

THERMOHALINE STAIRCASES IN THE ARCTIC OCEAN: DETECTION, EVOLUTION,
AND INTERACTION

by

Mikhail Schee

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy

Department of Physics
University of Toronto

© Copyright 2024 by Mikhail Schee

Mikhail Schee
Doctor of Philosophy
Department of Physics
University of Toronto
2024

Abstract

Thermohaline staircases consist of a series of horizontal, well-mixed layers, each on the order of a meter thick, separated by thin interfaces, across which temperature and salinity make abrupt jumps. While they have been consistently observed several hundred meters below the surface of the Arctic Ocean for over fifty years, little is known about their long-term evolution. Such stratification structures affect the propagation of internal waves and, because of an effect called internal gravity wave tunnelling, interactions between internal waves and staircases can be complex. This thesis presents a novel method of detecting thermohaline staircase layers in observations, analyzes their evolution on a decadal scale, and examines their interactions with internal waves.

Inspired by the patterns made by observations of thermohaline staircases in temperature-salinity space, I develop a novel detection method. Using the Hierarchical Density-Based Spatial Clustering of Applications with Noise algorithm, I find I can detect and connect staircase layers across datasets of hydrographic profiles from the Canada Basin in the Arctic Ocean. This offers an advantage over previous detection methods which treat each profile individually as, here, the sprawling horizontal nature of the layers can be analyzed.

Using this clustering method, I identify layers in the Beaufort Gyre Region which span over 1000 km horizontally and persist for nearly two decades. In addition to reproducing many results from previous studies, I find the layers to be evolving in time. The layers are sinking at approximately the same rate as the overall downwelling in the region. I also find that layers in the upper staircase are warming while layers near the bottom are cooling.

I develop a set of numerical experiments to examine the interactions between internal waves and idealized staircase stratification structures. For structures with one layer, I find the transmission of waves decreases monotonically as the layer thickness gets larger relative to the wavelength. With multiple layers present, I find peaks in transmission for particular ratios of thickness to wavelength, the patterns of which become more complex as more layers are added. I also reproduce the results of a laboratory experiment, finding the same pattern of reflection and transmission of waves.

Acknowledgements

I would like to acknowledge my supervisor, Prof. Nicolas Grisouard for invaluable help and guidance throughout my graduate studies. This research was supported in part by the National Science Foundation under Grant No. NSFPHY-1748958 through my participation in the Layering in Atmospheres, Oceans and Plasmas program in 2021, hosted by the Kavli Institute for Theoretical Physics at the University of California Santa Barbara. I wish to acknowledge fruitful discussions with Prof. Maike Sonnewald about machine learning clustering algorithms, with Dr. Carine van der Boog about her thermohaline staircase detection algorithm, and with Prof. Varvara E. Zemsanova about heat flux and temperature change in the ocean. I would like to acknowledge Prof. Erica Rosenblum and Dr. Jonathan M. Lilly for their invaluable insight on the projects in Chapters 2 and 3 on oceanographic processes, data collection, statistics, and colormaps for plots. I would like to acknowledge Prof. Peacock et al. at ENDLab, MIT with collaborators at ENS de Lyon for enlightening discussions about the interactions of internal waves with stratification structures and the chance to see in person the wave tank which I recreate in my numerical experiments in Chapter 4. Prof. Nicolas Grisouard and I were supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) (funding reference numbers RGPIN-2015-03684 and RGPIN-2022-04560). Dr. Jonathan M. Lilly was supported by grant number 2049521 from the Physical Oceanography program of the United States National Science Foundation. Prof. Erica Rosenblum was supported by the NSERC Postdoctoral Fellowship award, the NSERC Canada-150 Chair (Award G00321321), and the NSERC Discovery Grant (funding reference number RGPIN-2024-05545).

To my family, I appreciate your support and encouragement.

Contents

1	Introduction	1
1.1	Background	1
1.2	Ice-Tethered Profiler Data	8
1.3	The HDBSCAN clustering algorithm	9
1.4	Overview	11
2	Unsupervised clustering identifies thermohaline staircases in the Canada Basin of the Arctic Ocean	13
2.1	Introduction	14
2.2	Data	16
2.3	Methods	17
2.3.1	Applying the HDBSCAN clustering algorithm	17
2.3.2	Selecting values of input parameters	18
2.3.3	Evaluating the clustering algorithm results	20
2.4	Results	22
2.4.1	Properties of detected layers	22
2.4.2	Outliers and splitting / merging layers	23
2.4.3	Remnant intrusions revealed by temperature	25
2.4.4	Dependence of R_L on pressure	27
2.5	Discussion	28
3	Tracking Beaufort Gyre staircase layers across 17 years	31
3.1	Introduction	31
3.2	Data	33
3.3	Methods	34
3.3.1	Defining the bounds of the water column	34
3.3.2	Correcting for spatial distributions	34
3.3.3	Clustering each time period	37
3.3.4	Connecting across time periods	39
3.3.5	Defining outlier clusters	40
3.4	Results	40
3.4.1	Properties of detected layers	41
3.4.2	Values of IR_{S_A} and R_L	42
3.4.3	Spatial variability	43

3.4.4	Trends in time	44
3.4.5	Layer thicknesses	46
3.4.6	Net heat flux	48
3.5	Discussion	49
4	Modeling internal wave propagation through stratified fluids	54
4.1	Introduction	54
4.1.1	Previous Work	55
4.1.2	Overview	56
4.2	Theory	57
4.2.1	The Boussinesq Equations of Motion	57
4.2.2	Dispersion Relation	58
4.2.3	Complex Demodulation	60
4.2.4	Calculating Vertical Energy Flux	61
4.3	Dedalus	63
4.4	Numerical Experiments in One Spatial Dimension	63
4.4.1	Experimental Setup	63
4.4.2	Methods	67
4.4.3	Results	70
4.5	Numerical Experiments in Two Spatial Dimensions	71
4.5.1	Experimental setup	72
4.5.2	Methods	77
4.5.3	Results	82
4.6	Discussion	85
5	Conclusions	88
5.1	Summary of results	88
5.2	Future work	91
A	Supplementary Materials for: Unsupervised clustering identifies thermohaline staircases in the Canada Basin of the Arctic Ocean	94
A.1	Practical salinity S_P vs. absolute salinity S_A	94
A.2	Varying the value of ℓ	96
A.3	Optional parameters for HDBSCAN	98
A.4	Clusters from ITP3	99
A.5	Clusters across time for ITP2	99
A.6	Differences from T08	100
A.7	Differences from L22	102
A.8	Subsampling profiles	102
B	Supplementary Materials for: Tracking Beaufort Gyre staircase layers across 17 years	105

C	Supplementary equations and derivations	117
C.1	The Boussinesq Equations of Motion	117
C.1.1	Conservation of Momentum	117
C.1.2	Conservation of Mass	120
C.1.3	Thermodynamic Equation	120
C.1.4	The Boussinesq Equations of Motion	123
C.2	The Polarization Relation	124

List of Figures

1.1	A map of the Arctic region. Adapted from CIA World Factbook (Public Domain). Retrieved June 6, 2024 from https://commons.wikimedia.org/w/index.php?curid=531491 .	2
1.2	An example profile, number 1257 from Ice-Tethered Profiler 1 (see Section 1.2) in (a) conservative temperature, (b) absolute salinity, (c) potential density anomaly referenced to 1000 dbar, and (d) buoyancy frequency. This profile was taken on Monday, June 26th, 2006 at 136.1525°W, 77.2870°N.	5
1.3	The two double diffusive regimes. The pink arrows indicate the direction in which the temperature gradient acts to increase density and the black arrows indicate in which direction salinity acts to increase density. Arrows pointing downwards represent stabilizing gradients while upwards pointing arrows represent destabilizing gradients.	6
1.4	The global distribution of thermohaline staircases in the salt-fingering regime (red) and the diffusive-convection regime (blue), with data between 2001 and 2020. The shade of each dot indicates the number of steps detected in the profile, either from an Argo float or Ice-Tethered Profiler. Reproduced from van der Boog, Dijkstra, et al. (2021).	6
1.5	(a) A schematic of an Ice-Tethered Profiler, showing how the surface buoy on the ice is connected to the wire upon which the profiler moves up and down while taking measurements. Reproduced from Toole et al. (2011). (b) A map of ITP profiles taken between 2004–2023. The red box designates the Canada Basin as defined by Peralta-Ferriz & Woodgate (2015).	9
1.6	The potential temperature and salinity values between 200–300 m depth from profiles taken by ITP2 spanning around 200 km laterally with lines indicating constant potential density anomaly (referenced to the surface). The thin (thick) line represents the westernmost (easternmost) profile. Reproduced from Timmermans et al. (2008).	10
2.1	A map showing the locations of all profiles used from ITPs 2 and 3, showing the whole Arctic in (a) and a zoomed-in view in (b). The red box designates the Canada Basin as defined by Peralta-Ferriz & Woodgate (2015).	17

2.2	Results from the clustering algorithm with $m_{pts} = 170$ and $\ell = 25$ dbar run on 53042 data points in the salinity range 34.05–34.75 g/kg from all up-going ITP2 profiles. (a) The data in Θ – S_P space with dashed lines of constant potential density anomaly (kg m^{-3}) referenced to the surface. The red box bounds the clusters marked in panels (b) and (d). (b) Profiles 183, 185, and 187 from ITP2 in a limited pressure range to show detail. Each profile is offset in S_P for clarity. (c) The spatial arrangement used as input for the algorithm where the gray points are noise and each color-marker combination indicates a cluster. The same color-marker combinations are used in each panel and the markers in panels (c) and (d) are at the cluster average for each axis. (d) A subset of the data in $\alpha\Theta$ – βS_P space with the linear regression line and inverse slope (R_L) noted for each individual cluster and with dashed lines of slope $\alpha\Theta/\beta S_P = 1$	19
2.3	A parameter sweep showing the number of clusters found (solid lines) and DBCV (dashed lines) in ITP2 as a function of (a) 27 different values of ℓ with $m_{pts} = 170$ and (b) 44 different values of m_{pts} with $\ell = 25$ dbar.	20
2.4	The average (a) pressure, (b) Θ , and (c) S_P for the points within each cluster for each profile (profile cluster average, PCA) across time. The clustering algorithm was run with $m_{pts} = 580$ and $\ell = 25$ dbar on 678575 data points in the salinity range 34.21–34.82 g/kg from all up-going ITP3 profiles.	22
2.5	The value of each cluster’s normalized inter-cluster range for salinity IR_{S_P} in (a) and (c) and the lateral density ratio R_L in (b) and (d) as a function of the cluster’s average pressure. The colors and markers for ITP2 in (a) and (b) are the same as the clustering shown in Figure 2.2 and for ITP3 in (c) and (d), they are the same as shown in Figure 2.4. Markers circled in red indicate outliers with a z-score greater than 2. In (b), the solid curve is a 2nd-degree polynomial fit (equation given by the annotation) for the non-outlier points from ITP2 and the dashed curve is the same for ITP3. In (d), the solid curve is a 2nd-degree polynomial fit (equation given by the annotation) for the non-outlier points from ITP3 and the dashed curve is the same for ITP2.	24
2.6	Individual S_P profiles from ITP2, specifically chosen to show the examples of outlier clusters in IR_{S_P} and R_L highlighted by the bands of color. (a) Profiles 67, 69, 73, 75, 81, 83, 91, 93, 97, and 99, collected between August 31–September 5, 2004. (b) Profiles 87, 89, 95, 97, 99, 101, 103, 105, 109, and 111, collected between September 3–7, 2004. The colors and markers are the same as the clustering shown in Figure 2.2. The gray dots are noise points and the black lines show the profiles. Each profile is offset in S_P for clarity.	25
2.7	The average Θ and S_P for the 34 layers found by Lu et al. (2022) in black dots and for the 43 layers found in our study using data from ITP3 and the same colors and markers as in Figure 2.4. Clusters circled in red are outliers in either IR_{S_P} or R_L	26

2.8	Individual profiles 313, 315, 317, 319, and 321 from ITP3, collected between November 10–12, 2005, specifically chosen to show the example of a temperature inversion highlighted by the bands of color. The colors and markers of the individual points are the same as the clustering shown in Figures 2.4 and 2.7. The gray dots are noise points. The black lines on the left of each pair are the S_P profiles while the red lines on the right are for Θ . Each profile is offset in both S_P and Θ for clarity.	27
3.1	(a) A map of the Arctic with a red bounding box designating the Beaufort Gyre Region (BGR) as defined by ST22. (b) A zoomed-in view of the BGR with each of the 15,202 profiles used in this study marked in a color indicating the date on which it was taken. The dotted-dashed lines indicate the bathymetry contours.	33
3.2	Maps showing the pressures at which the (a) LHW and (b) AW cores occur for each of the 15,247 profiles in the BGR. The color bars are oriented such that values physically higher in the water column are closer to the top of the plot.	35
3.3	The top row shows the LHW core and the bottom row shows the AW core. (a,e) Histograms of the pressures. (b,f) Maps of the pressure values of the core for each profile overlaying a heat map of the <code>polyfit2d</code> of the pressure. The red stars mark the location of the highest pressure value in the <code>polyfit2d</code> . (c,g) Maps of the residuals, pressure values minus the <code>polyfit2d</code> . (d,h) Histograms of the residuals. All vertical axes and color bars are oriented such that values physically higher in the water column are closer to the top of the plot.	36
3.4	Trends in pressure over time for the (a,b) LHW and (c,d) AW cores. Panels (a,c) show the trends before and panels (b,d) show the trends after subtracting their respective <code>polyfit2d</code> . The blue lines are the 30-day moving average. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. . .	37
3.5	Parameter sweeps in m_{pts} for (a) BGR1415 and (c) BGR1516. The results of clustering (b) BGR1415 and (d) BGR1516 with the chosen values of m_{pts} 1050 and 500, respectively, before relabeling the cluster ID's.	38
3.6	The salinity across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.	41
3.7	The value of each cluster's (a) normalized inter-cluster range for salinity IR_{S_A} and (b) lateral density ratio R_L as a function of the cluster's average pressure. The colors and numbers for each cluster are the same as in Figure 3.6. The circles mark clusters that are outliers in the respective variable. Note that the vertical domain is zoomed in to show detail and, while there are clusters with average pressures below 355 dbar, all of those clusters have been disregarded in further calculations.	42
3.8	Maps of (a) pressure, (b) salinity, and (c) temperature for Cluster 63 over the corresponding <code>polyfit2d</code> for each variable. The red star in (a) is at the location of the minimum of the <code>polyfit2d</code> : 74.22°N, 149.58°W. Histograms of (d) pressure, (e) salinity, and (f) temperature for Cluster 63, showing both the original and the residual values for each variable.	44

3.9	For all clusters, the trends in time of the <code>polyfit2d</code> residual for (a) pressure, (b) salinity, and (c) temperature with respect to the cluster average of salinity. A blue “+” denotes clusters that are outliers in IR_{SA} and a red “x” denote clusters that are outliers in R_L , as was indicated in Figure 3.7. All outliers were ignored when calculating the fit lines which have (from left to right) R^2 values of -0.328, 0.071, and -0.789. The blue and pink triangle markers denote the values for the LHW and AW, respectively.	45
3.10	The thicknesses (in dbar) of the detected layers. (a) The non-zero thickness values against salinity of each regular cluster for every profile in the BGR. While there are thicknesses of up to 20 dbar, axes bounds are limited here to show detail. For all clusters, the (b) cluster average of layer thickness and (c) trends of percent change in layer thickness over time with respect to the cluster average of salinity. The values in (c) are calculated by taking the trend in time for the cluster’s non-zero thicknesses then dividing by that cluster’s average thickness shown in (b), however are highly uncertain, see text for explanation. A blue “+” and a red “x” denote clusters that are outliers in IR_{SA} and R_L , respectively. All outliers were ignored when calculating the fit lines which have R^2 values of 0.609 for (b) and 0.041 for (c).	47
3.11	(a) The average net heat flux for each cluster against salinity. All outliers were ignored when calculating the fit line which has $R^2 = -0.677$. (b) The cumulative net heat flux of regular clusters, from the top to the bottom of the staircase.	49
4.1	(a) A typical density stratification profile in the Canadian Basin of the Arctic Ocean (Rainville & Winsor, 2008). (b) A magnified section of the profile showing the staircase structure. The red line is a 5 m bin-averaged smoothed density profile. (c) The corresponding profile for the buoyancy frequency $N(z)$. Reproduced from Figure 1 of Ghaemsaidi et al. (2016).	55
4.2	A schematic of a vertically oscillating cylinder experiment where (a) the oscillation frequency ω is less than the stratification frequency N_0 , making the “St. Andrew’s Cross” pattern with wave beams propagating at an angle θ to the vertical, (b) ω is less than but almost equal to N_0 , and (c) $\omega > N_0$ where the waves are evanescent. Reproduced from Figure 3.3 in Sutherland (2010).	59
4.3	(a) The input wavefield in equation (4.6). The upward (c) and downward (d) propagating components of the input wavefield, isolated using complex demodulation. (d) The output wavefield, which is the sum of (c) and (d), and is identical to the input.	61
4.4	A depiction of the domain of the experiment. The two black horizontal dashed lines show the top and bottom of the measurement extent. The orange and teal dashed lines show calculation domains for the amplitudes of the incoming and transmitted wave, respectively. (left) An example background $N(z)$ profile, showing a double-layer stratification structure in the purple solid line and the frequency ω used to force the waves in the blue dotted line. (right) The placement and shape of the windows for the boundary forcing $W_{bf}(z)$ (solid) and sponge layer $W_{sp}(z)$ (dotted).	65

4.5	The total wavefield of an example experiment where $k_x L = 1.55$ and $\theta = 45^\circ$. The solid red line represents the wave propagation speed c_{gz} . The two black horizontal dashed lines show the top and bottom of the measurement extent. The black vertical dashed line marks the time at which steady state is assumed. The orange and teal dashed lines show calculation domains for the amplitudes squared of the incoming and transmitted wave, respectively. After steady state was reached, $\mathbb{T} = 0.16$	68
4.6	The downward propagating wavefield in the measurement domain of an example experiment where $k_x L = 1.55$ and $\theta = 45^\circ$. The wave is incident on the layer from above and the lower amplitude of the transmitted wave can be seen. After steady state was reached, $\mathbb{T} = 0.16$. The orange and teal dashed lines show calculation domains for the amplitudes of the incoming and transmitted wave, respectively.	69
4.7	The transmission coefficient \mathbb{T} for the analytical solution of 1 mixed layer (4.26) and experiments of 1, 2, 3, 4, and 5 mixed layers. Each marker represents one experiment.	70
4.8	The transmission coefficient \mathbb{T} for the analytical solution of one mixed layer (4.26) and experiments of one and two mixed layers as calculated with Dedalus in this study or by the inviscid MATLAB code of G16. Each marker represents one experiment.	71
4.9	(Left) A schematic of the experimental setup used by G16 showing the dimensions of the tank and the measurement area as well as the positioning of the wave generator. (Right) Measurements of the density stratifications, $\Delta\rho = \rho - 1000$ for the (a) single- and (b) double-mixed-layer experiments. Reproduced from Figure 4 in G16.	73
4.10	(a) The profile of the sponge layer given by equation (4.30). The (b) single- and (c) double-mixed-layer stratification profiles in solid lines used to emulate those of G16 where the dotted lines represent the value of the forcing frequency ω used in each case. The black dashed lines represent the bottom of the measurement domain.	74
4.11	Total vertical energy flux for a constant stratification profile. (a) $\langle \mathcal{F}_z \rangle$ across the z domain and across time where the red dotted line shows the vertical group speed c_{gz} . (b) The running average of $\langle \mathcal{F}_z \rangle$ with a one oscillation period window through the top ($z = 0$) and bottom ($z = -0.5$) surfaces over time. This experiment had $\omega = 0.7071 \text{ s}^{-1}$, $\hat{A} = 2 \times 10^{-4}$, and was run for 20 oscillation periods.	80
4.12	The components of vertical energy flux from (4.43) through the top boundary for a constant stratification profile: (a) $\langle \mathcal{F}_a \rangle$, (b) $\langle \mathcal{F}_p \rangle$, and (c) $\langle \mathcal{F}_\nu \rangle$. Note the different vertical axis scales between subplots and that these are the running averages with a one oscillation period window. This experiment had $\omega = 0.7071 \text{ s}^{-1}$, $\hat{A} = 2 \times 10^{-4}$, and was run for 20 oscillation periods.	81
4.13	The wave tank results from G16 for the (a) single and (c) double mixed layer experiments, adapted from their Figure 5. The results of my numerical experiments for the (b) single- and (d) double-mixed-layer cases run with $\omega_1 = 0.67 \text{ s}^{-1}$ and $\omega_1 = 0.72 \text{ s}^{-1}$, respectively. The snapshot in (b) was taken at $t/T = 12.407$ and the snapshot in (d) was taken at $t/T = 12.476$, both with colorbars saturated at $\pm 1.3 \times 10^{-3} \text{ m/s}$	82

4.14	Total vertical energy flux for the single (a,c) and double (b,d) layer stratification profiles. (a,b) $\langle \mathcal{F}_z \rangle$ across the z domain and across time. (c,d) The running average of $\langle \mathcal{F}_z \rangle$ with a one oscillation period window through the top ($z = 0$) and bottom ($z = -0.5$) surfaces over time. These experiments had (a,c) $\omega = 0.67 \text{ s}^{-1}$ or (b,d) $\omega = 0.71 \text{ s}^{-1}$, $\hat{A} = 2.3 \times 10^{-4}$, and were run for 10 oscillation periods.	84
4.15	The components of vertical energy flux from (4.43) through the top boundary for the single (a,c,e) and double (b,d,f) layer stratification profiles: (a,b) $\langle \mathcal{F}_a \rangle$, (c,d) $\langle \mathcal{F}_p \rangle$, and (e,f) $\langle \mathcal{F}_\nu \rangle$. Note the different vertical axis scales between subplots and that these are the running averages with a one oscillation period window. These experiments had (a,c,e) $\omega = 0.67 \text{ s}^{-1}$ or (b,d,f) $\omega = 0.71 \text{ s}^{-1}$, $\hat{A} = 2.3 \times 10^{-4}$, and were run for 10 oscillation periods.	85
A.1	Results from the clustering algorithm with $m_{pts} = 170$ and $\ell = 25$ dbar run on 53042 data points in the S_P range 34.05–34.75 g/kg from all up-going ITP2 profiles. This is the same as Figure 2.2, but the clustering algorithm was run on absolute salinity S_A instead of practical salinity S_P . (a) The data in Θ – S_A space with dashed lines of constant potential density anomaly (kg m^{-3}) referenced to the surface. (b) Profiles 183, 185, and 187 from ITP2 in a limited pressure range to show detail. Each profile is offset in S_A for clarity. (c) The spatial arrangement used as input for the algorithm where the gray points are noise and each color-marker combination indicates a cluster. The same color-marker combinations are used in each panel and the markers in panels (c) and (d) are at the cluster average for each axis. (d) A subset of the data in $\alpha\Theta$ – βS_A space with the linear regression line and inverse slope (R_L) noted for each individual cluster and with dashed lines of slope $\alpha\Theta/\beta S_A = 1$	95
A.2	The value of each cluster’s normalized inter-cluster range for salinity IR_{S_A} in (a) and the lateral density ratio R_L in (b) as a function of the cluster’s average pressure for ITP2. This is the same as Figure 2.5(a,b), but the clustering algorithm was run on absolute salinity S_A instead of practical salinity S_P . The colors and markers are the same as the clustering shown in Figure A.1. Markers circled in red indicate outliers with a z-score greater than 2.	96
A.3	The original profile 185 from ITP2 and the smoothed version of that profile for 4 different values of ℓ : (a) 2.5 dbar, (b) 12.5 dbar, (c) 25 dbar, and (d) 37.5 dbar. The Θ axes on each panel are offset slightly to show both lines.	97
A.4	Clustered data from ITP2 in the salinity range 34.05–34.75 g/kg plotted in Θ' – S_P space for 4 different values of ℓ : (a) 2.5 dbar, (b) 12.5 dbar, (c) 25 dbar (the same as in Figure 2.2(c)), and (d) 37.5 dbar, where the gray points are noise and each color-marker combination indicates a cluster.	98

A.5	Results from the clustering algorithm with $m_{pts} = 580$ and $\ell = 25$ dbar run on 678575 data points in the S_P range 34.21–34.82 g/kg from all up-going ITP3 profiles. (a) The data in Θ – S_P space with dashed lines of constant potential density anomaly (kg m^{-3}) referenced to the surface. (b) Profiles 313, 315, and 317 from ITP3 in a limited pressure range to show detail. Each profile is offset in S_P for clarity. (c) The spatial arrangement used as input for the algorithm where the gray points are noise and each color-marker combination indicates a cluster. The same color-marker combinations are used in each panel and the markers in panels (c) and (d) are at the cluster average for each axis. (d) A subset of the data in $\alpha\Theta$ – βS_P space with the linear regression line and inverse slope (R_L) noted for each individual cluster and with dashed lines of slope $\alpha\Theta/\beta S_A = 1$	99
A.6	A parameter sweep showing the number of clusters found (solid lines) and DBCV (dashed lines) in ITP3 as a function of 72 different values of m_{pts} with $\ell = 25$ dbar.	100
A.7	The average (a) pressure, (b) Θ , and (c) S_P for the points within each cluster for each profile (profile cluster average, PCA) across time. The clustering algorithm was run with $m_{pts} = 170$ and $\ell = 25$ dbar on 53042 data points in the salinity range 34.05–34.75 g/kg from all up-going ITP2 profiles.	101
A.8	Examples of clusterings when subsampling the ITP3 data to every (a,d) second, (b,e) third, or (c,f) fourth point. The top row shows the resulting clusterings and the bottom row shows a set of five example profiles to show the effect of subsampling on individual profiles.	103
A.9	A section of the temperature and salinity profile taken at Camp Snowbird, Station 1, on 16 May 1975. The step structure can be observed, although distorted by an eddy. Reproduced from Figure 20 of Bauer et al. (1980).	104
A.10	Examples of clusterings when subsampling the ITP3 data to every (a,d) second, (b,e) third, or (c,f) fourth point. The top row shows the resulting clusterings and the bottom row shows a set of five example profiles to show the effect of subsampling on individual profiles.	104
B.1	The dates on which profiles were taken within the BGR by each instrument in the time period analyzed in this study. This shows only the up-going profiles which extend below 400 dbar. Note that, while the same colors are used here as in the clustering plots, the colors for each ITP are only to make them easier to distinguish and do not correspond to certain clusters with the same color.	107
B.2	Maps showing, for each of the 15,247 profiles in the BGR, the (a) salinity and (b) temperature of the LHW core as well as the (c) salinity and (d) temperature of the AW core. The color bars are oriented such that values physically higher in the water column are closer to the top of the plot. Similar plots for the pressure of the LHW and AW cores are in Figure 3.2.	108

B.3	The salinity across all time periods, colored by cluster, before relabeling to create Figure 3.6. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.	109
B.4	The pressure across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.	110
B.5	The temperature across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.	111
B.6	The density anomaly, referenced to 1000 m, across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.	112
B.7	Maps of temperature for clusters 6, 23, 35, 52, 63, 69, 79, 90, and 96 over the corresponding <code>polyfit2d</code> . The cluster number increases with depth and none of the clusters here are outliers in IR_{S_A} or R_L	113
B.8	Trends in time of pressure for Cluster 63 both (a) before ($R^2 = 0.638$) and (b) after correcting with the <code>polyfit2d</code> ($R^2 = 0.632$). The blue lines are the 30-day moving average. This includes data from the 66,395 points within Cluster 63.	113
B.9	For all clusters, the trends in time of the <code>polyfit2d</code> residual for (a) pressure, (b) salinity, and (c) temperature with respect to the cluster average of salinity. This is similar to Figure 3.9, but with each cluster marked with it's ID and color. Blue and red circles denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have (from left to right) R^2 values of -0.328, 0.071, and -0.789. The blue and pink triangle markers denote the values for the LHW and AW, respectively.	114

B.10	The thicknesses (in dbar) of the detected layers. (a) The non-zero thickness values against salinity of each cluster that is neither an outlier in IR_{S_A} nor R_L for every profile in the BGR, colored by cluster. For all clusters, the (a) cluster average of layer thickness and (b) trends of percent change in layer thickness over time with respect to the cluster average of salinity. This is similar to Figure 3.10, but with each cluster marked with it's ID and color. The values in (c) are calculated by taking the trend in time for the cluster's non-zero thicknesses and dividing by that cluster's average thickness shown in (b), however are highly uncertain, see text for explanation. Blue and red circles denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have R^2 values of 0.609 for (b) and -0.041 for (c).	114
B.11	(a) The non-zero thicknesses (in dbar) of the detected layers for just ITP13 against pressure for each cluster that is neither an outlier in IR_{S_A} nor R_L . The linear trend has an R^2 value of 0.693. The red box denotes the domain bounds of the inset. (b) Reproduced from Shibley & Timmermans (2019). The thickness of layers they detected for ITP13. The blue dots represent the average depth and thickness for the data binned to 0.5-m increments and the linear fit line through them has a slope of 0.12 ± 0.01 m/m.	115
B.12	The mean of the non-zero thicknesses (in dbar) per profile of the detected layers for all clusters that are neither outliers in IR_{S_A} nor R_L . (a) A map of these mean layer thicknesses per profile on top of the <code>polyfid2d</code> . (b) A histogram of the mean layer thicknesses per profile.	115
B.13	The cluster average of each component of heat flux against salinity. (a) The cluster average layer thickness. (b) The cluster average isobaric heat capacity A blue “+” and a red “×” denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have R^2 values of 0.609 for (a) and -0.107 for (b).	116

List of Tables

2.1	Details of the Ice-Tethered Profilers (ITPs) used in this study.	16
2.2	The values of parameters used to run the clustering algorithm over both datasets.	20
2.3	The median normalized inter-cluster ranges and differences between average values of adjacent clusters for ITP2 and ITP3, calculated after removing outliers with z-score > 2 in the respective variable.	23
3.1	The coefficients for the <code>polyfit2d</code> of pressure, following Equation (3.1).	36
3.2	For each time period, summaries of the data used and clustering parameters.	39
3.3	A summary of median values for each variable v . The $\hat{\cdot}$ symbol indicates the median across each regular cluster (neither an outlier in IR_{SA} nor R_L), Δv represents the distance between the average values in v for neighboring clusters, v_{span} is the distance between v_{max} and v_{min} for a cluster, and $\underline{\sigma}_v$ is the standard deviation of a cluster in v	42
3.4	Comparing values of the lateral density ratio R_L from this study to that of Chapter 2. The means and standard deviations of R_L along with the second degree polynomial fits of R_L in pressure p from Figures 3.7(b) and 2.5(b,d), as well as the coefficients of determination for those fits.	43
3.5	The properties of the histograms for Cluster 63 shown in Figure 3.8. The means of all <code>polyfit2d</code> versions were zero.	44
3.6	A summary of median trends for each variable v across all regular clusters, before and after taking the residuals with <code>polyfit2d</code>	45
4.1	Values of parameters for the constant N , single layer, and double layer experiments, as well as the magnitude of the coefficients of vertical energy flux per unit mass $\hat{\mathcal{F}}$ predicted by (4.44) and the steady state measured values of these $\langle \mathcal{F} \rangle$ found by integrating (4.43) across the top boundary. Note that the values of $\langle \mathcal{F}_a \rangle$ for the single- and double-layer experiments had variations on the same order of magnitude as the reported values.	79
B.1	The compound trends for each variable v across the cluster average of each variable $d_v(d_t v)$ for both the original variables and the residuals with the corresponding <code>polyfit2d</code>	106

Chapter 1

Introduction

1.1 Background

Climate change and Arctic Amplification

In the past 50 years, global mean surface temperature has been rising at a rate of around 0.2°C per decade (NOAA, 2022), faster than in any other period in at least the last 2000 years. Similarly, the current rate of increase in global ocean heat content is faster than in any other period since the last deglaciation, which was between 11–18 thousand years ago (Gulev et al., 2021). In the previous decade (2010–2019), global sea surface temperatures rose at a rate of 0.28°C per decade, 4.5 times faster than the rate from 1900–2019, which was of 0.062°C per decade (Garcia-Soto et al., 2021). Below the ocean’s surface, these rates are slower, however still increasing. In the upper 2000 m, global ocean temperatures rose at the rate of 0.047°C per decade from 1993–2000, increasing to 0.064°C per decade from 2000–2020 (Su et al., 2022). Compared to the rest of the world, the mean surface air temperature of the Arctic has been warming at a rate around four to five times faster (Gulev et al., 2021; Rantanen et al., 2022), a phenomenon known as Arctic amplification (Cohen et al., 2014; Timmermans & Marshall, 2020). This is related to recent changes in the sea ice (Li et al., 2024) and the Arctic Ocean is projected to warm at a rate of up to three times that of the global ocean average by the end of the 21st century (Shu et al., 2022). This makes the Arctic a particularly crucial region to monitor for the effects of climate change.

The North Atlantic Current sends water through the Fram Strait between Greenland and Svalbard into the Arctic Ocean (Carmack et al., 1997; Wells & Wettlaufer, 2007). This interacts with the cold and fresh Arctic surface layer, which mainly comes from river runoff, sea ice melt, and low-salinity water from the North Pacific (Ekwurzel et al., 2001; Lincoln et al., 2016). The warm and salty Atlantic Water (AW) subducts beneath these surface waters to occupy depths between 250–800 meters (Timmermans et al., 2008). This AW layer is the most significant heat reservoir in the Arctic Ocean (Richards et al., 2022) and has warmed markedly since the mid-twentieth century (Carmack et al., 1997; Polyakov, Pnyushkov, & Timokhov, 2012). Moreover, many authors have noted that, while the AW is stable at depth, it contains enough heat to melt all Arctic sea ice if it were somehow able to reach the surface (Maykut & Untersteiner, 1971; Aagaard & Greisman, 1975; Steele & Boyd, 1998; Ekwurzel et al., 2001; Timmermans et al., 2008; Turner, 2010; Fer, 2014; Shaw & Stanton, 2014; Polyakov et al., 2017; Stranne et al., 2017; Shibley et al., 2020; Lu et al., 2022;



Figure 1.1: A map of the Arctic region. Adapted from CIA World Factbook (Public Domain). Retrieved June 6, 2024 from <https://commons.wikimedia.org/w/index.php?curid=531491>.

Brown & Radko, 2022; Richards et al., 2022). Indeed, Arctic sea ice is steadily disappearing with summer sea ice extent declining at a rate of around 10% per decade (Comiso et al., 2008; Garcia-Soto et al., 2021). Exposure to heat from below, in addition to the warming atmosphere above, would accelerate this decline (Turner, 2010).

Properties of Seawater

The density ρ of seawater is determined by three factors: pressure, temperature, and salinity. Overall, ρ is fairly constant around a reference value $\rho_0 = 1.027 \times 10^3 \text{ kg m}^{-3}$ (Vallis, 2017). The biggest variation in density is the background vertical profile $\bar{\rho}$, though this only amounts to a few percent change in total density from top to bottom (Sutherland, 2010). With a density stratification, a fluid parcel displaced vertically will undergo oscillations as buoyancy forces work to return it to its original depth. The frequency of these oscillations N is known as the Brunt-Väisälä, stratification, or buoyancy frequency and its square is proportional to the gradient of the background density profile.

When $\rho_0 \approx \rho$, then $N^2 = -(g/\rho_0)\partial_z \hat{\rho}$ to a good approximation (Sutherland, 2010; Kundu et al., 2015).

In certain cases, it can be useful to employ a quantity called potential density anomaly. Potential density is the density a particular fluid parcel would have if brought to a reference pressure p_r through an isentropic and isohaline process (IOC, 2010). That is, without transferring heat, mass, or salinity to its surroundings and without the dissipation of mechanical energy. In this thesis, I use a reference pressure of $p_r = 1000$ dbar. The potential density anomaly σ is simply the potential density minus 1000 kg m^{-3} (IOC, 2010).

The majority of pressure differences are accounted for by the background profile, which is in hydrostatic balance with $\hat{\rho}$ (see Section 4.2.1) (Vallis, 2017) and, except for in the abyssal ocean, does not have a significant effect on density (Sutherland, 2010). For small fluctuations in both temperature and salinity, the density of seawater can be expressed in the following linear equation of state (Sutherland, 2010; Cushman-Roisin & Beckers, 2011; Vallis, 2017)

$$\rho = \rho_0[1 - \alpha(T - T_0) + \beta(S - S_0)] \quad \Rightarrow \quad N^2 = g(-\alpha\partial_z T + \beta\partial_z S), \quad (1.1)$$

where T_0 , and S_0 are reference values of temperature and salinity, respectively; α (units of K^{-1}) is the thermal expansion coefficient (Timmermans et al., 2008; Cushman-Roisin & Beckers, 2011), and β (unitless) is the haline contraction coefficient (Bebieva & Timmermans, 2019; van der Boog, Otto Koetsier, et al., 2021).

