Two Interpolation Methods Using Multiply-Rotated Piecewise Cubic Hermite Interpolating Polynomials

PAUL M. BARKER AND TREVOR J. MCDUGALL

School of Mathematics and Statistics, University of New South Wales, Sydney, New South Wales, Australia

(Manuscript received 16 December 2019, in final form 6 February 2020)

ABSTRACT

Two interpolation methods are presented, both of which use multiple Piecewise Cubic Hermite Interpolating Polynomials (PCHIPs). The first method is based on performing 16 PCHIPs on 8 rotated versions of the plot of the data versus an independent variable (such as pressure or time). These 16 PCHIPs are then used to form 8 interpolations of the original data, and finally, these 8 are averaged. When the original data are unevenly spaced with respect to the independent variable, we show that it is best to perform the Multiply-Rotated PCHIP (MR-PCHIP) method using the “data index” as the independent variable, and then to subsequently perform one last PCHIP of the data index with respect to the original independent variable. This MR-PCHIP method avoids the flat spots that are a feature of the PCHIP method when the data have multiple values approximately equal to a local extreme value. The MR-PCHIP interpolated data have continuous first derivatives at the data points. This method also avoids the unrealistic overshoots that can occur when using the standard cubic spline interpolation procedure. The second interpolation method is designed specifically for hydrographic data with the aim of minimizing the formation of unrealistic water masses by the interpolation procedure. This is achieved by applying a Piecewise Cubic Hermite Interpolating Polynomial to each of 8 rotations of the salinity versus temperature plot (Multiply-Rotated Salinity–Temperature PCHIP, MRST-PCHIP) with bottle number (that is, data index) as the vertical interpolating coordinate, thereby making the MRST-PCHIP method independent of the height of a water column. This method is equivalent to interpolating in the salinity–temperature diagram, and MRST-PCHIP proves very effective at avoiding the production of anomalous water masses that otherwise occur when interpolating temperature and salinity separately.

1. Introduction

Interpolating a time series using cubic splines can lead to unrealistic overshoots, especially when the increments of the independent variable (e.g., time) are not constant between successive samples. A significant advantage of the Piecewise Cubic Hermite Interpolating Polynomial (PCHIP; [Fritsch and Carlson 1980](#)) method is that it does not display these unrealistic overshoots; rather a PCHIP interpolant is strictly bounded by the data points. The cubic polynomial between a pair of data points in the PCHIP method is determined by knowledge of the data values at these points as well as the chosen values of the derivatives at these data points. The value of the derivative at a data point is chosen by considering the data point in question, along with the data point to the left and the right of it. These three points define two straight-line slopes

between successive data points. If these linear slopes have different signs, or if both of them are zero, then the slope at the point is set to zero; otherwise, the slope at the point is set equal to the harmonic mean of the two slopes.

While avoiding unrealistic overshoots, the PCHIP interpolant displays “flat spots” near data extrema where there are data points that are approximately equal to a local extreme value. Here we describe interpolation methods that are designed to avoid these flat spots (see the [appendix](#) for information on algorithm availability). The first method uses several rotations of the data $C(\tau)$ as a function of the independent variable τ . That is, several rotated versions [we use eight rotations (the original data plus seven rotations)] of the original data on the C – τ diagram are formed, and the PCHIP-interpolating algorithm is applied to each of the 16 data series. The aim of the rotations is to avoid or minimize the occurrences of multiple extreme values in the data of the 16 rotated axes, and so minimize the flat-spot

Corresponding author: Paul M. Barker, paul.barker@unsw.edu.au

DOI: 10.1175/JTECH-D-19-0211.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](#) (www.ametsoc.org/PUBSReuseLicenses).

nature of the PCHIP procedure at such multiple extrema. Finally, linear combinations of these 16 interpolants are taken to reconstruct the interpolated $C(\tau)$.

When the data are unevenly spaced in the independent variable we find that it is preferable to perform the rotated PCHIP interpolations as a function of a variable that increments by one with each data pair (the “data index” variable) and to separately perform a single PCHIP of this data index variable as a function of the original independent variable.

Oceanic data often consist of salinity and temperature measurements made at discrete irregular pressures in the water column. To facilitate their use oceanographers typically interpolate these samples to regular pressure intervals. To date they have typically applied either splines, or the [Reiniger and Ross \(1968, hereinafter RR68\)](#) interpolation method, or simple linear interpolation, and each of these methods are applied with the independent variable being pressure. The [RR68](#) method uses a weighted mean of values obtained from the linear interpolation of the two nearest data points, and the extrapolation of the pairs of data above and below to construct a “reference curve.” The function to compute the reference curve contains a rational function, where the denominator can and does sometimes go to zero. It calculates two fits to the adjacent values uses a weighted combination of two parabolas with weights that were determined empirically based on mismatches to the reference curve. These existing methods treat salinity and temperature as independent variables and interpolate them against pressure without taking into consideration any damage done to the structure of the salinity–temperature relationship; for example, we show that using splines and [RR68](#) in pressure space sometimes leads to unphysical loops on the salinity–temperature diagram. The second interpolation method of this paper applies the PCHIP interpolation method to eight differently rotated versions of data on the salinity–temperature diagram. Like the first method, it is advantageous to have the independent variable of these 16 PCHIPs to be the data index (i.e., the “bottle number”) and to perform one last PCHIP to relate the pressure and the bottle number of the data. With the introduction of modern CTD’s, hydrographic data are now recorded at regular fine intervals so that simple interpolation techniques such as linear interpolation are often sufficiently accurate for describing the water-mass structure, although, with linear interpolation, the first derivatives in the vertical are discontinuous. However, not all modern data are available at finescale intervals such as those used on marine animals where the data are transmitted at low resolution due to satellite transmission constraints ([Siegelman et al. 2019](#)) or GEOTRACES data (www.geotraces.org).

2. The Multiply-Rotated PCHIP interpolation method

Here the first interpolation method is introduced, motivated by the aim of (i) avoiding the extreme overshooting that cubic splines sometimes exhibit while also (ii) avoiding or minimizing the occurrences of flat spots in the interpolant when the PCHIP method is used on data that contain multiple extreme values. The idea is to perform PCHIP interpolations on several different linear combinations of the dependent and independent variables, C and τ , to use these interpolations to find multiple different interpolated estimates of $C(\tau)$, and then to average these to form the final version. The linear combinations of C and τ are chosen as the axes of rotated versions of the data on the C – τ diagram; we have used eight such rotations in our tests of the method. The idea is that by PCHIP interpolating 16 differently rotated versions of the original data, a flat spot near a particular data point will only appear on at most 1 or 2 of these 16 versions. After averaging eight differently interpolated estimates of $C(\tau)$, the influence of a flat spot on one or two of the eight interpolants will not unduly influence the average. This benefit of rotating the data on the C – τ diagram is illustrated for a single rotation by 45° in [Fig. 1](#). If computational time is a concern, then a user can consider reducing the number of rotations.

