

MANUAL OF MATLAB CODE

Nonlinear Inversion of Thorium and Particle Data

Paul Lerner

plerner@whoi.edu

MIT-WHOI Joint Program

Woods Hole Oceanographic Institution

Disclaimer	3
Introduction.....	4
User Inputs	5
1. Number of Thorium Isotopes.....	5
2. Boundary Values.....	6
3. Data Errors.....	6
4. Prior Estimate of the State Vector	7
5. Particle availability.....	7
6. Data input.....	8
FMINCON iterative display.....	8
Example and Graphical Output	9
1. Numerical Model	9
2. Input with idealized dataset	10
3. Graphical Output	10
Acknowledgement	12
References	12

Disclaimer

While the present code has been through repeated tests, there is no guarantee that it is error free. Thus, the user is encouraged to notify the author of any obscurities or errors in the code or in this manual. In addition, the current iteration of this manual does not provide a complete description of the various subroutines used in this algorithm. The purpose of this manual is to ensure that the user is able to run the routine from the command line, with no exploration or editing of the subroutines required. Finally, this code should be flexible enough to use any dataset that is a combination of ^{228,230,234}Th in the dissolved and particulate phases, the associated parent radionuclides for the isotopes of interest, and particle concentration as input. However, the current state of the algorithm cannot take as input any other trace elements or isotopes, even if the dynamic equations for such elements or isotopes are similar (i.e. they are described by 1-d vertical single particle class reversible exchange model similar to the one described here).

Introduction

This document details how to fit a simple, one-dimensional (vertical) model of thorium and particle cycling in the oceanic water column to radionuclide data, including $^{228,230,234}\text{Th}$, $^{234,238}\text{U}$, ^{228}Ra , and particle concentration data. The radionuclide data are in units of dpm m^{-3} , while the particle concentration data are in units of mg m^{-3} . The algorithm used to this end is a nonlinear programming technique as implemented in Matlab's FMINCON, a constrained optimization routine used to solve nonlinear problems (see Matlab's user manual, reference herein). For details on the algorithm, the reader is invited to consult the papers of Byrd et al. (2000) and Waltz et al. (2006). The model itself is described in detail in Lerner et al. (2016). For convenience, however, a brief description of the model follows.

The model considers balance equations for dissolved and particulate thorium activities, as well as a balance equation for the concentration of particles in a single particle class. The balance equations account for the production of $^{228,230,234}\text{Th}$ from their radioactive parents, the radioactive decay of $^{228,230,234}\text{Th}$, the adsorption of thorium onto particles, the desorption of thorium from particles, particle degradation, and particle sinking. The model equations are

$$0 = A_\pi - (k_1 + \lambda)A_d + (k_{-1} + \beta_{-1})A_p \quad (1)$$

$$0 = k_1A_d - (k_{-1} + \beta_{-1} + \lambda)A_p - w \frac{\partial A_p}{\partial z} \quad (2)$$

$$0 = \beta_{-1}P - w \frac{\partial P}{\partial z} \quad (3)$$

Here, A_π is the production from the radioactive parent, A_d (A_p) is the activity of ^{228}Th , ^{230}Th , or ^{234}Th in dissolved (particulate) form, P is particle concentration, λ is the radioactive decay constant for ^{228}Th , ^{230}Th , or ^{234}Th , k_1 is the rate constant for Th adsorption onto particles, k_{-1} is the rate constant for Th desorption from particles, β_{-1} is the rate constant for particle degradation, w is the particle sinking speed, and z is depth in the water column. The model assumes steady state, as well as negligible contributions from water transport by advection and diffusion.

The model (1-3) is a set of first-order ordinary differential equations that requires two boundary conditions (one for (2) and another for (3)) for its solution. In the model, the boundary conditions are imposed at the upper boundary of the domain, which is set by the user (see section 3.1). The partial derivatives in (2-3) are approximated with finite differences on an irregular grid using a backward difference scheme. The grid levels

coincide with sampling depths and are thus generally different for different oceanographic stations.

The algorithm used to fit the model (1-3) to the data places the radionuclide activities, particle concentration, and the rate constants, all at different depths of the model grid, in a state vector \mathbf{x} . FMINCON seeks a vector \mathbf{x} which both perfectly satisfies the model (that is, the rhs of equations (1-3) equal 0 when a solution is found), and fits radionuclide and particle concentration data, as well as prior estimates of rate parameters, in a least squares sense, with due consideration for the errors in the data and in the estimates (for further details, see eq. 5 in Lerner et al., 2017).