In this thesis, I use the quantity called conservative temperature Θ , the temperature measure recommended by TEOS-10 (McDougall & Barker, 2011). This is defined as the potential enthalpy h^0 divided by a fixed heat capacity (McDougall & Barker, 2011), where h^0 is the enthalpy (internal thermal energy plus pressure-volume potential energy) a parcel of fluid would have after being brought to sea level pressure through an adiabatic and isohaline process (Kundu et al., 2015). Many other studies use the similar quantity, potential temperature θ . For the range of temperatures and salinities of data analyzed in this thesis, the difference $|\theta - \Theta|$ is a systematic shift that is always less than 0.05°C (McDougall & Barker, 2011).

Seawater is a solution of ions such as chloride, sodium, sulfate, magnesium, etc. However, as first reported in 1819, the ratios between these different components is nearly constant throughout the ocean and the term *salinity* refers to the amount of these materials in a particular sample of water (Millero, 2010; Vallis, 2017). Specifically, absolute salinity S_A is defined as the mass fraction of dissolved non- H_2O material in seawater and is reported in g/kg (Millero, 2010; IOC, 2010). However, as measuring the dissolved material in a seawater sample is cumbersome, most salinity observations use the practical salinity S_P scale based on the measured conductivity (Millero, 2010). While S_P is easier to measure, different samples with the same conductivity from different areas around the ocean can have different compositions, and so conversions are made from S_P to S_A based upon the location of the measurement (McDougall & Barker, 2011).

Double Diffusion and Thermohaline Staircases

Following the widely-used naming conventions of equation (1.1), two thirds of the world's oceans are *alpha* oceans which are stratified by temperature, with the warmest waters at the surface and colder waters at depth (Stewart & Haine, 2016). This is in contrast to the 15% that are *beta* oceans, the only

regions where sea ice can form over deep water, such as the Arctic which are primarily stratified by salinity; therefore the warm, salty AW is stable at depth (Carmack, 2007). This distinction arises from the non-linear thermal contraction of seawater, whose influence on density becomes negligible compared to that of salinity when temperatures approach 0°C. Beta oceans are found at high latitudes, alpha oceans at lower latitudes, and the remaining oceans are called “transition zone oceans” (Stewart & Haine, 2016). The border between these two domains has been moving poleward in the northern hemisphere, part of a process called “Atlantification,” and can be seen in the Greenland Sea which transitioned from the beta to alpha domain over the past 40 years (Gjelstrup & Stedmon, 2024).

The AW core is defined to be at the location of the subsurface temperature maximum, generally centered around depths of 300-500 m, just below the pycnocline where density increases rapidly with depth (Aagaard et al., 1981; Wells & Wettlaufer, 2007). This density stratification above the AW core has historically been strong enough to insulate surface waters from the warmth below (Aagaard et al., 1981; Fer, 2014; Shibley et al., 2017). Heat diffusion in the ocean is very slow (Turner, 2010), so heat from the AW could only rise at a significant rate due to vertical mixing from external sources (Guthrie et al., 2013).

The top of the pycnocline usually coincides with the lower halocline waters, the core of which is defined as the depths of salinity $S = 34.1$ (Lu et al., 2022). Between the cores of the lower halocline waters and the AW, the water’s density changes suddenly every few meters forming a “staircase” (Sutherland, 2016). This characteristic pattern can be seen in Figure 1.2 with layers that are several meters thick, have well defined and homogeneous temperature, salinity, and density, and are separated by sharp gradient interfaces (Sutherland, 2016; Shibley & Timmermans, 2019). Because these stratification structures are found where temperature and salinity are changing rapidly, they are often referred to as *thermohaline staircases*. However, they fall under the more general category of *double-diffusive staircases*.

When a fluid’s density increases monotonically downwards, it is considered stably stratified. However, if a stably stratified fluid has two components that affect its density and the vertical gradient of one component is statically unstable, this gives rise to double diffusion (van der Boog, Dijkstra, et al., 2021). Potential energy is stored in the unstable gradient which, if not for the stabilizing gradient of the other component, would cause the density to increase upwards. While initially documented in the ocean, double diffusion is now studied in many contexts such as saline lakes (Newman, 1976), lakes (Xue et al., 2017), and in stars and gas giants such as Jupiter and Saturn where it is referred to as *semi-convection* (André et al., 2017; Pontin et al., 2021). Double diffusion can occur when there are two density-affecting components that diffuse at significantly different rates.

In the ocean, the two components are temperature, which diffuses at a rate $\kappa_T \approx 1.4 \times 10^{-7} \text{m}^2/\text{s}$, and salinity, which diffuses at a rate $\kappa_S \approx 1.1 \times 10^{-9} \text{m}^2/\text{s}$ (Radko, 2013; Shibley & Timmermans, 2019). Seawater could have a stabilizing gradient of temperature decreasing with depth as well as a destabilizing gradient of salinity decreasing with depth. This case, where the more slowly diffusing component is unstably stratified, is referred to as the *salt fingering* regime. The case where unstable stratification is in the faster diffusing component (i.e., both temperature and salinity increasing with depth) is called *diffusive-convection* (Linden & Shirtcliffe, 1978). These two regimes are graphically depicted in Figure 1.3.

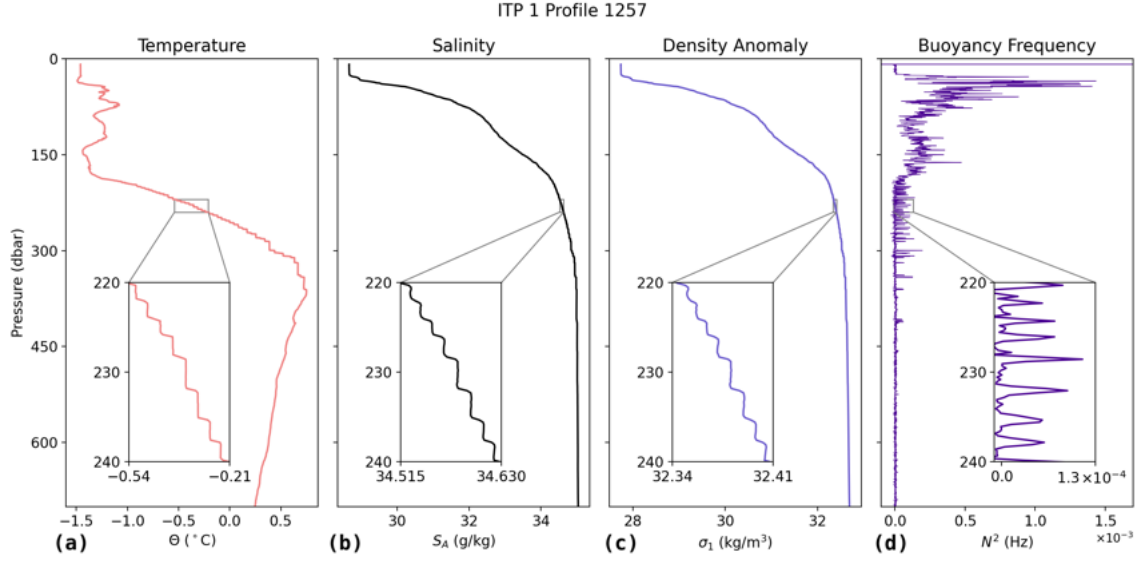


Figure 1.2: An example profile, number 1257 from Ice-Tethered Profiler 1 (see Section 1.2) in (a) conservative temperature, (b) absolute salinity, (c) potential density anomaly referenced to 1000 dbar, and (d) buoyancy frequency. This profile was taken on Monday, June 26th, 2006 at 136.1525°W, 77.2870°N.

The layers that make up the staircases are known to be laterally coherent across hundreds of kilometers, either through tracing individual layers across different profiles or through features in temperature–salinity space, as discussed in Section 1.3 (Schmitt et al., 1987; Timmermans et al., 2003, 2008; Toole et al., 2011; Lu et al., 2022). They have aspect ratios as large as 10^5 – 10^6 and, for this reason, are sometimes referred to as “sheets.” The exact mechanisms which form staircases have been debated for more than fifty years (Radko, 2013). A recent study presented a method in which the formation has less to do with double-diffusive instabilities and more with shear instabilities (Y. Ma & Peltier, 2022). However, it is accepted that having opposing stability gradients in two components is a necessary component (Sutherland, 2016; Shibley & Timmermans, 2019; Y. Ma & Peltier, 2022). Occurring at such small scales, these processes are difficult to represent in climate models and therefore difficult to predict their impacts on the global ocean (Ménèsqueun et al., 2022). In this thesis, I am not concerned with the processes behind the formation of thermohaline staircases, only on their existing presence in the Arctic Ocean.

Using data from autonomous instruments, van der Boog, Dijkstra, et al. (2021) mapped locations of thermohaline staircases throughout the world’s oceans, with the salt fingering regime corresponding to alpha oceans and diffusive-convection to beta oceans (see Figure 1.4). In particular, the highest concentration of diffusive-convection staircases is in the Arctic, a region where such structures have been particularly well-studied (Timmermans et al., 2008; Lu et al., 2022; Ménèsqueun et al., 2022). The first recorded observation of thermohaline staircases in the Arctic Ocean was made in 1969 at Ice Island T-3 (Neal et al., 1969). This iceberg, located somewhat northeast of the Canada Basin (often defined as 72–84°N, 130–155°W, see e.g. Peralta-Ferriz & Woodgate (2015)) was the site of dozens of hydrographic profiles that contain clearly visible staircases (Neshyba et al., 1971, 1972; Neal & Neshyba, 1973). Subsequent observations indicate that the staircases have been a consistent feature of the Canada Basin, including from data collected during the Arctic Internal

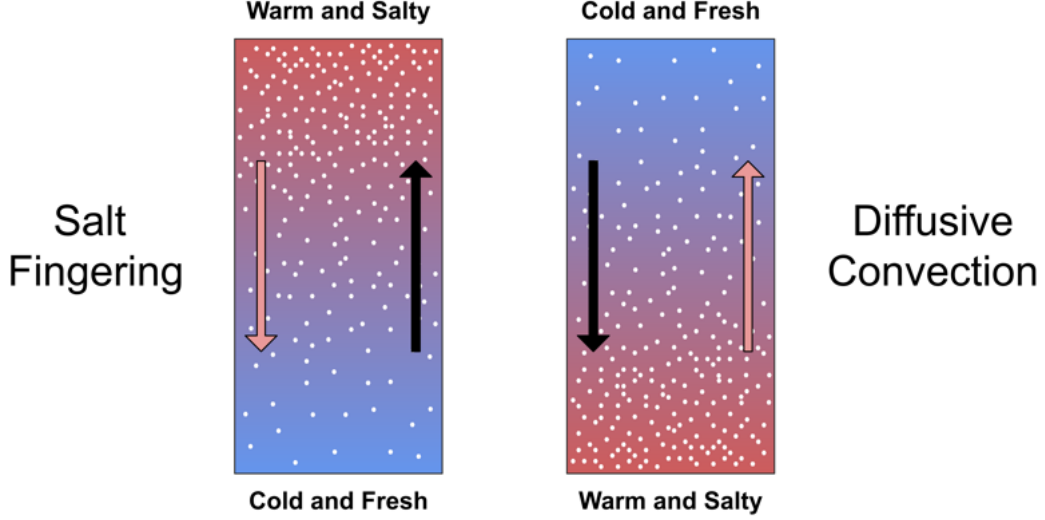


Figure 1.3: The two double diffusive regimes. The pink arrows indicate the direction in which the temperature gradient acts to increase density and the black arrows indicate in which direction salinity acts to increase density. Arrows pointing downwards represent stabilizing gradients while upwards pointing arrows represent destabilizing gradients.

Wave Experiment (AIWEX) (Padman & Dillon, 1987, 1988, 1989) and the Surface Heat Budget of the Arctic (SHEBA) experiment (Shaw & Stanton, 2014). The frequency of such observations increased dramatically in 2004, when the introduction of autonomous Ice-Tethered Profilers made possible the continuous, year-round sampling of the Arctic Ocean water column (Toole et al., 2011).

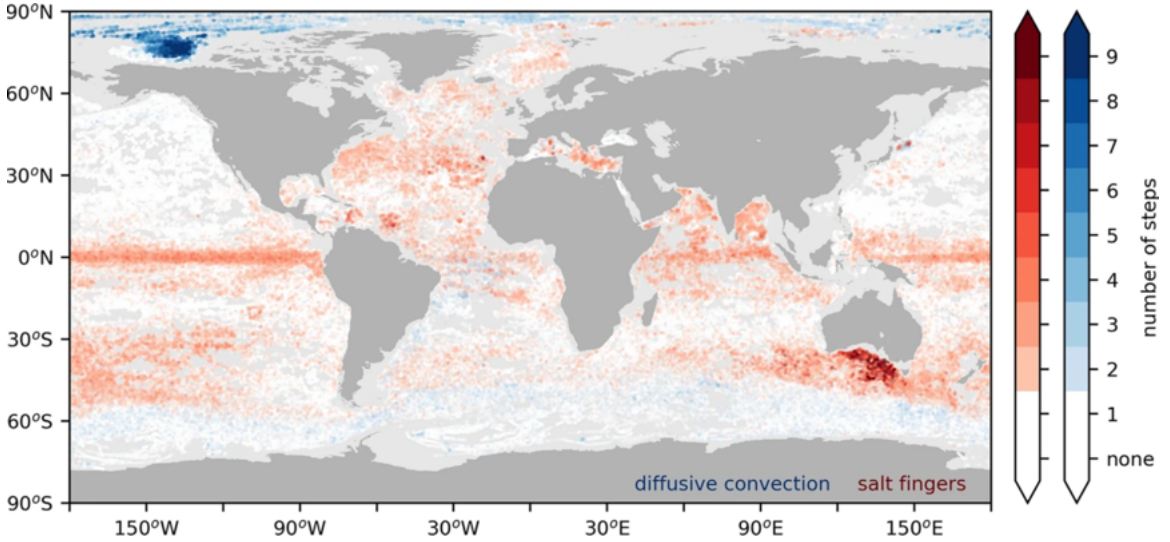


Figure 1.4: The global distribution of thermohaline staircases in the salt-fingering regime (red) and the diffusive-convection regime (blue), with data between 2001 and 2020. The shade of each dot indicates the number of steps detected in the profile, either from an Argo float or Ice-Tethered Profiler. Reproduced from van der Boog, Dijkstra, et al. (2021).

Since the start of high-resolution observations, thermohaline staircases have been persistent and wide-spread features of the Arctic Ocean (Timmermans et al., 2003; Polyakov, Pnyushkov, Rembe, et al., 2012; Lu et al., 2022). However, the lower sections of Arctic staircases have been disappearing

in recent years. While the mechanisms behind this transition are as of yet unclear, it coincides with changes to the large scale circulation and the decline in sea ice cover, which plays a major role in regulating energy input to the Arctic Ocean (Ménèsguen et al., 2022).

Internal Waves in the Arctic Ocean

While surface waves occur at the boundary between the ocean and the atmosphere, internal waves propagate within the ocean itself (Sutherland, 2010). Because these waves depend on the influence of buoyancy, their interactions with intricate stratification structures such as thermohaline staircases can be complex (Sutherland, 2016). Compared to most regions of the global ocean, the Arctic’s internal wave field is less energetic (Guthrie et al., 2013; Ghaemsaidi et al., 2016; Lincoln et al., 2016). This is mainly due to three factors: weak tidal forcing, limited interactions with other oceans, and sea ice.

Tidal forcings from the sun and the moon oscillate the ocean across bottom topography which generates internal waves (Simmons et al., 2004). As detailed in Section 4.2.2, internal waves in a strongly stratified fluid cannot propagate when their frequency is less than the absolute value of the Coriolis frequency, defined as $f = 2\Omega_e \sin(\phi)$ at a latitude ϕ where $\Omega_e = 2\pi/(24 \text{ hours})$ is the rotation rate of the Earth (Vallis, 2017). The barotropic M_2 tide is forced by the Moon and has a period of 12.4 hours (Nash et al., 2012; Morozov & Paka, 2010). In the Arctic, there is a critical latitude of 74.5°N at which the M_2 tidal frequency equals the Coriolis frequency, i.e., $\omega_{M_2} = f$ (Simmons et al., 2004). Therefore, in the region north of this latitude, which encompasses the majority of the Arctic Ocean, waves generated by the M_2 tide cannot propagate (Fer, 2014). This weak tidal forcing leads to lower internal-tide energies in the Arctic compared to oceans at lower latitudes (Guthrie et al., 2013).

The majority of the region south of the critical latitude for the M_2 tide at 74.5°S is covered by Antarctica. Therefore the same weak tidal forcing does not apply to the Southern Ocean. However, the Arctic Ocean is also less energetic because of its relatively restricted interactions with the global ocean. The Southern Ocean is surrounded on all sides by other oceans and, while some robust staircases have been observed, many are seasonally disrupted by turbulent mixing (Bebieva & Speer, 2019). In contrast, Arctic Ocean is primarily surrounded by land with only limited interactions with the Pacific and Atlantic Oceans through relatively narrow straits (Hassol, 2004).

Persistent ice cover in the Arctic limits the energy which can be transferred to the ocean from wind-driven shear and also acts to dissipate internal wave energy at the ocean-ice boundary (Guthrie et al., 2013; Fer, 2014). Where the ocean is exposed, wind blowing across the surface can generate internal waves leading to a more energetic internal wave field when wind speeds are high (Dosser & Rainville, 2016; Ghaemsaidi et al., 2016). These internal waves can act as a conduit between the surface and the ocean interior, exerting forces on the fluid and causing turbulence and mixing where they break (Sutherland, 2010; Klymak & Legg, 2010). While staircases have been shown to be notably resistant to shear forces (Brown & Radko, 2022), they are less often observed in turbulent areas of the Arctic, such as those with minimal ice cover (Shibley & Timmermans, 2019).

The tides, geography, and sea ice cover lead to the Arctic Ocean having an environment conducive to the stable, year-round presence of thermohaline staircases. However, as Arctic sea ice continues to decline, there is growing concern over the potential impacts of increased interactions between the atmosphere and the Arctic Ocean (Rainville et al., 2011). Indeed, increases of up to two orders of

magnitude in kinetic energy have been seen in response to ice breakup (Lincoln et al., 2016). This has lead many to speculate that an increase in wind-generated internal waves could lead to deeper vertical mixing (Rainville et al., 2011; Fer, 2014; Shibley & Timmermans, 2019). As vertical mixing increases, heat from the AW will more easily be transferred to the surface leading to further ice loss which will allow more wind energy to transfer into the ocean in a positive feedback loop (Turner, 2010). This is in addition to the well-known ice-albedo positive feedback where the ocean, being much less reflective than sea ice, absorbs more solar radiation, warming up, melting more sea ice, and exposing more ocean to the sun (Comiso et al., 2008; Cohen et al., 2014; Sledd & L’Ecuyer, 2021). Continued loss of sea ice in the Arctic will have major consequences for both the resident ecosystems and large scale circulation (Comiso et al., 2008; Ghaemsaidi et al., 2016).

While an increase in internal wave activity has been seen in correlation with decreases in ice cover, observations have not shown a similar increase in mixing from the AW compared to ice-covered regions (Lincoln et al., 2016). This is often attributed to the sharp stratification of the pycnocline. While internal waves could potentially bring energy that could be used for vertical mixing, thermohaline staircases, where the water’s density changes suddenly every few meters, present obstacles to their propagation at each step (Sutherland, 2016). As shown in Section 4.2.2, internal waves cannot propagate in regions where their frequency ω is greater than the buoyancy frequency N which is proportional to the derivative of density. As can be seen in Figure 1.2(d), N is effectively zero within each well-mixed layer, creating barriers to any internal wave with non-zero frequency. However, if these well-mixed layers are thin enough in comparison to the wavelength, the internal waves will propagate out the other side, a phenomenon known as *internal gravity wave tunnelling* (Eckart, 1961; Sutherland & Yewchuk, 2004; Ghaemsaidi et al., 2016). Many authors have studied this phenomenon under particular circumstances (Sutherland & Yewchuk, 2004; Ghaemsaidi et al., 2016; Wunsch, 2018; Supekar & Peacock, 2019; Boury et al., 2019), however, it is still unclear how much internal waves can transmit through stratification structures in the Arctic (Sutherland, 2016).

1.2 Ice-Tethered Profiler Data

The hydrographic profile data which I use in this thesis come from the Ice-Tethered Profilers (ITPs). These are automated, vertically profiling instruments that are connected via a wire cable to a surface buoy on an ice floe as depicted in Figure 1.5(a) (Toole et al., 2011). The wire extends from the surface, through the ice, down to depths of 500–800 m. The profiler travels up and down the wire roughly two to three times a day collecting high-resolution (~ 25 cm) measurements of salinity, temperature, and pressure with a salinity precision of ± 0.005 g/kg and a temperature precision of $\pm 0.001^\circ\text{C}$ (Timmermans et al., 2008; Shibley et al., 2017; Bebieva & Timmermans, 2019). As opposed to ship-based measurements, ITPs necessarily record profiles under areas of high ice concentration. ITP data are archived at the NOAA National Centers for Environmental Information repository (Toole et al., 2016) (accessed September 21, 2023). In order to take full advantage of the ~ 25 cm resolution, I use the Level III Matlab-format data files as opposed to the Level III “final” format which bin-averages the data to a 1-dbar vertical resolution (Krishfield et al., 2008).

Each traverse, alternating up or down, is recorded as a separate profile. ITPs have recorded over 160,000 profiles across the entire Arctic Ocean (see Figure 1.5(b)), however the highest concentration of observations is in the Canada Basin. ITP data have been used in numerous studies to analyze

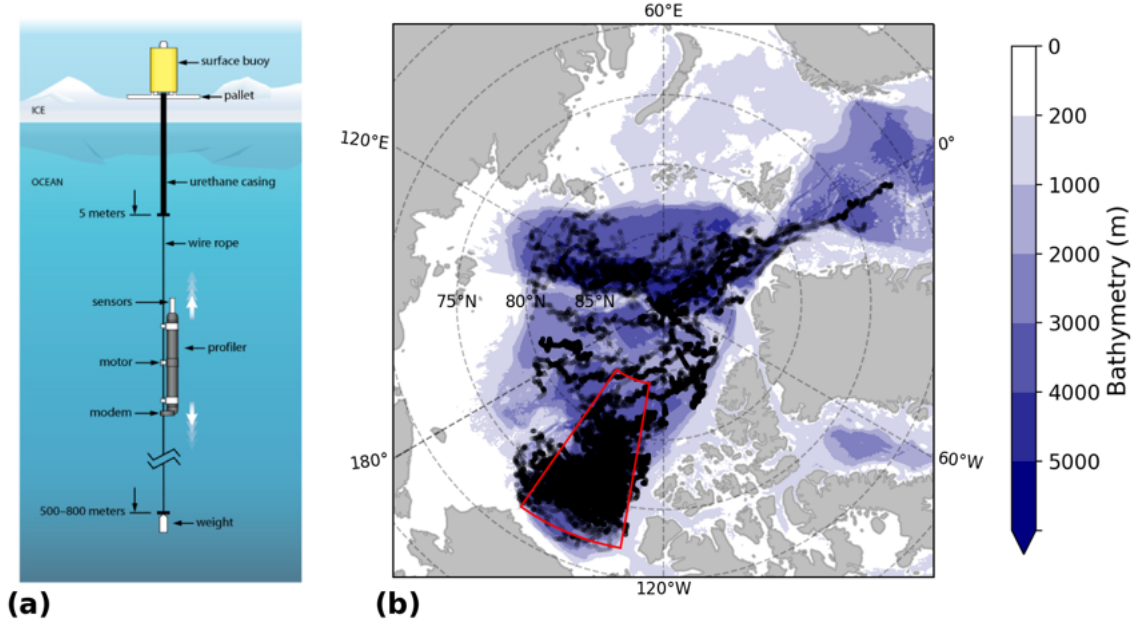


Figure 1.5: (a) A schematic of an Ice-Tethered Profiler, showing how the surface buoy on the ice is connected to the wire upon which the profiler moves up and down while taking measurements. Reproduced from [Toole et al. \(2011\)](#). (b) A map of ITP profiles taken between 2004–2023. The red box designates the Canada Basin as defined by [Peralta-Ferriz & Woodgate \(2015\)](#).

thermohaline staircases ([Timmermans et al., 2008](#); [Shibley et al., 2017](#); [Bebieva & Timmermans, 2017](#); [Shibley & Timmermans, 2019](#); [Bebieva & Timmermans, 2019](#); [van der Boog, Otto Koetsier, et al., 2021](#); [Ménèsguen et al., 2022](#); [Lu et al., 2022](#); [Shibley & Timmermans, 2022](#)), and many have noted the wake problem. Because the measuring instruments are located at the top of the profiler, the down-going profiles contain distortions caused by the wake of the profiling unit itself. While these distortions do not change the general shape of the profile, they do disrupt fine-scale features. Therefore, studies that analyze thermohaline staircase structure generally limit analysis to only the up-going profiles ([Timmermans et al., 2008](#); [Shibley et al., 2017](#); [Lu et al., 2022](#)).

1.3 The HDBSCAN clustering algorithm

Many studies, notably [Timmermans et al. \(2008\)](#) in Figure 1.6, have observed that a collection of profiles that includes well-mixed layers, such as thermohaline staircases, is associated with clustered patterns when plotted in temperature-salinity (T - S) space ([Schmitt et al., 1987](#); [Toole et al., 2011](#); [Yu et al., 2017](#); [Bebieva & Timmermans, 2019](#)). This clustering indicates that layers are thin sheets, extending laterally for 100's of kilometers ([Timmermans et al., 2008](#); [Lu et al., 2022](#)). These patterns occur where a staircase is present because all data points within a particular layer have approximately the same temperature and salinity values as other observations from the same layer, regardless of their vertical position within the layer. This fact suggests that staircases could be detected by directly identifying clusters in T - S space. Clustering algorithms are a type of unsupervised machine learning, and have been previously used in a variety of different oceanographic applications ([Sonnewald et al., 2021](#)). Examples include grouping observations from freely-drifting instruments in the Nordic Seas

(Koszalka & LaCasce, 2010), studying surface wave variability in the North Atlantic (Espejo et al., 2014), classifying the heat content of hydrographic profiles (Maze et al., 2017), resolving trapping in mesoscale eddies (C. Ma et al., 2019), defining spatial regions for Southern Ocean temperature profiles (Jones et al., 2019), detecting ENSO events (Houghton & Wilson, 2020), finding hot spots for mixing in the Southern Ocean (Rosso et al., 2020), and identifying shifts in the North Atlantic circulation (Desbruyères et al., 2021). However, to my knowledge, they have never before been used to identify thermohaline staircases.




Figure 1.6: The potential temperature and salinity values between 200–300 m depth from profiles taken by ITP2 spanning around 200 km laterally with lines indicating constant potential density anomaly (referenced to the surface). The thin (thick) line represents the westernmost (easternmost) profile. Reproduced from Timmermans et al. (2008)

In Chapters 2 and 3, I apply the Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) algorithm to data from ITPs. While Campello et al. (2013) present the algorithm in full detail, here I briefly review its general principles.

In the context of density based clustering algorithms, the term “density” refers to a measure of the relative number of data points in a certain region of parameter space. HDBSCAN estimates the density of a region based on the distances between points and a number of their nearest neighbors, creating a hierarchy of clusterings from which it chooses the most prominent. First, it calculates the “core distance” ε for each point as the minimum radius of a circle needed to encompass its m_{pts} nearest neighbors. The inverse of ε represents density; when it is small, points are close together and when it is large, points are spread out. HDBSCAN then creates a hierarchy of clusterings, starting with the largest ε and working down. For each value ε_0 , it first detects the level-set of non-noise points, where $\varepsilon \leq \varepsilon_0$, then connects points together in the same cluster if they are within a distance of ε_0 .

As the value of ε_0 decreases, a particular cluster may change by shrinking, splitting, or disappearing entirely. Having built this hierarchy tree, HDBSCAN can then select clusters at the ends of branches even though different branches may terminate at very different values of ε_0 . However, when the data are particularly noisy, just before crossing the ε_0 threshold where a cluster would disappear, it may split into many small, spurious clusters. To avoid including these artifacts in the final result, the algorithm ignores all clusters with fewer than m_{clSize} number of points. The recommendation of the authors of HDBSCAN and the default behavior of the “hdbscan” Python package is to set $m_{clSize} = m_{pts}$ (Campello et al., 2013; Moulavi et al., 2014), giving one hyperparameter which controls both how the core distance is calculated and the minimum points per cluster. The term hyperparameter is used to describe parameters given as input before running the algorithm. While there are other optional hyperparameters for HDBSCAN (see Section A.3), the results I obtained using the default settings were satisfactory and so I did not investigate their effects.

I choose HDBSCAN over other types of clustering algorithms for several reasons. Many previous oceanographic studies used partitioning algorithms such as k -means (Koszalka & LaCasce, 2010; Espejo et al., 2014; Houghton & Wilson, 2020) or Gaussian mixture models (Maze et al., 2017; Jones et al., 2019; Rosso et al., 2020; Desbruyères et al., 2021); however, partitioning algorithms require one to specify the number of clusters *a priori*, which is not known for my application. The first reason for choosing HDBSCAN is that it does not have this requirement (Ester et al., 1996; Jones et al., 2019). Second, HDBSCAN allows points that lie outside a cluster to be categorized as “noise.” In my case, this is important because points in an interface between layers should not be assigned to any cluster. Third, unlike DBSCAN, which uses the same threshold ε_0 for all clusters throughout the domain (Ester et al., 1996; Campello et al., 2013), HDBSCAN creates a hierarchy of clusterings with different ε_0 and can therefore correctly identify clusters that vary significantly in both the number of points per cluster and the densities of points within the clusters (McInnes et al., 2017). Fourth, HDBSCAN can correctly identify arbitrarily shaped clusters, whereas partitioning algorithms like k -means generally find center-defined clusters, which, because points are assigned to clusters based on their distance from the cluster’s center point, are only equipped to find globular, or convex, clusters (Hinneburg & Keim, 2003; Ester et al., 1996). This is important because the shapes of clusters associated with thermohaline staircases are not necessarily globular. Lastly, HDBSCAN requires only one hyperparameter (m_{pts}) to be specified, reducing the number of choices to be made before each run of the algorithm. Further, I determine the value of m_{pts} systematically, as explained in Section 2.3.2.

1.4 Overview

The overall objective of this thesis is to examine thermohaline staircases and their interactions in the Arctic Ocean, particularly in the Canada Basin where they have been most consistently observed. Given their potential role in modulating the Arctic climate, it is important to be able to accurately identify thermohaline staircases in observations to monitor changes.

Many previous studies have analyzed thermohaline staircases in the Arctic Ocean (Timmermans et al., 2008; Shibley et al., 2017; Bebieva & Timmermans, 2019; Lu et al., 2022). However, most large-scale efforts to detect staircases have relied on methods which treat each hydrographic profile individually. While such methods can provide information such as the locations of staircases, layer

thicknesses, fine-scale structure of temperature and salinity, etc., they do not allow for the analysis of the staircases as the sprawling horizontal structures they are. In Chapter 2, I describe a novel technique to detect and connect staircase layers across many hydrographic profiles using an unsupervised clustering method. I compare the properties of the layers this clustering method detects to those of previous studies using different detection methods on the same data sets and find close agreement. I describe how this method could be extended to identify layers in much larger datasets. The work described in Chapter 2 is published in *Environmental Data Science* (Schee et al., 2024).

Thermohaline staircases have been a well-known feature of the Arctic Ocean for over fifty years and their layers have been shown to extend horizontally for hundreds of kilometers (Neal et al., 1969; Timmermans et al., 2003, 2008; Lu et al., 2022). However, few studies have described how the properties of individual layers change over more than a few years (Polyakov, Pnyushkov, Rembe, et al., 2012; Ménesguen et al., 2022; Lu et al., 2022). These structures have potentially significant impacts on vertical transport of energy, however the impact of their stratification is not well-represented in climate models (Timmermans & Pickart, 2023). It is important to analyze in detail how staircase layers change over time in order to better predict the future evolution of the Arctic Ocean. In Chapter 3, I apply the unsupervised clustering method to a dataset of hydrographic profiles. I establish new coherency scales for staircase layers, finding individual ones which span across 17 years and the entire Canada Basin, on the order of 1000 km. I also find layers changing in temperature across time, with layers near the top of the staircase warming and layers near the bottom cooling. An article based on the work described in Chapter 3 is in preparation.

As Arctic sea ice declines, there is concern that an increase in wind-generated internal waves could provide the energy needed to release heat from the underlying AW (Sutherland, 2010; Turner, 2010; Rainville et al., 2011; Fer, 2014). The recent increase in internal wave activity has not coincided with an increase in mixing from the AW, a fact attributed to the strong stratification of the pycnocline (Lincoln et al., 2016). However, according to theory and confirmed by experiments, internal gravity wave tunnelling can occur, allowing internal waves to pass through certain types of stratification. (Eckart, 1961; Sutherland & Yewchuk, 2004; Ghaemsaidi et al., 2016; Sutherland, 2016). As the Arctic continues to change, the question remains as to under what conditions might internal waves be able to propagate through thermohaline staircases. In Chapter 4, I present numerical experiments of internal wave interactions with stratification structures resembling those of thermohaline staircases. I solve the Boussinesq equations of motion for experiments in both one and two spatial dimensions and, in both cases, find results that match those of previous studies. This agreement suggests that such numerical experiments could be extended to make predictions of how interactions between internal waves and thermohaline staircases may look in the future. In Chapter 5, I provide a summary of the salient results from this thesis and describe promising avenues for future exploration on these topics.

Chapter 2

Unsupervised clustering identifies thermohaline staircases in the Canada Basin of the Arctic Ocean

Abstract Thermohaline staircases are a widespread stratification feature that impact the vertical transport of heat and nutrients and are consistently observed throughout the Canada Basin of the Arctic Ocean. Observations of staircases from the same time period and geographic region form clusters in temperature-salinity (T - S) space. Here, for the first time, we use an automated clustering algorithm called the Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN), to detect and connect individual well-mixed staircase layers across profiles from Ice-Tethered Profilers (ITPs). Our application only requires an estimate of the typical layer thickness and expected salinity range of staircases. We compare this method to two previous studies that used different approaches to detect layers, and reproduce several results including the mean lateral density ratio R_L and that the difference in salinity between neighboring layers is a magnitude larger than the salinity variance within a layer. We find that we can accurately and automatically track individual layers in coherent staircases across time and space between different profiles. In evaluating the algorithm’s performance, we find evidence of different physical features, namely splitting or merging layers and remnant intrusions. Further, we find a dependence of R_L on pressure, whereas previous studies have reported constant R_L . Our results demonstrate that clustering algorithms are an effective and parsimonious method of identifying staircases in ocean profile data. The contents of this chapter are adapted from the corresponding article in *Environmental Data Science* (Schee et al., 2024). Replication code can be found on Zenodo: <https://zenodo.org/doi/10.5281/zenodo.8029947>.

Impact Statement Clustering algorithms are unsupervised machine learning methods that are used across many areas of data science. The use of such methods can automate the identification of certain features, thus allowing for analysis of very large datasets. Here, we show that a particular clustering algorithm called HDBSCAN can be used to automatically identify thermohaline staircases in hydrographic profiles from the Arctic Ocean. Compared to previous detection methods,

HDBSCAN has the advantages of requiring minimal prior knowledge and of automatically connecting individual staircase “steps” across different hydrographic profiles. We expect that this method could be applied to many similar datasets, offering a straightforward way to identify and track layers in thermohaline staircases across the world’s oceans.

2.1 Introduction

Thermohaline staircases, formed by differential diffusion rates of heat and salt, appear as a series of vertically well-mixed horizontal layers each separated by thin, strongly stratified interfaces. These structures have been observed throughout the world’s oceans (van der Boog, Dijkstra, et al., 2021) and in particular, they are well known to occur in the Arctic Ocean (Timmermans et al., 2008; Lu et al., 2022; Ménesguen et al., 2022). As described in Section 1.1, around 250 to 800 meters below the surface of the Arctic Ocean, there is a layer of water originating from the Atlantic, which is warmer and saltier than the topmost layer that is in contact with the sea ice above (Timmermans et al., 2008). The core of the Atlantic Water (AW) is defined as the maximum subsurface temperature and is generally at a depth of 400 meters in the Canada Basin. Above this depth, in a region in which both temperature and salinity increase downwards, lies the pycnocline where the staircases are found (Lu et al., 2022).

While Arctic sea ice has been steadily disappearing (Comiso et al., 2008), the density stratification above the AW core has historically been strong enough to insulate surface waters from the warmth at depth (Shibley et al., 2017). However, the lower sections of Arctic staircases have been disappearing in recent years (Ménésguen et al., 2022). Many works have noted that the AW contains enough heat to melt all Arctic sea ice, if it were somehow able to reach the surface (Maykut & Untersteiner, 1971; Turner, 2010; Stranne et al., 2017; Shibley et al., 2020). Given their potential role in modulating the Arctic climate, it is important to be able to accurately identify thermohaline staircases in observations to monitor changes.