The first task is to choose a scaling factor between the units (and hence dimensions) of the two axes C and τ . This scaling factor R that we choose is based on the maximum magnitude of the slope between pairs of data on the C – τ plot,

$$\tilde{R} = \max \left| \frac{C_{i+1} - C_i}{\tau_{i+1} - \tau_i} \right|. \quad (1)$$

In practice we have found that the value $R = 0.33\tilde{R}$ is preferred over $\tilde{R}$ itself, but this multiplying factor of 0.33 is best thought of as an adjustable parameter that a user might wish to change based on their particular data. With the oceanographic data that we have tested we have found that multiplying factors between 0.1 and 1 gave good interpolants. After the scaling is applied to τ , the two variables are

$$q_1 = C \quad \text{and} \quad v_1 = R\tau, \quad (2)$$

and the rotated versions (q_n, v_n) of the discrete data, at an angle of rotation of ϑ_n , are [the original axes, given by [Eq. \(2\)](#), correspond to $\vartheta_1 = 0$]

$$\begin{pmatrix} q_n \\ v_n \end{pmatrix} = \begin{pmatrix} \cos \vartheta_n & -\sin \vartheta_n \\ \sin \vartheta_n & \cos \vartheta_n \end{pmatrix} \begin{pmatrix} q_1 \\ v_1 \end{pmatrix}. \quad (3)$$

FIG. 1. (a) A sketch illustrating how the rotation of data on this plot of the data vs the independent variable, “data index,” can avoid data extremes at successive data points. The original data (blue axes) have two data points with the same Absolute Salinity, and the straightforward PCHIP interpolation (continuous blue line) of this data has an unrealistic flat spot between these two data points. When the data are rotated by 45° small red axes), the data are not necessarily monotonic, but importantly, the two data points that have the same value of Absolute Salinity now have different lengths measured along both red axes as can be seen in (b) and (c). When PCHIPs are performed on the data along both the red axes (q and v), and the values of Absolute Salinity reconstructed, the smooth red curve is found, and this does not exhibit the “flat spot” that arises when using the straightforward PCHIP interpolation method. The MR-PCHIP curve (black dashed line) is the average of the blue and red curves.

The interpolated time series $q_n^{\text{interp}}(\tau)$ and $v_n^{\text{interp}}(\tau)$ are formed for the eight angles of rotation ϑ_n , $0, \pi/16, \pi/8, 3\pi/16, \pi/4, 5\pi/16, 3\pi/8$, and $7\pi/16$, by performing a PCHIP interpolation as a function of the interpolating variable τ between the discrete values $q_n(\tau_i)$ and $v_n(\tau_i)$, thus requiring 16 PCHIPs. The eight desired estimates of $C(\tau)$ are obtained from the eight pairs, $q_n^{\text{interp}}(\tau)$ and $v_n^{\text{interp}}(\tau)$, by rotating them by the angle $-\vartheta_n$, so that

$$\begin{pmatrix} q_1^{\text{interp},n}(\tau) \\ v_1^{\text{interp},n}(\tau) \end{pmatrix} = \begin{pmatrix} \cos\vartheta_n & \sin\vartheta_n \\ -\sin\vartheta_n & \cos\vartheta_n \end{pmatrix} \begin{pmatrix} q_n^{\text{interp}}(\tau) \\ v_n^{\text{interp}}(\tau) \end{pmatrix}. \quad (4)$$

We are interested only in the eight estimates of $q_1(\tau)$, namely $q_1^{\text{interp},n}(\tau)$, not in the eight estimates of $v_1(\tau)$, namely $v_1^{\text{interp},n}(\tau)$. These eight estimates of $q_1(\tau)$, which we might call $C_n(\tau)$, are [from Eq. (4)]

$$C_n(\tau) = q_n^{\text{interp}}(\tau) \cos\vartheta_n + v_n^{\text{interp}}(\tau) \sin\vartheta_n. \quad (5)$$

The final step is to simply average these eight estimates to obtain $C(\tau)$.

In the next section we show some examples of the method for data that are uniformly spaced in the independent variable τ , and in a later section we discuss the best approach for interpolating data that are nonuniformly spaced in τ .

3. Examples of Multiply-Rotated PCHIP with uniform increments of the independent variable

In Fig. 2a we show three interpolation methods applied to a traditional approximation to a delta function, while Fig. 2b is an example of a dataset that has two consecutive equal values of the data at an extremum of the function. The cubic spline interpolation method displays the well-known ringing (Gibbs phenomenon), which does not seem desirable where the data are constant over several data points. The regular PCHIP method has the distinctive feature that it does not have extrema between successive data points, and so in Fig. 2 the PCHIP interpolant is exactly horizontal (a flat spot) between pairs of data points when the two function values are identical such as in an oceanographic mixed layer. If the data represent the output of some physical system that obeys a differential equation, this exactly horizontal flat-spot feature of Fig. 2b is unlikely to be physically relevant. So while this feature of the PCHIP method is an advantage for some datasets, it is not particularly desirable for some datasets from the natural world.

The motivation for this work has been to avoid both 1) the ringing phenomena and the well-known overshooting nature of the cubic spline interpolation method and 2) the flat spotting of the PCHIP interpolation method. Both these features are evident in Fig. 2. Results for the Multiply-Rotated PCHIP (MR-PCHIP) method of this paper are also shown in Fig. 2. This interpolation method (Fig. 2a) exhibits smaller overshoots than those of the cubic spline and (Fig. 2b) these overshoots do not persist beyond the first interval over which there are multiple data points having the same data values.

FIG. 2. Interpolation of two data structures using the usual PCHIP method, cubic splines, and using our Multiply-Rotated PCHIP (MR-PCHIP) method.