The documentation of this nonlinear inverse procedure follows. First, the user-defined inputs to the algorithm are described. Next, a brief description of the output of FMINCON is provided. Finally, the graphical output of the algorithm is given along with sample figures generated from an idealized radionuclide and particle dataset.

User Inputs

1. Number of Thorium Isotopes

To initiate the inversion procedure, run “mainThP.m.” When running this routine, you will be asked to load the data. The data should contain some combination of $^{228,230,234}\text{Th}$ activity data, with accompanying parent isotope and particle concentration data. While any number of isotopes may be chosen for use, each isotope must have both dissolved and particulate data, and particle concentration data must be provided. The radionuclide and particle concentration data should be contained within a single “.mat” file.

The user will be prompted to input the number of thorium isotopes he (she) would like to include. The input should be one number between 1-7 included. The numbers are conventional and correspond to the following combinations of Th isotopes: 1 to ^{234}Th , 2 to ^{230}Th , 3 to ^{228}Th , 4 to $^{230,234}\text{Th}$, 5 to $^{228,234}\text{Th}$, 6 to $^{228,230}\text{Th}$, and 7 to $^{228,230,234}\text{Th}$ (Fig 1). The routine will then proceed with the combination of Th isotopes defined by the user.

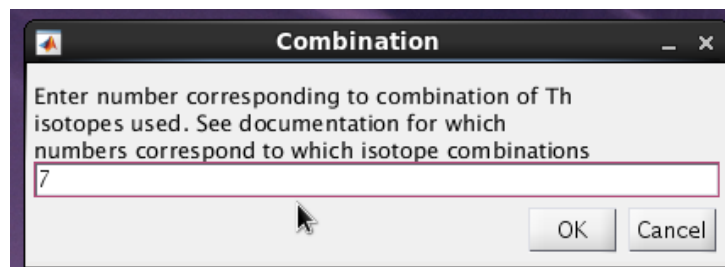


Fig. 1: Prompt for the combination of thorium isotopes to be used. For example, entry 7 corresponds to a study including $^{228,230,234}\text{Th}$, their associated parents ($^{234,238}\text{U}$ and ^{228}Ra), and the particle concentration.

2. Boundary Values

The user will be asked whether, at the depth of the boundary value of his (her) particulate data (particulate $^{228,230,234}\text{Th}$ and P), there are measurements of the dissolved data (dissolved $^{228,230,234}\text{Th}$, ^{228}Ra , and $^{234,238}\text{U}$). If this is indeed the case, then each particulate variable should contain one value more than each dissolved variable (that is, the particulate arrays should be one element longer than the dissolved arrays). An example of such a situation is given in section 5.2. If the dissolved data do not contain measurements at the depth of the domain boundary, type 0 in the prompt. Otherwise, type 1. Note that if the number of dissolved data equals the number of particulate data, the algorithm will remove the shallowest dissolved radionuclide values, and treat the particulate values at this depth as the boundary values.

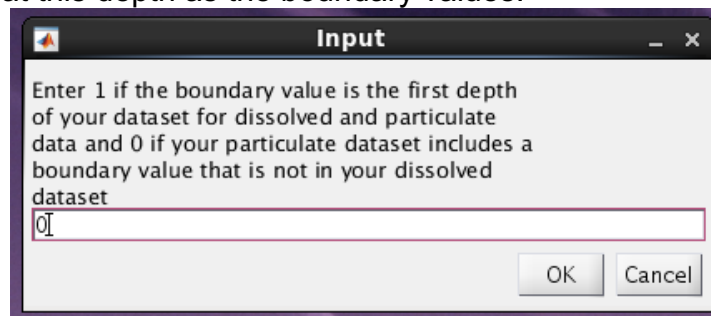


Fig 2: Prompt for setting the boundary value of the particulate data. Entry 0 indicates that the shallowest depth of your particulate data set (i.e., the depth where the boundary condition is imposed) does not include data of dissolved Th activities. Entry 1 indicates that the shallowest depth of your particulate data set includes data of dissolved Th activities.