Although many early studies identified staircases by visual inspection, the recent increase in available data since the deployment of autonomous Ice-Tethered Profilers in 2004 has spurred researchers to turn to algorithmic approaches. Here, we detail several previous studies whose foundational work was critical in the development of the current study. All of these studies used a similar approach to detect data points that fall within some well-mixed layer (henceforth for brevity simply “layer”) on a profile-by-profile basis. Timmermans et al. (2008) defined a point in a hydrographic profile to be detected within a layer when the local vertical potential temperature gradient $\partial\theta/\partial z$ is below $0.005^\circ\text{C m}^{-1}$, roughly an order of magnitude smaller than the overall gradient for a typical profile. Shibley et al. (2017) extended the automated detection method of Timmermans et al. (2008) by including two additional conditions: (1) For each experiment, they visually determined different threshold values on both $\partial\theta/\partial z$ and on the temperature difference between neighboring points and (2) after then running the detection method, those authors only considered staircases that consisted of at least three layers. Van der Boog et al. (van der Boog, Otto Koetsier, et al., 2021) developed a similar staircase detection algorithm, but chose threshold values of vertical density gradients that were expected to be applicable to staircases in all of the world’s oceans, not just the Arctic. Specifically, they set a vertical gradient threshold of $\partial\sigma_1/\partial p \leq 0.0005 \text{ kg m}^{-3} \text{ dbar}^{-1}$ where σ_1 is the density anomaly referenced to 1000 dbar. Then, from the subset of data that meets that condi-

tion, only mixed layers with a maximum variation of density anomaly of 0.005 kg m^{-3} or less, and whose neighboring interfaces had thicknesses of 30 dbar or less, were considered. [Lu et al. \(2022\)](#) defined the intersections between layers and interfaces as locations where the difference in the potential temperature gradient between two neighboring points is greater than $0.003 \text{ }^{\circ}\text{C m}^{-1}$, then disregarded the points within the interfaces. After performing this detection, [Lu et al. \(2022\)](#) then made cross-profile connections with the layer points that remained, following the work of [Padman & Dillon \(1988\)](#) and used valleys in histograms of salinity to guide their choice of boundaries between layers, some of which were manually adjusted. We find it useful to distinguish between these two steps: the detection of points within layers for each profile and the connection of these points to those in other profiles that are within the same layer. Many previous studies did not describe a connection method because it was not needed for their ends; we use this framing to call attention to the novel aspects of our contribution.

All of the above techniques first detect a subset of the data points that are likely within layers based upon vertical gradient thresholds. This approach requires both sufficient knowledge of the staircases properties in order to select appropriate thresholds, as well as data with sufficient vertical resolution to accurately estimate the gradients. The connection algorithm [Lu et al. \(2022\)](#) describe, after its completion, still requires some amount of manual intervention to produce a final dataset of staircase layers. A commonality of all of these approaches is that they are purpose-built for the task of detecting staircases in specific contexts. These factors motivate the search for a more general approach to detect and connect thermohaline layers across different profiles which could offer advantages such as greater scalability and applicability as well as more ready reproducibility, all of which would accelerate the pace of research on these important structures.

As mentioned in [Section 1.3](#), previous authors have noted that data from sets of profiles that contain thermohaline staircases form clusters in temperature-salinity (T - S) space, suggesting that staircase layers could be detected by a clustering algorithm. Such algorithms have been used to identify many features in the ocean, however, to our knowledge, have never before been applied to thermohaline staircases. Here, we apply a method based on the Hierarchical Density-Based Spatial Clustering of Applications with Noise or HDBSCAN algorithm ([Campello et al., 2013](#)) to both detect and connect thermohaline staircases across Arctic hydrographic profiles. This method has several advantages. It detects and automatically connects staircase layers across large hydrographic datasets in one step. In the past, HDBSCAN has been successfully applied to datasets with a number of points that were an order of magnitude larger than in this study, suggesting the application we present here could be scaled to accommodate more data ([Logan & Fotopoulou, 2020](#)). Also, it does not consider profiles individually and therefore does not require that each profile have fine vertical resolution. Our implementation of this algorithm does require knowledge of the typical layer thickness and the expected global salinity range of the staircase; however, as we will show, this is a more flexible requirement compared to determining threshold gradients. Most importantly, it exhibits excellent performance producing a final, connected dataset of layers to which no subsequent adjustments need to be made, and is thus suitable for application to datasets with a large number of points as well as being more easily reproducible. The purpose of this study is to present this method, and so we will note possible implications of our results but will not explore them in great detail.

The structure of this chapter is as follows. First, in [Section 2.2](#), we introduce and explain our choice of the datasets used in our analysis. In [Section 2.3](#), we introduce HDBSCAN and describe how

ITP	Start Date	Days in Operation	Starting Position	Along-path Distance	Longest Span ^a	Total Profiles	Up-going Profiles
2	2004/8/19	40	77 10.4'N, 141 13.0'W	390 km	191 km	244	121
3	2005/8/23	382	74 15.0'N, 135 9.0'W	2541 km	441 km	1532	766

^aThe maximum distance between any two, not necessarily consecutive, profiles.

Table 2.1: Details of the Ice-Tethered Profilers (ITPs) used in this study.

we use it to identify staircases and how we choose the input parameters. Then, in Section 2.4, we apply this method to data from two different Ice-Tethered Profiler experiments, focusing on results that reproduce those from the studies of Timmermans et al. (2008) and Lu et al. (2022), hereafter denoted as T08 and L22, respectively. Finally, in Section 2.5, we use the comparison between our results and those of T08 and L22 to evaluate the performance of the clustering algorithm and give recommendations on when and how it can be best used to identify staircases.

2.2 Data

As described in Section 1.2, Ice-Tethered Profilers are automated instruments which take vertical profiles of salinity, temperature, and pressure down to depths of 500–800 m with a vertical resolution of ~ 25 cm (Timmermans et al., 2008; Shibley et al., 2017). In this study, data from two different ITP experiments, ITP2 and ITP3, are analyzed; see Table 2.1 for a summary. Because the measuring instruments are located at the top of the profiler, we use only the up-going profiles in order to avoid the known distortion caused by the wake of the profiling unit in the down-going profiles (Shibley et al., 2017). Figure 2.1 shows the locations of all up-going profiles for each of the two ITP experiments. We choose to analyze ITP2 and ITP3 in particular in order to reproduce several results of T08 and L22, respectively.

At this point we note some choices in the study relative to those of T08 and L22. Here we choose to work with pressure, a directly measured quantity, while both of those studies use depth, which is derived from pressure; as the conversion between pressure and depth is linear, this difference does not affect our results or comparisons. Also note that although both of those earlier studies use potential temperature θ , we choose instead to use conservative temperature Θ , as recommended by TEOS-10 (McDougall & Barker, 2011). For the range of temperatures and salinities in the data we analyze, the difference $|\theta - \Theta|$ is a systematic shift that is always less than 0.05°C . Because this shift is relatively constant, our results are not sensitive to this choice. Furthermore, TEOS-10 recommends using absolute salinity S_A over practical salinity S_P . Our results do not change significantly using one versus the other (see Section A.1), so we choose to use S_P to make direct comparisons to the salinity ranges used by T08 and L22.

Staircases are only found in certain vertical ranges of the water column. However, as noted by L22, the salinity values of the layers are much more consistent across different profiles than their pressure or temperature characteristics, so we focus our analyses on a specific salinity range for each ITP. An appropriate range can be estimated by inspection of a few individual profiles, or from general knowledge of the salinity range of staircases in the region. The bounds of this range need not be precise, but choosing a range that is too small will potentially miss layers. Running the algorithm on a liberally large range to capture all potential layers will take longer and may also return some

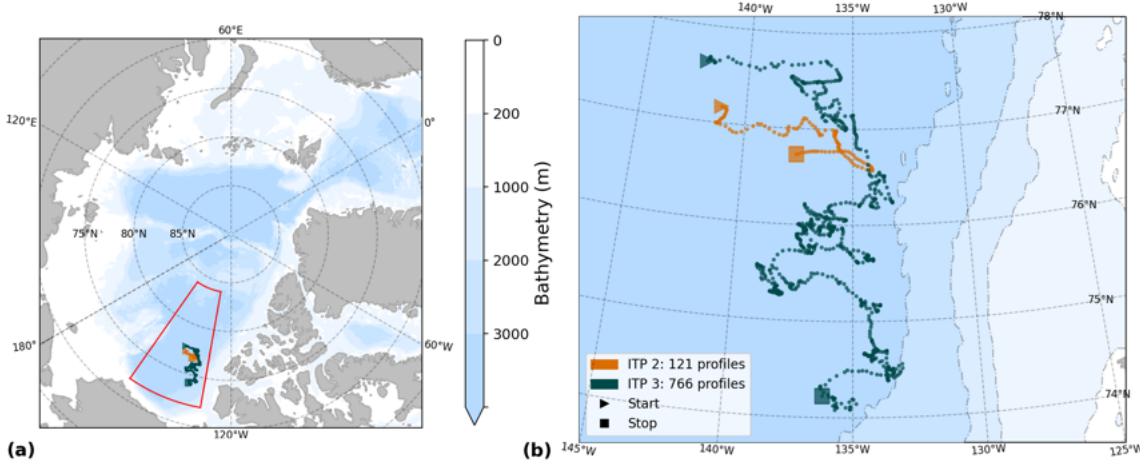


Figure 2.1: A map showing the locations of all profiles used from ITPs 2 and 3, showing the whole Arctic in (a) and a zoomed-in view in (b). The red box designates the Canada Basin as defined by [Peralta-Ferriz & Woodgate \(2015\)](#).

clusters that do not correspond to meaningful thermohaline staircase layers. However, in Section 2.3.3, we detail a process of identifying erroneous clusters. For easier comparison, we choose salinity ranges used in past work, namely, 34.05–34.75 g/kg for ITP2 following T08 and 34.21–34.82 g/kg for ITP3 following L22.

As discussed in more detail below, the clustering algorithm does not distinguish between the different times and locations at which different profiles were taken. It follows that a dataset to be analyzed should not span larger temporal or spatial scales than the scales the staircases are known to be coherent across. T08, who analyzed much of the same data we do, found staircases that spanned the entire Canada Basin (approximately 800 km) and lasted at least 2 years. We define the longest span as the maximum distance between any two, not necessarily consecutive, profiles. We find the longest span (≤ 441 km) and duration (≤ 382 days) of the two ITPs we analyzed, see Table 2.1, are indeed smaller than the known staircase coherence scales.

2.3 Methods

2.3.1 Applying the HDBSCAN clustering algorithm

To identify T – S clusters as evidence for staircases within the ITP data, we use the HDBSCAN algorithm which, as detailed in Section 1.3, clusters data based on the relative densities in different regions with a single hyperparameter, m_{pts} , the minimum number of points per cluster. Algorithm-identified clusters in the ITP data are expected to correspond to staircase layers, with some exceptions as detailed in Section 2.3.3.

Having chosen the HDBSCAN clustering algorithm, we now turn to specifying the two-dimensional space within which the clustering algorithm will operate. Figure 2.2(a) shows data from ITP2 in Θ – S_P space, where discrete groups of points, associated with individual layers and spanning multiple profiles, are apparent; these are colored according to their eventual partitioning into clusters. Note the occurrence of occasional gaps in Θ values, seen in the salinity ranges 34.16–34.30 and 34.65–34.74 g/kg. These result from the uneven spatial coverage of the meandering drift path of the ITPs

(see Figure 2.1) together with the tendency for the temperature to vary more horizontally within a particular layer than salinity (Lu et al., 2022). In order to avoid HDBSCAN splitting such groups of points into multiple clusters for the same layer, for each temperature profile $\Theta(z)$, we define the local anomaly as

$$\Theta'(z) = \Theta(z) - \frac{1}{\ell} \int_{z-\ell/2}^{z+\ell/2} \Theta(z') dz' \quad (2.1)$$

where ℓ is the width of a rectangular moving average window. Presenting the ITP2 data in $\Theta'-S_P$ space rather than $\Theta-S_P$ space, as in Figure 2.2(c), leads to groups that are more centered around zero along the temperature axis, without notable gaps. We choose to work in this space as it will allow HDBSCAN to group points more accurately. We also confirmed the difference between using Θ or θ and S_P or S_A did not significantly affect the results. Finally, we mention that HDBSCAN is sensitive to the aspect ratio of the axes. Because of the relative scaling between the temperature and salinity axes in the ranges we investigate, the algorithm puts more emphasis on salinity when determining clusters. However, as noted above, previous studies have shown that the salinity is the most relevant variable by which to identify layers and because the results were found to be satisfactory, we did not investigate how scaling the axes might affect the resulting clusters.

2.3.2 Selecting values of input parameters

It remains to choose values for the method parameters. The HDBSCAN algorithm is deterministic; that is, given the same arrangement of input data and the same value of the m_{pts} parameter, it will find the same clusters every time. The exact arrangement of data in $\Theta'-S_P$ space that we feed into the algorithm depends on three factors: (1) the set of profiles that we include, (2) the salinity range that we decide to analyze, and (3) the window width ℓ used to calculate the local anomaly of conservative temperature Θ' in Equation (2.1). We discussed our method of selecting the profiles and salinity range previously in Section 2.2. Here, we explain how we select values for ℓ and m_{pts} .

The results of a clustering algorithm can be judged on the basis of either external or internal validation. External validation methods involve comparing the clustering results to an external “ground truth,” while internal validation methods use the data themselves to provide a measure of quality for the clustering (Moulavi et al., 2014). In this application, external validation would require detailed labeling, indicating to which layer, if any, each individual point belongs. Such labels could be determined by a separate method; however, we aim for this method to be broadly applicable to more than just reproducing previous results. Therefore, we tune our selection of ℓ and m_{pts} using Density-Based Clustering Validation (DBCV) (Moulavi et al., 2014) as an internal validation. DBCV considers good clustering solutions to be those in which the lowest-density regions within the clusters are still denser than the highest-density regions of the surrounding noise points. It bases the density estimates off the so-called “mutual reachability distance,” defined for each pair of points to be the maximum ε of either point or the distance between the two, whichever is largest. For more details, see Moulavi et al. (2014).

To evaluate algorithm performance, we performed a parameter sweep through different values of ℓ and m_{pts} , and present the number of clusters found together with the DBCV scores in Figure 2.3. For the ℓ dependence, see Figure 2.3(a), we find a downward trend in the number of clusters as ℓ increases. DBCV scores tend to be larger in the middle of the ℓ range, with the highest score

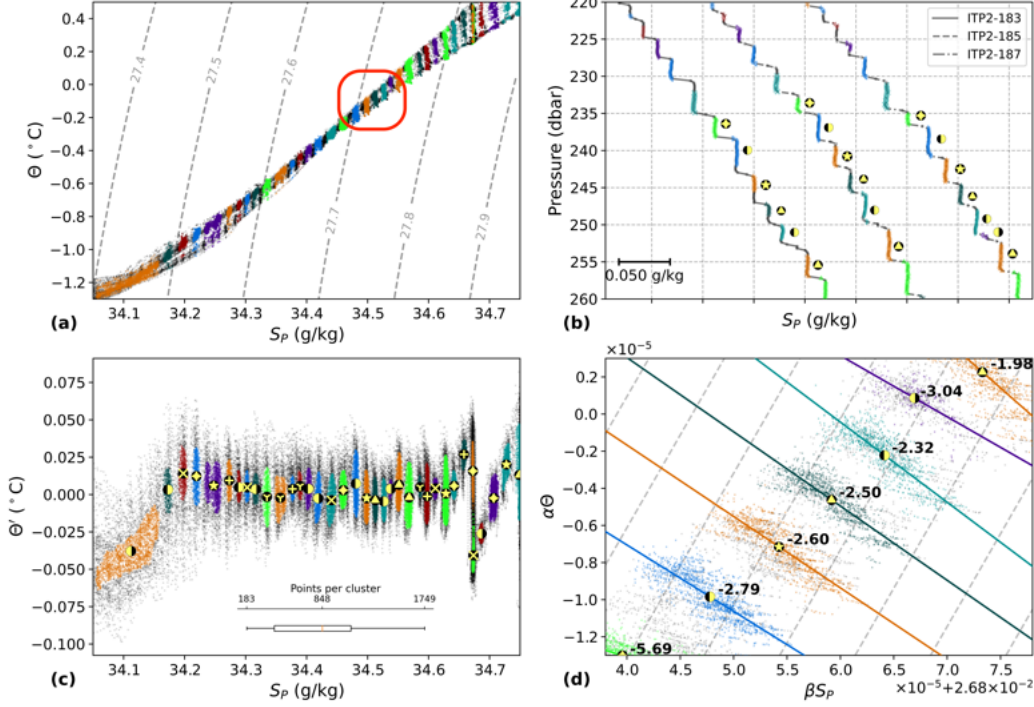


Figure 2.2: Results from the clustering algorithm with $m_{pts} = 170$ and $\ell = 25$ dbar run on 53042 data points in the salinity range 34.05–34.75 g/kg from all up-going ITP2 profiles. (a) The data in Θ – S_P space with dashed lines of constant potential density anomaly (kg m^{-3}) referenced to the surface. The red box bounds the clusters marked in panels (b) and (d). (b) Profiles 183, 185, and 187 from ITP2 in a limited pressure range to show detail. Each profile is offset in S_P for clarity. (c) The spatial arrangement used as input for the algorithm where the gray points are noise and each color-marker combination indicates a cluster. The same color-marker combinations are used in each panel and the markers in panels (c) and (d) are at the cluster average for each axis. (d) A subset of the data in $\alpha\Theta$ – βS_P space with the linear regression line and inverse slope (R_L) noted for each individual cluster and with dashed lines of slope $\alpha\Theta/\beta S_P = 1$.

occurring for $\ell = 25$ dbar. We may also note that the choice of ℓ shapes a large-scale feature of the Θ' – S_P plots. In Figure 2.2(c), a zig-zag pattern of increasing, rapidly decreasing, and then increasing again Θ' is seen in the range of $S_P \approx 34.63$ –34.72 g/kg. This pattern is due to the presence of the AW subsurface temperature maximum in Θ profiles; it disappears when ℓ is small while becoming more exaggerated for larger ℓ (see Section A.2). Based upon previous studies of staircases in the Canada Basin during this time period (Timmermans et al., 2008; Lu et al., 2022), we estimated the typical layer thickness to be 5 m in height, or 5 dbar in pressure, though we found similar results for estimates of 0.5–7.5 dbar (see Section A.2). The choice $\ell = 25$ dbar, where the largest DBCV score occurs, thus corresponds to approximately five times the typical layer thickness. This value is found to be large enough that the staircases are completely smoothed out yet small enough that the features outside the analyzed pressure range do not significantly affect the moving average.

We now turn now to m_{pts} , which under the default settings of HDBSCAN sets the minimum number of points in a cluster (Campello et al., 2013). If the value of m_{pts} is too small, the algorithm may erroneously split a cluster that represents one layer into multiple, smaller clusters, while a too-large value of m_{pts} would lead to the incorrect grouping of multiple discrete layers into a single




Figure 2.3: A parameter sweep showing the number of clusters found (solid lines) and DBCV (dashed lines) in ITP2 as a function of (a) 27 different values of ℓ with $m_{pts} = 170$ and (b) 44 different values of m_{pts} with $\ell = 25$ dbar.

ITP	Study to Reproduce	Input Parameters				Results	
		Number of Profiles	Salinity Range (g/kg)	ℓ (dbar)	m_{pts}	Number of Clusters ^a	DBCV
2	Timmermans et al. (2008)	121	34.05–34.75	25	170	36 (31)	0.3034
3	Lu et al. (2022)	766	34.21–34.82	25	580	43 (40)	0.3862

^aThe first number is the total number of clusters found by the algorithm. The second is the number of clusters that were neither outliers in IR_{SP} nor R_L .

Table 2.2: The values of parameters used to run the clustering algorithm over both datasets.

cluster. Note that this upper bound on a reasonable m_{pts} depends greatly on the number of data points given to the algorithm; the more data points the algorithm is given, the higher the value of m_{pts} can reasonably be set. In the parameter sweep of Figure 2.3(b), we find the number of clusters decreases rapidly until $m_{pts} \approx 60$, then decreases at a much slower rate, while the highest DBCV scores occur for intermediate values of m_{pts} . As with ℓ , we choose the value of m_{pts} having the highest DBCV score. For ITP2, this led to our selection of $m_{pts} = 170$.

Running HDBSCAN on ITP2 using the procedure outlined above with the choices $\ell = 25$ and $m_{pts} = 170$ leads to the clusters presented in Figure 2.2. Following the same parameter selection process for ITP3, we obtain the values $\ell = 25$ dbar and $m_{pts} = 580$ (see Section A.4). Table 2.2 summarizes our input parameter choices and the resulting number of clusters and DBCV values for both ITP2 and ITP3.

2.3.3 Evaluating the clustering algorithm results

The DBCV score gives a measure of quality for the clusters in terms of their densities of points relative to the surrounding noise. However, DBCV does not take into account the properties of the clusters that we expect from the physical situation of staircases, such as their spans in Θ and S_P or how far they are from neighboring clusters. We therefore present two metrics to help predict whether each cluster will accurately represent what we expect from layers within staircases: the lateral density ratio R_L and the normalized inter-cluster range IR .

The relative strength of horizontal variations in salinity and temperature along the i th layer is described by the lateral density ratio

$$R_L^i = \frac{\beta \Delta S_P}{\alpha \Delta \Theta}, \quad (2.2)$$

where $\beta = \rho^{-1} \partial \rho / \partial S_P$ is the haline contraction coefficient, $\alpha = -\rho^{-1} \partial \rho / \partial \Theta$ is the thermal expansion coefficient, and ΔS_P and $\Delta \Theta$ are the variations in salinity and temperature, respectively, along a particular layer (Radko, 2013; Bebieva & Timmermans, 2019). We estimate R_L by finding the inverse slope of the best-fit line through each cluster in $\alpha \Theta$ – βS_P space (see Figure 2.2(d)) (Timmermans et al., 2008; Chen, 1995). These lines are found using Orthogonal Distance Regression, which is more suitable than ordinary least squares in our case due to the presence of variability along both the $\alpha \Theta$ and βS_P axes (Winton, 2011); however, both methods yield similar results (not shown).

R_L quantifies the relative importance of S_P and Θ for the density of that layer (Timmermans et al., 2008; Bebieva & Timmermans, 2019) and is known to be directly related to the ratio of the vertical fluxes of salinity and heat within a staircase (Bebieva & Timmermans, 2019). Note that the lateral density ratio R_L is distinct from the density ratio, R_ρ , which is defined using the same equation (2.2) but with ΔS_P and $\Delta \Theta$ taken in the vertical direction (Shibley et al., 2017). The relative constancy of R_L values across time and space has been interpreted as reflecting the remarkable degree of lateral coherence of staircase layers (Toole et al., 2011). The values of R_L have also been shown to be remarkably similar across neighboring layers (Timmermans et al., 2008). Therefore, if a value of R_L lies significantly outside of the general range of its neighbors, the cluster either reflects the erroneous grouping of multiple layers into a single cluster, or else a physical merging or splitting of layers as discussed in Section 2.4.2.

We also define a measure of the relative spread of a variable, such as temperature or salinity, within a cluster in comparison with the differences between adjacent clusters. Ordering the clusters sequentially in density, the normalized inter-cluster range for the i th cluster is given by

$$IR_v^i = \frac{v_{max}^i - v_{min}^i}{\min(|\bar{v}^i - \bar{v}^{i-1}|, |\bar{v}^i - \bar{v}^{i+1}|)}, \quad (2.3)$$

where $i - 1$ and $i + 1$ denote the adjacent clusters to either side, v is the variable of interest (i.e., pressure, Θ , or S_P), v_{max}^i and v_{min}^i are the maximum and minimum values of the variable v within cluster i , and \bar{v}^i denotes the mean value of v for cluster i . The numerator is the span between the maximum v_{max}^i and minimum v_{min}^i within cluster i . The denominator is the span between the mean of that cluster \bar{v}^i and the mean of either the cluster above or below, whichever is smaller. For clusters at either end of the variable space, we take the denominator to be the span between the mean of the i th cluster and that of its single neighbor.

The inter-cluster range IR_v therefore quantifies the range of a given variable, v , within a cluster in comparison with the range to the nearest neighboring cluster. For a staircase, the salinity values within one layer are generally well separated from the salinity values of the neighboring layers (Lu et al., 2022). Therefore, we expect that the clusters with large IR_{S_P} could represent a part of a single layer that was erroneously divided into multiple clusters by the algorithm, or entirely different physical features, as discussed in Section 2.4.3.

In order to evaluate clusterings, we choose a method to detect outliers in both IR_{S_P} and R_L . We define outliers as points more than two standard deviations from the mean, or, equivalently, with a z-score greater than two (approximately corresponding to a p-value of 0.05 for a two-tailed test). More sophisticated outlier detection methods exist and, while this approach is not guaranteed




Figure 2.4: The average (a) pressure, (b) Θ , and (c) S_P for the points within each cluster for each profile (profile cluster average, PCA) across time. The clustering algorithm was run with $m_{pts} = 580$ and $\ell = 25$ dbar on 678575 data points in the salinity range 34.21–34.82 g/kg from all up-going ITP3 profiles.

to find all erroneous clusters, outliers in IR_{S_P} and R_L give us an indication, based on simple and measurable physical characteristics, of particular clusters which may not represent a single, full layer and which therefore require closer inspection. Although such outliers could be manually adjusted to better capture single, complete layers, we take the approach of disregarding them when calculating statistics and trends.

2.4 Results

Now having presented the HDBSCAN clustering algorithm, our process of selecting ℓ and m_{pts} , and metrics to identify erroneous clustering, we apply this method to data from ITP2 and ITP3.

2.4.1 Properties of detected layers

As a starting point, we examine the average value of pressure, Θ , or S_P for each cluster found from profiles collected by ITP3. Inspired by L22, we plot those values over time in Figure 2.4. We find that the clustering algorithm is capable of tracking individual layers across hundreds of profiles collected along the 2541 km-long track traced by ITP3 over 382 days. A pattern emerges, where the pressures of individual layers appear to have more variability than temperature and salinity, consistent with T08. Moreover, we find that salinity variations within a layer are smaller than the salinity differences between two neighboring layers, while the opposite is true for pressure and temperature (similar to L22; see their Figure 3). We find similar results for ITP2 (see Figure A.7).

More quantitatively, we find differences in salinity between clusters are approximately seven times

	\widehat{IR}_p	\widehat{IR}_Θ	\widehat{IR}_{S_P}	$\widehat{\Delta\Theta}$ (°C)	$\widehat{\Delta S_P}$ (mg/kg)
ITP2	8.394	4.087	0.707	0.048	16.5
ITP3	16.641	7.141	0.927	0.050	15.4

Table 2.3: The median normalized inter-cluster ranges and differences between average values of adjacent clusters for ITP2 and ITP3, calculated after removing outliers with $z\text{-score} > 2$ in the respective variable.

larger than variations within a cluster. That is, we compute the standard deviation of salinity within each cluster and find that the median is 2.2 mg/kg, while the median absolute difference between the average salinity of adjacent clusters is 15.4 mg/kg. This is in agreement with L22 who found inter-layer salinity differences to be an order or magnitude larger than variations within a layer. Furthermore, the median normalized inter-cluster ranges (Equation 2.3) in Table 2.3 quantitatively confirm the qualitative patterns noted by T08 and L22 as $\widehat{IR}_p, \widehat{IR}_\Theta > 1$ and $\widehat{IR}_{S_P} < 1$ for both ITPs, where $\widehat{\cdot}$ indicates the median. Table 2.3 also contains the median of the differences between neighboring cluster averages of temperature, $\widehat{\Delta\Theta}$, and salinity, $\widehat{\Delta S_P}$, for both ITPs. For ITP3, these values of 0.048°C and 15.4 mg/kg match those of L22, which are 0.05°C and 17 mg/kg. They also agree with T08, who reported the the difference in temperature and salinity across interfaces to be $\delta\theta \approx 0.04^\circ\text{C}$ and $\delta S \approx 14$ mg/kg.

2.4.2 Outliers and splitting / merging layers

As discussed previously in Section 2.3.3, outliers have been identified in the inter-cluster salinity range IR_{S_P} as well as the lateral density ratio R_L . Figure 2.5 shows IR_{S_P} and R_L for both ITPs with outliers indicated by red circles. We find that these outliers can either be due to erroneous clustering or indicate the presence of different physical features such as the splitting or merging of thermohaline staircase layers, which is a well-known phenomenon (Neshyba et al., 1972; Padman & Dillon, 1988; Kimura et al., 2015).

We can learn more about these features from the illustrative sets of particular salinity profiles from ITP2 presented in Figure 2.6. Both sets of profiles span less than a week and we narrow the displayed pressure ranges so that the individual steps are visible. This figure shows that overall, the algorithm captures the layered structure very well, marking points within interfaces as noise. Nevertheless, it is imperfect. Near layer boundaries, the algorithm sometimes includes points from an interface within a cluster, and sometimes neglects to include points within a layer. Additionally, as seen in the first profile of Figure 2.6(b) around 236 dbar, the algorithm can also miss layers entirely, especially when the layer is particularly thin and only present in a small number of profiles. Occasional issues such as these are to be expected with any automated detection method.

Focusing on ITP2 as an example, we find two outliers in IR_{S_P} marked by an orange 4-pointed star and a green “×” in Figure 2.5(a). We can track the same clusters in Figure 2.2(c), indicated by orange or green dots with a 4-pointed star or a “×” at the center of the cluster. They both have an average $S_P \approx 34.67$ g/kg but are separated in Θ' . A series of individual salinity profiles associated with these outliers over approximately six days is shown in Figure 2.6(a). We find that these outliers correspond to a relatively thick single layer that was erroneously split across two clusters by the algorithm. We attribute this erroneous splitting to the zig-zag pattern in Θ' - S_P space mentioned in Section 2.3.2 and note that it could be eliminated with a different selection of

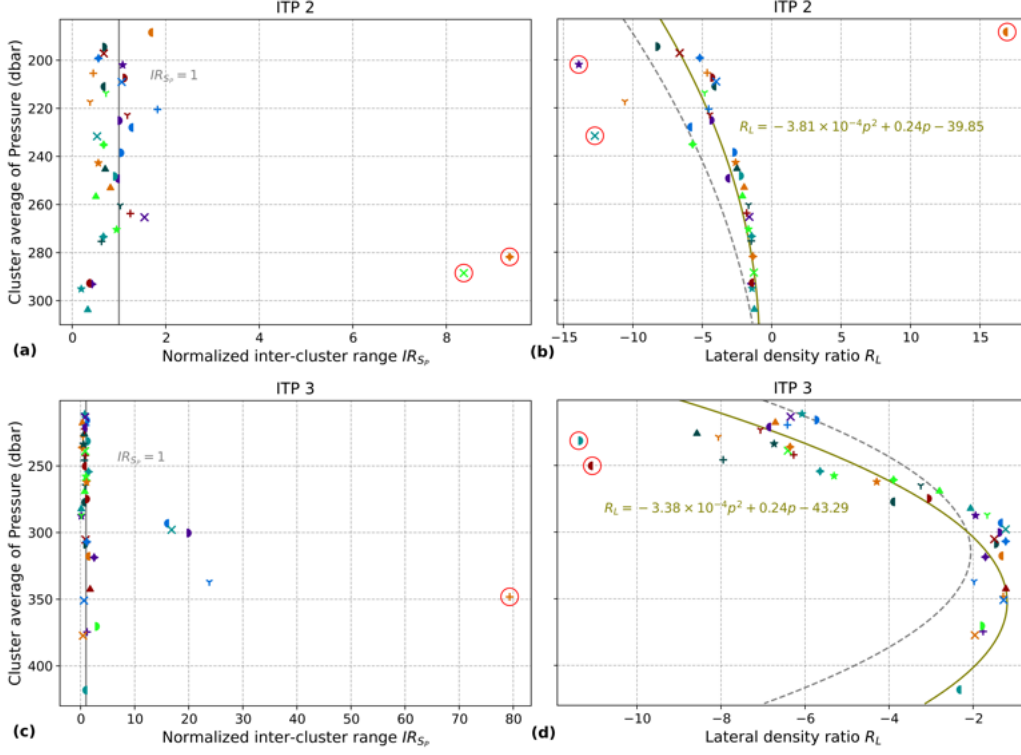


Figure 2.5: The value of each cluster’s normalized inter-cluster range for salinity IR_{S_p} in (a) and (c) and the lateral density ratio R_L in (b) and (d) as a function of the cluster’s average pressure. The colors and markers for ITP2 in (a) and (b) are the same as the clustering shown in Figure 2.2 and for ITP3 in (c) and (d), they are the same as shown in Figure 2.4. Markers circled in red indicate outliers with a z-score greater than 2. In (b), the solid curve is a 2nd-degree polynomial fit (equation given by the annotation) for the non-outlier points from ITP2 and the dashed curve is the same for ITP3. In (d), the solid curve is a 2nd-degree polynomial fit (equation given by the annotation) for the non-outlier points from ITP3 and the dashed curve is the same for ITP2.

ℓ . Although we can attribute these outliers in IR_{S_p} to an artifact of the method, there are other instances where such outliers indicate the presence of physical features, as we will discuss in Section 2.4.3.

Next, we examine outliers in R_L , defined in Equation 2.2, for ITP2, marked by red circles in Figure 2.5(b). These outliers correspond to clusters that appear to have multiple layers grouped together. For example, the outlier cluster marked by an orange half-circle in Figure 2.5(b) can be seen in Figure 2.2(c) spanning $S_P = 34.054$ – 34.159 , a much wider range than any other cluster. Similarly, in Figure 2.2(c), the outlier cluster marked by a purple star spans $S_P = 34.233$ – 34.261 and clearly encompasses what should be two distinct clusters. Some outliers in R_L are the result of erroneously clustering multiple layers together. On the other hand, the particular feature of the outlier cluster marked by a teal “x”, centered around 232 dbar with $R_L = -12.7$ in Figure 2.5(b), indicates splitting or merging. As highlighted in Figure 2.6(b), it typically spans multiple stair steps, but the last few profiles it only spans a single, larger step. This illustrates that we can use outliers in R_L to identify clusters that are not single, complete layers and to find instances of potential splitting or merging.




Figure 2.6: Individual S_P profiles from ITP2, specifically chosen to show the examples of outlier clusters in IR_{S_P} and R_L highlighted by the bands of color. (a) Profiles 67, 69, 73, 75, 81, 83, 91, 93, 97, and 99, collected between August 31–September 5, 2004. (b) Profiles 87, 89, 95, 97, 99, 101, 103, 105, 109, and 111, collected between September 3–7, 2004. The colors and markers are the same as the clustering shown in Figure 2.2. The gray dots are noise points and the black lines show the profiles. Each profile is offset in S_P for clarity.

2.4.3 Remnant intrusions revealed by temperature

Direct comparisons with other layer detection methods can also be used to identify interesting physical features. To demonstrate this, here we directly compare average layer characteristics computed by HDBSCAN to those reported by L22, both using data from ITP3. L22 used 758 profiles from ITP3 while we used all 766 available up-going profiles. Based on the gaps in Figure 3 of L22, we believe that the eight missing up-going profiles are from July 2006. Figure 2.7 shows the average Θ and S_P for each cluster found in our study and by L22 based on the values in their Table A1 after converting θ to Θ . We initially find 43 clusters and, after eliminating outliers in IR_{S_P} and R_L as described in Section 2.3.3, we find 40 clusters. While L22 identified only 34 thermohaline layers, we find close agreement between those and the clusters we found for $S_P \lesssim 34.74$ g/kg. Below where this salinity occurs in the water column, we find five more clusters than L22.

The differences between these two results appears to be related to the presence of *remnant intrusions*, which display features of both staircase layers and intrusions, and are thought to represent an intermediate stage in staircase formation. Such features are known to appear near the bottom of staircases around the AW core and have been analyzed in detail by [Bebieva & Timmermans \(2019\)](#). They are characterized by homogeneous salinity, which leads the method of L22 to treat them as single layers. However, they have a temperature structure that is inverted (warmer above colder) compared to the typical gradient within a thermohaline staircase. The warm and cold sections are distinct enough for the clustering algorithm to split the structure into multiple clusters, each of which is homogeneous in salinity and relatively homogeneous in temperature. Figure 2.8 highlights an example of a remnant intrusion where the method of L22 gives results that differ from ours;




Figure 2.7: The average Θ and S_P for the 34 layers found by Lu et al. (2022) in black dots and for the 43 layers found in our study using data from ITP3 and the same colors and markers as in Figure 2.4. Clusters circled in red are outliers in either IR_{S_P} or R_L .

the layer spanning approximately 320–340 dbar is constant in S_P , but decreases in Θ with depth and is divided into two clusters using the clustering method. These two clusters can be found in Figure 2.5(c), indicated by the blue “Y” centered around 337 dbar with $IR_{S_P} = 23.8$ and the orange “+” centered around 348 dbar with $IR_{S_P} = 79.3$.

Such remnant intrusions exist between the active intrusions of the AW core and the staircase layers of the lower halocline waters. Staircases may be formed from intrusions but such features have distinctly different patterns of heat and salt flux than double-diffusively driven staircases (Bebieva & Timmermans, 2019). The disagreement between our results and those of L22 highlights the difficulty of detecting remnant intrusions. While neither method is designed to automatically distinguish between these and staircase layers, the method presented here offers the opportunity to identify them when evaluating outlier clusters. Moreover, although Table 2.3 supports the suggestion by L22 that salinity is the most appropriate variable by which to identify staircase layers, having shown that a layer identified by L22 is in fact a remnant intrusion illustrates why it remains important to consider temperature as well.




