


**Fig. 6.** Plots of the length of a certain run – a set of consecutive observations that can be definitely classified as belonging to the same part of an annual cycle – against depth within the core, for four different values of the threshold $\nu$ for classifying individual points. Runs representing peaks (labelled P in main text) are represented by red stars, descending runs (D) by orange circles, troughs (T) by blue stars, and ascending runs (A) by green circles. **(a)** Runs defined by threshold $\nu = 0.5$. **(b)** Runs defined by threshold $\nu = 0.6$. **(c)** Runs defined by threshold $\nu = 0.7$. **(d)** Runs defined by threshold $\nu = 0.8$.





CPD
8, 2477–2502, 2012



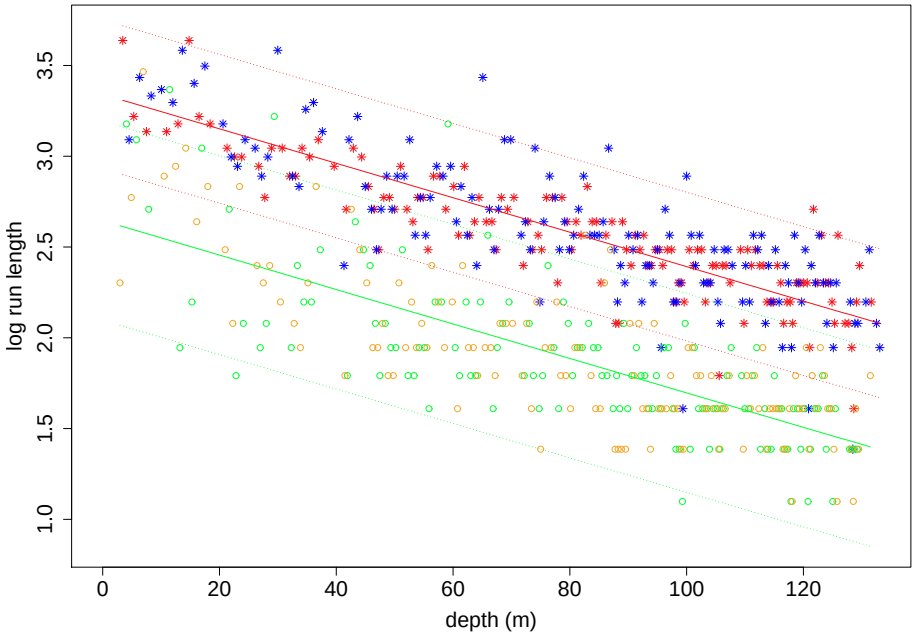


**Fig. 7.** A linear model for the logarithm of the length of a certain run – a set of consecutive observations that can be definitely classified as belonging to the same part of an annual cycle – against depth within the core. Runs representing peaks (labelled P in main text) are represented by red stars, descending runs (D) by orange circles, troughs (T) by blue stars, and ascending runs (A) by green circles. The solid lines indicate the fitted values for peak/trough runs (red) and ascending/descending runs (green), and the red and green dotted lines represent the respective 95% predictive intervals for individual runs.





CPD
8, 2477–2502, 2012



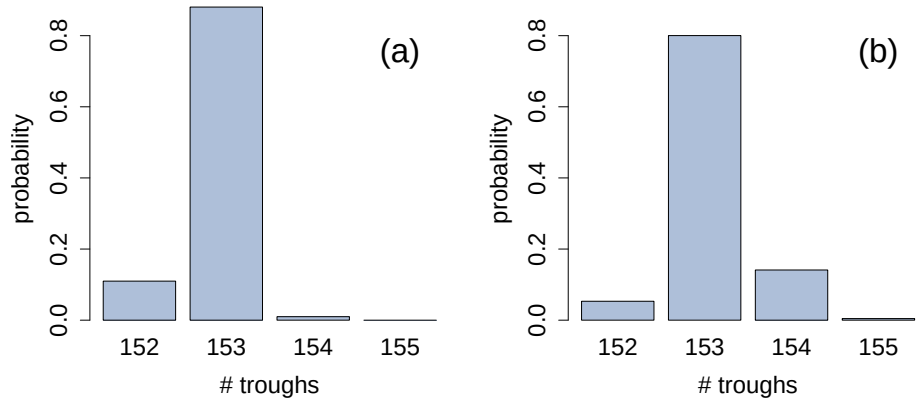


**Fig. 8.** Bar plots showing the probability distributions of the number of annual troughs – a summary of the chronology – in the dating of the Gomez core, based on two possible values of the threshold $\nu$ for classifying individual points, and dividing the core into $\beta = 6$ sub-sections for the initial smoothing process. **(a)** Threshold $\nu = 0.5$. **(b)** Threshold $\nu = 1/\sqrt{2} \approx 0.707$.