3. Data Errors

The statistics of the errors in a multivariate data set (e.g., ^{228}Th , ^{230}Th , etc.) are typically included in an error covariance matrix. Error covariance matrices are $N \times N$ matrices, where N is the number of data. These square symmetric matrices include the error variances (squared errors or uncertainties) along the diagonal and the error covariances outside the diagonal. If an error covariance matrix is available to you, type 0 in the prompt. If the errors in your data are available in the form of 1-d arrays with one measure of uncertainty for each type of data but with no error covariances, type 1. No other option is provided, with the understanding that data errors could (in fact, should!) always be estimated.

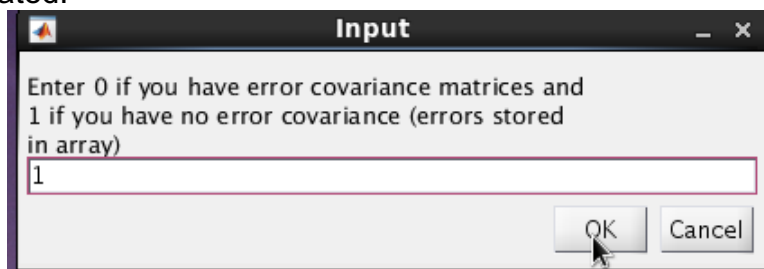


Fig 3: Prompt for setting the type of errors to the user is supplying. Entry 1 indicates that the errors are in 1-d arrays and that there is no error covariance in the measurements.

4. Prior Estimate of the State Vector

The state vector \mathbf{x} is the vector of radionuclide activities, particle concentrations, and rate parameter values, which is sought by the inverse procedure. FMINCON relies on a prior estimate of this vector, \mathbf{x}_{k0} . Here, the subscript “k0” denotes that this is the value of \mathbf{x} before the first iteration of FMINCON. Type 0 in the prompt if you want this prior estimate of the state vector to equal the data provided. Otherwise, type 1, in which case you will be asked to provide prior estimates of each radionuclide activity (dissolved and particulate) and particle concentration, at all depths of the model grid. These prior estimates should be within the .mat file you initially loaded.

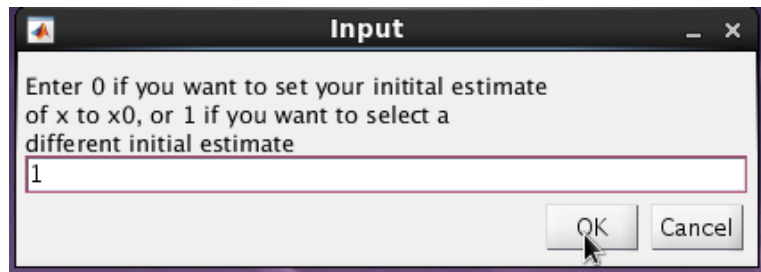


Fig 4: Prompt for setting the prior values of the elements in the state vector. Entry 1 indicates that a prior estimate different from the data, will be used.

5. Particle availability

The routine can proceed in the absence of particle concentration data. In that case, the model does not consider equation (3) in section 1, and only the balance equations for dissolved and particulate thorium (eq. 1-2, section 1) are satisfied. This option should be used with caution, because equation (3) provide an additional constraint on $\beta_{.1}$. In the limit where the uncertainties in the radionuclide and particle concentration are small, this allows for unique solutions for $\beta_{.1}$ and $k_{.1}$. Without equation (3), only the sum ($\beta_{.1}+k_{.1}$) can be resolved, and not the individual values of these rate parameters. Therefore, particle concentration data is necessary if individual estimates of particle degradation and thorium desorption are sought.

To run the routine without particle concentration data, type “0” in the prompt below. Otherwise, type “1.”

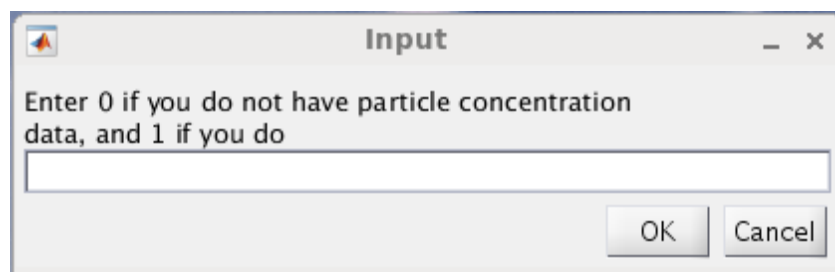


Fig 5: Prompt for the inclusion of particle concentration data. Entry 1 indicates that particle concentration data is available, while entry 0 indicates this data is unavailable and will not be considered in the model.