Figure 2.8: Individual profiles 313, 315, 317, 319, and 321 from ITP3, collected between November 10–12, 2005, specifically chosen to show the example of a temperature inversion highlighted by the bands of color. The colors and markers of the individual points are the same as the clustering shown in Figures 2.4 and 2.7. The gray dots are noise points. The black lines on the left of each pair are the S_P profiles while the red lines on the right are for Θ . Each profile is offset in both S_P and Θ for clarity.

2.4.4 Dependence of R_L on pressure

Lastly, we revisit Figure 2.5(b,d) where the distribution of points suggest a dependence of R_L on pressure. This is in contrast to previous studies which have found a constant R_L for the depth range we analyze (Timmermans et al., 2008; Toole et al., 2011; Bebieva & Timmermans, 2019). For ITP2 in panel (b), after removing the five outliers indicated across panels (a) and (b), we use a second-degree polynomial fit to find $R_L = -3.81 \times 10^{-4}p^2 + 0.24p - 39.85$, where p is pressure, with a coefficient of determination $R^2 = 0.67$. For ITP3 in panel (d), we remove the three outliers in panels (c) and (d) then find quite the similar dependence of $R_L = -3.38 \times 10^{-4}p^2 + 0.24p - 43.29$ with a coefficient of determination $R^2 = 0.84$. Comparing these two curves in Figure 2.5(b,d), we find most of the difference can be explained by a downward shift of roughly 20 dbar in the upper water column to 50 dbar in the lower water column from ITP2 to ITP3. Note that a second-degree polynomial was chosen here as a simple parametric model to capture the apparent non-linear dependence of that R_L on pressure.

By contrast, T08 concluded there was no vertical dependence of R_L . They analyzed data from ITP1 through ITP6, which sampled the Canada Basin during the period from 2004 to 2007, and found a constant value of $R_L = -3.7 \pm 0.9$. Their Figure 6(a) shows five values of R_L for ITP2 that range from -3.5 to -3.0. While Bebieva & Timmermans (2019) found that R_L changes below the depth of the temperature maximum, those authors also found in their Figure 3(b) that R_L is constant in the depth range we consider in this study. Additionally, the presence of remnant intrusions does not explain the difference in our results as the pressure dependence of R_L is evident in the upper part of the water column where they are absent. Although we find that R_L depends on pressure, our results agree with those of T08 on the magnitude of R_L , as we find the mean value for all non-outlier clusters of ITP2 to be -3.55 with a standard deviation of 2.24, and so we compare more of our results to those of T08.

We reproduce several of the figures from T08 in Figure 2.2. The combinations of color and markers for the clusters are the same in Figures 2.2, 2.5, and 2.6. In Figure 2.2(a), we plot data from ITP2 in Θ - S_P space, reproducing Figure 5(a) from T08. Those authors noted that points from each particular layer clustered along lines in Θ - S_P space that cross isopycnals and when we mark the clusters found by the algorithm, those layers become visually distinct. Figure 2.2(b) shows

profile 185 in the same depth range as Figure 4 from T08, plus the up-going profiles taken by ITP2 immediately before and after. This illustrates that the clustering algorithm is indeed marking points within those particular layers as in the same cluster and points within interfaces as noise. It also shows that the clustering algorithm tracks the same layers across neighboring profiles even though the pressure at which those layers are found varies. Figure 2.2(c) shows the ITP2 data on the axes used by the clustering algorithm. In Figure 2.2(d), we plot the clusters bounded by the box in Figure 2.2(a) in $\alpha\Theta-\beta S_P$ space, reproducing Figure 6(a) from T08. The dashed lines of slope 1 correspond to isopycnals as $\Delta\rho = -\alpha\Delta\Theta + \beta S_P$. Overall, the panels in Figure 2.2 are close visual matches for those particular figures in T08.

After examining the differences between the two in detail (see Section A.6), it remains unclear why the values of R_L we found for clusters in ITP2 data differ from T08. However, we believe that the quantitative agreement in magnitude for the values of R_L and the qualitative match between the clustering results and the features in Figure 2.2 show that the clustering method reproduces those results from T08. Overall, these findings suggest that the clustering algorithm is indeed revealing a dependence of R_L on pressure in the data we analyzed from ITP2 and ITP3. Since R_L is directly related to the ratio of the vertical derivatives of the vertical fluxes of salinity and heat within a staircase (Bebieva & Timmermans, 2019), this could indicate a pressure dependence of that ratio.

2.5 Discussion

In this study, we have presented a method based on the HDBSCAN clustering algorithm to both detect and connect well-mixed layers in thermohaline staircases across Arctic Ocean hydrographic data. HDBSCAN has previously been successfully applied to sets with millions of data points (Logan & Fotopoulou, 2020). As the comparison with results from previous studies was favorable, this suggests that this method may produce trustworthy results when applied to datasets with larger numbers of points. This study contributes to a growing set of examples of clustering algorithms being used in oceanography (e.g., Koszalka & LaCasce, 2010; Espejo et al., 2014; Maze et al., 2017; C. Ma et al., 2019; Jones et al., 2019; Houghton & Wilson, 2020; Rosso et al., 2020; Desbruyères et al., 2021; Sonnewald et al., 2021). The continued and extended use of clustering algorithms in oceanography and related fields is an important development, since discovering and detecting features in all manner of datasets becomes more challenging as they grow ever larger.

Special attention was given herein to the identification of outliers using parameters output by the method itself. We found that clusters which are outliers in the lateral density ratio R_L often indicate multiple layers that are erroneously clustered together, but can also highlight potential instances of layer splitting or merging. By introducing the normalized inter-cluster range IR_v , we quantitatively showed that the pressure and temperature values vary more within a layer than the difference between the values of neighboring layers, while the opposite is true for salinity. Because we know that the practical salinity S_P values in a particular layer are well separated from those of neighboring layers, we identified clusters with notably large values of IR_{S_P} as likely to be either only part of a layer that was erroneously split by the algorithm or a remnant intrusion. This study also suggests that there may be a pressure dependence of the lateral density ratio R_L , seen in both ITP2 and ITP3. The reasons for this dependence are unclear and are worthy of further study. Following the model presented by Bebieva & Timmermans (2019), the ratio of vertical fluxes of heat

and salt could be recalculated using this more complex vertical dependence of R_L . This could then be compared to simulations (Yang et al., 2022) and parameterizations (Radko et al., 2014; Shibley & Timmermans, 2019) of the flux ratio in other models. However, in order to verify these with observations, vertical resolution fine enough to resolve the interfaces between layers, that is, finer than provided by ITPs, would be needed.

The method we presented here has certain advantages in particular circumstances. For example, all of the other methods referred to herein (Timmermans et al., 2008; Shibley et al., 2017; van der Boog, Otto Koetsier, et al., 2021; Lu et al., 2022) require setting one or multiple thresholds on gradients in temperature, salinity, or density which are used to identify sections of profiles where layers may be present. However, choosing reasonable values for these thresholds requires precise prior knowledge of the staircase properties, which may not be available before identifying staircases in that particular region and time period. The method presented here requires the selection of a salinity range in which to search and an estimate of the typical layer thickness which can be determined from a brief look at a dataset or from previous studies. To evaluate the clusters, a selection of an outlier detection method is required. In addition, while the other techniques require resolution high enough to resolve small-scale vertical gradients, the method presented here—which does not distinguish between individual profiles—only requires that a sufficient number of data points be available in order to detect clusters. Therefore, this technique could potentially find staircases in datasets that have too low a resolution to resolve the steps in any particular profile. While, in this study, we use all data points available after filtering, we found that subsampling profiles to every second, third, or fourth point yielded similar results (see Figure A.8). Finally, the method automatically connects layers across profiles. While L22 used an automated connection method, it still required subsequent adjustments to the layers to be made manually. Consequently the method presented here could be readily applied to datasets with larger numbers of points and used to examine the large-scale, lateral properties of coherent layers in thermohaline staircases.

The clustering algorithm has several limitations. It cannot be used to examine sets of profiles spanning temporal or spatial scales larger than those on which stairs are known to be coherent. This method considers all profiles simultaneously, and it is not applicable to identifying stair steps in datasets of independent individual profiles. While this method could be applied on a profile-by-profile basis, each profile would require a separate selection of m_{pts} which itself becomes highly sensitive with a small number of points to cluster. Moreover, as the clustering algorithm does not consider profiles individually, it may miss layers which are only present in a small number of profiles, as may occur especially for layers that split or merge. Finally, although the clustering method captures the overall structure of the staircase layers well, it sometimes miscategorizes points, especially near the boundaries between layers and interfaces.

The method presented here could be used in conjunction with other staircase detection methods. For example, if a different detection method were used to identify which data points in a collection of profiles are in layers, the clustering algorithm could be run on just the layer points to automatically connect the layers across profiles. Additionally, the clustering algorithm could be used on a large collection of datasets to identify which subsets contain staircases; then a more specifically tuned staircase detection method could be used on just that subset, reducing the amount of time-consuming analysis.

Because the method can be scaled to datasets with a larger number of points, a natural extension

would be to apply it to a dataset with profiles from many ITPs. Furthermore, although we have focused on identifying upper-ocean Arctic thermohaline staircases, it could also be useful in other oceanic regimes. Staircases with steps on the order of 50 m thick have been observed in the Arctic between depths of 2000–3000 m (Timmermans et al., 2003), much deeper than ITPs measure, while the Argo network of autonomous profiling floats has captured staircases in regions all over the world (van der Boog, Dijkstra, et al., 2021). Since this method does not assume a consistent vertical resolution across the profiles, it could be run on a mix of data from different types of instruments, assuming they operated within the same region and time period. This method could also be adjusted to identify other structures that appear in Θ – S space such as different types of layers or water masses, or even structures that can appear in different spaces such as Θ – O_2 (Rosso et al., 2020). On the technical side, another topic of future study would be to adapt this method to automatically distinguish between well-mixed layers and intrusions. However, when using this clustering algorithm to specifically search for staircase layers, we recommend avoiding the bottom of the thermocline around the AW core where remnant intrusions are known to appear (Bebieva & Timmermans, 2019). Finally, there may be methods other than the DBCV validation process used herein that may be better able to guide the choice of parameters, potentially improving the detection fidelity.

Chapter 3

Tracking Beaufort Gyre staircase layers across 17 years

Abstract In the Beaufort Gyre Region of the Arctic Ocean, thermohaline staircases exist in the pycnocline with water from the Pacific Ocean above and from the Atlantic below. Observations show that both the Pacific and Atlantic waters in this region have warmed significantly in the past several decades. However, to the best of my knowledge, no studies have investigated the evolution of individual steps in thermohaline staircases for a period longer than four years. Here, I use an unsupervised clustering algorithm to identify staircase layers in 17 years of hydrographic profiles from 2005–2022. I find that these layers are sinking at a rate of 2.24 ± 1.97 dbar per year, a rate which matches the net downwelling for the Beaufort Gyre, suggesting that the layers passively follow the evolution of isopycnals. I also find that the layers are warming at a rate on the order of 0.1°C per decade and that there is a robust trend of this rate decreasing from the top to the bottom of the staircase. From this, I calculate the average net heat flux for each layer which is positive in the upper part of the staircase and negative near the bottom, summing to a net 0.054 W/m^2 , matching the magnitude of estimates from previous studies. I find the average layer thickness to be 1.55 dbar and that thicknesses increase with depth at a rate of $(2.53 \pm 0.40) \times 10^{-2}$ dbar per dbar, which are also consistent with previous studies. However, I also find that the thickness of layers are likely either slightly thinning or not changing in time. While thermohaline staircases have been known to be present in the Arctic Ocean for over 50 years, these results show that individual staircase layers can remain coherent over long periods and how they change in time. Replication code can be found at https://github.com/scheemik/Staircase_Clustering_Detection_Algorithm.v2.

3.1 Introduction

As was introduced in Section 1.1, the Canada Basin of the Arctic Ocean contains a layer of warm, salty water from the Atlantic Ocean which occupies depths between 250–800 m (Timmermans et al., 2008). This mass of Atlantic Water (AW) is the most significant reservoir of heat in the Arctic Ocean (Richards et al., 2022) and, as many have noted, contains enough heat to melt all Arctic sea ice if it were somehow able to reach the surface (Maykut & Untersteiner, 1971). It has also been warming significantly since the mid-twentieth century (Carmack et al., 1997; Polyakov, Pnyushkov,

& Timokhov, 2012). Above the AW, and above the pycnocline, exists a layer of water originating from the Pacific Ocean, typically between depths of 30–100 m (Timmermans et al., 2018). The heat contained in this layer is also significant, estimated to be able to melt around a meter of sea ice if it were able to rise to the surface (Timmermans et al., 2018; MacKinnon et al., 2021). The Pacific water is warming as well, with its heat content doubling from 1987 to 2017 (Timmermans et al., 2018; Dosser et al., 2021).

Between these two water masses, lie the thermohaline staircases. As mentioned in Section 1.1, staircases have appeared in profiles of the pycnocline for decades including those taken at Ice Island T-3 (1969–1971) (Neshyba et al., 1971, 1972; Neal & Neshyba, 1973), during AIWEX (1985) (Padman & Dillon, 1987, 1988, 1989), and during SHEBA (1997–1998) (Shaw & Stanton, 2014).

Many studies have commented on this persistent presence of staircases in the Arctic Ocean. From a set of 12 hydrographic profiles spanning 1990–2001, Timmermans et al. (2003) found staircases to be persistent feature of the deep Canada Basin (below 2400 m). Both Shibley et al. (2017) and van der Boog, Dijkstra, et al. (2021) show the presence or absence of staircases in each ITP profile taken in the Arctic between 2004–2013 and 2004–2020, respectively. However, these studies do not connect specific, individual layers across different profiles, either because such an effort was irrelevant to the study or because there was a lack of temporal resolution (Timmermans et al., 2003).

Several studies have noted the remarkable lateral coherence of layers. Using data from ITPs 1–6, Timmermans et al. (2008) argued that clusters which appear in Temperature–Salinity indicate that individual layers appear across the profiles analyzed which span approximately 800 km. A similar argument was made by Bebieva & Timmermans (2019). However, neither study analyzed how specific layers changed over time.

While there have been studies that have tracked the evolution of specific layers in the Arctic Ocean, all have only been across relatively short time spans. From a 30-hour series of temperature profiles taken every four minutes at Ice Island T-3, Neshyba et al. (1972) were able to trace specific layers and analyze their fine-scale temporal evolution. Using temperature profiles from AIWEX, Padman & Dillon (1988) were able to track layers along a 20-hour time series of 96 profiles spanning 600 m. However, for a set of 25 profiles spanning 9 days, they were unable to track individual layers due to the temperature variation between profiles. Polyakov, Pnyushkov, Rembe, et al. (2012) used high-resolution profiles from moorings to track six interfaces between staircase layers in the pycnocline across a year-long period between 2003–2004. Similarly, Ménesguen et al. (2022) used moorings to trace layers across a year-long period from 2006–2007. Lu et al. (2022) analyzed thousands of profiles from ITPs between 2005–2009, tracking 34 layers throughout that four year period. They also found the average depth of staircases to be between 200–350 m and noted that this is similar to the average depths Neal et al. (1969) found in 1969 (220–340 m), suggesting that the staircase layer depth range may have remained stable across those decades.

Here, I use a clustering method based on HDBSCAN, which I introduced in Section 2.3.1, to automatically detect thermohaline layers in the Canada Basin of the Arctic Ocean and track their evolution across 17 years of hydrographic profiles. Being able to analyze the evolution of staircase layers in the Arctic Ocean is key as stratification of this region is not well-represented in global climate models (Timmermans & Pickart, 2023). This chapter is organized as follows. First, in Section 3.2, I describe how I selected and filtered the data for this study. In Section 3.3, I analyze the top and bottom boundaries I define for the section of the water column I analyze, introducing

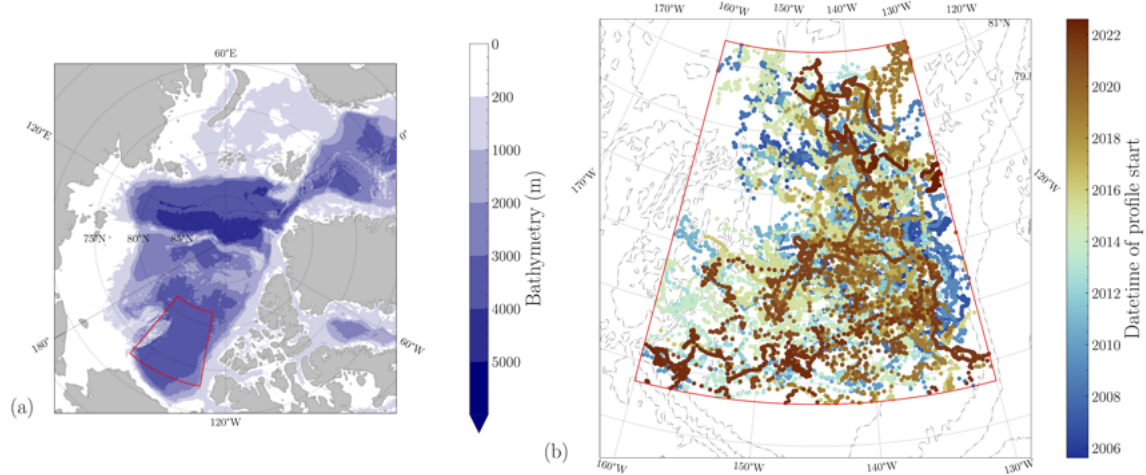


Figure 3.1: (a) A map of the Arctic with a red bounding box designating the Beaufort Gyre Region (BGR) as defined by ST22. (b) A zoomed-in view of the BGR with each of the 15,202 profiles used in this study marked in a color indicating the date on which it was taken. The dotted-dashed lines indicate the bathymetry contours.

the methods I use to correct for spatial variations and find trends in time. I then apply HDBSCAN to year-long time periods and connect layers across time. With this dataset of staircase layers, in Section 3.4 I present results of the properties of the detected layers and how they evolve in time. Finally, in Section 3.5, I discuss the implications of my results, evaluating the performance of the methods used, and offer possible directions for future studies.

3.2 Data

In this study, I analyze ITP profiles within the Beaufort Gyre Region (BGR) which, following Shibley & Timmermans (2022) (hereafter referred to as ST22), I define to span 130–160°W and 73–81.5°N (see the red box in Figure 3.1). The maximum distance any two points could span in this region is 1168 km. To avoid the wake problem discussed in Section 1.2, I keep only the up-going profiles.

The first ITP to take measurements in the BGR was ITP2 on 2004/8/19. However, ITP2 was only in operation for 40 days leaving a gap of 320 days until ITP1 took its first measurement in the BGR on 2005/8/15. Since the deployment of new ITPs has generally occurred around August 15th of each year (see Figure B.1), I choose to divide the time series into year-long periods, each starting and ending on August 15th. At the time of analysis, there was insufficient data available after 2022/8/15, so I choose to analyze only the profiles within the 17 years between 2005/8/15 and 2022/8/15.

Similar to ST22, I filter out profiles that do not cover a sufficient depth range, keeping only profiles which extend past 400 dbar. This significantly reduced the number of profiles to analyze because some ITPs employed a particular repeating pattern of observation, taking a number of shallow profiles, then a deep profile, then taking a number of shallow profiles again. For example, ITPs 105, 113, and 114 had approximately every 1 in 8 profiles extending down to around 800 dbar while the others stopped around 250 dbar. Figure 3.1 shows the locations of the 15,202 profiles analyzed in this study that are within the BGR and the chosen 17-year period, after filtering out

profiles which do not extend deeper than 400 dbar.

As discussed in Section 2.2, the clustering method of finding staircase layers should not be used on a dataset which spans larger temporal or spatial scales than those which staircase layers are known to be coherent. As I will discuss in Section 3.5, this study establishes new maximum known coherence scales for these layers. However, as detailed in Section 3.3.3, I only perform clustering on subsets of the data that respect the previously-known coherence scales.

Similar to Section 2.2, I choose to use slightly different variables than the studies to which I compare my results, notably ST22. Here, I work with the directly measured quantity of pressure, as opposed to depth which is derived from pressure. Again, the conversion between the two is linear, and does not affect my results or comparisons. As recommended by TEOS-10 (McDougall & Barker, 2011), I use conservative temperature Θ and absolute salinity S_A , as opposed to potential temperature θ and practical salinity S_P . As the results from Chapter 2 did not change significantly using S_P vs. S_A (see Section A.1), it is reasonable to expect the results I present here would not change significantly if using different versions of the variables. After selecting the set of profiles and variables I will use, I then define the bounds of the section of the water column in which I will detect staircase layers.

3.3 Methods

3.3.1 Defining the bounds of the water column

Only certain depths of the water column contain staircases. Lu et al. (2022) (hereafter referred to as L22) noted that this section of the pycnocline is bounded above by the Lower Halocline Water (LHW) and below by the Atlantic Water (AW), the cores of each defined to be, respectively, the depths at which salinity is 34.1 and at which the subsurface temperature is maximum. These are similar to the bounds used by ST22 who noted that some care must be taken to avoid selecting erroneous subsurface temperature maximums near the surface, and so I first mask out all temperature values that occur at pressures less than 100 dbar.

With the data that remains, for each profile, I set the AW core to $p(\Theta_{max})$, the pressure at which the temperature is maximum, and the LHW core to $p(S_A \approx 34.1)$, the pressure at which the salinity is closest to 34.1. When searching for the LHW core, I disregard pressures below the AW core to ensure the LHW core is located physically higher in the water column. When detecting and analyzing the staircase layers, I use only the data whose pressures occur between $p(S_A \approx 34.1) < p < p(\Theta_{max})$.

3.3.2 Correcting for spatial distributions

Figure 3.2 shows maps of the values of $p(S_A \approx 34.1)$ and $p(\Theta_{max})$, the pressure contours of the LHW and AW cores, for the 15,247 profiles within the BGR and my chosen 17-year period. Both contain a similar bowl-shaped geographic distribution to that noted by L22. Similar plots for salinity and temperature are in Figure B.2 and, for the LHW core, there is obviously no spatial distribution in salinity nor is there any discernible spatial pattern in temperature. In the AW core, both salinity and temperature increase from the southeast to northwest, again matching the patterns seen by L22. Because the clustering method described in Section 2.3 takes only temperature and salinity

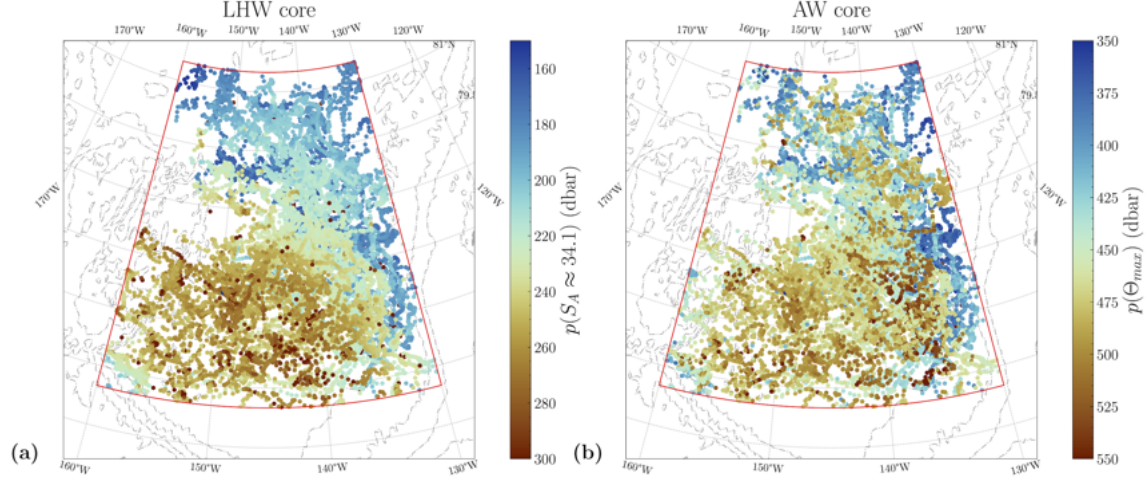


Figure 3.2: Maps showing the pressures at which the (a) LHW and (b) AW cores occur for each of the 15,247 profiles in the BGR. The color bars are oriented such that values physically higher in the water column are closer to the top of the plot.

into account, the spatial distributions in pressure will not affect its results. They may, however, affect the calculation of temporal trends of the clusters.

The BGR has seasonal wind forcing due to ice cover, causing downwelling in the summer and upwelling in the winter (Timmermans & Toole, 2023). However, over the past few decades, these oscillations coexisted with a cumulative downwelling of around 2–5 m/yr (Meneghello et al., 2018; Proshutinsky et al., 2019; Timmermans & Toole, 2023). The spatial distributions in Figure 3.2 span hundreds of dbar with 17 years of data. Because the overall trend in time is orders of magnitude smaller than the trends across space, I first perform a correction for the horizontal variations, implicitly assuming the general spatial distributions do not change in time.

To correct for spatial distributions, I fit a two-dimensional polynomial over latitude and longitude for a given variable v (e.g., p , S_A , Θ , etc.) using the least squares function from the linear algebra module of the Numpy Python library. In order to capture the large-scale features of the spatial distribution in Figure 3.2 while remaining relatively simple, I compute the 10 coefficients up to third power in latitude (ϕ) and longitude (λ), fitting the general formula

$$v_{fit} = A + B\lambda + C\lambda^2 + D\lambda^3 + E\phi + F\lambda\phi + G\lambda^2\phi + H\phi^2 + I\lambda\phi^2 + J\phi^3. \quad (3.1)$$

I hereafter refer to this process as finding the `polyfit2d` for a set of data.

In order to correct the trends in pressure across time, I found the `polyfit2d` on pressure for the LHW and AW cores. Figure 3.3 panels (a,e) show histograms of the pressures of the LHW and AW cores, respectively. Panels (b,f) show the same distribution of pressures for each core as seen in Figure 3.2, but on top of the `polyfit2d` heat maps, the coefficients for which are in Table 3.1. Both use the same color maps such that the fit can be evaluated by noting the difference in color between the heat maps and the overlying data points. Evaluations can also be made with the maps of residuals in panels (c,g) which are the differences for each point between its value and the value of the `polyfit2d` at the same latitude and longitude. Panels (d,h) show the histograms of the residuals for the LHW and AW, respectively. The limits on the colorbars in (c,g) and the vertical axes in

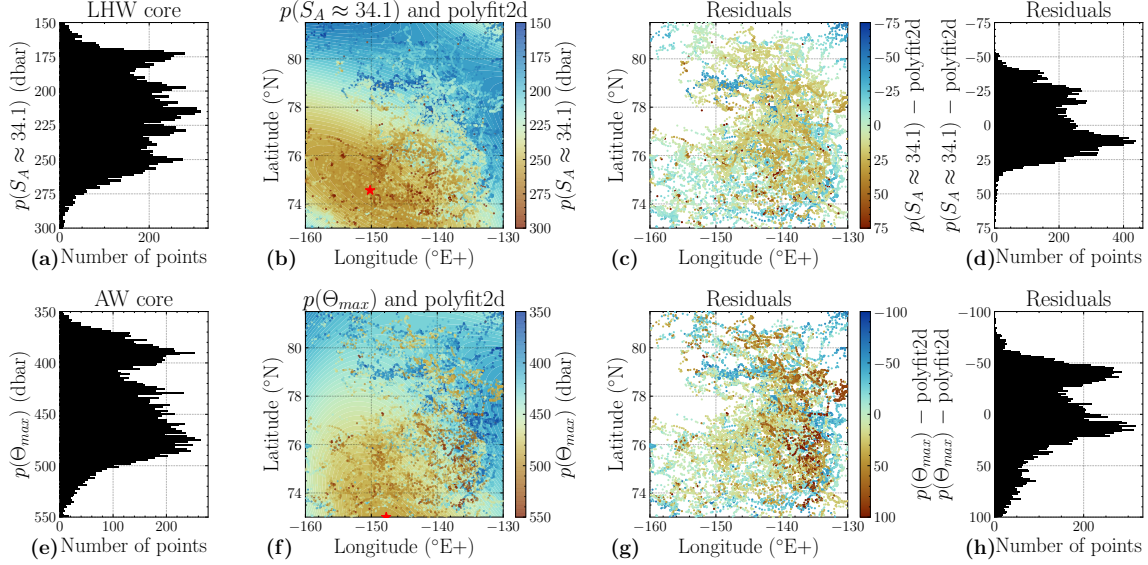


Figure 3.3: The top row shows the LHW core and the bottom row shows the AW core. (a,e) Histograms of the pressures. (b,f) Maps of the pressure values of the core for each profile overlaying a heat map of the `polyfit2d` of the pressure. The red stars mark the location of the highest pressure value in the `polyfit2d`. (c,g) Maps of the residuals, pressure values minus the `polyfit2d`. (d,h) Histograms of the residuals. All vertical axes and color bars are oriented such that values physically higher in the water column are closer to the top of the plot.

Core	A	B	C	D	E	F	G	H	I	J
LHW	-4.67×10^4	220.46	-2.26	1.88×10^{-3}	2184.32	-14.60	0.04	-41.59	0.17	0.28
AW	4.76×10^4	181.66	0.35	6.79×10^{-3}	-1527.95	-4.02	0.03	15.95	0.09	-0.01

Table 3.1: The coefficients for the `polyfit2d` of pressure, following Equation (3.1).

(d,h) have been chosen such that they cover the same number of dbar as the corresponding panels in (a,b,e,f), as well as oriented such that the values towards the top of the plots represent values that are physically higher in the water column (e.g., a residual of -60 is 60 dbar above the `polyfit2d` at the same latitude and longitude).

L22 found that the deepest region of this shape corresponds to the center of the Beaufort Gyre, which they defined as 74.3°N , 144.2°W . Within the BGR, I found the deepest part of these `polyfit2d` spatial distributions to be 74.6°N , 150.2°W for the LHW and 73.0°N , 147.7°W for the AW, as marked by red stars in Figure 3.3(b,f). These are, respectively, 181.86 km and 181.32 km from the center defined by L22.

As can be seen in Figure 3.3(c,g), the residuals have little to no discernible spatial distributions. Between panels (a) and (d), the standard deviation of the pressures of the LHW core went from 38.32 dbar to 30.54 dbar while between (e) and (h), the standard deviation of the AW went from 45.19 dbar to 40.98 dbar. These reductions in the standard deviations and the lack of a spatial pattern in the residuals implies that the `polyfit2d` is indeed removing the overall spatial distributions in pressure.

Figure 3.4 shows the trends in pressure and residuals from the `polyfit2d` on pressure for both the LHW and AW cores. I calculated the trends using the ordinary least-squares regression. It is evident that the variation in pressure decreases when finding the residuals. The slope of the LHW

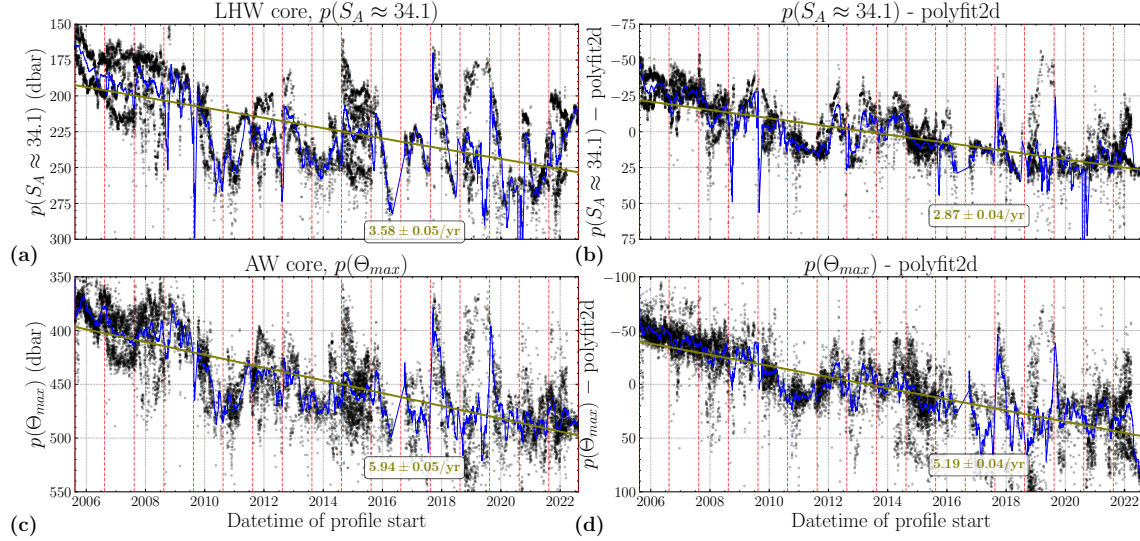


Figure 3.4: Trends in pressure over time for the (a,b) LHW and (c,d) AW cores. Panels (a,c) show the trends before and panels (b,d) show the trends after subtracting their respective `polyfit2d`. The blue lines are the 30-day moving average. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods.

(AW) core across time for pressure, 3.42 ± 0.05 dbar/yr (5.87 ± 0.05 dbar/yr), decrease slightly for the residuals, 2.91 ± 0.04 dbar/yr (5.32 ± 0.04 dbar/yr), where the \pm values are the standard errors of the estimated slope. The R^2 values also similar between the slopes of the pressures and the residuals, being 0.480, 0.512, 0.699, and 0.698 for panels (a), (b), (c), and (d), respectively.

While the trends in time do not change much after correcting with the `polyfit2d`, the trends are different between the LHW and AW cores. Both the LHW and AW cores have a cumulative downwelling rate that agrees with previous studies, however the AW core is sinking faster than the LHW core. This difference could be happenstance or indicative of some general behavior of the water column. One possibility is that deeper layers are sinking faster than shallower layers, stretching the pycnocline. Another possibility is that new layers are forming at the bottom, just above the AW core. If, at the same time, layers are disappearing at the top, this could be thought of as an “escalator” motion. By identifying individual staircase layers and analyzing their trends, I will be able to evaluate whether the sinking pycnocline contains a stretching staircase, an escalator, or something in between.

3.3.3 Clustering each time period

As previously mentioned, the clustering method does not distinguish between individual profiles and should not be applied to a set of data that span larger scales than across which staircases have been known to be coherent. The longest any previous study has tracked an individual thermohaline staircase layer in the Arctic is four years (Polyakov, Pnyushkov, Rembe, et al., 2012; Lu et al., 2022). For consistency, I split the data into consecutive year-long periods, of which there are 17, each starting and ending on August 15, which is generally around when new ITPs have been deployed (see Figure B.1). This not only guarantees each dataset respects the known temporal coherence scale, but also reduces the maximum size of any particular dataset to a manageable quantity. The

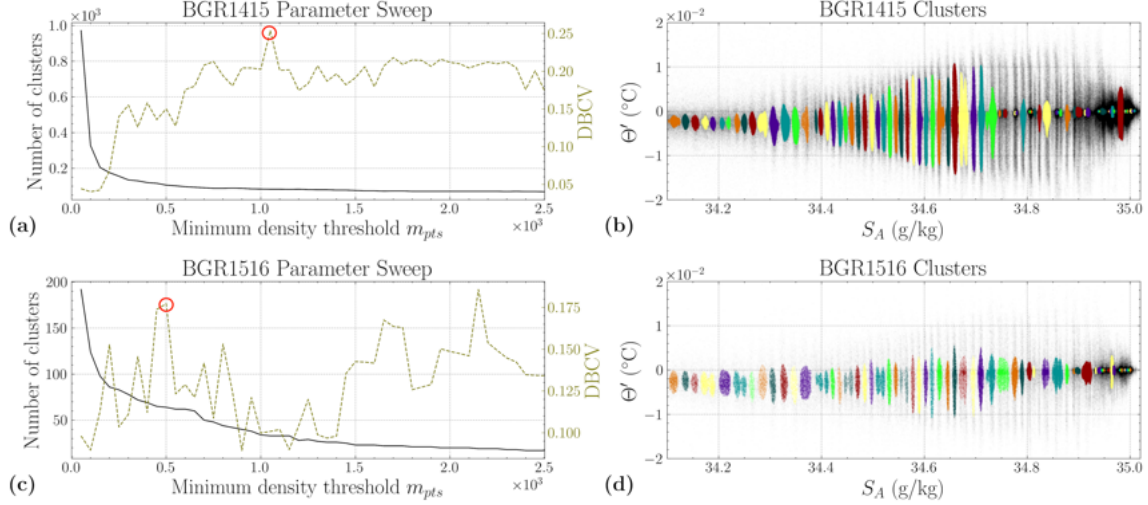


Figure 3.5: Parameter sweeps in m_{pts} for (a) BGR1415 and (c) BGR1516. The results of clustering (b) BGR1415 and (d) BGR1516 with the chosen values of m_{pts} 1050 and 500, respectively, before relabeling the cluster ID's.

dataset for each period contains every profile taken within the BGR during that time filtered to pressures between $p(S_A \approx 34.1) < p < p(\Theta_{max})$. Note that this eliminates the 45 profiles where $p(S_A \approx 34.1) = p(\Theta_{max})$, reducing the total count to 15,202 profiles. I name each period BGRXXYY where XX and YY are the last two digits of the years the period spans. For example, BGR0506 contains the profiles taken between 2005/8/15 and 2006/8/15. On all plots for this work with time on the horizontal axis, I mark each August 15 with vertical red dashed lines. Table 3.2 summarizes Figure B.1, listing the ITPs, the number of profiles, and the number of data points for each period.