CPD
8, 2477–2502, 2012



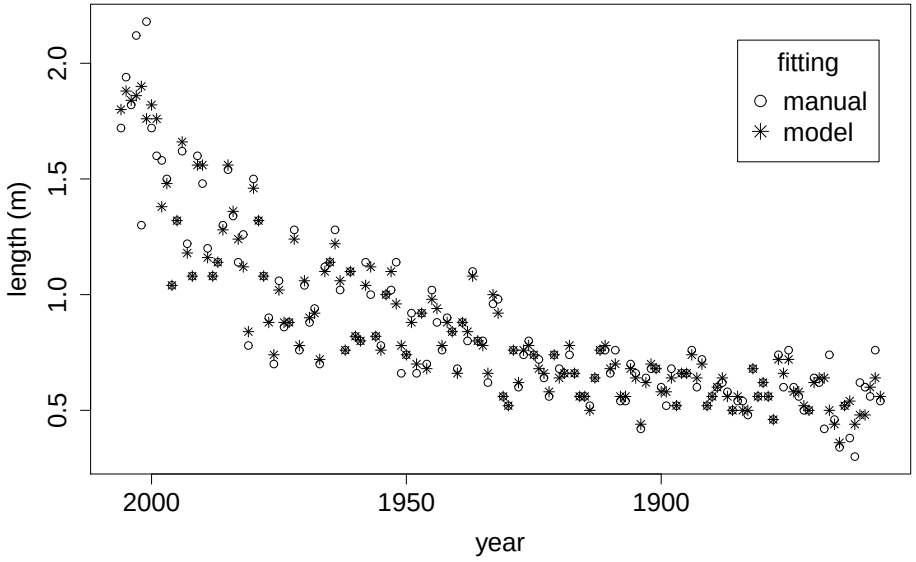


**Fig. 9.** Thickness of estimated annual layers in the Gomez core, based on manual counting and on the automated approach defined here, as a function of calendar year. Circles indicate the thicknesses resulting from manual counting; stars indicate the corresponding thicknesses based on the most likely reconstruction (probability greater than 0.8) using automated layer-counting of the $H_2O_2$ signal.





CPD
8, 2477–2502, 2012



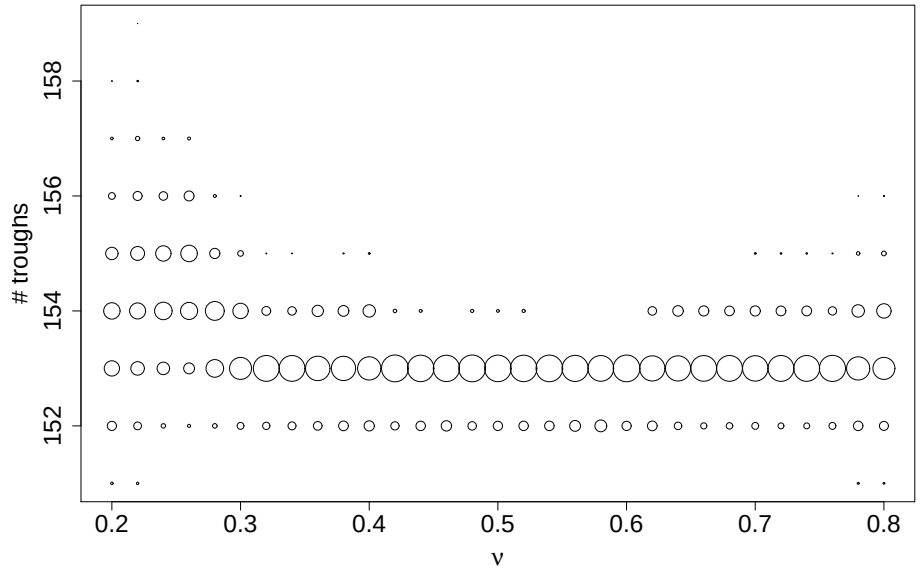


**Fig. 10.** The effect of the threshold $\nu$, for classifying individual points, on the probability distribution obtained for the number of annual troughs – a summary of the chronology – in the dating of the Gomez core, in each case dividing the core into $\beta = 10$ sub-sections for the initial smoothing process. Each circle has area proportional to the probability of a particular number of troughs, when the reconstruction uses a particular value of the threshold $\nu$.





CPD
8, 2477–2502, 2012



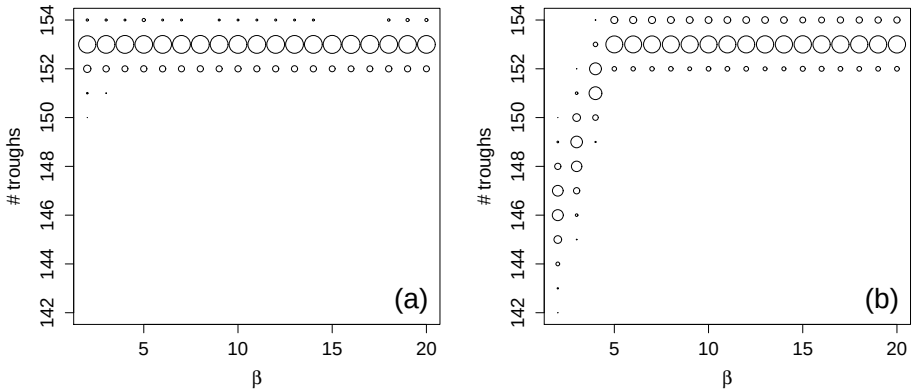


**Fig. 11.** The effect of the number of sub-sections $\beta$ used in the initial smoothing of the $H_2O_2$ signal on the probability distribution obtained for the number of annual troughs – a summary of the chronology – in the dating of the Gomez core, with two possible values of the threshold $\nu$ for classifying individual points. Each circle has area proportional to the probability of a particular number of troughs, when the reconstruction uses particular values of $\beta$ and $\nu$. **(a)** $\beta = 2, \ldots, 20$, with $\nu = 0.5$. **(b)** $\beta = 2, \ldots, 20$, with $\nu = 1/\sqrt{2} \approx 0.707$.





CPD
8, 2477–2502, 2012