The bulk of our testing has been done on oceanographic data where historically the data were sampled at discrete points throughout the water column, usually at irregular depth intervals. These samples typically consist of three properties, being salinity, temperature, and pressure. Oceanographers use the word “cast” to describe a vertical profile of data and the word “bottle” to describe a data point at a particular pressure on such a cast. In Fig. 3 we illustrate these interpolation methods for oceanographic data of Conservative Temperature Θ and Absolute Salinity S_A (IOC et al. 2010) as functions of the observation number (data index) down a particular vertical cast of data. The spline interpolation method results in unrealistic excursions to low values of temperature and salinity in the deepest part (largest data index) of the vertical cast. The linear interpolation method involves infinite changes of slopes at the data points and also is unrealistic where the data exhibit obvious curvature in the vertical direction. The MR-PCHIP method appears to perform well on these vertical casts.

4. Examples of MR-PCHIP with nonuniform increments of the independent variable

The data presented in Fig. 3 were actually values of Conservative Temperature Θ and Absolute Salinity S_A from 21.85°N, 201.64°E in the Pacific Ocean measured at nonuniform values of pressure down a vertical cast. In forming Fig. 3 we ignored the pressure data and interpolated the temperature and salinity data as though they were observations that were evenly spaced, calling the interpolating variable τ data index.

To interpolate these vertical casts of Θ and S_A as a function of pressure we begin by applying the method of section 2 with the interpolating variable being pressure, p , rather than the data index. The results are shown in

Fig. 4 where it is clear that the cubic spline method behaves very poorly, now exhibiting very large positive excursions of both of Θ and S_A in the deepest bottle pair, to be compared with the negative excursions of both Θ and S_A between this same bottle pair when done with τ as the independent variable (see Fig. 3). The straightforward PCHIP method seems to perform the best in Fig. 4, but we know that this method exhibits unrealistic flat spots when data values near data extrema are similar to the extreme value. The straightforward application of the method of section 2 of this paper does not perform at all well, with undesirable negative excursions between the deepest pair of bottles [see the MR-PCHIP(p) lines in Fig. 4]. We conclude that using a nonuniformly distributed independent variable (in our case, pressure) as the interpolating variable is quite unsatisfactory for both splines and the method of section 2 of this paper.

We have found that a far better way of interpolating data that are measured at nonuniform increments of the interpolating variable p is to first ignore these nonuniform increments and perform the interpolation method of section 2 with respect to “observation number” or data index τ_i , and then to subsequently stretch this data index axis back to pressure space using an additional straightforward PCHIP interpolation $\tau(p)$. To place the interpolated values of $S_A(\tau)$ and $\Theta(\tau)$ of section 2 at a given value of pressure on the water column, it would seem sensible to perform another PCHIP interpolation of pressure as a function of τ , $p(\tau)$. This would then treat all three observed variables, namely Absolute Salinity, Conservative Temperature, and pressure equally, with $S_A(\tau)$, $\Theta(\tau)$, and $p(\tau)$ being interpolated as a function of the same interpolating variable, τ . This would work, but in the most common use of interpolation the user specifies values of pressure at which the interpolated

FIG. 3. Examples of the interpolation methods using data index τ as the interpolation variable. (a) Absolute Salinity interpolated with data index, (b) Conservative Temperature interpolated with data index, (c) zoom in on three bottles of the interpolated Absolute Salinity, and (d) zoom in on three bottles of the interpolated Conservative Temperature.

variable is required, and because of this, the final PCHIP we recommend is not $p(\tau)$ but $\tau(p)$. This is done merely as a numerical expediency in order to avoid the need to perform a Newton–Raphson iteration each time we

enter the algorithm with a known value of pressure. By performing the one-dimensional PCHIP of $\tau(p)$ instead of the more easily justified $p(\tau)$ we are assuming that these two functional relationships are almost identical.

FIG. 4. Examples of the interpolation methods using both pressure p and τ as interpolation variables. The MR-PCHIP interpolation method that uses τ as the interpolation variable is plotted against pressure using the seventeenth PCHIP $\tau(p)$.

Since pressure is a monotonic function of the data index, we do expect this to be the case. We have tested these two ways of interpolating between pressure and data index and find that the difference in pressure in doing the interpolation one way versus the other is indeed very small.

In this way our combined MR-PCHIP interpolation method based on [section 2](#) arrives at the interpolation $\Theta(p)$ after doing a total of 17 PCHIP interpolations. This procedure is a multiply rotated 17 PCHIP interpolation method, MR-PCHIP that yields a piecewise cubic polynomial for both $\Theta(\tau)$ and $\tau(p)$. These piecewise cubic polynomials can be combined in the function-of-function form $\Theta[\tau(p)]$, and we note that this can be regarded as a sixth-order polynomial of pressure.

The results of this way of dealing with the nonuniform spacing of the independent variable is also shown in [Fig. 4](#) for the same vertical cast of Absolute Salinity, Conservative Temperature and pressure. The MR-PCHIP(τ) method of [section 2](#), performed with the interpolating variable being data index τ and with the subsequent stretching to pressure using the seventeenth PCHIP $\tau(p)$, performs well, closely approximating the Multiply-Rotated Salinity–Temperature MRST-PCHIP interpolation method that we develop below.

5. The Multiply-Rotated Salinity–Temperature PCHIP Method (MRST-PCHIP)

The key motivation in developing this second interpolation method of this paper was to have the interpolated data display realistic water-mass structure on the Absolute Salinity–Conservative Temperature (S_A – Θ) diagram. That is, the interpolated data, when plotted on the S_A – Θ diagram should be as realistic as possible. It was found that when interpolating temperature and salinity separately as a function of pressure using existing interpolation methods, the resulting S_A – Θ diagram of the interpolated data was often quite unrealistic, sometimes displaying erroneous loops and other features. Moreover, while the MR-PCHIP method developed above shows better S_A – Θ characteristics than splines, the straightforward PCHIP and the [RR68](#) method, we thought that further improvements might be possible. This motivated us to develop the following method that is based directly on the data plotted on the S_A – Θ diagram.

We first scale the temperature and salinity axes of the S_A – Θ diagram to reflect the relative variation of temperature and salinity in the World Ocean. Thus, we form the scaled salinity variable $\tilde{S}_A$,

$$\tilde{S}_A = q_1 = RS_A, \quad (6)$$

and we have found that a suitable value of the scaling parameter R is $9 \text{ K (g kg}^{-1}\text{)}^{-1}$. This scaling factor was chosen to reflect the global variations in temperature and salinity in the World Ocean. That is, we scale the S_A – Θ diagram so that the length along the x axis represented by a difference of Absolute Salinity of 1 g kg^{-1} is the same as the length along the y axis of a difference of Conservative Temperature of 9°C . The interpolation technique performs well when this scaling factor is between 3 and $27 \text{ K (g kg}^{-1}\text{)}^{-1}$.