In Section 2.3.2, I describe how I selected the parameters ℓ , the length of the running average window in dbar, and m_{pts} , the minimum points per cluster. As described in Section A.2, the value of ℓ does not significantly affect the results of the clustering, except for near the bottom of the staircase where the influence of intrusions is greatest. While choosing a smaller value of ℓ may negate the ability to identify the presence of remnant intrusions in the results, I do not analyze such features in this work. Therefore, I select a smaller value of ℓ than in Chapter 2 to reduce the zig-zag pattern across the high-salinity region of Θ' - S_A space. In this chapter, I use a value of $\ell = 2.5$ dbar for every period.

As was in the case in Section 2.3.2, I use density-based clustering validation (DBCV) as a metric to guide my selections of m_{pts} . For each dataset, I performed a parameter sweep, running HDBSCAN across a range of m_{pts} from 100 to 2500, recording the DBCV and number of clusters for each clustering result. Examples of such parameter sweeps for BGR1415 and BGR1516 are shown in Figure 3.5(a,c). I see the same general trend in number of clusters as in Figure 2.3(b), rapidly decreasing for small values of m_{pts} then decreasing much more slowly as m_{pts} grows larger. Generally, I followed the same procedure in selecting m_{pts} as before, choosing the value with the highest DBCV as long as it does not occur below the values at which the number of clusters is rapidly decreasing. For BGR1415, this led me to choose $m_{pts} = 1050$ where DBCV was maximum at 0.2537; this particular clustering is shown in Figure 3.5(a). However, for several periods, this approach led to clusters which aligned poorly with staircase layers. In these cases, I chose a value

Period	ITPs present	Profiles used	Total points	Noise points	m_{pts}	Number of clusters	DBCV
BGR0506	1, 3	1,425	1,152,866	616,986	650	75	0.2869
BGR0607	1, 3, 4, 5, 6, 8, 13	1,518	1,267,893	739,831	1100	68	0.2539
BGR0708	4, 5, 6, 8, 13, 18, 21, 30	1,454	1,126,688	623,000	1300	56	0.3172
BGR0809	8, 11, 18, 21, 25, 30	573	509,054	509,054	450	65	0.1615
BGR0910	21, 32, 33, 34, 35	774	628,730	400,559	300	80	0.1797
BGR1011	33, 41, 42, 43	743	621,887	446,750	500	78	0.1193
BGR1112	41, 52, 53, 54, 55	880	709,249	406,165	1300	54	0.2632
BGR1213	41, 62, 64, 65	597	473,227	278,529	1100	49	0.2256
BGR1314	64, 68, 69, 70, 77, 78, 79, 80	611	493,043	228,962	1350	47	0.2910
BGR1415	77, 79, 80, 81, 82, 84, 85, 86, 87	1,907	1,600,412	947,319	1050	79	0.2537
BGR1516	82, 85, 86, 88, 89	377	290,354	170,627	500	64	0.1772
BGR1617	97, 99	397	318,621	194,903	200	82	0.2130
BGR1718	97, 100, 101, 108	470	344,856	190,093	550	67	0.1973
BGR1819	103, 104, 105, 107, 109, 110	684	615,743	410,198	700	68	0.1433
BGR1920	105, 113, 114, 117, 118	488	443,676	268,088	750	61	0.1791
BGR2021	113, 114, 120, 121	1,063	967,756	707,763	1150	68	0.1719
BGR2122	120, 121, 122, 123	1,241	1,180,791	722,087	1150	63	0.1823
Total	60 different ITPs	15,202	12,744,846	—	—	—	—

Table 3.2: For each time period, summaries of the data used and clustering parameters.

of m_{pts} that was lower than that with the maximum DBCV, yet a prominent value of DBCV with respect to the m_{pts} , using the fact that the optimal value of m_{pts} is roughly proportional to the number of data points. For example, the maximum DBCV of the parameter sweep for BGR1516 in Figure 3.5(c) was 0.1857 at $m_{pts} = 2150$. However, I chose to use $m_{pts} = 500$ which gave a prominent value DBCV of 0.1772 compared to nearby points and was at a similar order of magnitude in m_{pts} compared to other periods with a similar order of magnitude number of data points (see Table 3.2). The chosen clustering for BGR1516 is shown in Figure 3.5(d).

3.3.4 Connecting across time periods

After clustering each time period, I then began the process of connecting the clusters between different time periods which represent the same layer. When running the HDBSCAN algorithm, each cluster is assigned an integer, starting at 0, as an ID. However, the order in which these IDs are assigned is based upon the processes of building the hierarchy tree described in Section 2.3.1, not upon the clusters' positions in $\Theta'-S_A$ space. As was shown in Section 2.4.1, the salinity of the layers is more stable than the pressure or temperature. Therefore, I first reassign the cluster IDs to be in order of the clusters' average S_A . This is the ordering shown in the coloring of the clusters in Figure 3.5(b,d). While I use only seven distinct colors, each distinct cluster has a distinct ID.

However, as noted in Section 2.4.2, the clustering algorithm can sometimes erroneously group multiple layers into a single cluster or erroneously divide one layer into multiple clusters. There is no guarantee these errors will be consistent across different time periods and therefore there is no guarantee that, after ordering the cluster IDs in each period with increasing average S_A , the cluster IDs for a particular layer will match across time periods. For example, the purple cluster in Figure 3.5(d) around 34.21 g/kg does not correspond to the purple cluster in panel (b) around 34.19 g/kg because the yellow cluster in panel (d) around 34.18 g/kg clearly encompasses two layers.

To correct for this, I manually relabeled the cluster IDs using several guiding principles. I ensured that clusters with the same ID across different periods had approximately the same average value of S_A . When a cluster in a particular period clearly encompassed more than one layer, I gave that cluster its own distinct ID if the majority of other periods did, in fact, distinguish those layers. I avoided relabeling two clusters within the same period to the same ID unless it was abundantly clear that the clustering algorithm divided one layer into multiple clusters. For values of S_A above approximately 34.9, the number of clusters increased and their spans in S_A decreased, often overlapping with each other. This part of the staircase is very near the AW core, a region known to have intrusions (Bebieva & Timmermans, 2019). Because I am not analyzing intrusions in this study, I did not make an attempt to connect across time periods for clusters with an average $S_A \gtrsim 34.9$ g/kg, simply assigning them IDs with values of 120 or larger. The result of this relabeling can be seen in Figure 3.6. A version of the same plot before the relabeling can be seen in Figure B.3.

3.3.5 Defining outlier clusters

In 2.3.3, I describe how I evaluate the validity of individual clusters based on their lateral density ratio R_L and normalized inter-cluster range in salinity IR_{S_A} . I determine that outliers in R_L represent either clusters where multiple layers have erroneously been grouped together or a physical splitting or merging of layers. Outliers in IR_{S_A} represent either parts of a single layer erroneously split into multiple clusters or intrusion features. In this study, I wish to analyze the properties of single, full layers and therefore chose to disregard such clusters when calculating statistics and trends.

Here, I mark clusters as outliers in a different manner than in Chapter 2. First, I consider all clusters with IDs 120 or higher (i.e., clusters with an average $S_A \gtrsim 34.9$) to be outliers as I did not attempt to relabel these clusters to connect them across time periods. With those remaining, I define outliers in either R_L or IR_{S_A} based upon the modified z-score

$$M_i = \frac{b(x_i - \hat{x})}{\text{MAD}} \quad (3.2)$$

where $b = 0.6745$ is a normalization parameter assuming a normal distribution (Rousseeuw & Croux, 1993), x_i represents the value of R_L or IR_{S_A} for a particular cluster, \hat{x} is the sample median of all the clusters, and MAD is the median of the absolute deviations about the median

$$\text{MAD} = \text{median}_i\{|x_i - \hat{x}|\} \quad (3.3)$$

where median_i represents taking the median across all values of $|x_i - \hat{x}|$ (Rousseeuw & Leroy, 1987). I use the modified z-score here as it is based upon estimators that are robust to small sample sizes and I follow the recommendation of Iglewicz & Hoaglin (1993), marking points as outliers when $|M_i| > 3.5$. In Figure 3.7, I mark such outliers in IR_{S_A} with blue and in R_L with red. Hereafter, I refer to clusters that are outliers in neither as “regular clusters.”

3.4 Results

Having used the clustering algorithm to identify layers in the 17 year-long periods and subsequently connected the layers across those time periods, I can analyze their various properties and trends.




Figure 3.6: The salinity across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.

3.4.1 Properties of detected layers

By connecting the clustering results of each year-long time period, I find that I can track individual layers across 17 years and the entire BGR, spanning up to 1168 km. In Figure 3.6, I plot these results in salinity against time. Blue and pink triangle markers denote the values of the LHW and AW cores, respectively, for each profile. I use seven other colors to distinguish between the different clusters and mark the IDs for clusters in BGR0506 (BGR2122) with numbers in the corresponding colors on the left (right) of the plot. Noise points are marked in black and the vertical red dashed lines mark every August 15th, the boundaries between time periods.

Similar plots with vertical axes of pressure, temperature, and density anomaly can be seen in Figures B.4, B.5, and B.6. From these, it can be seen that individual layers vary most in pressure and then, in decreasing order, temperature, density anomaly, and salinity. This pattern is reflected quantitatively in the median values of the normalized inter-cluster ratio as $\widehat{IR}_p > \widehat{IR}_\Theta > \widehat{IR}_{\sigma_1} > \widehat{IR}_{S_A}$ (see Table 3.3). In Section 2.4.1, I found the same relative magnitudes for pressure, temperature, and salinity, which matches with the patterns noted in previous studies (Timmermans et al., 2008; Lu et al., 2022). For density anomaly, it follows that $\widehat{IR}_\Theta > \widehat{IR}_{\sigma_1} > \widehat{IR}_{S_A}$ as σ_1 is

Variable	Pressure (p)	Salinity (S_A)	Temperature (Θ)	Density Anomaly (σ_1)
\widehat{IR}_v	365	4.49	80.7	18.8
$\widehat{\Delta v}$	18.3 dbar	9.80×10^{-3} g/kg	5.74×10^{-2} °C	7.60×10^{-3} kg/m ³
$\widehat{v}_{\text{span}}$	176 dbar	2.51×10^{-2} g/kg	0.77 °C	6.25×10^{-2} kg/m ³
$\widehat{\sigma}_v$	26.9 dbar	4.50×10^{-3} g/kg	7.72×10^{-2} °C	6.46×10^{-3} kg/m ³

Table 3.3: A summary of median values for each variable v . The $\widehat{\cdot}$ symbol indicates the median across each regular cluster (neither an outlier in IR_{S_A} nor R_L), Δv represents the distance between the average values in v for neighboring clusters, v_{span} is the distance between v_{max} and v_{min} for a cluster, and σ_v is the standard deviation of a cluster in v .




Figure 3.7: The value of each cluster’s (a) normalized inter-cluster range for salinity IR_{S_A} and (b) lateral density ratio R_L as a function of the cluster’s average pressure. The colors and numbers for each cluster are the same as in Figure 3.6. The circles mark clusters that are outliers in the respective variable. Note that the vertical domain is zoomed in to show detail and, while there are clusters with average pressures below 355 dbar, all of those clusters have been disregarded in further calculations.

influenced by both Θ and S_A . I once again conclude that salinity is the most appropriate variable by which to identify any particular layer in the staircase, in agreement with L22.

3.4.2 Values of IR_{S_A} and R_L

As described in Section 3.3.5, I mark outliers in IR_{S_A} and R_L , following a similar procedure as was used in Chapter 2. However, I use the modified z-score defined in (3.2) and the recommended threshold of 3.5 (Iglewicz & Hoaglin, 1993). In Figure 3.7, I mark outliers in IR_{S_A} with blue and outliers in R_L with red. I keep this same coloring scheme for the other plots in this chapter. As mentioned above, I disregard clusters with IDs ≥ 120 when calculating statistics and trends. In plots, I mark these clusters as outliers in both IR_{S_A} and R_L , however few of these clusters appear in Figure 3.7 as I have restricted the vertical domain to show detail.

When a cluster has $IR_{S_A} < 1$, this means that its range of salinity values does not overlap with those of the clusters above and below. In this study, only Cluster 91 has $IR_{S_A} < 1$, as can

Dataset	\overline{R}_L	σ_{R_L}	Polynomial fit of $R_L(p)$	R^2
BGR all	-10.62	4.15	$-2.39 \times 10^{-4}p^2 + 0.11p - 22.08$	0.04
ITP2	-3.55	2.24	$-3.81 \times 10^{-4}p^2 + 0.24p - 39.85$	0.67
ITP3	-3.86	2.41	$-3.38 \times 10^{-4}p^2 + 0.24p - 43.29$	0.84

Table 3.4: Comparing values of the lateral density ratio R_L from this study to that of Chapter 2. The means and standard deviations of R_L along with the second degree polynomial fits of R_L in pressure p from Figures 3.7(b) and 2.5(b,d), as well as the coefficients of determination for those fits.

be seen at 238 dbar in Figure 3.7(a). This is in contrast to Figure 2.5(a,c) where roughly half the clusters have no overlap with their neighbors. I attribute this to the manner in which I relabeled the cluster IDs, described in Section 3.3.4. When it was clear that a cluster for a particular time period clearly encompassed more than one layer which were clustered separately in most other time periods, I would give that cluster its own unique ID. For example, in Figure 3.7(a) with a cluster average pressure of 232 dbar, Cluster 2 has $IR_{S_A} = 154$, indicating a high degree of overlap in S_A with at least one of its neighboring clusters. This overlap can be seen in Figure 3.6 as Cluster 2 is the topmost layer in BGR1112–BGR1314, colored dark red. In each of those three periods, a single cluster encompassed a significant range of S_A and I assigned it an ID of 2. This is in contrast to BGR1415 where the clusters which I assigned IDs of 0, 4, 6, and 7 cover approximately the same range in S_A , but were separated by the algorithm. These clusters can also be seen in BGR0506–BGR0607. This extreme example of a phenomenon common across the clusters which illustrates why there are so few clusters with $IR_{S_A} < 1$. However, it is a common occurrence and because the outlier detection method I employ seeks relatively extreme values, I conclude that outliers in IR_{S_A} still correspond to clusters that do not represent single layers.

In Section 2.4.4, I discuss how results in Chapter 2 suggest a dependence of R_L on pressure. Following that section, I plot a second-degree polynomial fit through the clusters that are not outliers in R_L in Figure 3.7(b) to find $R_L = -2.39 \times 10^{-4}p^2 + 0.11p - 22.08$. While the coefficients in this equation are close to those shown in Figure 2.5(b,d), the coefficient of determination, $R^2 = 0.04$, is much lower (see Table 3.4). The mean lateral density ratio for the 76 regular clusters, $\overline{R}_L = -10.62$, is also significantly different from the results of Chapter 2 and from Timmermans et al. (2008), who reported a mean value of -3.7 ± 0.9 . It is unclear whether this supports or negates the conclusion from Chapter 2 that R_L has a dependence on pressure. The similarity in the coefficients of the polynomial fit is striking, however it has a very low R^2 . The difference in magnitude of \overline{R}_L could be, similar to IR_{S_A} , attributed to my particular method of connecting across time periods.

3.4.3 Spatial variability

Before I find the trends in time for each cluster, I first correct for spatial distributions, as explained in Section 3.3.2. As an example, Figure 3.8 shows maps and histograms of the points within Cluster 63 for pressure, salinity, and temperature. The maps in panels (a–c) show the individual points plotted on top of the contour of the `polyfit2d` for the corresponding variable. The general bowl-shaped pattern in pressure matches very closely with that of the LHW (see Figure 3.2(b)) with the location of the minimum pressure in the `polyfit2d` only 46.14 km away at 74.22°N , 149.58°W . The direction of increase of salinity shown in Figure 3.8(b), from the southwest to the northeast, matches that of the AW (see Figure B.2(c)), albeit with a gradient about one third as steep. Note, however,

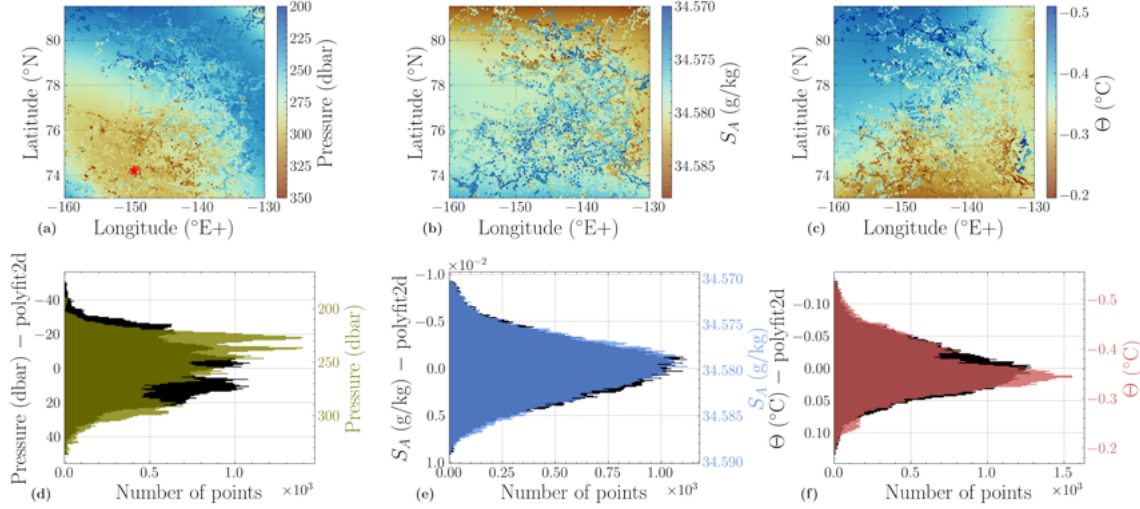


Figure 3.8: Maps of (a) pressure, (b) salinity, and (c) temperature for Cluster 63 over the corresponding `polyfit2d` for each variable. The red star in (a) is at the location of the minimum of the `polyfit2d`: 74.22°N, 149.58°W. Histograms of (d) pressure, (e) salinity, and (f) temperature for Cluster 63, showing both the original and the residual values for each variable.

For Cluster 63	Pressure (dbar)	Salinity (g/kg)	Temperature (°C)
Mean	256	37.5798	-0.3610
Standard Deviation	27.4	0.0032	0.0591
<code>polyfit2d</code> Standard Deviation	17.1	0.0031	0.0453

Table 3.5: The properties of the histograms for Cluster 63 shown in Figure 3.8. The means of all `polyfit2d` versions were zero.

that the range of variations of salinity is extremely narrow around the mean value.

In Figure 3.8(c), there is a clear pattern of increasing temperatures from the northwest to the southeast, opposite in direction yet equal in magnitude to that of the AW (see Figure B.2(d)). I attribute this to what L22 identified as the influences of the different water masses above and below, the LHW and AW. That is, layers closer to a particular water mass will tend to resemble those masses in spatial variability. Indeed, I do find the patterns of temperature across each layer change from top to bottom, most notably around $S_A \approx 34.7$ g/kg (see Figure B.7). The reason for why this influence is more evident in temperature than in salinity I attribute to temperature’s higher rate of diffusion.

The histograms in Figure 3.8(d-f) show both the distributions of the original points overlaid on the distributions of the residuals for each variable. Upon making these spatial corrections, the standard deviations of the distributions in each variable decreased, although only by a significant amount for pressure (see Table 3.5). This again implies that taking the `polyfit2d` is removing the overall spatial distribution.

3.4.4 Trends in time

After using `polyfit2d` to correct for spatial distributions in the clusters, I then find the trends in time using the ordinary least squares method. Note that I did not account for the fact that different clusters cover different time spans. Not all clusters appear in every time period and, especially for

Variable	$\widehat{d_t p}$ (dbar/yr)	$\widehat{d_t S_A}$ (g/kg/yr)	$\widehat{d_t \Theta}$ ($^{\circ}\text{C}/\text{yr}$)	$\widehat{d_t \sigma_1}$ ($\text{kg}/\text{m}^3/\text{yr}$)
Original	3.43	-3.39×10^{-4}	1.18×10^{-2}	-1.04×10^{-3}
Residuals	2.24	-2.42×10^{-4}	5.73×10^{-3}	-6.24×10^{-4}

Table 3.6: A summary of median trends for each variable v across all regular clusters, before and after taking the residuals with `polyfit2d`.

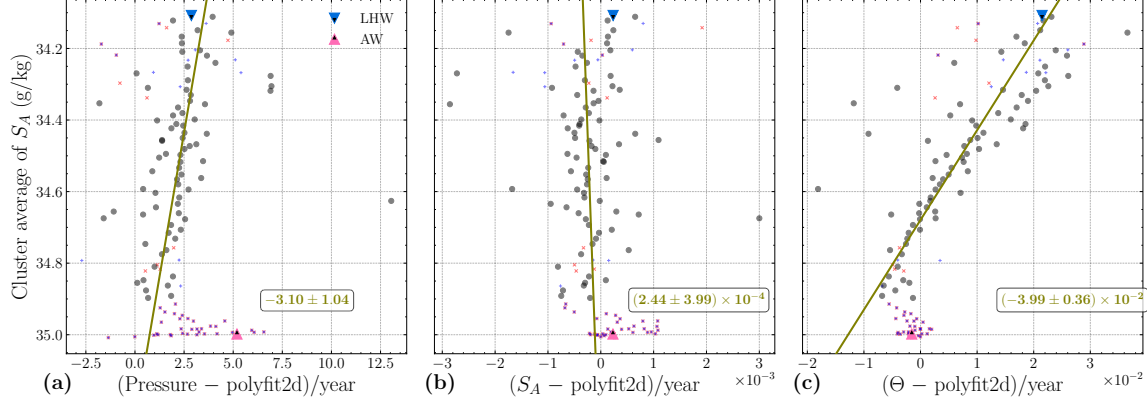


Figure 3.9: For all clusters, the trends in time of the `polyfit2d` residual for (a) pressure, (b) salinity, and (c) temperature with respect to the cluster average of salinity. A blue “+” denotes clusters that are outliers in IR_{S_A} and a red “x” denote clusters that are outliers in R_L , as was indicated in Figure 3.7. All outliers were ignored when calculating the fit lines which have (from left to right) R^2 values of -0.328, 0.071, and -0.789. The blue and pink triangle markers denote the values for the LHW and AW, respectively.

clusters that only appear in a few periods, the trends are much more susceptible to influence from extreme values. An example of these trends can be seen for Cluster 63 in Figure B.8 where, just as it was for the LHW and AW, the slopes decrease after finding the residuals. This pattern holds across all the clusters as can be seen in Table 3.6 where the median trend over time for each variable decreases after taking the residuals.

In Figure 3.9 I plot the trends in time, corrected by `polyfit2d`, for each cluster in pressure, salinity, and temperature against the cluster average of salinity. I have marked clusters with a blue “+” and a red “x” to denote clusters that are outliers in IR_{S_A} and R_L , respectively, as was shown in Figure 3.7. Here, I have given all clusters the same markers for clarity however, a version of this figure where the markers denote which cluster is which can be seen in Figure B.9. I mark the trends in the respective variables for the LHW and AW in blue and pink triangles, respectively.

Ignoring all clusters that are outliers in either IR_{S_A} or R_L , I found the linear fit of these trends as a function of the cluster’s average salinity to be $d_{S_A}(d_t p) = -3.10 \pm 1.04$ dbar/yr/(g/kg), $d_{S_A}(d_t S_A) = (2.44 \pm 3.99) \times 10^{-4}$ (g/kg)/yr/(g/kg), and $d_{S_A}(d_t \Theta) = (-3.99 \pm 0.36) \times 10^{-2}$ $^{\circ}\text{C}/\text{year}/(\text{g}/\text{kg})$. The R^2 values are -0.328, 0.071, and -0.789, respectively. I refer to these as compound trends as they are trends in a particular variable (noted on the vertical axis) of the trends in time for a particular variable (noted on the horizontal axis). I do not include the LHW or AW when calculating compound trends. A summary of the compound trends for each combination of variables, before and after taking the residuals can be found in Table B.1.

For the regular clusters, the pressure trends in time decrease going down the water column, meaning the deeper layers are sinking more slowly than the upper layers. This can be seen in

Figure 3.9(a) and holds true no matter which variable is on the vertical axis. This is the opposite of what one might assume from solely the trends for the upper and lower boundaries; the AW is sinking more quickly than the LHW. Indeed, there is a transition around $S_A = 34.9$ g/kg (or 355 dbar, or 0.5°C) between the trends in the regular clusters and the LHW decreasing with depth to the trends in the outlier clusters with IDs ≥ 120 and the AW increasing with depth. This transition does not necessarily represent a physical phenomenon as these particular outlier clusters were not carefully connected across time periods. However, it is striking to note that the trend of the AW nevertheless falls along the same line.

As can be seen in Figure 3.9(b), the trends of salinity over time are scattered around zero and the compound trend is insignificant. This holds true for compound trends in salinity across vertical axes of pressure, temperature, and density anomaly. In each case, both before and after taking the residuals, the standard error is larger than the slope and the R^2 value is less than 0.085 (see Table B.1). This is as expected because salinity is the main factor for identifying layers within each period and connecting them across different periods. By virtue of this identification process, it would be unlikely to find the salinity within clusters to be changing over time. The same expectation holds with the LHW, which is identified as the point in the water column with salinity closest to 34.1 g/kg. However, it is notable that the AW, which is identified as the point with the maximum temperature in each profile, has a negligible trend in salinity over time as well.

In Figure 3.9(c), there is a significant compound trend in the trends of temperature over time with warming trends near the top decreasing and even transitioning to cooling near the bottom. The significance of this trend will be discussed in context of heat flux in Section 3.4.6. It should be noted that the trend for temperature over time for the AW indicated in Figure 3.9 is only for the point at which temperature is maximum, and does not represent the trend in the temperature of the Atlantic Water layer as a whole. In general, the AW has been warming significantly over the past half-century (Carmack et al., 1997; Polyakov, Pnyushkov, & Timokhov, 2012).

3.4.5 Layer thicknesses

The vertical thickness of layers in Arctic thermohaline staircases have long been observed to be on the order of meters (Padman & Dillon, 1987; Timmermans et al., 2008). While many laboratory studies have seen the formation of layers (Huppert & Linden, 1979; Fernando, 1989; Guo et al., 2016), to my knowledge, no studies have noted whether the thickness of an individual layer changes over long periods of time. Here, I calculated layer thickness for each cluster on a per-profile basis. That is, for each cluster, I analyzed each profile, finding the maximum and minimum value of pressure within that cluster and within that profile, and taking the thickness to be the difference between the two.

In Figure 3.10(a), I plot all these values for all regular clusters as a function of salinity. I do not include values equal to zero as these simply indicate there was only one point for a particular cluster in a particular profile. While I do find thicknesses of up to 20 dbar for the lower layers, I have restricted the horizontal axis to show detail. A version of this plot spanning a larger horizontal domain and with the points colored by cluster can be found in Figure B.10. There is a structure to the distribution of layer thicknesses that comes from the manner in which the data was recorded. The vertical resolution of ITPs is ~ 25 cm, and so values of layer thickness tend to fall around integer multiples (i.e. 0.25 m, 0.5 m, 0.75 m, etc.).




Figure 3.10: The thicknesses (in dbar) of the detected layers. (a) The non-zero thickness values against salinity of each regular cluster for every profile in the BGR. While there are thicknesses of up to 20 dbar, axes bounds are limited here to show detail. For all clusters, the (b) cluster average of layer thickness and (c) trends of percent change in layer thickness over time with respect to the cluster average of salinity. The values in (c) are calculated by taking the trend in time for the cluster’s non-zero thicknesses then dividing by that cluster’s average thickness shown in (b), however are highly uncertain, see text for explanation. A blue “+” and a red “x” denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have R^2 values of 0.609 for (b) and 0.041 for (c).

Figure 3.10(b) shows the average thickness for each cluster. The mean of these cluster average thicknesses is 1.55 dbar with a standard deviation of 1.53 dbar, matching the ranges of thicknesses reported in previous studies for this region (Padman & Dillon, 1987; Timmermans et al., 2008; Shibley et al., 2017; Lu et al., 2022). As these layers span the entire BGR, approximately 1000 km across, they have an aspect ratio on the order of 10^6 , again in agreement with previous studies (Timmermans et al., 2008; Toole et al., 2011; Lu et al., 2022).

There is also a trend of increasing thickness for layers deeper in the water column. In Figure 3.10(b), I find a trend in cluster average thickness of 4.43 ± 0.67 dbar/(g/kg) ($R^2 = 0.609$). When plotted against pressure, this becomes a trend of $(2.53 \pm 0.40) \times 10^{-2}$ dbar/dbar ($R^2 = 0.591$). If I instead fit a line to all thicknesses from all clusters against pressure shown in Figure 3.10(a), I find a slope of $(1.431 \pm 0.006) \times 10^{-2}$ dbar/dbar ($R^2 = 0.475$) which is the same order of magnitude. In the ocean, one dbar is almost exactly equivalent to one meter. Therefore, this rate of increase in layer thickness with respect to depth matches that of Padman & Dillon (1987), who reported a change of around 1.1×10^{-2} m/m using data from AIWEX in 1985. Shibley & Timmermans (2019) reported a trend of 0.12 ± 0.01 m/m from analyzing data from ITP13 from August 2007 to August 2008. That slope is an order of magnitude larger than $(4.14 \pm 0.04) \times 10^{-2}$ dbar/dbar, the slope I find with a linear trend through the layer thicknesses I detected from just ITP13. I attribute this difference mainly to the fact that I include layer thicknesses from deeper in the water column, but also to the fact that Shibley & Timmermans (2019) found their reported slope by first binning the data into 0.5 m increments (see Figure B.11). While the magnitude of the slope I found may not agree in this case, the overall trend does follow the consensus that layer thickness generally increases with depth (Neal et al., 1969; Guo et al., 2016; Shibley et al., 2017; Lu et al., 2022). I do find a similar geographic distribution to ST22, with the highest mean layer thicknesses in the southeast of the region (see Figure B.12).

There are predictions that, as the AW continues to warm, staircase layers will be thinner (Shibley & Timmermans, 2022). However, as seen in Figure 3.10(c), the trends of percent change in layer

thickness over time have no significant pattern. These trends are calculated in the same manner as other trends shown above, but have been normalized for each cluster by that cluster’s average thickness. Note that these trends are more uncertain than trends in other variables. For each cluster, there is only one value of thickness per profile, meaning the sets of points are smaller and therefore more susceptible to extreme values. This is in addition to the issue mentioned above, the influence from extreme values for clusters that do not span the entire time range. Indeed, every regular cluster with an absolute value of percent change in thickness per year greater than 15% appears in five or fewer year-long time periods. That is to say, the trend in layer thickness for any particular cluster in Figure 3.10(c) should not be taken as precise. I present this set of trends to comment on their overall characteristics. The compound trend of percent change in layer thickness over time against salinity is $1.86 \pm 5.45 \text{ \%}/\text{yr}/(\text{g}/\text{kg})$ with $R^2 = 0.041$, implying no pattern along the staircase. For the identified layers, approximately two-thirds of the trends are negative. From this I conclude that their thicknesses are somewhere between not changing over time, or becoming slightly thinner overall.

This does not exclude the possibility that newly forming layers are becoming thinner over time. However, the methods used in this study exclude the possibility of seeing new layers form. Because layers are stable in salinity over time and I define the LHW ($S_A \approx 34.1$) as my top bound, I cannot determine whether or not new layers are forming at the top of the staircases. While the lower bound of the AW is defined in terms of maximum temperature, I did not connect clusters across time periods above $S_A \approx 34.9$. In Figure 3.6, I find layers with salinities just less than this point that are stable throughout the entire 17 years and so I also cannot determine whether or not new layers are forming at the bottom bound of the staircase.

3.4.6 Net heat flux

Measuring the heat flux through each layer in a naturally-occurring thermohaline staircase is a difficult problem, one which may require specialized observational efforts (Polyakov, Walsh, & Kwok, 2012; Stranne et al., 2017). However, with the extensive dataset developed in this study, I can make estimates of the net heat flux for individual layers. Lacking internal heat sources, a change in temperature of a layer must be due to the net heat flux through its boundaries. Using the layer thicknesses and trends in temperature I found for each cluster, I can compute the cluster average net heat flux with the following

$$\overline{\Delta \mathcal{F}_H} = \overline{\Delta h} \cdot \overline{C_p} \cdot \overline{d_t \Theta} \cdot \overline{\rho} \quad (3.4)$$

where $\overline{\Delta h}$ is the cluster average layer height (taking 1 dbar \approx 1 m), $\overline{C_p}$ is the cluster average isobaric heat capacity [$\text{J kg}^{-1} \text{ K}^{-1}$], $\overline{d_t \Theta}$ is the cluster average trend in temperature over time [K s^{-1}], and $\overline{\rho}$ is the cluster average density [kg m^{-3}]. I plot $\overline{\Delta \mathcal{F}_H}$ against salinity in Figure 3.11. The individual components can be seen in Figure B.13 where $\overline{d_t \Theta}$, which is also shown in Figure 3.9(c), is the only component to change sign. Values of net heat flux calculated in this way do not indicate in which direction heat is moving. However, they do allow estimates of the change in heat of individual layers which can be compared to estimates of heat flux across staircases made in previous studies.

There is a significant trend of decreasing net heat flux from top to bottom of $d_{S_A}(\overline{\Delta \mathcal{F}_H}) = (-5.06 \pm 0.64) \times 10^{-3} \text{ W/m}^2/(\text{g}/\text{kg})$ ($R^2 = -0.677$). This holds for vertical axes of pressure and

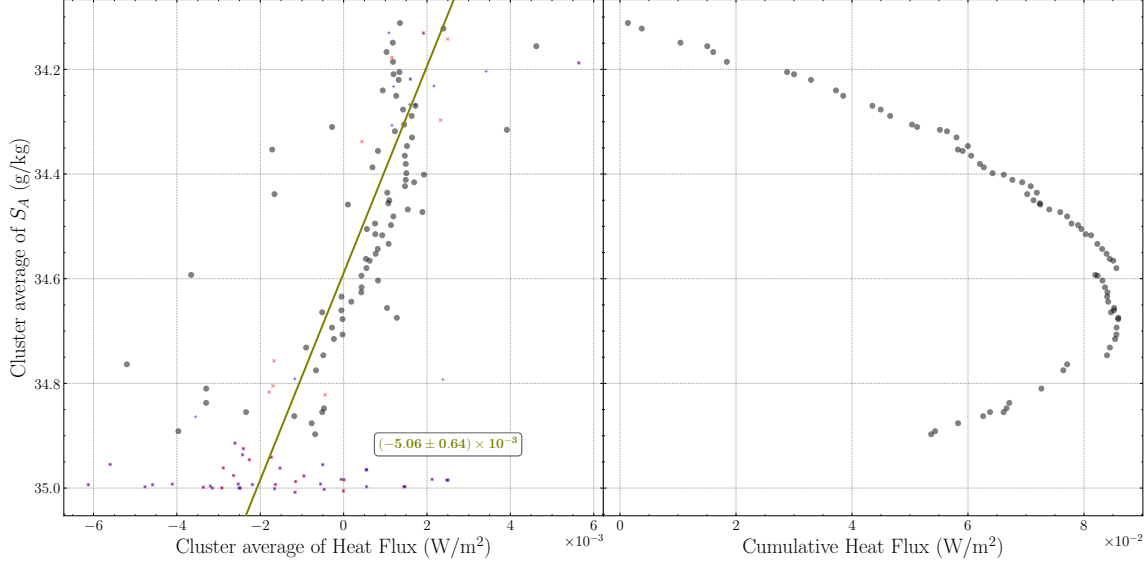


Figure 3.11: (a) The average net heat flux for each cluster against salinity. All outliers were ignored when calculating the fit line which has $R^2 = -0.677$. (b) The cumulative net heat flux of regular clusters, from the top to the bottom of the staircase.

temperature, with $d_p(\overline{\Delta \mathcal{F}_H}) = (-2.82 \pm 0.39) \times 10^{-5} \text{ W/m}^2/\text{dbar}$ ($R^2 = -0.641$) and $d_\Theta(\overline{\Delta \mathcal{F}_H}) = (-2.43 \pm 0.28) \times 10^{-3} \text{ W/m}^2/^\circ\text{C}$ ($R^2 = -0.713$). There is a zero crossing at $S_A = 34.6 \text{ g/kg}$ (or 278 dbar, or -0.28°C) where clusters go from having positive net heat fluxes above to having negative net heat fluxes below. This shift can also be seen in the inflection point of the cumulative net heat flux in Figure 3.11(b) which sums the net heat flux of each regular cluster from the top to the bottom of the staircase. Here, I have made the assumption that the net heat flux for a layer is entirely due to the heat flux through its top and bottom surfaces because the extreme aspect ratio of layers ($O(10^6)$) makes the surface area of the edges negligible. The total sum of the regular clusters' net heat fluxes is 0.054 W/m^2 . This is on the same order of magnitude of estimates of heat flux through staircases from previous studies (Timmermans et al., 2008; Turner, 2010; Guthrie et al., 2015; Shibley & Timmermans, 2019). However, this is only a small fraction of the estimates of total ocean-to-sea-ice heat flux (Timmermans et al., 2008).