6. Data input

The user will be asked to type a variable name for (i) the activity of each radionuclide ($^{228,230,234}\text{Th}$, ^{228}Ra , $^{234,238}\text{U}$), (ii) their associated error variances or error covariances, (iii) the concentration of particles, and (iv) their errors or error covariances, which are to be considered in the analysis. For example, the user will be asked to “Please input dissolved Th234 for x0,” at which point the user will need to type the name of the variable correspond to the dissolved ^{234}Th data. These data should be in the form of Nx1 arrays (1 array for each variable, which are Th isotope activity in the dissolved phase, for each Th isotope activity in the particulate phase, for each parent isotope, and particle concentration), where N is the number of datapoints for each variable, and the errors should be either in the form of arrays if the errors for a given isotope or particle concentration at different depths do not covary, or error covariance matrices if they do covary. The user will also be asked to input the variable name associated with the depths of the measurements of (i) dissolved and particulate activities, and (ii) particle concentration.

If the prior state vector, \mathbf{x}_{k0} , is not set equal to data, then the user will also be asked to type a variable name for the prior estimate of each radionuclide activity (dissolved and particulate) and particle concentration. For example, the user will be asked “Please input dissolved Th234 for xk0,” at which point the user will need to type the name of the variable (Nx1 array, where N is the number of data points) corresponding to the dissolved ^{234}Th data prior values.

FMINCON iterative display

The code is constructed so that when Matlab’s optimization algorithm (FMINCON) is initiated, the following values should appear on screen.

Iter	F-count	f(x)	Feasibility	First-order optimality	Norm of step
0	565	3.674389e+02	1.961e+00	3.134e+02	
1	1135	3.161486e+02	8.494e-01	5.360e+02	1.910e+00
2	1703	2.784580e+02	8.476e-01	6.781e+02	3.435e-01
3	2270	2.572123e+02	8.469e-01	7.120e+02	2.934e-01
4	2837	2.389758e+02	8.414e-01	3.064e+02	1.169e+00
5	3403	2.313158e+02	8.387e-01	3.665e+02	8.388e-01
6	3970	2.187779e+02	8.355e-01	4.827e+02	8.269e-01
7	4537	1.951501e+02	8.321e-01	6.817e+02	1.090e+00
8	5104	1.698125e+02	8.305e-01	8.411e+02	7.691e-01
9	5672	1.579525e+02	8.269e-01	9.250e+02	5.057e-01

Fig 6: Example output of FMINCON on command window.

These values will continue appearing until a solution \mathbf{x} is found. The left column, *Iter*, is the iteration number of the routine. *F-count* is the cumulative number of times the

function being optimized (the cost function) has been evaluated. $f(\mathbf{x})$ is the value of the cost function. *Feasibility* describes the maximum constraint violation, the constraint here being that equations (1-3) should equal 0. It is thus equal to the maximum of the right-hand sides of equations (1-3) at each iteration, and the feasibility constraint is satisfied when this value is very close to 0 (in this routine, the constraint is assumed to be satisfied if this value is less than 10^{-4}). *First-order optimality* measures how close \mathbf{x} is to optimizing the function $f(\mathbf{x})$. The first optimality measure should be 0 if \mathbf{x} reaches a stationary point. The optimality tolerance, at which point a stationary point is assumed to have been reached, is set internally to 10^{-3} . Finally, *Norm of step* is the Euclidean norm of the change in \mathbf{x} between subsequent iterations. For further details about FMINCON output, the user is encouraged to examine Matlab's documentation:

<https://www.mathworks.com/help/optim/ug/fmincon.html>

Example and Graphical Output

1. Numerical Model

An idealized $^{228,230,234}\text{Th}$, $^{234,238}\text{U}$, ^{228}Ra , and particle concentration dataset is used to illustrate some of the results that may be generated from the inverse procedure. In order to derive an idealized dataset, the model equations (1-3) are solved, and the $^{228,230,234}\text{Th}$ activities and particle concentrations obtained from the solution are contaminated with noise. This model solution is obtained as follows. The rate parameters k_1 , k_{-1} , β_{-1} , and w are set to, respectively, 0.7 yr^{-1} , 3 yr^{-1} , 1 yr^{-1} , and 600 m yr^{-1} . The $^{234,238}\text{U}$ activities are set to constant values along the water column, with $^{238}\text{U} = 2400 \text{ dpm m}^{-3}$ and $^{234}\text{U} = 2750 \text{ dpm m}^{-3}$. In contrast, ^{228}Ra is set as depth-dependent, using a linear combination of two exponential functions (Marchal and Lam 2012),