The general expression for the rotated variables (q_n, v_n) in terms of the original axes $(\tilde{S}_A, \Theta)$ for an angle of rotation of ϑ_n is

$$\begin{pmatrix} q_n \\ v_n \end{pmatrix} = \begin{pmatrix} \cos\vartheta_n & -\sin\vartheta_n \\ \sin\vartheta_n & \cos\vartheta_n \end{pmatrix} \begin{pmatrix} \tilde{S}_A \\ \Theta \end{pmatrix}. \quad (7)$$

The original unrotated axes are recovered with $\vartheta_n = 0$, so that, for example, $q_1 = \tilde{S}_A$ and $v_1 = \Theta$. Recovering the $\tilde{S}_A$ and Θ from knowledge of (q_n, v_n) involves rotation by the angle $-\vartheta_n$ so that

$$\begin{pmatrix} \tilde{S}_A \\ \Theta \end{pmatrix} = \begin{pmatrix} \cos\vartheta_n & \sin\vartheta_n \\ -\sin\vartheta_n & \cos\vartheta_n \end{pmatrix} \begin{pmatrix} q_n \\ v_n \end{pmatrix}. \quad (8)$$

Writing these equations in component form, we have [Eq. (7)]

$$q_n = \tilde{S}_A \cos\vartheta_n - \Theta \sin\vartheta_n \quad \text{and} \quad v_n = \tilde{S}_A \sin\vartheta_n + \Theta \cos\vartheta_n, \quad (9)$$

while Eq. (8) is

$$\tilde{S}_A = q_n \cos\vartheta_n + v_n \sin\vartheta_n \quad \text{and} \quad \Theta = -q_n \sin\vartheta_n + v_n \cos\vartheta_n. \quad (10)$$

We choose the angles $0, \pi/16, \pi/8, 3\pi/16, \pi/4, 5\pi/16, 3\pi/8$, and $7\pi/16$, and for each of these eight angles we have a rotated version of the $\tilde{S}_A$ – Θ diagram and we perform a PCHIP interpolation on the data along each of these 16 axes. That is, we form eight axes of both q_n and v_n from Eq. (9) and a PCHIP is performed for each of these. Having formed these 16 interpolants, they are rotated back to the direction of the original temperature and salinity axes using Eq. (10). At this stage there are 8 different interpolating cubics of both $\tilde{S}_A$ and Θ to each of which we give an equal weight of one-eighth.

So far in this description we have not discussed our choice of the independent variable that is chosen when performing each of the 16 PCHIPS. We first performed these 16 PCHIPS using pressure as the interpolating variable, and found a huge improvement over the existing interpolation methods of (i) cubic splines, (ii) linear

interpolation, and (iii) the straightforward use of a PCHIP separately for each of $S_A(p)$ and $\Theta(p)$. This 16 rotated PCHIP method avoids the extreme wiggles of splines, it avoids the biased nature of linear interpolation, and it also avoids the strict flat-spot behavior of the standard PCHIP method at property extrema. Our Multiply-Rotated PCHIP method has the character of a smooth spline-like interpolation, but without the sometimes extreme overshooting of spline interpolation. This new method has continuous first derivatives but discontinuous second derivatives at the data points.

However, it would be preferred that the interpolated S_A – Θ diagram be completely independent of pressure, so that a simple vertical heaving of the water column does not lead to any change whatsoever in the interpolated S_A – Θ curve. To achieve this it is necessary to change the independent variable (the interpolating variable) from pressure to something that is a function of Absolute Salinity and Conservative Temperature only. With this in mind we tried the independent variable to be the cumulative sum down the cast of

$$\Delta\tau = \sqrt{0.01 + (\Delta\Theta/K)^2 + (\Delta\tilde{S}_A/K)^2}, \quad (11)$$

where $\Delta\tilde{S}_A$, $\Delta\Theta$, and $\Delta\tau$ are the differences between adjacent bottle pairs. This choice of τ produced good results, but then we found that the interpolated S_A – Θ diagram improved slightly by increasing the leading term (the 0.01 number above) to the point where $\Delta\tau$ approached a constant. It was decided to adopt this simpler approach, so that the interpolating variable τ is simply put equal to the bottle number (or data index), with τ in increments of by one between each bottle pair. In this way the plot of the interpolated S_A and Θ fields on the S_A – Θ diagram is completely independent of the pressure of the original data.

An advantage of this choice of interpolating coordinate is that when the observed bottles have very similar S_A and Θ values and so are close to being collocated on the S_A – Θ curve but are widely separated in depth, interpolating with respect to the new coordinate will naturally limit the variations of Absolute Salinity and Conservative Temperature between the bottles. In contrast, if depth or pressure is the interpolating variable, there can be much larger variations of salinity and temperature between the bottles in this situation. Another example of the benefit of the data index over pressure as the independent variable is when the pressure increment between successive bottles is very small. This causes very large temperature and salinity gradients, which strongly impact the interpolation between adjacent bottle pairs. This can occur in ocean mixed layers, causing unwanted temperature and salinity variations

in an otherwise nearly well-mixed region, and this behavior is greatly diminished by using data index as the independent variable.

So, at this stage the “multiply rotated 16 PCHIP method” produces a piecewise cubic interpolation of Absolute Salinity and Conservative Temperature as a function of τ , which at the data, is equal to the bottle number of the data. Between these integer values of τ are noninteger values of τ that, at this stage, serve simply as a parameter to give $S_A(\tau)$ and $\Theta(\tau)$, which is functionally equivalent to interpolating $S_A(\Theta)$ or $\Theta(S_A)$. If all we wanted was an interpolated S_A – Θ diagram we could perhaps stop here. But we also want to find interpolated vertical profiles of Absolute Salinity and Conservative Temperature as a function of pressure down the cast. To place the interpolated values of $S_A(\tau)$ and $\Theta(\tau)$ at a given value of pressure on the water column we adopt the same procedure as used with our MR-PCHIP method, as described in section 4. That is, we perform 16 PCHIPS using τ as the interpolating variable, as well as a final (seventeenth) PCHIP $\tau(p)$.