3.5 Discussion

In this chapter, I have demonstrated how a method based on the HDBSCAN clustering algorithm can be used to analyze the long-term evolution of thermohaline staircase layers in the BGR. The clustering method I used was shown in Chapter 2 to automatically detect and connect well-mixed staircase layers across hydrographic profiles. Many previous studies have remarked on the persistent presence of thermohaline staircases in the Arctic Ocean since they were first observed in 1969 (Neal et al., 1969). In this study, I found that individual staircase layers were coherent over the entire 17 years analyzed. Having shown the stability of these layers on the decadal scale, it could be speculated that individual layers have remained coherent since those first measurements, or even longer. These structures are not only long-lived, but also large-scale, stretching 1000 km horizontally with a remarkable aspect ratio which I found to be on the order of 10^6 , as others have previously

(Timmermans et al., 2008; Toole et al., 2011; Lu et al., 2022). Recently, Ménesguen et al. (2022) has shown evidence that the lower sections of these staircases have been disappearing. While, in this study, I did not analyze down to the depths at which these disappearances were noted and I did not see indications that the layers I identified were disappearing, I did find that layers are changing, especially in pressure and temperature. As the Arctic continues to warm and the sea ice continues to retreat, this raises the question as to whether this will also result in the loss of these unique stratification structures.

Using the normalized inter-cluster range IR_p , I confirmed qualitatively that values of pressure and temperature vary much more within a layer than salinity. In agreement with L22, I conclude that salinity is the most appropriate variable by which to identify individual layers. To analyze spatial distributions, I fit different layers to a two-dimensional polynomial of up to third order in latitude and longitude. Using these `polyfit2d`'s, I identified the same bowl-shaped distribution in pressure noted by L22. I also note that, for the upper staircase, temperatures generally increase from north to south, from northwest to the southeast in the middle, and from west to east near the bottom of the staircase, I attribute this to the differing influences of the water masses above (LHW) and below (AW). By taking the residuals as the original points minus the `polyfit2d`, I find that the standard deviations in each layer's pressure, salinity, and temperature decrease and showed this process removes the overall spatial distribution.

With a dataset of staircase layers tracked over 17 years and after removing the geospatial variations with `polyfit2d`, I was able to note how individual layers changed over time and how that contributed to the evolution of the staircase as a whole. I found that the LHW is sinking 2.91 ± 0.04 dbar/yr, a rate slower than the AW at 5.32 ± 0.04 dbar/yr. This is in contrast to how I found that layers at the top of the staircase were sinking at higher rates than those near the bottom, implying a compression of the staircase. On average, I found layers to be sinking at a rate of 2.24 ± 1.97 dbar per year. This is the same magnitude as the approximately 2–5 m/yr cumulative downwelling rate in the Beaufort Gyre reported by Meneghello et al. (2018) and Proshutinsky et al. (2019). Their rates are based on the Ekman transport derived from surface stress fields and I note the agreement with the rates I found from measurements taken beneath the surface. This implies that, overall, the layers are passively following the movements of the whole water column.

When calculating the average cluster thicknesses, I found a mean value of 1.55 dbar with a standard deviation of 1.53 dbar, which matches the thicknesses reported in previous studies (Padman & Dillon, 1987; Timmermans et al., 2008; Shibley et al., 2017; Lu et al., 2022). I also confirmed the well-reported consensus that layer thickness increases with depth (Neal et al., 1969; Padman & Dillon, 1987; Guo et al., 2016; Shibley et al., 2017; Shibley & Timmermans, 2019; Lu et al., 2022). However, the trends in thickness over time that I found did not indicate a definitive pattern and I conclude that, overall, either layers are not changing in thickness or are becoming slightly thinner. This leaves open several possibilities to reconcile with the overall compression of the staircase. If layer thicknesses are stable over time, then the compression could imply that the interfaces between the well-mixed layers must be decreasing in thickness. This seems unlikely given the interfaces are generally on the order of centimeters (Padman & Dillon, 1987; Neshyba et al., 1971) and also the vertical resolution of ITPs is not fine enough to be able to confirm this (Shibley et al., 2020). Another possibility is that layers are merging, a common process seen in many laboratory studies of layer formation (Turner, 1965; Huppert & Linden, 1979). However, I do not see evidence of layers

systematically merging over time. If the layers are indeed thinning over time, then this immediately reconciles the compression of the staircase.

Previous studies have posited that new staircase layers are formed by intrusions near the bottom of the staircase (Bebieva & Timmermans, 2019; Boury et al., 2022). While I defined the AW core as the maximum temperature in each profile, my effective lower bound was the $S_A \approx 34.9$ g/kg cutoff where I stopped connecting clusters across time periods. I found layers around this cutoff to be stable across the entire 17-year period and therefore cannot determine whether or not new layers were forming. However, at the bottom of the staircase, there is a gap between the layer represented by the lowest regular cluster and the AW core. This lowest regular cluster is sinking at a rate 0.56 ± 0.04 dbar/yr, a magnitude slower than that of the AW; the gap between these two layers is expanding which could leave room for new layers. If layers are disappearing at the top at the same time, this forms an “escalator” motion. At the top of my domain I have a bound of the LHW at 34.1 g/kg. While the layers near that bound were generally not nearly as well-captured by the clustering algorithm compared to lower in the column, I do not see clear indications of layers disappearing over time. While these results do not exclude the possibility of an escalator motion, they do show that such boundary processes, if present, are not taking place within the salinity range 34.1–34.9 g/kg.

Estimates of heat flux through thermohaline staircases have often used parametrizations that depend on the jumps in temperature and salinity between neighboring layers (Kelley, 1990; Shibley et al., 2017; Lu et al., 2022). Making direct measurements of heat flux in this part of the ocean is difficult (Polyakov, Walsh, & Kwok, 2012; Stranne et al., 2017). Using the trends in time I found for the temperature of each cluster, I was able to estimate the net heat flux for each of the detected layers across the 17-year period. I found a significant trend of $d_p(\overline{\Delta\mathcal{F}_H}) = (-2.82 \pm 0.39) \times 10^{-5}$ W/m²/dbar with layers near the top of the staircase having positive net heat flux going to layers near the bottom having negative net heat flux. The point at which the sign changes is around 278 dbar, or 34.6 g/kg, which is also reflected in the inflection point of the cumulative net heat flux. The reason why this inflection point occurs at this particular location is unclear. It could possibly be related to the notion put forward by L22 that layers are more prominently influenced by whichever water mass, LHW or AW, to which they are closest. However, this inflection point does not match the range in which I saw the lateral spatial pattern of temperature in layers reverse ($S_A \approx 34.7$ g/kg). The cumulative net heat flux for all regular clusters amounts to 0.054 W/m². This matches with the magnitude of the estimates from previous studies (Guthrie et al., 2015; Shibley & Timmermans, 2019), however the layer-by-layer analysis reveals a structure in net heat flux within the staircase. Note that these estimates were made by combining the temporal averages of each component. A better estimate of the net heat flux could be obtained by combining the components at each point in the time series, and then evaluating the rate of change.

There are several technical aspects in this study that could be improved upon to employ more automated decision making which would increase this method’s applicability. These include the usage of DBCV to guide the selection of m_{pts} for each year-long time period, connecting clusters across each period, and taking into account the fact that not all clusters appeared in all periods when calculating trends in time. While each of these issues could be eliminated or minimized by clustering all 17 years of data at once instead of in year-long periods, I detail the implications of each below.

Once the dataset to be clustered is determined, HDBSCAN requires only one hyperparameter,

m_{pts} (Logan & Fotopoulou, 2020). For most time periods, I simply chose the value of m_{pts} which resulted in the largest DBCV, following the procedure described in Section 2.3.2. However, for a couple time periods, I instead chose a lower, yet still prominent, value of DBCV to determine m_{pts} as the clustering with the highest DBCV did not identify staircase layers well. Manually inspecting the clustering results for different values of m_{pts} adds time and complexity to the layer detection process, particularly in this case where 17 different datasets were independently clustered. One could reduce the effort needed to select m_{pts} values by restricting the selected salinity range. With less data, there are fewer reasonable values of m_{pts} to consider. Because salinity values tend to plateau around the temperature maximum, a small restriction to the maximum salinity bound could greatly reduce the amount of data to cluster. If only one cohesive dataset needs to be clustered, the extra time needed to manually select an appropriate value of m_{pts} may not be a significant hindrance. Based on the results of this study establishing long coherence times for staircases, it would be reasonable to apply the clustering algorithm once to the entire dataset rather than to each year-long period separately. However, for applications where many independent datasets would be clustered separately, a more algorithmic approach to selecting appropriate values of m_{pts} should be sought.

The manual process I used to connect clusters across time periods does not lend itself to reproducibility and cannot be easily scaled to a dataset with more time periods to connect. For some time periods, certain clusters clearly encompassed more than one layer while in other periods, those layers were clustered separately. In these cases, I took the approach to give different cluster IDs to the erroneously grouped layers in one period and the clusters covering the same layers in other periods. For this reason, I attribute the fact that I found values of IR_{S_A} and R_L larger in magnitude than expected based upon previous studies to this manual connection process. Most clusters had $IR_{S_A} > 1$, compared to Chapter 2 which had the opposite. While the polynomial fit of R_L with respect to pressure had similar coefficients to the fit from Chapter 2, the coefficient of determination of the fit was very low and the mean value $\overline{R_L}$ for the clusters was significantly larger in magnitude. Employing a different process of connecting across time periods, by splitting clusters that erroneously cover multiple layers for example, could bring those values more in line with expectations. However, if only one dataset is being clustered or if all datasets to be clustered are unrelated, then there is no need for a process to connect clusters across time.

Because of the manual process I used to connect clusters across time periods, there were a number of clusters that appeared in very few of the 17 year-long periods. This had an effect on all the trends in time I calculated. Clusters which spanned relatively short amounts of time were particularly susceptible to extreme values unduly influencing the calculated trend. This was particularly notable when calculating trends in layer thickness as the process of calculating layer thickness reduced the number of data points for each cluster, increasing the susceptibility to extreme values. This could be accounted for by using a different method of regression than ordinary least squares which is influenced less by extreme values. Another way would be to consider the amount of time spanned by each cluster when calculating statistics, perhaps marking cluster which cover less than a certain threshold time span as outliers. As this issue is a result of the manual process of connecting clusters across time periods, it would also be negated by clustering the entire time span at once.

In this chapter, I have shown that individual layers in thermohaline staircases can remain coherent across 17 years. With this new temporal coherence scale established, a future study could apply the clustering method to a single dataset which spans a similarly long time period, negating the need to

connect layers across different time periods and eliminating the resulting issue of calculating trends in time for clusters that only span a fraction of the time period. This, along with a more robust method of selecting an appropriate value of m_{pts} would make the clustering algorithm a promising way to detect thermohaline staircases or other similar features in large sets of oceanographic data.

Chapter 4

Modeling internal wave propagation through stratified fluids

Abstract The warm layer below the pycnocline of the Arctic Ocean could dramatically increase the rate of sea ice loss if it were somehow able to rise to the surface. Internal waves, which can be generated by wind and ice floes on the surface, can propagate downwards, providing a possible energy source for heat to mix upward. As ice cover in the Arctic declines, it will be important to predict how internal waves' interactions with stratification profiles will change. In this chapter, I develop numerical experiments to solve the Boussinesq equations of motion using spectral methods. The results match predictions from theory that, for stratification structures with more than one mixed layer, there are particular values of the ratio between layer thickness and wavelength where the wave transmission is much higher than for just one layer. Following [Foran \(2017\)](#), I also use numerical experiments reproduce the results of laboratory wave tank experiments by [Ghaemsaïdi et al. \(2016\)](#) showing the transmission and reflection of internal waves from either one or two mixed layers. For these experiments, I measure the vertical energy flux through the stratification and find that the relative magnitudes of the components match those predicted by theory.

4.1 Introduction

As described in [Section 1.1](#), Arctic sea ice is declining, leaving more of the ocean's surface exposed to wind. With this, there is growing concern over how wind-generated internal waves may affect vertical mixing ([Fer, 2014](#)). The majority of the heat in the Arctic Ocean is in the Atlantic Water (AW) layer, which contains enough heat to melt all Arctic sea ice, if it were to rise to the surface ([Aagaard & Greisman, 1975](#)). Between the cold surface waters and the warm AW below, the intricate stratification structures of the thermohaline staircases pose barriers to vertically propagating internal waves.

As shown by the dispersion relation (see [Section 4.2.2](#)), when an internal wave encounters a region where its frequency ω is greater than the buoyancy frequency N , that wave becomes evanescent and

cannot propagate. As shown in the example of a staircase’s N structure in Figure 4.1(c), $N = 0$ within each well-mixed layer. Therefore, any internal wave with $\omega > 0$ would encounter many evanescent layers upon traversing a staircase. Ray theory predicts that an internal wave incident on an evanescent region will reflect entirely. This prediction is valid when the thickness of the evanescent region L is large compared to the vertical wavelength λ_z , however, it becomes unreliable when $L \sim \lambda_z$ as the theory’s underlying assumption of slowly-varying media no longer holds (Sutherland, 2010). If a region in which the wave can propagate again exists beyond the evanescent one and $L \sim \lambda_z$, then the wave will partially reflect and partially transmit through the layer of uniform density in a process called *internal gravity wave tunnelling* (Eckart, 1961; Sutherland & Yewchuk, 2004).




Figure 4.1: (a) A typical density stratification profile in the Canadian Basin of the Arctic Ocean (Rainville & Winsor, 2008). (b) A magnified section of the profile showing the staircase structure. The red line is a 5 m bin-averaged smoothed density profile. (c) The corresponding profile for the buoyancy frequency $N(z)$. Reproduced from Figure 1 of Ghaemsaidi et al. (2016).

4.1.1 Previous Work

There is a body of previous work on the interactions between internal waves and double-diffusive staircases. Many of these previous studies make differing assumptions and it is still unclear to what extent internal waves can penetrate such stratification structures in particular circumstances (Sutherland, 2016).

Sutherland & Yewchuk (2004) analytically solved the transmission rate of small-amplitude internal waves through one well-mixed layer for the case of a two-dimensional, inviscid, non-rotating, Boussinesq fluid. They also conducted a wave tank experiment which roughly matched their predictions of transmission, even when neglecting the effects of dispersion and dissipation.

While Sutherland & Yewchuk (2004) were mainly concerned with applications of wave tunnelling to the atmosphere, Ghaemsaidi et al. (2016) considered the context of internal waves in the Arctic Ocean where the presence of double-diffusive staircases make it necessary to consider the impact of multiple evanescent layers. Using a weakly viscous linear model, they directly solved for the transmission rate numerically using `bvp4c`, a MATLAB implementation of a fourth-order Runge-Kutta method of solving boundary value problems. Their results show, for particular combinations of the wave’s angle of incidence θ and the ratio between the layer thickness and horizontal wavelength

$L/\lambda_x = k_x L$, a two-layer system exhibits significantly higher internal wave transmission than the one-layer case. This is due to interference patterns of waves between layers (Wunsch, 2018). Additionally, they conducted a laboratory wave-tank experiment with which they found remarkable agreement with the weakly viscous model for both the one and two-layer cases.

Sutherland (2016) extended the theory on transmission through multiple layers, introducing the effects of rotation but neglecting viscosity entirely. They analytically solved for the transmission rate of small-amplitude internal waves through an arbitrary number of mixed layers with equal vertical extents separated by infinitely thin interfaces, establishing bounds on $k_x L$ outside which transmission is negligible. Similar to Ghaemsaidi et al. (2016), Sutherland (2016) implemented their model to be solved numerically. They used this to show that the qualitative aspects of their analytical predictions for steps of uniform height L still hold when the steps are uneven, but with a mean height L .

The models presented by Ghaemsaidi et al. (2016) and Sutherland (2016) are powerful, allowing for the transmission rate to be calculated directly. However, solving directly for the transmission rate becomes more and more difficult as more realistic complexity is added to the system. A different approach is evident from laboratory wave tank experiments; the system is set up and allowed to evolve in time. Here, the transmission rate can be calculated from measurements of the velocity field. While laboratory experiments are an important aspect of gaining insight, numerical experiments offer a flexible avenue to investigate these systems without the need for large and specialized equipment. Foran (2017) used a Fourier spectral Boussinesq code developed by Kraig Winters at the Scripps Institution of Oceanography (Winters et al., 2004) to create a direct numerical experiment (DNS) of the wave tank experiment conducted by Ghaemsaidi et al. (2016). They qualitatively reproduced the results for both the one and two layer cases.

4.1.2 Overview

In this chapter, I detail my progress toward understanding how the transmission of internal waves through stratification structures depends on various factors in situations that more closely represent the reality of the Arctic Ocean. In Section 4.2, I describe the theoretical basis for the equations of motion I use in the numerical experiments as well as the processes of complex demodulation and calculating vertical energy flux which I employ in my analysis. In the codes for this work (available at https://github.com/scheemik/1D_measure-T and https://github.com/scheemik/2D_boundary_forced_waves.v2), I use the Dedalus framework for solving partial differential equations which I describe in Section 4.3. In Section 4.4, I describe the numerical experiments I performed with the equations collapsed to one dimension, similar to Sutherland (2016). I describe the setup of the numerical experiments, how I used complex demodulation to calculate the transmission ratios, and how those ratios compare to the predictions of previous studies. Section 4.5 presents the numerical experiments I performed with the full, two dimensional equations of motion. I describe how, following Foran (2017), I recreated the setup of a previous laboratory wave tank experiment by Ghaemsaidi et al. (2016), hereafter referred to as G16, and compare their results to those of my numerical experiments as well as analyzing the vertical energy flux. In Section 4.6, I discuss the implications of the numerical experiments, comparing and contrasting the results from the two different sets as well as summarizing my conclusions from these experiments and noting possible avenues for future work.

4.2 Theory

In this section, I explain the theoretical basis of my numerical experiments in this chapter including the equations of motion and the dispersion relation. I also detail the processes of how I separate the wavefield by the direction of motion using complex demodulation and how I calculate vertical energy flux.

4.2.1 The Boussinesq Equations of Motion

I base my numerical experiments of internal wave motion on the following formulation of the Boussinesq equations of motion

$$\frac{D\vec{v}}{Dt} = -2\vec{\Omega} \times \vec{v} + \nu \nabla^2 \vec{v} - \nabla p^* + b\hat{z} \quad (4.1a)$$

$$\nabla \cdot \vec{v} = 0 \quad (4.1b)$$

$$\frac{Db}{Dt} = \kappa \nabla^2 b - N^2(z)w \quad (4.1c)$$

where $D/Dt = \partial_t + \vec{v} \cdot \nabla$ is the material derivative, $\vec{v} = (u, v, w)$ is the fluid velocity measured in a Cartesian coordinate system $(\hat{x}, \hat{y}, \hat{z})$, $\vec{\Omega}$ is the rotation rate of the Earth, $\nabla^2 = \partial_x^2 + \partial_y^2 + \partial_z^2$ is the Laplacian, $\nabla = (\partial_x, \partial_y, \partial_z)$ is the gradient, $p^* = p'/\rho_0$ is the specific pressure where the prime represents perturbations in all dimensions (t, x, y, z) and ρ_0 is a constant reference density, $b = -g\rho'/\rho_0$ is the buoyancy, κ is the buoyancy diffusivity, and $N^2(z) = -\frac{g}{\rho_0} \partial_z \hat{\rho}$ is the background stratification frequency where $\hat{\rho}(z)$ represents the background density variations which are taken to only vary in the vertical direction. Equation (4.1a), an approximation of the Navier-Stokes equation, is the momentum equation, (4.1b) is the conservation of mass, and (4.1c) is the thermodynamic equation (Sutherland, 2010; Cushman-Roisin & Beckers, 2011; Lautrup, 2011; Kundu et al., 2015; Vallis, 2017). While I present a full derivation of these equations in Section C.1, here I briefly highlight the pertinent underlying assumptions.

I assume the four relevant forces acting on fluid parcels in the system are due to the Coriolis effect, viscosity, pressure, and gravity. I take gravity to act purely in the vertical direction. I then take the system to cover less than 10° in latitude. This leads to the ‘ f -plane’ approximation where the Coriolis frequency is constant ($2\vec{\Omega} = f\hat{z}$) (Sutherland, 2010). I also assume that seawater is an isotropic, Newtonian fluid whose viscosity is constant across the relevant temperature range.

Generally, the density of water in the ocean only changes by a few percent between the surface and the bottom and by even less in the horizontal directions (Sutherland, 2010). Therefore, I define a constant density ρ_0 as being much greater than the background density profile $\bar{\rho}$ which is in hydrostatic balance with the background profile of pressure, $\partial_z \bar{p} = -\bar{\rho}g$ (Vallis, 2017). This also leads to the Boussinesq approximation which considers density variations ρ' to be negligible unless they affect buoyancy forces because g is relatively large (Vallis, 2017). It follows that seawater is nearly incompressible, and so the conservation of mass (4.1b) states that the divergence of the velocity vanishes (Cushman-Roisin & Beckers, 2011).

I assume there are no internal sources of heat and that the changes to the internal energy are

only due to the diffusion of heat and salt (Kundu et al., 2015). Because of the incompressibility condition, I can assume density has a linear dependence on heat and salt if fluctuations in both are small (Cushman-Roisin & Beckers, 2011). As detailed in Section 1.1, the difference in the magnitudes of diffusivities of heat and salt, κ_T and κ_S , is what gives rise to double diffusion. However, I do not attempt to represent the double diffusive process in my numerical experiments. Rather, I impose a fixed background stratification and make the simplifying assumption that $\kappa_T = \kappa_S = \kappa$ (Cushman-Roisin & Beckers, 2011).

In terms of the streamfunction

The above assumptions lead to the form of the Boussinesq equations in (4.1). However, for certain applications, these equations can be expressed more simply in terms of the streamfunction, ψ . First, I describe a few additional simplifications I employ for the numerical experiments in Section 4.4.

Assuming the amplitude of the waves is small enough such that the terms proportional to amplitude squared (in this case, the nonlinear advection terms) can be neglected, the material derivatives collapse to be just time derivatives, $D/Dt \rightarrow \partial_t$. This is known as the *small amplitude assumption* (Sutherland, 2010; Cushman-Roisin & Beckers, 2011). Here, I maintain the viscous terms in order to compare to other studies. By assuming an isotropic fluid and taking the f -plane assumption, there are no distinctions between horizontal directions. This allows me to always orient the x -axis along the horizontal direction of wave motion and assume a two-dimensional system (Sutherland, 2010). Additionally, I take the length scale of my system to be small enough that the Coriolis force is negligible ($f = 0$). This gives

$$\partial_x u + \partial_z w = 0 \quad (4.2a)$$

$$\partial_t b = -N^2(z)w \quad (4.2b)$$

$$\partial_t u = \nu(\partial_x^2 u + \partial_z^2 u) - \partial_x p^* \quad (4.2c)$$

$$\partial_t w = \nu(\partial_x^2 w + \partial_z^2 w) - \partial_z p^* + b. \quad (4.2d)$$

Because this flow is incompressible and restricted to two dimensions, I can define the streamfunction ψ to be $(u, w) = (-\partial_z \psi, \partial_x \psi)$ (Sutherland, 2010). The system is then reduced (see Section C.1.4 for details) to

$$\partial_t^2 \nabla^2 \psi = \nu \partial_t \nabla^4 \psi - N^2(z) \partial_x^2 \psi \quad (4.3)$$

where $\nabla^4 = \partial_x^4 + 2\partial_x^2 \partial_z^2 + \partial_z^4$ is the biharmonic operator, the square of the Laplacian (Ghaemsaïdi et al., 2016).

4.2.2 Dispersion Relation

When discussing internal waves, it is common to make the assumption of plane waves to derive the dispersion relation. As mentioned above, this relation can give insight into the general features of internal wave propagation under certain circumstances. For illustration, I'll derive here the dispersion

relation for an inviscid fluid, setting $\nu = 0$ in (4.3). The plane waves assumption states

$$\psi(t, \vec{x}) = \psi_0 \exp[i(\vec{k} \cdot \vec{x} - \omega t)] \quad (4.4)$$

where ψ_0 is the amplitude, \vec{k} is the wave vector, and ω is the wave frequency (Sutherland, 2010). In two dimensions, waves will have horizontal and vertical wavenumbers, k_x and k_z . Therefore $\vec{k} \cdot \vec{x} = k_x x + k_z z$ and so the first derivatives of ψ become $\partial_t \psi = -i\omega \psi$, $\partial_x \psi = ik_x \psi$, and $\partial_z \psi = ik_z \psi$ and the second derivatives become $\partial_t^2 \psi = -\omega^2 \psi$, $\partial_x^2 \psi = -k_x^2 \psi$, $\partial_z^2 \psi = -k_z^2 \psi$. Plugging these in to the streamfunction equation allows ψ to be eliminated and results in the dispersion relation

$$k_z^2 = k_x^2 \left(\frac{N^2}{\omega^2} - 1 \right). \quad (4.5)$$

Two main conclusions can be drawn from this. The first is that changing the sign of k_x or k_z does nothing to the dispersion relation. Therefore, for every frequency, there are four possible wavevectors (Mercier et al., 2008). This is demonstrated by the classic oscillating cylinder experiment which creates the “St. Andrew’s Cross” (see Figure 4.2) (Mercier et al., 2008; Sutherland, 2010). The dispersion relation can also be written as $k_z^2 = k_x^2 \tan^2(\theta)$ where $\theta = \pm \arccos(\omega/N)$ is the angle from the vertical along which the phase lines of the waves (i.e., the arms of the cross) are oriented (Sutherland, 2010). The second conclusion is that the waves will be evanescent when $\omega > N$ as k_z would need to be imaginary, meaning that the amplitude of such waves would decay exponentially away from the source (Aguilar et al., 2006). This is analogous to the quantum tunnelling of a particle through a potential barrier (Eckart, 1961; Sutherland & Yewchuk, 2004; Ghaemsaïdi et al., 2016). When including the Coriolis term, the rotating streamfunction equation is $\partial_t^2 \nabla^2 \psi = -f^2 \partial_z^2 \psi - N^2(z) \partial_x^2 \psi$. This leads to the dispersion relation $k_z^2 = k_x^2 (N^2 - \omega^2)/(\omega^2 - f^2)$ which adds the condition that $\omega > f$ in order to not be evanescent. Generally in the ocean, $N > f$, and so propagating waves must have $f < \omega < N$ (Sutherland, 2016).




Figure 4.2: A schematic of a vertically oscillating cylinder experiment where (a) the oscillation frequency ω is less than the stratification frequency N_0 , making the “St. Andrew’s Cross” pattern with wave beams propagating at an angle θ to the vertical, (b) ω is less than but almost equal to N_0 , and (c) $\omega > N_0$ where the waves are evanescent. Reproduced from Figure 3.3 in Sutherland (2010).

4.2.3 Complex Demodulation

Complex demodulation, also sometimes known as Hilbert Transform, is a method that can be used to isolate certain frequencies or wavenumbers from some wavefield $\psi(\vec{x}, t)$ (Croquette & Williams, 1989; Mercier et al., 2008; Grisouard & Thomas, 2015). This process works by using Fourier transforms to construct some complex-valued field $\tilde{\psi}(\vec{x}, t)$ specifically such that taking its real part recovers the original field $\text{Re}[\tilde{\psi}(\vec{x}, t)] = \psi(\vec{x}, t)$. Then, filters on $\tilde{\psi}$ in Fourier space will extract particular components. For example, past studies have used a band-pass filter to isolate near-inertial waves and estimate their vertical displacement amplitudes (Dosser et al., 2014; Dosser & Rainville, 2016; Lincoln et al., 2016). Other studies have used filters to eliminate the inertial component from drift trajectories (Shaw et al., 2009; McPhee, 2012). Here, we filter by the sign of a component of the wavevector \vec{k} to extract just the waves moving in a particular direction.

Following Mercier et al. (2008), we present an example of complex demodulation to demonstrate how it works. While, from the two-dimensional dispersion relation (4.5), we know that a given set of (ω, N) allows for four possible wavevectors corresponding to four different directions of propagation, we consider a simpler wavefield

$$\psi(t, x, z) = A \cos(\omega t - k_x x - k_z z) + B \cos(\omega t - k_x x + k_z z) \quad (4.6)$$

where there are only two wave components: “A” propagating upwards and to the right and “B” propagating downwards and to the right, both with the same wavenumber and frequency. A snapshot of this input wavefield is shown in Figure 4.3(a).

We can rewrite ψ in terms of complex exponentials using Euler’s formula:

$$\psi(t, x, z) = \frac{1}{2}(Ae^{-ik_z z} + Be^{ik_z z})e^{i\omega t - ik_x x} + \frac{1}{2}(Ae^{ik_z z} + Be^{-ik_z z})e^{-i\omega t - ik_x x}. \quad (4.7)$$

The first step, called *demodulation*, obtains the complex-valued field $\tilde{\psi}$. We take a Fourier transform on ψ in time and apply a band-pass filter to isolate the positive frequency ω . When taking this Fourier transform across a finite time span using a simple rectangular filter, it is important that the field is in steady state and the time span is an integer number of wave oscillation periods $T = 2\pi/\omega$.

For this particular situation, a band-pass filter simply involves setting the second term in (4.7) to zero. This eliminates half the amplitude so we then multiply by 2 and take the inverse Fourier transform in time to get the complex-valued signal

$$\tilde{\psi}(t, x, z) = (Ae^{i\omega t - ik_x x})e^{-ik_z z} + (Be^{i\omega t - ik_x x})e^{ik_z z} \quad (4.8)$$

where $\text{Re}[\tilde{\psi}(t, x, z)] = \psi(t, x, z)$. This *demodulation* step is critical as it allows the subsequent filtering based on the wavenumbers (Mercier et al., 2008).

With $\tilde{\psi}$, separating the upward and downward waves is simple after taking a Fourier transform across the z domain: eliminating the negative k_z wavenumbers will leave the upward propagating wave while eliminating the positive k_z wavenumbers will leave the downward propagating wave. The same process can be applied in other spatial dimensions. Similar to before, when taking Fourier transforms across a finite spatial domain, it is important that the domain be an integer number of wavelengths.




Figure 4.3: (a) The input wavefield in equation (4.6). The upward (c) and downward (d) propagating components of the input wavefield, isolated using complex demodulation. (d) The output wavefield, which is the sum of (c) and (d), and is identical to the input.

Figure 4.3(c) shows the upward propagating field $\tilde{A}(t, x, z) = A \exp i(\omega t - k_x x - k_z z)$, the result of taking a Fourier transform in z of $\tilde{\psi}$, filtering to just the positive k_z , and taking the inverse Fourier transform. Figure 4.3(d) shows the downward propagating field $\tilde{B}(t, x, z) = B \exp i(\omega t - k_x x + k_z z)$, resulting from the same process, except filtering to just the negative k_z . The output wavefield shown in Figure 4.3(b) is the result of adding panels (c) and (d), where $\text{Re}[\tilde{A} + \tilde{B}] = \psi$, that is, the components isolated using complex demodulation reconstruct the original wavefield.

4.2.4 Calculating Vertical Energy Flux

I am interested in finding the vertical energy flux per unit mass \mathcal{F}_z^* through a horizontal surface at some depth z . For simplicity, I hereafter drop the asterisk superscript which indicates “per unit mass”, understanding that this vertical energy flux has units of m^2/s^{-3} . I start by taking the dot product of the momentum equation (4.1a) with \vec{v} (Sutherland, 2010)

$$\vec{v} \cdot \partial_t \vec{v} + \vec{v} \cdot (\vec{v} \cdot \nabla \vec{v}) = -\vec{v} \cdot [(f\hat{z}) \times \vec{v}] + \vec{v} \cdot \nu \nabla^2 \vec{v} - \vec{v} \cdot \nabla p^* + \vec{v} \cdot \vec{b} \quad (4.9)$$

where, even though I assume the Coriolis force to be negligible in my system, the Coriolis term would cancel out anyway. This is because $(f\hat{z}) \times \vec{v} = (-fv, fu, 0)$ and $\vec{v} \cdot (-fv, fu, 0) = -fuv + fuv + 0 = 0$.

On the left-hand side of (4.9), I expand the first term using the reverse of the product rule:

$$\vec{v} \cdot \partial_t \vec{v} = \frac{1}{2} \partial_t |\vec{v}|^2 = \partial_t E_k, \quad (4.10)$$

where I've defined kinetic energy per unit mass as $E_k = (1/2)|\vec{v}|^2$. The other term on the left-hand side represents advection and expanding this using the dot product rule gives

$$\vec{v} \cdot (\vec{v} \cdot \nabla \vec{v}) = \vec{v} \cdot \left[\frac{1}{2} \nabla |\vec{v}|^2 - \vec{v} \times (\nabla \times \vec{v}) \right]$$

where the second term in the brackets vanishes because $\vec{v} \cdot [\vec{v} \times (\nabla \times \vec{v})] = 0$. The other term in brackets is ∇E_k and, using a vector identity of divergence,

$$\vec{v} \cdot \nabla E_k = \nabla \cdot (E_k \vec{v}) - E_k (\nabla \cdot \vec{v})$$

where the second term goes to zero because the incompressibility condition (4.1b) states the flow is non-divergent. Therefore,

$$\vec{v} \cdot (\vec{v} \cdot \nabla \vec{v}) = \nabla \cdot (E_k \vec{v}) \quad (4.11)$$

On the right-hand side of (4.9), I expand the Laplacian in the viscosity term

$$\vec{v} \cdot \nu \nabla^2 \vec{v} = \vec{v} \cdot \nu (\partial_x^2 \vec{v} + \partial_y^2 \vec{v} + \partial_z^2 \vec{v})$$

and then, using a rearrangement of the product rule

$$\vec{v} \cdot (\partial^2 \vec{v}) = \partial \left(\frac{1}{2} \partial |\vec{v}|^2 \right) - \partial \vec{v} \cdot (\partial \vec{v}),$$

the viscosity term becomes

$$\vec{v} \cdot \nu \nabla^2 \vec{v} = \nu \nabla \cdot (\nabla E_k) - \epsilon, \quad (4.12)$$

where the first term on the right-hand side represents the diffusion of kinetic energy and $\epsilon = \nu [(\partial_x \vec{v})^2 + (\partial_y \vec{v})^2 + (\partial_z \vec{v})^2]$ is the dissipation of kinetic energy.

For the pressure term on the right-hand side, the product rule states

$$\vec{v} \cdot \nabla p^* = \nabla \cdot (p^* \vec{v}) - p^* \nabla \cdot \vec{v}$$

where, again, the second term vanishes because of the incompressibility condition (4.1b), therefore $\vec{v} \cdot \nabla p^* = \nabla \cdot (p^* \vec{v})$. For the buoyancy term, we have simply $\vec{v} \cdot (b\hat{z}) = wb$. Substituting these along with (4.10), (4.11), and (4.12) into (4.9) gives

$$\partial_t E_k + \nabla \cdot (E_k \vec{v}) = \nu \nabla \cdot (\nabla E_k) - \epsilon - \nabla \cdot (p \vec{v}) + wb,$$

which can be rearranged into

$$\partial_t E_k = -\nabla \cdot \vec{\mathcal{F}} + wb - \epsilon, \quad \text{where} \quad \vec{\mathcal{F}} = (E_K + p^*)\vec{v} - \nu \nabla E_k$$

is the energy flux per unit mass, whose vertical component would then be

$$\mathcal{F}_z = (E_K + p^*)w - \nu \partial_z E_k = \frac{1}{2}|\vec{v}|^2 w + p^* w - \nu \partial_z \left(\frac{1}{2}|\vec{v}|^2 \right), \quad (4.13)$$

where the three terms represent vertical energy flux per unit mass due to advection, pressure, and viscous diffusion, respectively.

4.3 Dedalus

Dedalus ([Burns et al., 2020](#)) is a framework that uses spectral methods to solve arbitrary systems of partial differential equations (PDEs) which I used to perform the numerical experiments in this chapter. While the majority of a Dedalus simulation is spent in optimized C libraries, the front-end is provided as an open-source Python package. A major advantage of Dedalus is that equations, boundary conditions, and other parameters of a problem can be entered in human-readable, plain-text strings which are then translated into solvers. This allows easy entry and modification of the equations for a system which, along with being implemented in an open-source, high-level language like Python, makes Dedalus an accessible framework for solving PDEs. It can also be run in parallel using MPI. Dedalus has been used in many studies across a wide variety of scientific fields. In oceanography, it has been used to model processes such as the restratification ([Callies, 2018](#)) and tracer transport ([Holmes et al., 2019](#)) in abyssal mixing layers, instabilities in the bottom boundary layer ([Wenegrat et al., 2018](#)), as well as mixing of flow under ice keels ([De Abreu et al., 2024](#)).