$$^{228}\text{Ra} = Ae^{-(z-ze)/l} + Be^{-(zb-z)/l} \quad (4)$$

Here, $A = 30 \text{ dpm m}^{-3}$ and $B = 5 \text{ dpm m}^{-3}$ are ^{228}Ra activity values, $ze = 100 \text{ m}$ and $zb = 4010 \text{ m}$ are the upper and lower boundaries of the model domain, and $l = 500 \text{ m}$ is a vertical scale. The model grid is regular with vertical spacing between grid points equal to 100 m , so the number of depths is 41. Upper boundary values are also needed for $^{228,230,234}\text{Th}_p$ and P , and are set to 0.8 dpm m^{-3} , $0.0005 \text{ dpm m}^{-3}$, 600 dpm m^{-3} , and 15 mg m^{-3} , respectively. The numerical model is supplied in "ThP_numerical_model.m" and the generated dataset is supplied in "Thmodel.mat."

The script "ThP_numerical_model.m" generates idealized radionuclide and particle data contaminated with noise. The noise is sampled from a normal distribution with a mean of 0 and a standard deviation that is 10% of the radionuclide activities or particle concentrations, except for $^{234,238}\text{U}$. The activities of these radioactive parents are

assumed to have smaller uncertainties and hence the standard deviation of the noise for $^{234,238}\text{U}$ is set at 5% of the $^{234,238}\text{U}$ activities.

2. Input with idealized dataset

As described in Section 3, when running “mainThp.m,” the user will encounter 4 input prompts. On the first prompt, the user simply types the number corresponding to the combination of thorium isotopes which is desired in the analysis. For example, the user types 7 if he (she) wishes to use all three thorium isotopes. On the second prompt, type 0, since the idealized dataset includes the boundary values of particulate $^{228,230,234}\text{Th}$ and the boundary value of particle concentration (i.e., at $z = 100$ m), but does not have values of dissolved $^{228,230,234}\text{Th}$, ^{228}Ra , $^{234,238}\text{U}$ at the depth of the boundary values. On the third prompt, type 1, since we assume that no error covariances are available. On the fourth prompt, type 1, since the prior estimate of the state vector, \mathbf{x}_0 , will contain the idealized dataset.

The data generated from the numerical model, as well as their uncertainties, to be used in the inversion routine are contained within the following arrays (note that the model used to generate the idealized data assumes no error covariance):

Dissolved ^{234}Th : “Th234dmodel”, “Th234dmodelerr”
 Particulate ^{234}Th : “Th234pmodel”, “Th234pmodelerr”
 ^{238}U : “U238model”, “U238modelerr”

Dissolved ^{230}Th : “Th230dmodel”, “Th230dmodelerr”
 Particulate ^{230}Th : “Th230pmodel”, “Th230pmodelerr”
 ^{234}U : “U234model”, “U234modelerr”

Dissolved ^{228}Th : “Th228dmodel”, “Th228dmodelerr”
 Particulate ^{228}Th : “Th228pmodel”, “Th228pmodelerr”
 ^{228}Ra : “Ra228model”, “Ra228modelerr”

Type the corresponding variable when asked to input a given radionuclide or particle concentration dataset for \mathbf{x}_0 . The depth array for the idealized dataset is named “z”. This array includes the boundary value depth, so the depths carrying the dissolved Th isotopes are in “z(2:end)”.

When asked to type the radionuclide and particle names of \mathbf{x}_{k0} , the variables are arrays named in the same manner as the data used for \mathbf{x}_0 , but appended with “rnd” (e.g., instead of typing “U238model”, type “U238modelrnd”).

3. Graphical Output

Figure 6 below illustrates some of the output of the computer code (Fig 6). Figure 6a is the profile of dissolved ^{230}Th , with ± 1 standard deviation, from both the idealized

dataset in \mathbf{x}_0 (black) and the solution \mathbf{x} (blue). Figure 6b is the profile of k_1 , as assigned in \mathbf{x}_0 (uniform value of 0.5 yr^{-1} ; black circles), and estimated from the inversion. Horizontal error bars are ± 1 standard deviation. Note that \mathbf{x}_0 does not include the value of $k_1 = 0.7 \text{ yr}^{-1}$ that has been used to generate the idealized dataset. This decision allows one to check that, if given such a dataset, the algorithm can converge upon a solution containing k_1 used to generate the data.