6. Results of the MRST-PCHIP method

Figure 5 shows the same vertical cast as Figs. 3 and 4, a cast taken by the R/V *Moana Wave* on 19 January 1994 in the North Pacific Ocean, adjacent to Hawaii, as part of the Joint Global Ocean Flux Study (JGOFS) Hawaii Ocean Time Series (HOTS) project. This vertical cast depicts an extreme case where no data were recorded between 500 and 2500 dbar, a situation that is encountered in historical data where the water column properties were measured with irregularly spaced Niskin bottles or from a mooring.

The cast was interpolated using splines, RR68, linear, and the usual PCHIP method, in each case with the independent variable being pressure, and with the interpolation method being applied to both $S_A(p)$ and $\Theta(p)$, separately. In addition, Fig. 5 shows the results of using both our MR-PCHIP and MRST-PCHIP methods, both of which use a total of 17 PCHIP calculations, and both of which use the bottle number τ as the independent variable for 16 of these 17 PCHIP calculations. Figures 5a and 5b depict the vertical S_A and Θ distributions, and the spline- and RR68-interpolated curves are completely unrealistic. In the water column between 500 and 2500 dbar, the spline produced temperatures above 40°C with salinities greater than 38 g kg^{−1}, whereas the temperatures and salinities of the observed bottles bounding this section of the water column ranged between 7° and 2°C and 34.2 and 34.7 g kg^{−1}, respectively. Over this same section of water the RR68 technique produced extremely fresh water with salinity levels down to 33.2 and

FIG. 5. Cast located at 21.85°N, 201.64°E adjacent to Hawaii in the North Pacific Ocean. (a) Absolute Salinity vs pressure and (b) Conservative Temperature vs pressure. Zoom in on the shallowest 150 dbar of (c) Absolute Salinity vs pressure and (d) Conservative Temperature vs pressure. (e) S_A - θ plot, (f) S_A - θ plot, zoomed in (g) S_A - θ plot, zoomed in on the shallowest portion of the water column. (h) S_A - θ plot, zoomed in on the deepest portion of the water column.

temperatures less than -12°C , which is significantly less than the freezing temperature. For both of these interpolation techniques the interpolated values were either much greater or much less than of the observations in the whole water column. Over this same section of the water column the linear, PCHIP, MR-PCHIP, and MRST-PCHIP methods all produced values that were more realistic.

In the shallowest 150 dbar (Figs. 5c,d), the spline, RR68 and MR-PCHIP curves contain a significant amount of structure between the observed bottles, with the spline exhibiting larger anomalies than any of the other techniques. Between 50 and 60 dbar the MR-PCHIP method shows a similar bump to higher salinities as does the spline method, and this does not seem particularly realistic. Similarly, between 5 and 25 dbar, the MRST-PCHIP method seems the most realistic method.

When the interpolated values are plotted on the S_A - Θ diagram (Figs. 5e-h) the unrealistic nature of some of the interpolated values becomes obvious. The spline fit produces the most unrealistic features, while the RR68 method gives a very cool (too cool by $\sim 10^{\circ}\text{C}$) and fresh water mass in this deep interpolated region. Eliminating these two clearly wrong methods from discussion we are able to get a better understanding of the performance of the remaining three methods at producing a realistic water column structure (Fig. 5f). A close inspection of the upper 500 dbar of the water column, on the S_A - Θ plot (Fig. 5g), reveals that the spline, RR68 and MR-PCHIP results are also unrealistic because the interpolated results from these methods contain complete loops or spiral-like features around the data points. The deepest portion of the water column (Fig. 5h) reveals that the PCHIP-interpolated values are the least like the observed data in the region of the cast, with the linear interpolation and the MRST-PCHIP method of the present paper being preferred.

Figure 6 shows a cast taken by the R/V *Melville* in 1971 in the central Pacific Ocean. The salinity versus pressure plot (Fig. 6a) reveals two reversals in the salinity concentration, the first occurs between 150 and 200 dbar, while the second is between 300 and 400 dbar. The first reversal (Fig. 6b) occurs between two bottles that have nearly identical salinities. Applying the standard historical interpolation techniques (spline, RR68, PCHIP, and linear) with pressure as the interpolating variable results in a wide range of estimates of the salinity. In the case of the first retroflexion (Fig. 6b), all the interpolation methods overshoot the data except for the linear and PCHIP methods. Usually an overshoot is a good feature as it is very unlikely that the data samples have recorded the maximum and minimum concentrations in the water column. The RR68 curve is

not smooth over this depth range; this feature is due to their use of a rational function in the calculation of the interpolated values, so that the denominator can tend toward zero. Like the spline and the RR68 curve, our MR-PCHIP and MRST-PCHIP overshoot the observed data points and like the spline they produce smooth curves. Of the six ways of interpolating the data, the present MR-PCHIP and MRST-PCHIP methods seems the most realistic in this depth range.

The second reversal (Fig. 6c), located a little deeper in the water column, occurs at a local minimum in the salinity concentration. The spline, PCHIP, and MRST-PCHIP curves are quite similar to each other. The linear interpolation scheme overestimates the salinity and the RR68 technique contains kinks or bumps in the interpolated values, while the MR-PCHIP curve, like the RR68 curve, significantly overshoots the known data values, producing low salinity concentrations that do not seem realistic.

In the shallowest section of the Absolute Salinity-pressure profile (0–75 dbar; see Fig. 6b where the uppermost four data points are in the ocean mixed layer) the interpolated spline curve contains oscillations between the data points; this is an undesirable feature that is also evident in the MR-PCHIP interpolant but is not a feature of the other interpolation techniques. The Conservative Temperature-pressure plots are not shown in Fig. 6 as they do not show much variation between the different interpolation techniques. This is due to the monotonic nature of the temperature distribution.

The S_A - Θ plot (Fig. 6d), like the salinity pressure plot, contains the two large reversals in the salinity distribution. The shallow reversal in the salinity concentration in the S_A - Θ curve (Fig. 6e), highlights the differences in the interpolation techniques. The linear interpolation unsurprisingly does not contain any overshoot between the data points. The spline produces a smooth curve that behaves well in this region. The magnitude of the overshoot is less in the MR-PCHIP and MRST-PCHIP interpolants than in the spline; however, all three are symmetric. Of the six interpolation methods the MRST-PCHIP method of the present paper looks the most realistic.