4.4 Numerical Experiments in One Spatial Dimension

I start developing numerical experiments in a moderately simplified system. The analytical solutions for transmission through multiple layers presented by [Sutherland \(2016\)](#) neglect viscosity. While this may be a reasonable assumption at large scales in the ocean, [Sutherland & Yewchuk \(2004\)](#) note that including viscosity is needed to accurately compare theoretical predictions with laboratory experiments. As I will be comparing my results to the wave tank experiments conducted by G16, I choose to include viscosity, but neglect the Coriolis force, which is negligible at those scales.

4.4.1 Experimental Setup

For these experiments, I represent the system using (4.3), the streamfunction representation of the Boussinesq equations in two dimensions

$$\partial_t^2 (\partial_x^2 \psi + \partial_z^2 \psi) = \nu \partial_t \nabla^4 \psi - N^2(z) \partial_x^2 \psi,$$

which neglects the Coriolis force and diffusion, assumes small amplitude waves, and where $\nabla^4 = \partial_x^4 + 2\partial_x^2 \partial_z^2 + \partial_z^4$ is the biharmonic operator, the square of the Laplacian. This is exactly the equation

used by G16.

Collapsing the horizontal dimension

For these experiments, I assume the flow is bounded above and below by infinite horizontal boundaries. Because of this, and because the stratification depends on z , I need to keep the z -dependence of the streamfunction undetermined at this point, but I can assume that the flow is periodic in the horizontal direction, namely,

$$\psi(t, x, z) = \hat{\psi}(t, z) \exp[ik_x x] \quad (4.14)$$

where k_x is the horizontal wave number (Sutherland, 2010). This is equivalent to the standard plane waves assumption (4.4), except only one spatial component of the wave is explicitly assumed. Note that, while $\hat{\psi}(t, z)$ may be complex valued, I assume the actual flow is the real part of (4.14) (Sutherland, 2010). With this form, I substitute for the derivatives of ψ with respect to x in the above streamfunction equation to give

$$\partial_t^2(\partial_z^2\psi - k_x^2\psi) = \nu(\partial_z^4\psi - 2k_x^2\partial_z^2\psi + k_x^4\psi) + N^2(z)k_x^2\psi \quad (4.15)$$

where now the only direction of consequence is the vertical z . I choose the same value of viscosity as G16, $\nu = 10^{-6} \text{ m}^2 \text{ s}^{-1}$. If viscosity were neglected, equation (4.15) becomes the same as used by Sutherland & Yewchuk (2004).

Domains and Bases

By collapsing the horizontal dimension, I now have only one spatial dimension z to define. In order to use complex demodulation to separate waves by their direction of motion, I need to have equally spaced grid points. I choose to accomplish this by using a Fourier basis in z , which also implicitly sets periodic boundary conditions. In this way, I cannot force waves directly from a boundary. However, I achieve the same effect by dedicating the top portion of the extent to generate waves. This not only creates the intended waves that propagate from the top to the bottom of the domain, but also creates waves that propagate upward, through the periodic boundary, and continue propagating from the bottom to the top of the domain. To dampen these unwanted waves, and to avoid the intended waves from looping back around once reaching the bottom, I dedicate the bottom portion of the extent to be a sponge layer. The spatial domain is then divided into three extents: for the boundary forcing, the measurement, and the sponge layer, as depicted in Figure 4.4.

The sizes of each extent are largely defined in terms of the generated vertical wavelength $\lambda_z = 2\pi/k_z$ which I set to be 1 m for these experiments. The boundary forcing extent starts at $z = 3\lambda_z$ and goes down to $z = 0$. The top of the measurement extent is always at $z = 0$ and the length of the measurement extent is always an integer number of λ_z . However, the exact number of λ_z , and therefore the measurement extent's bottom depth z_f , depends on the particular stratification structure. Similar to the boundary forcing extent, the sponge layer extent is $3\lambda_z$ in length, but extends down past z_f .

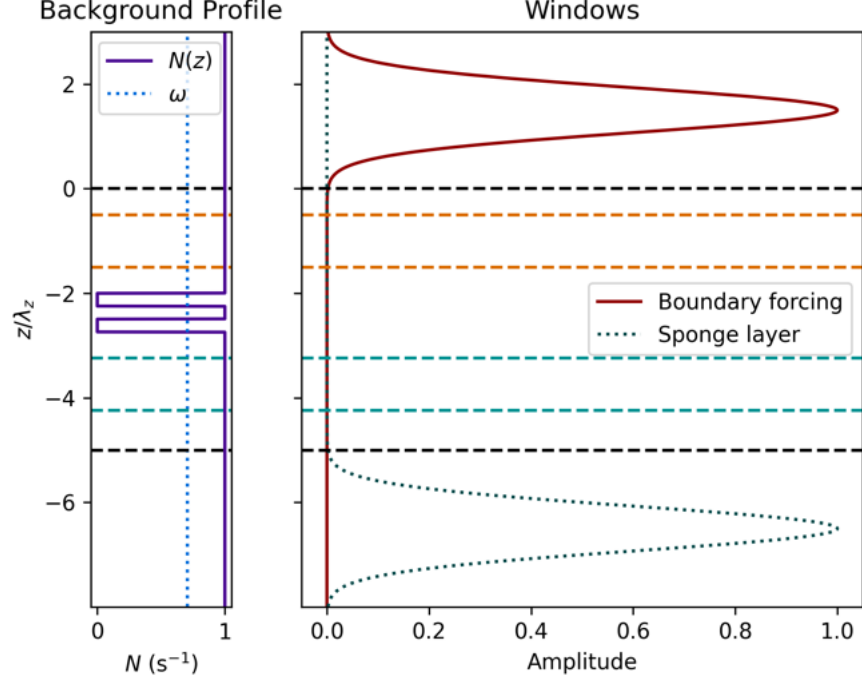


Figure 4.4: A depiction of the domain of the experiment. The two black horizontal dashed lines show the top and bottom of the measurement extent. The orange and teal dashed lines show calculation domains for the amplitudes of the incoming and transmitted wave, respectively. (left) An example background $N(z)$ profile, showing a double-layer stratification structure in the purple solid line and the frequency ω used to force the waves in the blue dotted line. (right) The placement and shape of the windows for the boundary forcing $W_{bf}(z)$ (solid) and sponge layer $W_{sp}(z)$ (dotted).

Stratification profiles

The stratification profiles are defined piece-wise, with $N(z) = 0$ within each of the n layers and $N(z) = N_0$ everywhere else:

$$N(z) = \begin{cases} 0 & z_{min}^i < z < z_{max}^i, \text{ for } i = 1, \dots, n \\ N_0 & \text{elsewhere} \end{cases} \quad (4.16)$$

where I set the reference stratification frequency to $N_0 = 1 \text{ rad s}^{-1}$ and define the top and bottom of each layer i , z_{max}^i and z_{min}^i to be

$$z_{max}^i = z_{min}^{i-1} - R_i L \quad (4.17)$$

$$z_{min}^i = z_{max}^{i-1} - L \quad (4.18)$$

where L is the layer thickness and R_i is the ratio of interface to layer thickness, which I set to be 1.

To give a buffer for waves incoming from the boundary forcing extent, I explicitly set $z_{max}^1 = -2\lambda_z$ as the top of the stratification structure. To provide a buffer for waves exiting the stratification structure before hitting the sponge layer, and to define the measurement extent to be an integer

number of wavelengths, the bottom of the measurement extent is defined to be

$$z_f = z_{min}^n - \lambda_z \left[1 - \left(\frac{z_{max}^1 - z_{min}^n}{\lambda_z} \right) + \left\lfloor \frac{z_{max}^1 - z_{min}^n}{\lambda_z} \right\rfloor \right] - 2\lambda_z. \quad (4.19)$$

where $\lfloor \cdot \rfloor$ represents rounding to the nearest integer. An example of this type of stratification profile with $n = 2$ layers can be seen in Figure 4.4 (left) where $z_f = 5\lambda_z$ is marked by a black dashed line.

Sponge layer forcing

To only dampen waves within the bottom $3\lambda_z$ of the domain, I define a Gaussian sponge layer window

$$W_{sp}(z) = \exp \left[-4 \ln(2) \left(\frac{z - z_{sp,c}}{\lambda_z} \right)^2 \right] \quad (4.20)$$

where $z_{sp,c} = z_f - 3\lambda_z/2$ is the center of the Gaussian function with a full-width at half max of λ_z . The shape and position of this window is depicted in Figure 4.4(right) and is set to provide a smooth transition of increasing strength as a sudden onset of a strong sponge layer can cause waves to reflect instead of being dissipated (Jablonowski & Williamson, 2011).

Inside this window, I apply a Rayleigh friction sponge forcing in the form

$$F_{sp}(\psi, z) = \frac{W_{sp}(z)}{\tau_{sp}} (\partial_z^2 \partial_t \psi - k^2 \partial_t \psi) \quad (4.21)$$

where I set the time constant to be $\tau_{sp} = 1$ s. This counteracts the wave field, effectively damping any motion that enters the sponge layer window. Rayleigh friction is used as a sponge layer at the top of the atmosphere in some Global Climate Models (GCMs) (Jablonowski & Williamson, 2011).

Boundary forced waves

In order to force waves only within the top $3\lambda_z$ of the domain, I define a Gaussian boundary forcing window in a similar way to the sponge layer

$$W_{bf}(z) = \exp \left[-4 \ln(2) \left(\frac{z - z_{bf,c}}{\lambda_z} \right)^2 \right] \quad (4.22)$$

where $z_{bf,c} = 3\lambda_z/2$ is the center of the Gaussian function, again with a full-width at half max of λ_z . The shape and position of this window is depicted in Figure 4.4(right).

Inside this window, I force plane waves using a similar assumption to (4.14) but in the vertical dimension

$$F_{bf}(\psi, z) = \frac{W_{bf}(z)}{\tau_{bf}} \left[\hat{A} \sin(k_z z - \omega t) - \psi \right] \quad (4.23)$$

where I set the time constant to be $\tau_{bf} = 1$ s, the forced wave amplitude to be $\hat{A} = 2 \times 10^{-4}$ (unitless), and the wave frequency to be $\omega = N_0 \cos(\theta)$ which, because I define $N_0 = 1 \text{ rad s}^{-1}$ and, as detailed below, $\theta = 45^\circ$, I have $\omega = 0.7071 \text{ s}^{-1}$.

Equation of motion

Adding in the forcing terms for the boundary forcing (4.23) and the sponge layer (4.21), the one dimensional equation of motion (4.15) becomes

$$\partial_t (\partial_z^2 \partial_t \psi - k_x^2 \partial_t \psi) - \nu (\partial_z^4 \psi - 2k_x^2 \partial_z^2 \psi + k_x^4 \psi) = k_x^2 N^2(z) \psi + F_{bf}(\psi, z) - F_{sp}(\psi, z) \quad (4.24)$$

which is the equation I use in the code for this work.

4.4.2 Methods

In this work, I am interested in quantifying how much of a wave that is incident on the top of a stratification structure transmits out the bottom for a range of different structures. This is calculated by comparing the amplitude squared of the incident wave $A_{\mathbb{I}}$ to the amplitude squared of the transmitted wave $A_{\mathbb{T}}$ to find the transmission coefficient

$$\mathbb{T} = A_{\mathbb{T}}/A_{\mathbb{I}} \quad (4.25)$$

which is equivalent to the fraction of the energy of the incident wave which is transmitted through the stratification. In the case of one layer, the analytical solution of (4.25) is

$$\mathbb{T}_{ana} = \left[1 + \left(\frac{\sinh(k_x L)}{\sin(2\theta)} \right)^2 \right]^{-1}, \quad (4.26)$$

where k_x is the horizontal wave number, L is the layer thickness, and $\theta = \arccos(\omega/N_0)$ is the incident angle (Sutherland & Yewchuk, 2004). From this, the largest values of \mathbb{T} occur when $\theta = 45^\circ$ and so I choose this angle for all of my experiments in order to have the largest theoretically possible signal. I choose to run my experiments over a range of $k_x L$ values to see how \mathbb{T} changes. For a lower limit, when $k_x L = 0$, there is no stratification structure and $\mathbb{T} = 1$. As $k_x L$ increases, \mathbb{T} asymptotically approaches zero, becoming less than 2×10^{-4} when $k_x L = 5$, and so I choose this as my upper limit.

Analytical solutions exist for stratification structures with more than one layer (Sutherland, 2016), which are complex enough that reproducing them here would not be illuminating. These solutions of \mathbb{T} for multiple layers with $\theta = 45^\circ$ follow the same edge conditions as one layer with $\mathbb{T}(k_x L = 0) = 1$ and reaching $\mathbb{T} \approx 0$ around $k_x L = 5$. However, unlike (4.26), they do not monotonically decrease in \mathbb{T} , instead having spikes of high transmission for particular values of $k_x L$. In order to have enough resolution to resolve these features, for each number of layers, I run 128 experiments evenly spaced across values of $k_x L$ from 0 to 5.

Selecting the timestep

In order to calculate $A_{\mathbb{I}}$ and $A_{\mathbb{T}}$, I extract just the downward propagating wavefield $\tilde{\psi}_{\downarrow}$ by using complex demodulation, a process detailed in Section 4.2.3. This process requires the spatial extent to be an integer number of that dimension's wavelengths. Due to the way I define the measurement extent in (4.19), this is always satisfied. Using complex demodulation also requires a wavefield in steady state that lasts an integer number of wave periods, $T = 2\pi/\omega$.

While any integer number of T would suffice, I consider that complex demodulation involves taking Fourier transforms. The implementation of the Fast Fourier Transforms (FFTs) used in the code expect the array size to be a power of two. Not only is this implementation of FFT most efficient when the data array length is already a power of two, but padding the array until it reaches the next power of two can sometimes produce errors in the calculation, especially for periodic functions which are not zero outside the chosen domain (Briggs & Henson, 1995). Therefore, I choose a time step size and number of wave periods such that the number of time steps over which I perform complex demodulation is a power of two. Specifically, I choose a time step $\Delta t = T/64$ and an interval of 8 wave periods so that I perform complex demodulation over 512 time steps for each experiment.

Measuring the transmission coefficient

Measuring the transmission coefficient \mathbb{T} for an experiment involves trimming the wavefield ψ to just the measurement extent in both time and space, performing complex demodulation (see Section 4.2.3) to extract just the downward propagating wavefield $\tilde{\psi}_\downarrow$, and calculating the amplitudes squared $A_\mathbb{I}$ and $A_\mathbb{T}$. As an example to illustrate this process, Figure 4.5 shows the total resulting wavefield ψ for a experiment where $\theta = 45^\circ$ and the stratification had one layer with $k_x L = 1.55$.




Figure 4.5: The total wavefield of an example experiment where $k_x L = 1.55$ and $\theta = 45^\circ$. The solid red line represents the wave propagation speed c_{gz} . The two black horizontal dashed lines show the top and bottom of the measurement extent. The black vertical dashed line marks the time at which steady state is assumed. The orange and teal dashed lines show calculation domains for the amplitudes squared of the incoming and transmitted wave, respectively. After steady state was reached, $\mathbb{T} = 0.16$.

The measurement extent, where I measure the amplitudes squared, should reach steady state after the first waves have been able to propagate through its bottom boundary z_f . In an abundance of caution, I allow a transient period of time it would take a wave moving at the vertical group velocity c_{gz} to be generated, travel down to the bottom of the measurement extent, then back up

to the top of the measurement extent. That is

$$t_{transient} = \frac{1}{c_{gz}}(|z_{bf,c}| + 2|z_f|) \quad (4.27)$$

where for the example experiment shown in Figure 4.5, $t_{transient}/T = 23$ and is marked by the vertical black dashed line. As I choose to perform complex demodulation over $8T$, the end of the experiment is then $t_f = t_{transient} + 8T$, where for the example experiment, $t_f/T = 31$. Figure 4.6 shows $\tilde{\psi}_\downarrow$.




Figure 4.6: The downward propagating wavefield in the measurement domain of an example experiment where $k_x L = 1.55$ and $\theta = 45^\circ$. The wave is incident on the layer from above and the lower amplitude of the transmitted wave can be seen. After steady state was reached, $\mathbb{T} = 0.16$. The orange and teal dashed lines show calculation domains for the amplitudes of the incoming and transmitted wave, respectively.

For both amplitudes squared of interest, $A_\mathbb{I}$ and $A_\mathbb{T}$, I choose a range in z over which to average that I define to be λ_z in length and $\lambda_z/2$ away from the stratification structure. Specifically, the orange dashed lines in Figure 4.6 are at $z_{min}^{A_\mathbb{I}} = z_{max}^1 + \lambda_z/2$ and $z_{max}^{A_\mathbb{I}} = z_{min}^{A_\mathbb{I}} + \lambda_z$ where z_{max}^1 is the top of the stratification structure. Similarly, the teal dashed lines are at $z_{max}^{A_\mathbb{T}} = z_{min}^n - \lambda_z/2$ and $z_{min}^{A_\mathbb{T}} = z_{max}^{A_\mathbb{T}} - \lambda_z$ where z_{min}^n is the bottom of the stratification structure.

I then calculate the amplitudes squared by

$$A = \langle |\tilde{\psi}_\downarrow| \rangle = \langle \tilde{\psi}_\downarrow \cdot \tilde{\psi}_\downarrow^{cc} \rangle \quad (4.28)$$

where $\tilde{\psi}_\downarrow^{cc}$ is the complex conjugate of $\tilde{\psi}_\downarrow$ and $\langle \rangle$ represents the average across both the time dimension between $t_{transient}$ and t_f and the spatial dimension between the applicable bounds; for $A_\mathbb{I}$, I average in the z direction between $z_{min}^{A_\mathbb{I}}$ and $z_{max}^{A_\mathbb{I}}$ and for $A_\mathbb{T}$, between $z_{min}^{A_\mathbb{T}}$ and $z_{max}^{A_\mathbb{T}}$. Having both $A_\mathbb{I}$ and $A_\mathbb{T}$, I can calculate the transmission coefficient which, for the example experiment in Figure 4.6, is

$\mathbb{T} = 0.16$.

4.4.3 Results

For stratification structures with 1, 2, 3, 4, and 5 layers, I perform 128 experiments with values of $k_x L$ ranging from 0 to 5. For one layer, every experiment with $k_x L > 2.5$ had $\mathbb{T} < 0.025$. For 2, 3, 4, and 5 layers, every experiment with $k_x L > 2.5$ had $\mathbb{T} < 0.01$. Therefore, Figure 4.7 shows these experiments for $k_x L$ from 0 to 2.5. The experiments for one layer are in excellent agreement with the values predicted by (4.26).

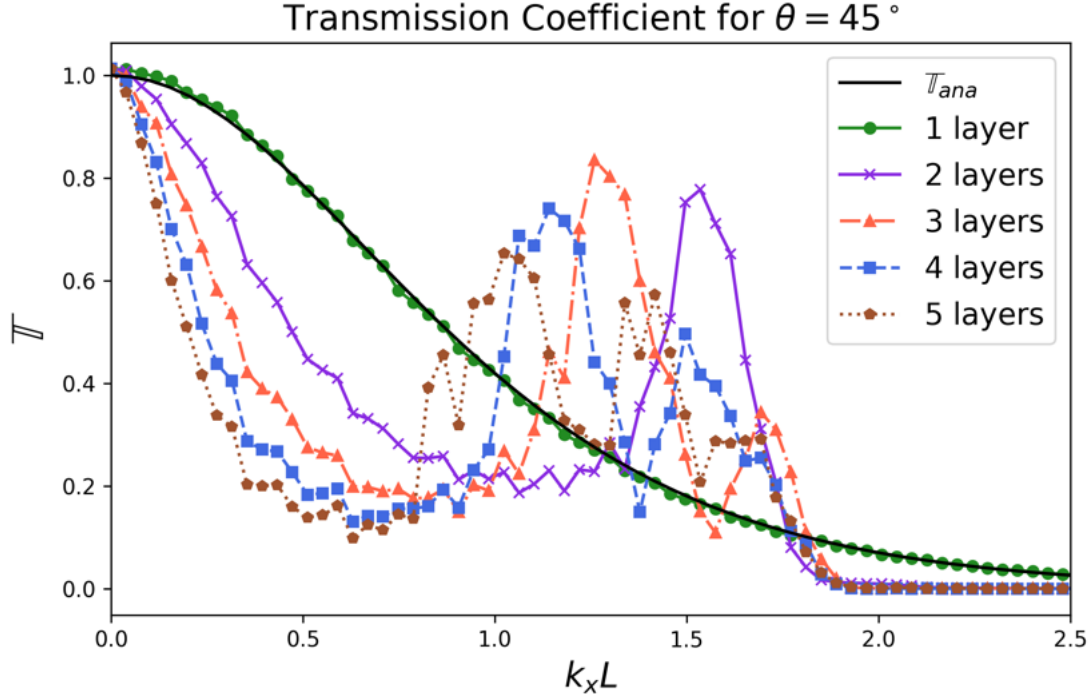


Figure 4.7: The transmission coefficient \mathbb{T} for the analytical solution of 1 mixed layer (4.26) and experiments of 1, 2, 3, 4, and 5 mixed layers. Each marker represents one experiment.

For the experiments with multiple layers, there are peaks in transmission at intermediate values of $k_x L$. This is a manifestation of the internal gravity wave tunnelling phenomenon (Sutherland & Yewchuk, 2004). The largest of these peaks for each number of layers decreases in $k_x L$ as the number of layers increases. These results cannot be directly compared to the predictions of Sutherland (2016) as I use finite-thickness interfaces. However, I compare to the results of G16, who modeled a two layer system with a finite-thickness interface.

G16 numerically solved both the inviscid and viscous equations directly for the transmission coefficient. While I was able to reproduce their results in the inviscid case with the codes they made available, I was unable to do so for the viscous case. Therefore, I make quantitative comparisons only to their inviscid results in Figure 4.8. For the one layer case, the results of my viscous numerical experiments as well as the inviscid results of G16 agree very well with the analytical prediction of Sutherland & Yewchuk (2004). For the two-layer case, the two methods agree in the value of $k_x L$ for which the peak in \mathbb{T} occurs and the approximate width of the resonance curve, but the G16 method finds a larger \mathbb{T} .

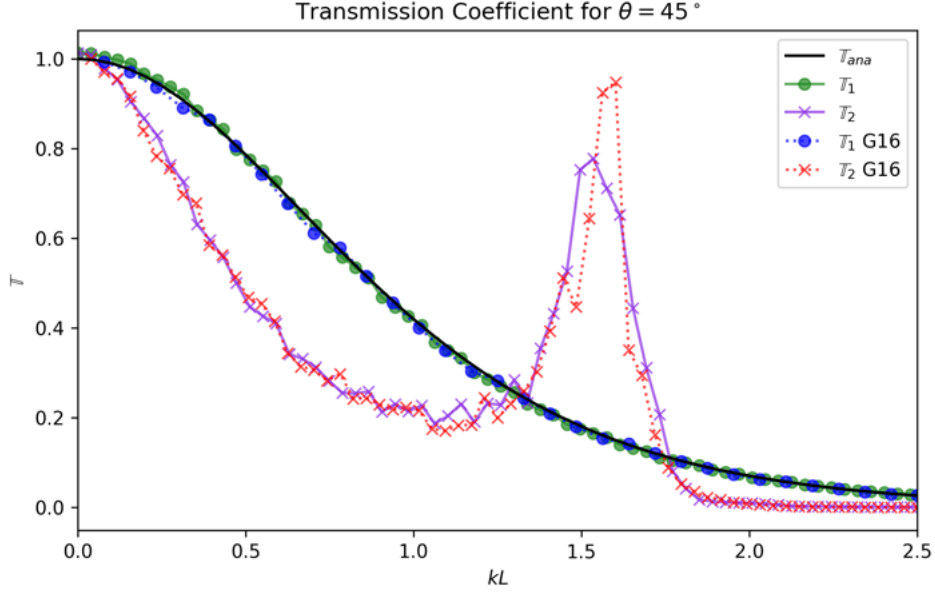


Figure 4.8: The transmission coefficient \mathbb{T} for the analytical solution of one mixed layer (4.26) and experiments of one and two mixed layers as calculated with Dedalus in this study or by the inviscid MATLAB code of G16. Each marker represents one experiment.

G16 ran both inviscid and viscous numerical experiments and found the magnitude in \mathbb{T} of the peaks to be lower in the viscous case. As my experiments included viscosity, I find agreement in this as the maximum \mathbb{T} I found in the transmission peak is lower than that for G16’s inviscid case. However, they also showed that the value of $k_x L$ for the maximum \mathbb{T} in the peak is noticeably lower in the viscous case compared to their inviscid experiments. This is opposed to how I found agreement in the $k_x L$ of the transmission peak in my viscous experiments compared to their inviscid case. While the reason for this discrepancy is unclear, the overall agreement suggests that the experiments in Dedalus correctly represent the theoretical predictions for this scenario. In the Dedalus framework, the equations of motion are explicitly written in to the code, offering a more transparent and more easily modifiable script than those used by G16.

For the case of more than two mixed layers, only qualitative comparisons can be made to the predictions of Sutherland (2016) as they assumed infinitely thin interfaces and plot the transmission coefficient on kL vs ω/N_0 . A similarity between Figure 4.7 and the results of Sutherland (2016) are that, with more layers, there are more peaks in transmission. The result that the largest peak for each number of layers occurs at smaller values of $k_x L$ as the number of layers increases remains to be confirmed.

4.5 Numerical Experiments in Two Spatial Dimensions

As demonstrated above, the transmission through stratification structures of well-mixed layers depends on the interference patterns of waves, creating peaks in transmission when that interference is constructive (Wunsch, 2018). In the previous section, I assumed infinite plane waves in the horizon-

tal, effectively collapsing the system to just one vertical spatial dimension. Allowing for horizontal variations, as more realistic waves would have, changes the manner of interference, and thus the transmission, when encountering layered stratification structures. In order to allow for these effects, and to accurately reproduce the results of wave tank experiments, I turn my attention to a fully two-dimensional system. G16 performed a laboratory experiment to investigate the propagation of internal waves through vertical stratification. Here, I explain how, following [Foran \(2017\)](#), I worked to recreate the conditions of those wave tank experiments within my Dedalus numerical experiments.

4.5.1 Experimental setup

Because Coriolis effects were not present in the laboratory experiment performed by G16 and they are not essential to the dynamics of interest, I choose to neglect rotation. Taking the Boussinesq equations of motion (4.1), expanding the material derivatives, and neglecting the Coriolis term gives the following system of equations:

$$\partial_t \vec{v} + \vec{v} \cdot \nabla \vec{v} = \nu \nabla^2 \vec{v} - \nabla p^* + b \hat{z} \quad (4.29a)$$

$$\nabla \cdot \vec{v} = 0 \quad (4.29b)$$

$$\partial_t \vec{b} + \vec{v} \cdot \nabla \vec{b} = \kappa \nabla^2 \vec{b} - N^2(z) w \quad (4.29c)$$

which, again, assumes a vertically density-stratified, nearly incompressible, and non-rotating fluid. These equations serve as a base for these experiments in Dedalus using a two-stage, second-order, diagonally implicit Runge-Kutta time stepping scheme with a constant $dt = 0.125$ s ([Ascher et al., 1997](#)). I use the same value of viscosity, $\nu = 10^{-6} \text{ m}^2 \text{ s}^{-1}$, as G16 which is the viscosity of pure water at 20°C. The diffusivity (κ) used was not specified, so I use the corresponding value for pure water at 20°C: $\kappa = 1.4 \times 10^{-7} \text{ m}^2 \text{ s}^{-1}$ ([Kundu et al., 2015](#)).

Domains and Bases

The wave tank experiments performed by G16 took place in a wave tank depicted schematically in Figure 4.9. The tank was 0.54 m deep and 5.46 m long. The width was 0.55 m, but a partition reduced the working section to a width of 0.2 m. Parabolic reflection barriers at either end directed unwanted waves generated in the working section around the partition into the 0.35 m wide dissipation section ([Echeverri, 2009](#)). G16 took measurements in a domain 0.5 m deep and 0.5 m long, centered along the length of the tank.

While the wave tank experiments were three dimensional, I choose to simplify the experiments down to two dimensions to reduce the required computing time and memory resources. This is a reasonable simplification for this study because, in the experiment performed by G16, the wave forcing in the tank did not vary along the width, the width of the tank was relatively small compared to the depth and the length, and they only recorded the wave motion in two dimensions. Therefore, I now assume $\nabla = (\partial_x, \partial_z)$ and $\vec{v} = (u\hat{x}, w\hat{z})$ in equations (4.29).

In my experiments, the full domain is 1.5×1.5 m, with the bounds $-0.5 \leq x \leq 1$ in the horizontal and $0 \leq z \leq -1.5$ in the vertical and with 512 grid points in each direction. I use the same 0.5×0.5 m size measurement domain as G16, with the top at $z_0 = 0$, the bottom at $z_f = -0.5$ and the horizontal measurement extent $0 \leq x \leq 0.5$. As the equations of motion (4.29) contain only up




Figure 4.9: (Left) A schematic of the experimental setup used by G16 showing the dimensions of the tank and the measurement area as well as the positioning of the wave generator. (Right) Measurements of the density stratifications, $\Delta\rho = \rho - 1000$ for the (a) single- and (b) double-mixed-layer experiments. Reproduced from Figure 4 in G16.

to quadratic nonlinearities, I use a dealiasing factor of $3/2$ in both spatial dimensions (Boyd, 2001; Burns et al., 2020). My primary interest for this project is vertical motion so I use a Chebyshev basis in the z direction which, within the Dedalus framework, allow me to force waves from the boundary and define non-constant coefficients (NCC's) that vary in the vertical, such as the background stratification $N(z)$ and sponge layer $S(z)$.

The vertical extent below the measurement domain is used as a sponge layer, extending from $z_{sp,0} = -0.5$ m down to $z_{sp,f} = -1.5$ m. Using a Chebyshev basis in the vertical also allows me to use a lower resolution in the sponge layer, saving computational resources. While I use a total of 512 grid points in the vertical, 472 were in the upper layer while 40 were in the sponge layer. In these experiments, I apply a diffusive sponge $S(z)$ to the viscosity ν , by which I increase the effects of viscosity near the bottom, artificially dampening downward propagating waves to prevent reflections caused by hitting the bottom boundary. Specifically, I define the sponge layer to be

$$S(z) = 1 + \frac{1}{2}(H_{sp} - 1) [\tanh(s_{sp}(z - z_{sp,c})) + 1] \quad (4.30)$$

where I set the amplitude of the sponge layer as $H_{sp} = 50$ (unitless), the slope of the sponge layer to be $s_{sp} = -5 \text{ m}^{-1}$, and $z_{sp,c} = z_{sp,f} - (2/3)(z_{sp,f} - z_{sp,0}) = -7/6$ m as the center of the sponge layer ramp. A depiction of this coefficient can be seen in Figure 4.10(a) and I apply this to the value of viscosity ν in the equations of motion. The values of the constants in (4.30) were chosen to give the sponge layer a high enough amplitude to stop downward waves with a gentle enough slope to avoid reflecting the waves back upwards (Jablonowski & Williamson, 2011).

For the horizontal direction, I use a Fourier basis which implicitly defines periodic boundary conditions. The waves were generated at the top left edge of the measurement domain. While most of the unwanted interference was eliminated by the sponge layer, the extra 1 m of simulated horizontal domain and the periodic boundary conditions in x allows enough time for the experiments




Figure 4.10: (a) The profile of the sponge layer given by equation (4.30). The (b) single- and (c) double-mixed-layer stratification profiles in solid lines used to emulate those of G16 where the dotted lines represent the value of the forcing frequency ω used in each case. The black dashed lines represent the bottom of the measurement domain.

to develop before unwanted reflections were able to reappear in the measurement domain.

Stratification profiles

G16 constructed two different stratification profiles in the wave tank: a single mixed layer approximately 7 cm thick as shown in Figure 4.9(a) and a double mixed layer with two layers approximately 7 cm thick separated by approximately 3.5 cm as shown in Figure 4.9(b). In both profiles, the stratification above and below the mixed layers was $N_0 = 0.95 \text{ rad s}^{-1}$ and $N_f = 1.24 \text{ rad s}^{-1}$, respectively, while $N = 0$ within the layers.

In the laboratory experiment, the stratification took a few centimeters to transition between values, making the layer edges smoother than the sharp transitions seen in Section 4.4. To reproduce these smooth and continuous layer edges, I form both the single- and double-layer stratification profiles using combinations of tanh functions. For the single layer case, I set

$$N_1(z) = 0.5N_0 [\tanh(s_{st}(z - z_{st,0})) + 1] + 0.5N_f [\tanh(-s_{st}(z - z_{st,f})) + 1], \quad (4.31)$$

where the top and bottom bounds of the stratification are $z_{st,0} = -0.22$ m and $z_{st,f} = -0.30$ m, respectively, and the stratification slope is $s_{st} = 200$. This is depicted in Figure 4.10(b). For the double layer case, I set

$$N_2(z) = 0.5N_0 [\tanh(s_{st}(z - z_{st,0})) + 1] + 0.5N_f [\tanh(-s_{st}(z - z_{st,f})) + 1] + N_{int}(z) \quad (4.32)$$

where the only differences from the one layer case are the bottom stratification bound $z_{st,f} = -0.38$ m and the stratification for the interface:

$$N_{int}(z) = 0.5N_i [\tanh(s_{st}(z - z_{int,0})) + 1] + 0.5N_i [\tanh(-s_{st}(z - z_{int,f})) + 1] \quad (4.33)$$

where the height of the interface is $N_i = 0.5(N_0 + N_f) = 1.095 \text{ rad s}^{-1}$ and the center of the interface is $z_{int,c} = z_{st,f} + (z_{st,0} - z_{st,f})/2 = -0.3$ m which makes the top and bottom bounds of the interface $z_{int,0} = z_{int,c} - 0.025$ m and $z_{int,f} = z_{int,c} + 0.025$ m, respectively, and so the thickness of the interface is 5 cm. This is depicted in Figure 4.10(c).

Boundary Conditions

The equations of motion for this problem (4.29) contain time derivatives of b , u , and w , making them state variables and leaving one diagnostic variable, namely, p^* (Cushman-Roisin & Beckers, 2011). Each of the state variables requires initial conditions as well as conditions on their values at the horizontal and vertical boundaries of the experiment at all times.

By choosing a Fourier basis in x , I implicitly set periodic boundary conditions in the horizontal direction. Using a Chebyshev basis in z allows me to directly set the state variables along the top and bottom boundaries, that is, I set six Dirichlet boundary conditions in the vertical direction. For the first three, I set the state variables equal to zero at the bottom of the domain, $u(z_f) = w(z_f) = b(z_f) = 0$, to guarantee no waves can originate from the bottom. At the top of the domain (z_0), I prescribe functions for the boundary conditions on u , w , and b to force downward-propagating internal waves.

In the wave tank, G16 vertically oscillated a 7.5 cm diameter cylinder which spanned the width of the working section. Oscillations had an amplitude of 7 mm and the cylinder was located just outside the top left corner of the measurement domain. This generated an internal wave beam which propagated through the measurement domain, downwards and to the right. To reproduce these waves in my experiments, I find the polarization relations between the variables starting with the plane wave assumption (4.4) in two spatial dimensions

$$\eta(t, x, z) = \eta_0 \exp[i(k_x x + k_z z - \omega t)], \quad (4.34)$$

where k_x and k_z are the horizontal and vertical wavenumbers, ω is the frequency of the boundary forcing, η represents any one of the variables u , w , p^* , or b , and η_0 is an arbitrary coefficient.

In the plane wave assumption (4.34) for buoyancy, I define the coefficient to take the form $b_0 = -\hat{A}g$. With this, I make the *small amplitude assumption*, that is, \hat{A} is small enough such that the nonlinear terms, which scale with the amplitude squared, can be neglected, linearizing the equations of motion (Sutherland, 2010; Cushman-Roisin & Beckers, 2011). With these linearized

equations, the polarization relation is then

$$u = -\hat{A}g \frac{\omega k_z}{N^2 k_x} \sin(k_x x + k_z z - \omega t) \quad (4.35a)$$

$$w = +\hat{A}g \frac{\omega}{N^2} \sin(k_x x + k_z z - \omega t) \quad (4.35b)$$

$$b = -\hat{A}g \cos(k_x x + k_z z - \omega t). \quad (4.35c)$$

$$p^* = -\hat{A}g \frac{\omega^2 k_z}{N^2 k_x^2} \sin(k_x x + k_z z - \omega t). \quad (4.35d)$$

where, in the experiments for this work, I use $\hat{A} = 2.3 \times 10^{-4}$ (unitless). A complete derivation of this polarization relation can be found in Section C.2.