Fig 7: (a) Dissolved ^{230}Th from the numerical model (black) and the inversion (blue). (b) Prior (black) and posterior (blue) values of k_1 . Error bars are ± 1 standard deviation.

As a note of caution, the algorithm is not guaranteed to perfectly to recover the rate parameters used to generate the idealized data in the numerical model. This is because both the data and the rate parameters are treated as unknowns in the routine. How well the model will fit the data vs. the prior estimates of the unknowns (which are not the same as the rate parameters used to generate the data) depends on the error covariance matrix, which contains errors for both the data and the prior estimates of the rate parameters. Data or rate parameters with large errors are not weighted as strongly as those with small errors. In the numerical model, the relative errors of the data is set to 20%, while the smallest relative error in the prior estimates of the rate parameters is 50% (for the relative error of w). One way to test whether the algorithm can exactly recover the rate parameters used in the numerical model is to reduce the errors in the idealized dataset so that they are negligible compared to the errors in the rate parameters. The user may do this by running the routine “mainThpsmallerror.m”, which performs the same operations as in “mainThp.m” but also reduces the errors in the data by 3 orders of magnitude.

Additionally, if the user decides to only use one isotope from the numerical model in the inversion, the inversion will not return the values of k_1 , k_{-1} , β_{-1} , and w used to generate the data. This is because the forward model uses only one value for each rate parameter to generate the radionuclide and particle data, while the routine attempts to estimate a value for each parameter at each depth. Therefore, even assuming the relative errors in the data are very small compared to the relative errors in the rate parameters, using only 1 isotope results in an underdetermined system with 120 equations (40 depths x (dissolved isotope equation + particulate isotope equation + particle equation)), and 160 unknowns (40 depth x 4 rate parameters).

The computer code also calculates the normalized residuals, i.e., the differences between the posterior values obtained by inversion and the prior values obtained from the data, normalized by the errors in the data. Normalized residuals are a measure of the ability of the model to reproduce the data, given the data errors. The normalized residuals associated with the fit of the model to the idealized dataset are shown in Figure 7a. All residuals are in the range $[-2, +2]$, showing that the posterior estimates differ from the prior estimates by less than 2 standard deviations.

The computer code also calculates normalized equation residuals, which are the rhs of equations (1-3) calculated from the solution, normalized by the largest term in each equation. Normalized equation residuals are a measure of the extent to which the model equations are satisfied. For the fit of the model to the idealized dataset, normalized equation residuals are all smaller than 10^{-13} in absolute magnitude (Fig. 7b), indicating that the model equations are satisfied to very high order.




Fig 8: (a) Normalized Residuals. The x-axis corresponds to the data point number, with 404 datapoints in total. The red horizontal lines cross the y-axis at ± 2 . (b) There are 280 equation residuals, one for each equation (1-3) at each depth.

Acknowledgement

This work has been supported by the US National Science Foundation (OCE-1232578).

References

1. Matlab (2016), Optimization Toolbox (r2016a), *User's guide*, The Mathworks, Inc., [http:// www.mathworks.com/help/optim/](http://www.mathworks.com/help/optim/), accessed: 22 July, 2016.
2. Waltz, R. A., Morales, J. L., Nocedal, J., and Orban, D. (2006). An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical programming*, 107(3), 391-408.
3. Byrd, R. H., Gilbert, J. C., & Nocedal, J. (2000). A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, 89(1), 149-185.
4. Lerner, P., Marchal, O., Lam, P. J., Anderson, R. F., Buesseler, K., Charette, M. A., Edwards, R. L., Hayes, C. T., Huang, K.F., Lu, Y, Robinson, L. F., and Solow, A. (2016). Testing models of thorium and particle cycling in the ocean using data from station GT11-22 of the US GEOTRACES North Atlantic section. *Deep Sea Research Part I: Oceanographic Research Papers*, 113, 57-79.
5. Lerner, P., Marchal, O., Lam, P. J., Buesseler, K., and Charette, M. (2017). Kinetics of thorium and particle cycling along the US GEOTRACES North Atlantic Transect. *Deep Sea Research Part I: Oceanographic Research Papers*, 125, 106-128.
5. Marchal, O., and Lam, P. J. (2012). What can paired measurements of Th isotope activity and particle concentration tell us about particle cycling in the ocean?. *Geochimica et Cosmochimica Acta*, 90, 126-148.