Farther down in the water column, at the salinity minimum (Fig. 6f), the RR68-interpolated values also exhibits wiggles in the S_A - Θ curve that do not appear in either the spline or the MRST-PCHIP curves, and the MR-PCHIP curve contains a bump that is of similar magnitude to the RR68 wiggles. The overestimate of the linear interpolant of salinity between the observed bottles is apparent.

The vertical interpolation scheme used to date by the National Oceanographic Data Center (NODC), described

FIG. 6. Data collected by the R/V *Melville* in 1971 in the central Pacific Ocean. (a) Absolute Salinity vs pressure, (b) zoom in on the upper 250 dbar of the water column, depicting Absolute Salinity vs pressure, and (c) zoom in on the deeper retroflexion in the salinity concentration. (d) An S_A - Θ plot for the whole water column, (e) zoom in on the shallowest retroflexion in the S_A - Θ plot, and (f) the S_A - Θ plot depicting the retroflexion that occurs deeper in the water column.

in Johnson et al. (2011), uses a combination of the RR68, linear, and three-point Lagrangian interpolation schemes, with the choice of scheme being based on the vertical spacing between data points. With the vertical spacing of data as on Fig. 6, the Johnson et al. (2011) scheme applies the RR68 method, and then checks to see if all the interpolated data between a pair of observations lie between those values. If a local extrema is detected then the

RR68 scheme is abandoned and linear interpolation is used. Hence with the data of Fig. 6, the NODC scheme displays the cusps of the RR68 scheme on Figs. 6c and 6f except at the reversals of salinity as a function of depth, where linear interpolation is used. In this way the Johnson et al. (2011) interpolation method chooses between two different interpolation methods, both of which we have demonstrated in Figs. 5 and 6 as being unphysical.

FIG. 7. Errors in (a) Conservative Temperature, (b) Absolute Salinity, and (c) water mass in the reconstructed cast at 28.8°S, 316.4°E, from Blue Earth Global Expedition 2003 (BEAGLE cruise), down to 3000 dbar due to vertical interpolation using linear [RR68](#), MR-PCHIP, and MRST-PCHIP methods. The water-mass error is defined between two points as the distance on the S_A – θ diagram perpendicular to the interpolated values and the data, measured in density units (kg m^{-3}).

To determine the effectiveness of the interpolation methods we smoothed and subsampled all 494 casts of the 2003 Blue Earth Global Expedition (BEAGLE cruise), in the Southern Hemisphere, then reconstructed them using linear, [RR68](#), MR-PCHIP and MRST-PCHIP interpolations. The smoothing was applied to remove spikes and any water-mass intrusions between the subsampled bottles; it consisted of a running mean where the span was set to be equal to the spacing of the subsampled bottles. We used observed pressures of station 83 from cruise 3, Gasconye 1963, being a randomly selected historical cast that extended down about 3000 dbar and contained 15 Niskin bottles at the subsampled bottle depths (5, 50, 100, 150, 200, 300, 500, 700, 900, 1100, 1300, 1500, 2000, 2500, 3000 dbar). The Conservative Temperature, Absolute Salinity and water-mass errors for each of the interpolation methods are shown in [Figs. 7a–c](#) for a BEAGLE cast, located at 28.8°S, 316.4°E, (southern Atlantic Ocean).

The nonuniform variation in the slope of the of [RR68](#)'s interpolated property is clearly shown in the error plots where they appear as sharp spikes; there are three such examples in the salinity and water-mass errors ([Figs. 7b,c](#)), about 650, 950, and 2100 dbar yet there is only one such

example in the temperature structure occurring at 1000 dbar. These features are due to the interpolation technique not the underlying data as they only occur in one of the interpolants.

The RMS and the mean depth-integrated errors are shown in [Tables 1](#) and [2](#) for Conservative Temperature and Absolute Salinity, while [Table 3](#) gives the RMS and the mean water-mass errors for all 494 casts. Two different pressure ranges are displayed: from the surface to 700 dbar and from the surface to 3000 dbar. These tables clearly indicate that linear interpolation produces the largest errors, which is also seen in [Figs. 7a–c](#). When linear interpolation is used to calculate the global heat content between the surface and 700 dbar, the error is 2.5 times as large as that of MRST-PCHIP and MR-PCHIP. Increasing the depth range to 3000 dbar results in the error being 5.4 times as large. Not surprisingly the linear interpolations also produces the largest errors when considering saline content and water-mass errors. The errors produced with [RR68](#) are considerably smaller than the linear errors, but they are larger than those produced by either MRST-PCHIP or MR-PCHIP.

To gain an understanding of the errors produced on a global scale by [RR68](#), linear and our interpolation

TABLE 1. Conservative Temperature errors ($^{\circ}\text{C}$) from all reconstructed casts of the Blue Earth Global Expedition 2003 (BEAGLE cruise), from 0 to 700 dbar and from 0 to 3000 dbar due to vertical interpolation using linear, [RR68](#), MR-PCHIP, and MRST-PCHIP methods. The RMS column is the root-mean-square of the difference between the interpolant and the data down the water column, while the mean heat content column is the average over all the casts of the vertical integral of the signed difference between the interpolated temperature and the data.

Pressure range (dbar)	RMS ($\times 10^{-3}$)		Mean heat content ($\text{J m}^{-2}; \times 10^6$)	
	0–700	0–3000	0–700	0–3000
Linear	183.44	109.49	72.8	289.0
RR68	116.09	68.23	32.9	72.8
MR-PCHIP	108.04	63.07	28.5	46.6
MRST-PCHIP	114.34	66.48	28.9	48.9

schemes we used the *World Ocean Atlas 2013* (WOA) ([Locarnini et al. 2013](#); [Zweng et al. 2013](#)) monthly composite for January. The vertical axis is supplied as depth, and this was converted to pressure. We then stabilized each of the profiles, which consisted of 102 levels, using the method of [Barker and McDougall \(2017\)](#). From this we interpolated to a regular 1 dbar in the vertical using the MRST-PCHIP method; we also performed this step using [RR68](#) and the linear interpolation methods and these options did not change our findings. Using the same observed pressures of station 83 from cruise 3, Gasconye 1963, down to 3000 dbars, we subsampled the high-vertical-resolution interpolated atlas and then reinterpolated each of the profiles back to 1 dbar using our MRST-PCHIP method. We defined the error in the vertical interpolation as the reinterpolated profile minus the initial constructed profile.

The global heat content error ([Fig. 8](#)) is defined as the difference in the integrated heat content over the top 700 dbar between the original and the reinterpolated temperature values on pressure surfaces. None of the methods are able to reproduce the original profiles without errors, all methods contain errors of different magnitude and varying location of the errors. In all of the discussion of the error estimates we do not include in our discussions of the waters in the Arctic, as the original WOA data there contained unrealistic S_A – Θ curves.

The [RR68](#) interpolation scheme ([Fig. 8a](#)) and our MRST-PCHIP ([Fig. 8c](#)) contain errors that have very similar magnitudes at similar horizontal locations. The most significant differences between the two occur between the equator and 30°S in the Pacific Ocean where [RR68](#) suffers its largest error. The linear interpolation contains larger errors, particularly in the equatorial Pacific Ocean ([Fig. 8b](#)). Unlike the [RR68](#) and MRST-PCHIP interpolations, there is a very clear spatial distribution to

TABLE 2. Absolute Salinity errors (g kg^{-1}) from all reconstructed casts of the Blue Earth Global Expedition 2003 (BEAGLE cruise), from 0 to 700 dbar and from 0 to 3000 dbar due to vertical interpolation using linear, [RR68](#), MR-PCHIP, and MRST-PCHIP methods. The RMS column is the root-mean-square of the difference between the interpolant and the data down the water column, while the mean salt content column is the average over all the casts of the vertical integral of the signed difference between the interpolated salinity and the data.

Pressure range (dbar)	RMS ($\times 10^{-3}$)		Mean salt content (kg m^{-2})	
	0–700	0–3000	0–700	0–3000
Linear	23.47	10.41	1.65	4.20
RR68	16.61	15.97	1.39	2.35
MR-PCHIP	16.31	6.41	0.98	1.61
MRST-PCHIP	15.50	6.99	0.97	1.59

the linear interpolation errors. These errors are biased to positive values that are concentrated in the subtropical and tropical latitudes (30°S–30°N) while the smaller negative errors are located in the southern temperate latitudes (60°–30°S) and in the western North Atlantic Ocean. When the errors in heat content per square meter are multiplied by the area of each of the grid boxes and summed, the linear interpolation error is equivalent to approximately 5 years worth of the current rate of warming of the ocean due to climate change.

The estimate of errors in the fresh water content that we have adopted is the global zonally averaged salinity error ([Fig. 9](#)), and this was calculated by comparing the difference between the interpolated values and the original data on pressure surfaces, then summing them zonally along lines of constant latitude. The white horizontal lines are where each of the pressures have data and so the errors are all zero as the interpolant pass through these points.

TABLE 3. Water-mass errors (kg m^{-3}) from all reconstructed casts of the Blue Earth Global Expedition 2003 (BEAGLE cruise), from 0 to 700 dbar and from 0 to 3000 dbar due to vertical interpolation using linear, [RR68](#), MR-PCHIP, and MRST-PCHIP methods. The water-mass error is defined as the difference between the interpolated values and the data in the direction perpendicular to a pair of data points on the S_A – Θ diagram, measured in density units (kg m^{-3}). The RMS column is the root-mean-square of the water-mass difference between the interpolant and the data down the water column, while the mean column is the average over all the casts of the magnitude of the depth-averaged water-mass differences.

Pressure range (dbar)	RMS ($\times 10^{-3}$)		Mean ($\times 10^{-3}$)	
	0–700	0–3000	0–700	0–3000
Linear	10.78	9.37	4.07	1.55
RR68	8.38	6.21	2.59	0.90
MR-PCHIP	7.25	5.98	1.644	0.52
MRST-PCHIP	6.90	5.55	1.69	0.51

FIG. 8. Global heat content error (J m^{-2}) due to (a) vertical interpolation using [RR68](#), (b) linear interpolation in the vertical, and (c) vertical interpolation using our MRST-PCHIP method.

The [RR68](#) interpolation scheme ([Fig. 9a](#)) and our MRST-PCHIP method ([Fig. 9c](#)) both contain errors that do not display a clear pattern. Excluding the northern latitudes where the original data are unrealistic, the maximum errors are in the upper section of the water column and in the Southern Ocean adjacent to the Antarctic coast. The linear interpolation error ([Fig. 9b](#)) contains clear and significant systematic errors, mostly caused by the overestimation of salinity due to the curvature of the S_A - Θ curve at the salinity minimum of Antarctic Intermediate Water (AAIW) (see [Fig. 6f](#)). The most obvious feature is the large positive error that stretches from 50°S to 40°N over the depth range of 200–1000 dbar. In the deeper waters (1500–3000 dbar) of the Southern Hemisphere (60° – 20°S), is a notable negative error due to the interpolation underestimating the higher salinity values that occur due to the North Atlantic Deep Water. The surface waters, which consist of a salty layer lying above the AAIW, also show a significant overestimation of the salinity.

FIG. 9. Global zonally averaged salinity error due to (a) vertical interpolation using [RR68](#), (b) linear interpolation, and (c) MRST-PCHIP interpolation.

The conclusion we draw from the figures, [Tables 1–3](#), and by examining many more vertical cast from historical ocean data, is that the most realistic method is MRST-PCHIP. Unlike any of the other methods presented here MRST-PCHIP has been designed to interpolate on the S_A - Θ diagram and as such it does not produce unrealistic water masses. Another advantage feature is that heave does not have any effect. When compared with MRST-PCHIP the water-mass errors are

1.5 and 2.4 times larger for [RR68](#) and linear interpolation, respectively, in the upper 700 dbar, they increase to 1.8 and 3 times larger, respectively, when the depth range is increased to 3000 dbar.

When one does not have access to both salinity and temperature profiles, then the MR-PCHIP method is recommended. The question then arises, What would be the best method for vertically interpolating another variable, such as nitrate concentration? One could use the MR-PCHIP method for this task, but we suspect that a better procedure would be to use MRST-PCHIP method with nitrate being input to the algorithm in place of Absolute Salinity, and with a differently chosen scaling factor R that is appropriate to the new variable. However, this would be dependent on the location of the data; in the polar regions, it may well be better to replace temperature with nitrate concentration keeping salinity, once again choosing a suitably chosen scaling factor R .

7. Conclusions

In this paper we have described two methods for interpolating hydrographic data. Both of these methods use 16 PCHIP calculations on rotated versions of the hydrographic data plots, along with a final PCHIP to provide the interpolated data with a pressure value. Both of these methods avoid the flat-spot nature of the usual PCHIP method, and neither of them exhibit the unrealistically large overshoots than are sometimes found with splines.

In developing these two multiply rotated interpolation methods we tried a variety of different methods and combinations of methods. We have shown that cubic splines sometimes give unrealistic overshoots, and this explains why cubic splines are no longer used in physical oceanography. Both our methods have PCHIP as their underlying interpolation technique, and we have found that it is advantageous to adopt the independent interpolating variable as the data index or bottle number rather than being the pressure of the data point. The eliminates issues that can arise when data points are very close together in pressure but the temperature or salinity gradient is large. This large gradient can have a strong influence on the interpolant in adjacent bottle pairs. The adoption of the data index as the independent variable ensures that such artificially large gradients do not arise, and one beneficial consequence is that the interpolated temperature and salinity values in the ocean mixed layer remain essentially well mixed even when there are small temperature or salinity differences that occur over very small differences in pressure.

Our Multiply-Rotated Salinity–Temperature MRST-PCHIP method effectively does its interpolation in

Absolute Salinity–Conservative Temperature space, and we have chosen a suitable scaling factor between temperature and salinity. In the ocean interior both salinity and temperature are subjected to the same physical processes of advection and turbulent mixing, and the different spatial structure of the salinity and temperature fields are due to the different boundary conditions for temperature and salinity at the ocean's boundaries. For this reason we chose the scaling factor to be the ratio of the ranges of the sea surface boundary conditions for temperature and salinity.

Of data we have examined, our MRST-PCHIP method behaves the best, closely followed by the MR-PCHIP method. Recall that our MRST-PCHIP method is specifically designed to respect the water-mass structure of the water column on the S_A – Θ diagram, as illustrated in [Figs. 5](#) and [6](#). The advantages of the MRST-PCHIP method over the MR-PCHIP method can be seen on [Fig. 5g](#) where the small artificial loops of the MR-PCHIP method on the S_A – Θ diagram do not occur with MRST-PCHIP. It is instructive to understand how our MRST-PCHIP method behaves in an ocean mixed layer where the temperature and salinity values are almost constant. In this case each of the 16 linear combinations of S_A and Θ are also almost constant in this region, and so each of the 16 PCHIP interpolants will also be constant in this region. In this way MRST-PCHIP has the desirable feature that mixed layers remain well mixed, in contrast to splines, which shows the Gibbs ringing phenomenon, and also to the MR-PCHIP method that allows nonconstancy in the deepest bottle pair of the mixed layer.

Acknowledgments. Dr. Simon Wotherspoon is thanked for introducing us to PCHIPs, and we gratefully acknowledge financial support from the Australian Research Council via Grant FL150100090.

APPENDIX

Software

The interpolation method detailed in this paper is included as functions in the Gibbs SeaWater (GSW) oceanographic toolbox of the Thermodynamic Equation Of Seawater—2010 (TEOS-10) ([IOC et al. 2010](#)). The toolbox is a collection of programs to compute the TEOS-10 properties and is available from www.TEOS-10.org (see [McDougall and Barker 2011](#)). The MR-PCHIP interpolation algorithm of this paper is available as the functions `gsw_tracer_interp` and `gsw_t_interp`; these programs are designed for use with in situ observations where salinity has not been recorded. The MRST-PCHIP is

available in the programs `gsw_SA_CT_interp` and `gsw_tracer_CT_interp`. Like the basic PCHIP algorithm our algorithms require the profile to consist of at least four bottles. If three bottles are provided, the software uses MATLAB's noncentered, shape-preserving three-point formulas, and when just two bottles are provided linear interpolation is applied.

In the MR-PCHIP method we have set the multiplying constant that relates $\hat{R}$ to R to be 0.33; however, this value can be changed.

REFERENCES

- Barker, P. M., and T. J. McDougall, 2017: Stabilizing hydrographic profiles with minimal change to the water masses. *J. Atmos. Oceanic Technol.*, **34**, 1935–1945, <https://doi.org/10.1175/JTECH-D-16-0111.1>.
- Fritsch, F. N., and R. E. Carlson, 1980: Monotone piecewise cubic interpolation. *SIAM J. Numer. Anal.*, **17**, 238–246, <https://doi.org/10.1137/0717021>.
- IOC, SCOR, and IAPSO, 2010: The International Thermodynamic Equation of Seawater—2010: Calculation and use of thermodynamic properties. Intergovernmental Oceanographic Commission Manuals and Guides 56, 220 pp., http://www.teos-10.org/pubs/TEOS-10_Manual.pdf.
- Johnson, D. R., T. P. Boyer, H. E. Garcia, R. A. Locarnini, O. K. Baranova, and M. M. Zweng, 2011: World Ocean Database 2009 documentation. NODC Internal Rep. 20, 175 pp., http://www.nodc.noaa.gov/OC5/WOD09/pr_wod09.html.
- Locarnini, R. A., and Coauthors, 2013: *Temperature*. Vol. 1, *World Ocean Atlas 2013*, NOAA Atlas NESDIS 73, 40 pp., http://data.nodc.noaa.gov/woa/WOA13/DOC/woa13_vol1.pdf.
- McDougall, T. J., and P. M. Barker, 2011: Getting started with TEOS-10 and the Gibbs SeaWater (GSW) oceanographic toolbox. SCOR/IAPSO Working Group 127 Rep., 28 pp., http://www.teos-10.org/pubs/Getting_Started.pdf.
- Reiniger, R., and C. Ross, 1968: A method of interpolation with application to oceanographic data. *Deep-Sea Res. Oceanogr. Abstr.*, **15**, 185–193, [https://doi.org/10.1016/0011-7471\(68\)90040-5](https://doi.org/10.1016/0011-7471(68)90040-5).
- Siegelman, L., F. Roquet, V. Mensah, P. Rivière, E. Pauthenet, B. Picard, and C. Guinet, 2019: Correction and accuracy of high- and low-resolution CTD data from animal-borne instruments. *J. Atmos. Oceanic Technol.*, **36**, 745–760, <https://doi.org/10.1175/JTECH-D-18-0170.1>.
- Zweng, M. M., and Coauthors, 2013: *Salinity*. Vol. 2, *World Ocean Atlas 2013*, NOAA Atlas NESDIS 74, 39 pp., http://data.nodc.noaa.gov/woa/WOA13/DOC/woa13_vol2.pdf.