The oscillating frequencies used by G16 in the scenarios I reproduce were not specified. Following Foran (2017), who successfully reproduced the experiment through different numerical experiments, I choose the frequencies $\omega_1 = 0.67 \text{ s}^{-1}$ and $\omega_2 = 0.72 \text{ s}^{-1}$ for the single and double mixed layer profiles, respectively. Both G16 and Foran (2017) used a characteristic total wavenumber of $k = 45 \text{ m}^{-1}$ for the boundary forcing. Through trigonometry and the dispersion relation of internal waves (4.5)

$$\theta = \arccos(\omega/N), \quad k_x = k \cos(\theta), \quad k_z = k \sin(\theta), \quad (4.36)$$

those frequencies correspond to an angle below the horizontal $\theta_1 = 45.1^\circ$, $k_{x,1} = 31.7 \text{ m}^{-1}$, and $k_{z,1} = 31.9 \text{ m}^{-1}$ for the single layer and $\theta_2 = 40.7^\circ$, $k_{x,2} = 34.1 \text{ m}^{-1}$, and $k_{z,2} = 29.4 \text{ m}^{-1}$ for the double layer. These values are summarized in Table 4.1.

The state variables also require initial conditions to be specified (Cushman-Roisin & Beckers, 2011), and I set them all to zero across the domain. Because because of this, using (4.35) directly as the top boundary forcing causes a transient oscillation that eventually dies out. To mitigate this transient, I apply a temporal ramp function to the boundary forcing

$$R_{bf}(t) = \frac{1}{2} \left(\tanh \left[\frac{4t}{nT} - 2 \right] + 1 \right), \quad (4.37)$$

where $T = 2\pi/\omega$ is the oscillation period and I set the ramp to last $n = 3$ oscillation periods because it is long enough to make the transient oscillation negligible but not long enough to make the experiment spin up time cumbersome.

To simulate the finite extent of the cylindrical wave generator, I apply a windowing function

$$W_{bf}(x) = \frac{1}{4} \{ \tanh[s_w(x - x_l)] + 1 \} \{ \tanh[-s_w(x - x_r)] + 1 \} \quad (4.38)$$

to the top boundary forcing where s_w is the slope of the window's sides, x_l is the window's left edge, and x_r is the window's right edge. I set $s_w = 20$ for a balance between edges too steep to run the code efficiently and edges too shallow to form well-defined wave beams. To have the waves originate from the top left corner of the display domain, I set $x_l = -\lambda_x/2$ and $x_r = \lambda_x/2$ where $\lambda_x = 2\pi/k_x$ is the horizontal wavelength.

With the parameters set as stated, I apply both the ramp (4.37) and windowing (4.38) to the polarization relation (4.35) for the state variables in order to get the following boundary conditions

at the top of the domain:

$$u(z_0) = -\hat{A} \frac{g\omega k_z}{N_0^2 k_x} \sin(k_x x + k_z z - \omega t) W_{bf}(x) R_{bf}(t) \quad (4.39a)$$

$$w(z_0) = +\hat{A} \frac{g\omega}{N_0^2} \sin(k_x x + k_z z - \omega t) W_{bf}(x) R_{bf}(t) \quad (4.39b)$$

$$b(z_0) = +\hat{A} g \cos(k_x x + k_z z - \omega t) W_{bf}(x) R_{bf}(t) \quad (4.39c)$$

This creates an internal wave beam that ramps up to its full amplitude after a few oscillation periods, propagating from the top left towards the bottom right of the measurement domain.

Condition for Linearity

In this project, I am only interested in investigating linearly stable flows. This could be ensured by monitoring the Richardson number $Ri = N^2/(\partial_z u)^2$ because if $Ri > 1/4$, then linear stability is guaranteed (Kundu et al., 2015). However, since Ri is proportional to N^2 and I set $N = 0$ explicitly in sections of the background profile, this is not a viable option.

To compare the relative magnitudes of linear versus nonlinear terms, I use the horizontal momentum equation (4.29a), defining a linearity criterion C_{lin} as the ratio of the nonlinear advection term $u\partial_x u + w\partial_z u$ and the linear term $\partial_t u$. Under the plane wave assumption (4.34) the derivatives become $\partial_x \rightarrow ik_x$, $\partial_z \rightarrow ik_z$, and $\partial_t \rightarrow -i\omega$. Therefore, I find the condition for linearity to be when

$$C_{lin} = \frac{k_x u + k_z w}{\omega} \ll 1 \quad (4.40)$$

is satisfied everywhere in the flow. For all experiments in this work, C_{lin} is on the order of 0.1 or smaller at all times.

Equations of motion

Using equations (4.29) in two spatial dimensions, I apply the sponge layer $S(z)$ given by (4.30), and define the background stratification $N(z)$ as either (4.31) (single layer) or (4.32) (double layer) to give

$$\partial_t u - S(z)\nu(\partial_x^2 u + \partial_z^2 u) + \partial_x p = -(u\partial_x u + w\partial_z u) \quad (4.41a)$$

$$\partial_t w - S(z)\nu(\partial_x^2 w + \partial_z^2 w) + \partial_z p - b = -(u\partial_x w + w\partial_z w) \quad (4.41b)$$

$$\partial_x u + \partial_z w = 0 \quad (4.41c)$$

$$\partial_t b - \kappa(\partial_x^2 b + \partial_z^2 b) = -N^2(z)w - (u\partial_x b + w\partial_z b) \quad (4.41d)$$

which, along with the boundary conditions, are the equations I use in the code for this work. The values of the constants are summarized in Table 4.1.

4.5.2 Methods

In the previous set of experiments, I focused my analysis on the transmission coefficient, measuring the amplitude squared of the wave that is transmitted through a stratification structure in comparison to the amplitude squared of the incident wave. I calculated the transmission coefficients

using complex demodulation. In the present experiments, I use a compound Chebyshev basis in the vertical direction. Because the Fourier transform of such a basis is not straight-forward (Fokas & Smitheman, 2012), I choose to focus my attention not on the transmission ratio, but on the vertical energy flux. As detailed Section 1.1, vertical energy flux due to internal waves could have implications for the Arctic Ocean system.

Measuring Vertical Energy Flux

Because the energy which internal waves can transport vertically can have a significant impact on the climate system (Ghaemsaidi et al., 2016), I choose to measure the vertical energy flux of these experiments. I would like to find the vertical energy flux per unit mass \mathcal{F}_z through some horizontal surface at a depth z . Taking the expression (4.13) for \mathcal{F}_z down to two dimensions and applying the product rule gives

$$\mathcal{F}_z = \frac{1}{2}(wu^2 + w^3) + p^*w - \nu(u\partial_z u + w\partial_z w), \quad (4.42)$$

where the three terms represent vertical energy flux per unit mass due to advection \mathcal{F}_a , pressure \mathcal{F}_p , and viscosity \mathcal{F}_ν , respectively all with units m^3/s^3 . Note that I set the signs in (4.42) such that a negative \mathcal{F}_z means the energy is propagating downwards.

To find \mathcal{F}_z through some depth z at a certain time t , I integrate across the horizontal domain from x_{min} to x_{max} for each term, starting with

$$\langle \mathcal{F}_a \rangle(z, t) = \int_{x_{min}}^{x_{max}} \frac{1}{2}(u^2 w + w^3) dx, \quad (4.43)$$

and similarly for $\langle \mathcal{F}_p \rangle$ and $\langle \mathcal{F}_\nu \rangle$, and where the total vertical energy flux is the sum of these three terms, $\langle \mathcal{F}_z \rangle = \langle \mathcal{F}_a \rangle + \langle \mathcal{F}_p \rangle + \langle \mathcal{F}_\nu \rangle$. I calculate each of these terms at all depths and times for each experiment.

Estimating Vertical Energy Flux Magnitude

In order to get an estimate for the relative contributions of each term I should expect, I substitute in the polarization relation (4.35) for u , w , p^* , and b for each term in (4.42):

$$\mathcal{F}_a = \frac{1}{2} \left(\hat{A} g \frac{\omega}{N^2} \right)^3 \left[\frac{k_z^2}{k_x^2} + 1 \right] \sin^3(k_x x + k_z z - \omega t) \quad (4.44a)$$

$$\mathcal{F}_p = - \left(\hat{A} g \frac{\omega}{N^2} \right)^2 \frac{\omega k_z}{k_x^2} \sin^2(k_x x + k_z z - \omega t) \quad (4.44b)$$

$$\mathcal{F}_\nu = -\nu k_z \left(\hat{A} g \frac{\omega}{N^2} \right)^2 \left[\frac{k_z^2}{k_x^2} + 1 \right] \sin(k_x x + k_z z - \omega t) \cos(k_x x + k_z z - \omega t) \quad (4.44c)$$

When integrating over an integer number of horizontal wavelengths, \mathcal{F}_a and \mathcal{F}_ν will go to zero. Therefore, in the ideal case of infinite plane waves, the total vertical energy flux per unit mass is expected to be entirely due to \mathcal{F}_p .

These experiments are not, however, the ideal case and so I can estimate the relative contributions of each term by calculating the coefficients $\hat{\mathcal{F}}_a$, $\hat{\mathcal{F}}_p$, and $\hat{\mathcal{F}}_\nu$ which amount to the magnitude of equations (4.44) without the trigonometric functions. I calculate these coefficients in Table 4.1 using

Parameter	Units	Constant N	Single layer	Double layer
g	m/s ²	9.81	—	—
κ	m ² /s	1.4×10^{-6}	—	—
ν	m ² /s	10^{-6}	—	—
\hat{A}	None	2.0×10^{-4}	2.3×10^{-4}	2.3×10^{-4}
N_0	rad/s	1.0	0.95	0.95
k	m ⁻¹	53.32	45	45
k_x	m ⁻¹	37.70	31.74	34.12
k_z	m ⁻¹	37.70	31.90	29.36
ω	rad/s	0.7071	0.67	0.72
θ	deg	45°	45.15°	40.72°
c_{gz}	m/s	9.38×10^{-3}	1.06×10^{-2}	1.04×10^{-2}
$n_{T,\text{lag}}$	None	6.00	5.05	5.49
$\hat{\mathcal{F}}_a$	m ³ /s ³	4.06×10^{-9}	4.72×10^{-9}	5.08×10^{-9}
$\langle \mathcal{F}_a \rangle$	m ³ /s ³	2.3×10^{-13}	1.43×10^{-13}	1.40×10^{-13}
$\hat{\mathcal{F}}_p$	m ³ /s ³	4.77×10^{-8}	5.95×10^{-8}	5.88×10^{-8}
$\langle \mathcal{F}_p \rangle$	m ³ /s ³	2.3×10^{-8}	4.21×10^{-8}	1.43×10^{-8}
$\hat{\mathcal{F}}_\nu$	m ³ /s ³	1.92×10^{-10}	1.80×10^{-10}	1.66×10^{-10}
$\langle \mathcal{F}_\nu \rangle$	m ³ /s ³	3.8×10^{-10}	2.04×10^{-10}	2.03×10^{-10}

Table 4.1: Values of parameters for the constant N , single layer, and double layer experiments, as well as the magnitude of the coefficients of vertical energy flux per unit mass $\hat{\mathcal{F}}$ predicted by (4.44) and the steady state measured values of these $\langle \mathcal{F} \rangle$ found by integrating (4.43) across the top boundary. Note that the values of $\langle \mathcal{F}_a \rangle$ for the single- and double-layer experiments had variations on the same order of magnitude as the reported values.

the parameters of each experiment. In all cases, $\hat{\mathcal{F}}_a = O(10^{-9})$, $\hat{\mathcal{F}}_p = O(10^{-8})$, and $\hat{\mathcal{F}}_\nu = O(10^{-10})$; the pressure term is indeed still expected to be the largest contributor.

Energy Flux for a Constant Stratification Profile

While the predictions of the sizes of the terms in the vertical energy flux above assume ideal plane waves, I expect the relative sizes of the terms to be comparable to the results of the numerical experiments. As a test, I use a profile with the constant stratification of $N = 1 \text{ s}^{-1}$ allowing free propagation through the domain. Without any stratification obstacles to cause reflections, I expect the vertical energy flux to be negative for all points in time and across all depths.

In order to bring this test as close as possible to ideal, infinite plane waves, I neglect using the windowing function $W_{bf}(x)$ on the boundary forcing. In order to satisfy the periodic horizontal boundary conditions, an integer number of horizontal wavelengths λ_x must fit within the horizontal experiment domain, which I set to be from 0 to 0.5 m. Therefore, I set $\lambda_x = 2\pi/k_x$ equal to one third of the horizontal extent. I choose $\theta = 45^\circ$ which results in $k_x = k_z = 37.7 \text{ m}^{-1}$ and $\omega = 0.7071 \text{ s}^{-1}$. The predictions of the energy flux coefficients for those values are displayed in Table 4.1.

Internal wave energy propagates vertically at the vertical group velocity, which is the derivative of ω with respect to k_z (Cushman-Roisin & Beckers, 2011; Kundu et al., 2015). Using the dispersion relation (4.5) with a constant N , this gives

$$c_{gz} = \frac{\partial \omega}{\partial k_z} = -\frac{\omega k_z}{(k_x^2 + k_z^2)} \quad (4.45)$$

which, with the chosen parameters, predicts a lag of $n_{T,\text{lag}} = 6$ oscillation periods between the energy flux through the top and bottom boundaries. Figure 4.11 shows the vertical energy flux through constant stratification over 20 oscillation periods. The red dashed line in panel (a) shows a slope corresponding to c_{gz} . This line is offset horizontally to accentuate how closely it matches the propagation of energy. In panel (b), the lag between the top and bottom boundaries in energy flux is close to the predicted 6 periods.

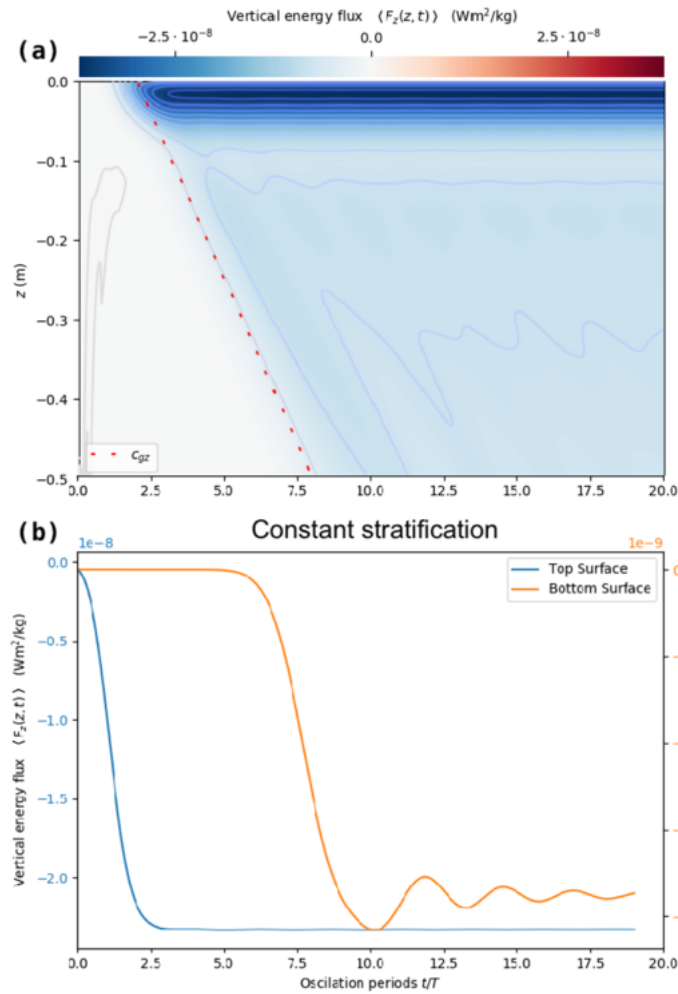


Figure 4.11: Total vertical energy flux for a constant stratification profile. (a) $\langle \mathcal{F}_z \rangle$ across the z domain and across time where the red dotted line shows the vertical group speed c_{gz} . (b) The running average of $\langle \mathcal{F}_z \rangle$ with a one oscillation period window through the top ($z = 0$) and bottom ($z = -0.5$) surfaces over time. This experiment had $\omega = 0.7071 \text{ s}^{-1}$, $\hat{A} = 2 \times 10^{-4}$, and was run for 20 oscillation periods.

The ramp in the forcing function is clearly evident in the energy flux through the top boundary. After approximately 3 periods, the top surface reaches a steady state. After a lag of around 6 periods, the ramp is also evident in the bottom boundary. While there is a small transient over the bottom boundary steady state, there is no evidence of interference of waves reflected from the bottom of the buffer in the simulated domain after 20 periods.

Figure 4.12 shows plots of the three energy flux terms from (4.43) integrated across the top

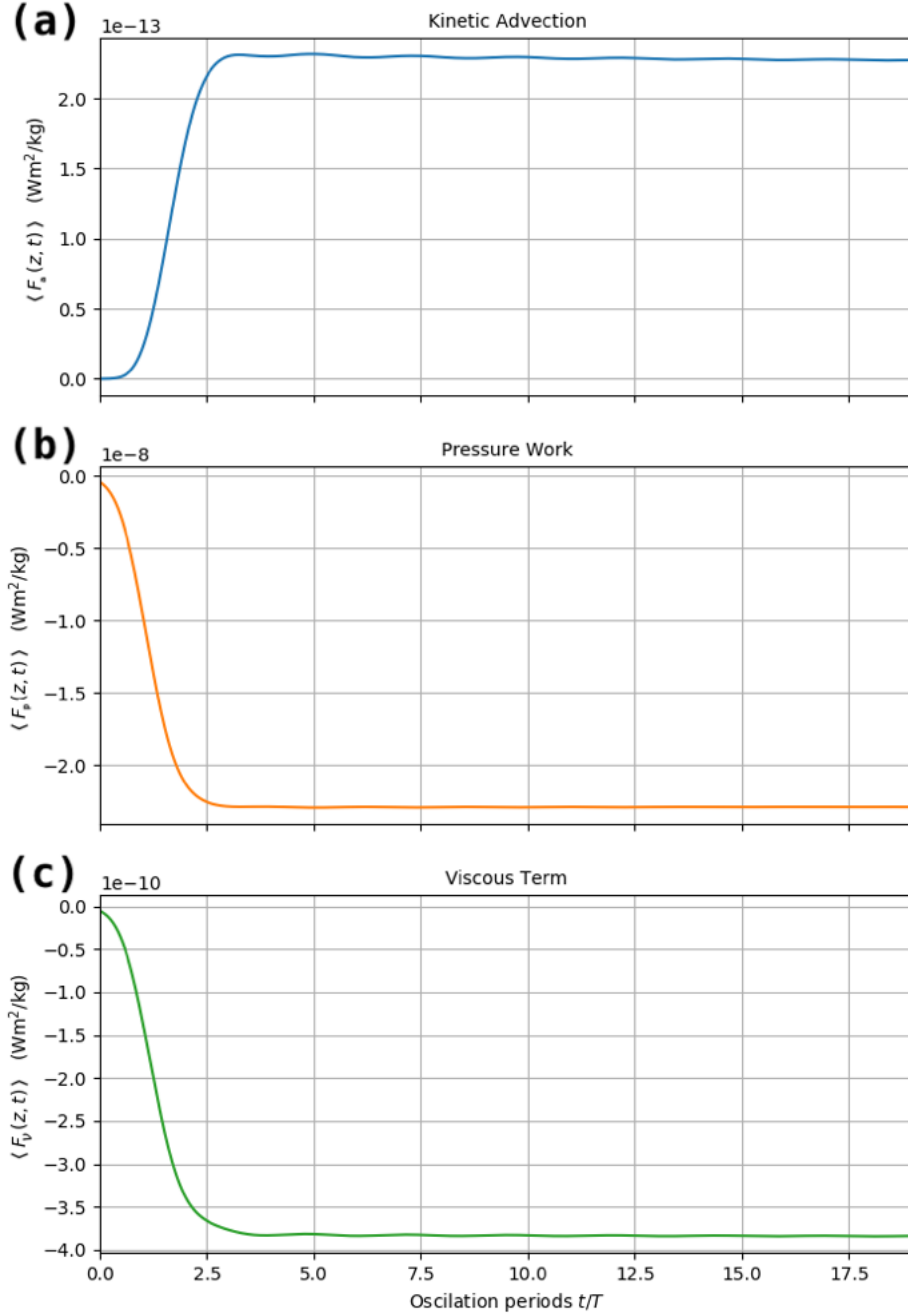


Figure 4.12: The components of vertical energy flux from (4.43) through the top boundary for a constant stratification profile: (a) $\langle \mathcal{F}_a \rangle$, (b) $\langle \mathcal{F}_p \rangle$, and (c) $\langle \mathcal{F}_\nu \rangle$. Note the different vertical axis scales between subplots and that these are the running averages with a one oscillation period window. This experiment had $\omega = 0.7071 \text{ s}^{-1}$, $\hat{A} = 2 \times 10^{-4}$, and was run for 20 oscillation periods.

boundary from the same experiment as Figure 4.11. Again, the ramp function over the first 3 periods is evident. In the ideal case, because the terms were integrated over 3 horizontal wavelengths λ_x , the only non-zero term should be pressure work. However, for these experiments I find steady state values of $\langle \mathcal{F}_{a,c} \rangle = 2.3 \times 10^{-13}$, $\langle \mathcal{F}_{p,c} \rangle = -2.3 \times 10^{-8}$, and $\langle \mathcal{F}_{\nu,c} \rangle = -3.8 \times 10^{-10}$ all with units m³/s³ or, equivalently, W m²/kg. From Table 4.1, the predicted magnitudes of the energy flux coefficients

are $\hat{\mathcal{F}}_{a,c} = 4.06 \times 10^{-9} \text{ W m}^2/\text{kg}$, $\hat{\mathcal{F}}_{p,c} = 4.77 \times 10^{-8} \text{ W m}^2/\text{kg}$, and $\hat{\mathcal{F}}_{\nu,c} = 1.92 \times 10^{-10} \text{ W m}^2/\text{kg}$. While the magnitudes of the measured $\langle \mathcal{F}_{p,c} \rangle$ and $\langle \mathcal{F}_{\nu,c} \rangle$ match the predictions of equation (4.44), the measured value $\langle \mathcal{F}_{a,c} \rangle$ is smaller than predicted, almost negligible compared to the other energy flux components. The reason for this discrepancy is unclear. However, the vertical energy flux through the top boundary due to the pressure work, shown in Figure 4.12, makes up the vast majority of the total energy flux, shown in Figure 4.11, as predicted by integrating (4.43).

4.5.3 Results

Comparison to previous work

In both laboratory experiments and theoretical experiments, G16 saw the majority of the wave beam reflect off the single layer and two distinct reflections and transmissions for the double layer. I reproduce these results from their Figure 5 below in Figure 4.13(a,c). All of the major features were successfully reproduced by Foran (2017) by recreating the stratification profiles as closely as possible and tuning ω until the correct angle was reached.




Figure 4.13: The wave tank results from G16 for the (a) single and (c) double mixed layer experiments, adapted from their Figure 5. The results of my numerical experiments for the (b) single- and (d) double-mixed-layer cases run with $\omega_1 = 0.67 \text{ s}^{-1}$ and $\omega_1 = 0.72 \text{ s}^{-1}$, respectively. The snapshot in (b) was taken at $t/T = 12.407$ and the snapshot in (d) was taken at $t/T = 12.476$, both with colorbars saturated at $\pm 1.3 \times 10^{-3} \text{ m/s}$.

By following Foran (2017), I also successfully reproduce the major features of the experiments

by G16, as seen in Figure 4.13(b,d). In Figure 4.13(b), the wave beam mostly reflects off the single layer, with a small amount transmitted. In Figure 4.13(d), the wave beam has both reflection and transmission from both mixed layers.

In Figure 4.13(a,c), G16 normalizes w in the colorbar by the “characteristic vertical velocity amplitude $A\omega$ ” where $A = 7.07$ mm is the amplitude and ω is the frequency of the wave-generating cylinder’s oscillations. Having no physical displacement distance of a wave-maker in my experiments, I do not normalize the vertical velocity. However, assuming the values of ω used by G16 were similar to those I use, the magnitude of the vertical velocities in their experiments and mine are both on the order of $O(10^{-3})$ m/s.

A reasonable estimate for the magnitude of vertical velocities I should expect can be found from the coefficient of the polarization relation for vertical velocity (4.35b) $\hat{A}g\omega/N_0^2$, either 1.68×10^{-3} m/s for the single layer case or 1.80×10^{-3} m/s for the double layer case. These match the magnitudes of vertical velocities in the numerical experiments however, I saturate the color bars in Figure 4.13(b,d) at $\pm 1.3 \times 10^{-3}$ m/s to more clearly show the transmitted wave beams, which are low in amplitude.

In the measurements taken by G16, they sampled the wave tank with a grid of 2042 by 2042 pixels across the measurement domain. This resolution was more than fine enough to resolve the features of the internal wave interactions in the tank. While I had 512 by 512 pixels total across my simulation domain, in the vertical, 40 were in the sponge layer below and only one third of the horizontal extent was used, giving the measurement domain a resolution of 170 by 472. However, even with relatively low resolution, the numerical experiments still reproduce the major qualitative features as well as match the magnitude of the vertical velocity field while remaining numerically stable.

Energy Flux for Single and Double Mixed Layer Profiles

Figure 4.14 shows the vertical energy flux over 10 oscillation periods through both the single- and double-layer stratification profiles. Because the boundary forcing is windowed in the x direction for these experiments, the vertical energy flux oscillates in time as can be seen in panels (a,b). These oscillations occur at twice the forcing frequency, as predicted by (4.44). In panels (c,d), I smooth out these oscillations by taking the running average over one oscillation period, which is why those plots stop at $t/T = 9$ instead of $t/T = 10$.

The ramp in the forcing function is again evident as the vertical energy flux through the top surface reaches its maximum magnitude around $t/T = 3$. Predictions of the lag between the top and bottom boundaries using (4.45) are $n_{T,\text{lag},1} = 5.05$ and $n_{T,\text{lag},2} = 5.49$ (see Table 4.1). These predictions, however, assume a constant stratification which is not the case for these experiments. The axis bounds for the vertical energy flux through the bottom surface covers an order of magnitude larger range of values for the constant stratification case in Figure 4.11(b) compared to the single- and double-layers cases in Figure 4.14(c,d). By contrast, the scales for the vertical energy flux through the top surface are close between the two. From this, I conclude that the presence of stratification has had an effect on the propagation of the vertical energy flux. Whether this effect was to dampen or merely delay the propagation is unclear as the experiments for the single and double layer cases were not run long enough for the vertical energy flux through the bottom surface to reach a steady state.

Figure 4.15 shows plots the running average with a one oscillation period window of the three

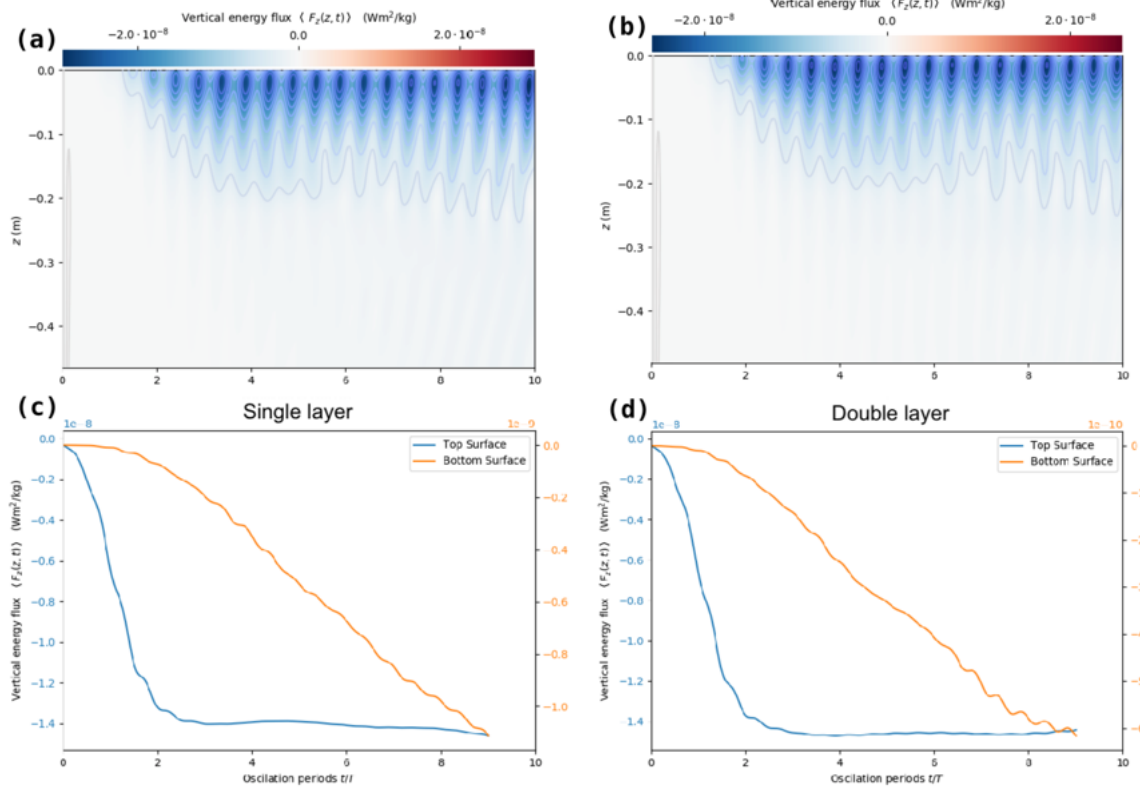


Figure 4.14: Total vertical energy flux for the single (a,c) and double (b,d) layer stratification profiles. (a,b) $\langle \mathcal{F}_z \rangle$ across the z domain and across time. (c,d) The running average of $\langle \mathcal{F}_z \rangle$ with a one oscillation period window through the top ($z = 0$) and bottom ($z = -0.5$) surfaces over time. These experiments had (a,c) $\omega = 0.67 \text{ s}^{-1}$ or (b,d) $\omega = 0.71 \text{ s}^{-1}$, $\hat{A} = 2.3 \times 10^{-4}$, and were run for 10 oscillation periods.

energy flux terms from (4.43) integrated across the top boundary from the same single- and double-layer experiments shown in Figure 4.14. Over the first three periods, the effect of the ramp function is evident for $\langle \mathcal{F}_p \rangle$ and $\langle \mathcal{F}_\nu \rangle$ in both cases. For $\langle \mathcal{F}_a \rangle$, there are significant oscillations within the ramping periods which I attribute to the windowing of the boundary forcing in the x direction. These oscillations are on the order of $10^{-11} \text{ W m}^2/\text{kg}$. However, the values after $t/T = 4$ are much smaller, with an average of $\langle \mathcal{F}_{a,1} \rangle = 1.4 \times 10^{-13} \text{ W m}^2/\text{kg}$ for the single-layer case in panel (a) and $\langle \mathcal{F}_{a,2} \rangle = 1.4 \times 10^{-13} \text{ W m}^2/\text{kg}$ for the double-layer case. Similar to the constant stratification case, these are much smaller values than predicted by (4.44). While these values match the magnitude of $\langle \mathcal{F}_{a,c} \rangle$, they have large standard deviations, $1.0 \times 10^{-13} \text{ W m}^2/\text{kg}$ for the single-layer case and $2.0 \times 10^{-13} \text{ W m}^2/\text{kg}$ for the double layer. Because the values and standard deviations of $\langle \mathcal{F}_{a,1} \rangle$ and $\langle \mathcal{F}_{a,2} \rangle$ are all on the order of $O(10^{-13}) \text{ W m}^2/\text{kg}$ and this is orders of magnitude less than the contributions to the overall vertical energy flux from pressure and viscosity, I conclude that the contributions from advection are effectively zero.

The steady state values of the vertical energy flux due to pressure and viscosity match in magnitude to both the constant stratification experiment and the predictions from (4.44). For the single-layer experiment, I find $\langle \mathcal{F}_{p,1} \rangle = -4.21 \times 10^{-8} \text{ W m}^2/\text{kg}$ and $\langle \mathcal{F}_{\nu,1} \rangle = -2.04 \times 10^{-10} \text{ W m}^2/\text{kg}$. For

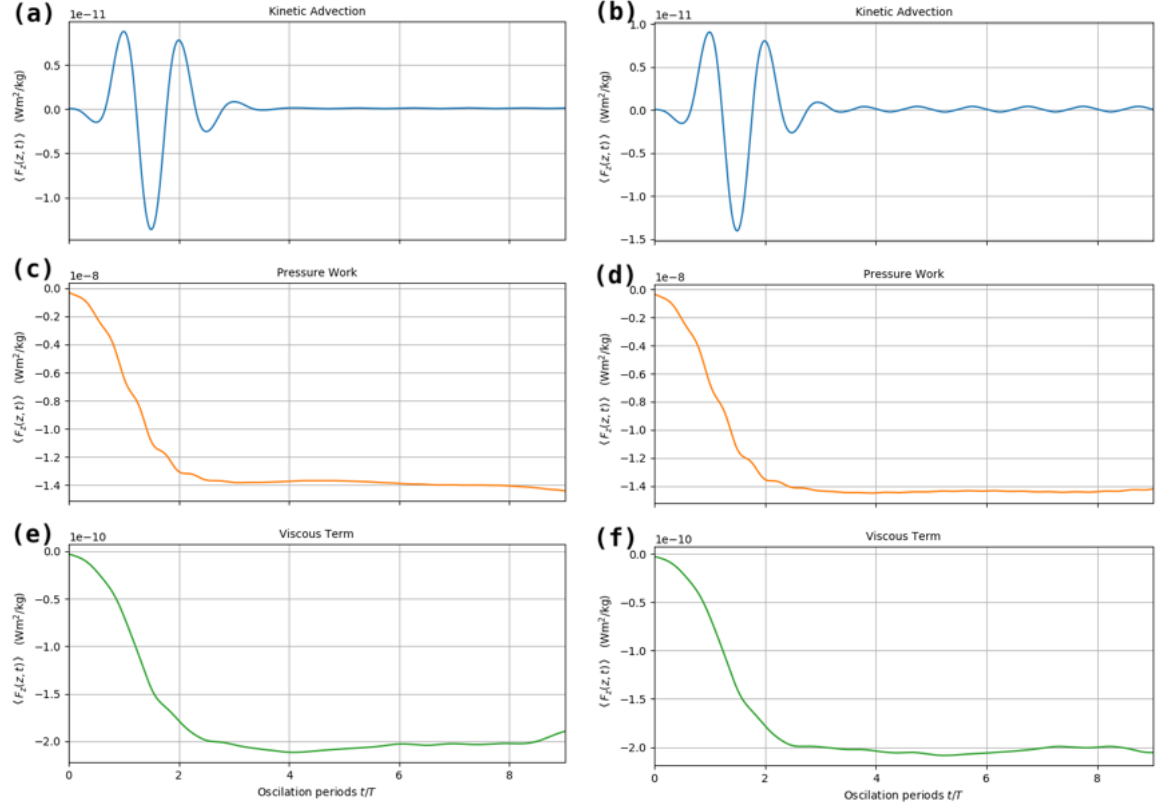


Figure 4.15: The components of vertical energy flux from (4.43) through the top boundary for the single (a,c,e) and double (b,d,f) layer stratification profiles: (a,b) $\langle \mathcal{F}_a \rangle$, (c,d) $\langle \mathcal{F}_p \rangle$, and (e,f) $\langle \mathcal{F}_\nu \rangle$. Note the different vertical axis scales between subplots and that these are the running averages with a one oscillation period window. These experiments had (a,c,e) $\omega = 0.67 \text{ s}^{-1}$ or (b,d,f) $\omega = 0.71 \text{ s}^{-1}$, $\hat{A} = 2.3 \times 10^{-4}$, and were run for 10 oscillation periods.

the double-layer experiment, I find $\langle \mathcal{F}_{p,2} \rangle = -1.43 \times 10^{-8} \text{ W m}^2/\text{kg}$, and $\langle \mathcal{F}_{\nu,3} \rangle = -2.03 \times 10^{-10} \text{ W m}^2/\text{kg}$. See Table 4.1 for a summary of these values. While these values do vary noticeably after $t/T = 4$, the magnitude of these variations are two orders of magnitude smaller than the steady state values. I find again, as I did for the constant stratification experiment, that the total vertical energy flux is almost entirely due to the vertical energy flux from pressure work, as predicted by integrating (4.43).

4.6 Discussion

In this chapter, I performed numerical experiments of internal waves interacting with stratification structures in both one and two spatial dimensions using the Dedalus framework. This was motivated by the question of how much internal waves might be able to transmit through thermohaline staircases in the Arctic. For the experiments in one spatial dimension, I found the transmission coefficient across values of $k_x L$, the ratio of layer thickness to horizontal wavelength, from 0 to 5 for waves propagating at an angle of 45° through stratification structures of one to five layers. I found the case of one mixed layer matches with analytical theory, monotonically decreasing in transmission

as $k_x L$ increased. For multiple layers, I found peaks in transmission at intermediate values of $k_x L$. These features are predicted by analytical theory as well, however I could only make qualitative comparisons as to their locations and magnitudes as theory assumes infinitely thin interfaces between layers (Sutherland, 2016). I found that the location of the largest peak in transmission for each number of layers decreased in $k_x L$ as the number of layers increased. However, I was unable to confirm whether this matches predictions from theory. My results agree quantitatively with those of G16 who used a finite-thickness interface when calculating transmission for a two-layer scenario.

With my experiments in two spatial dimensions, I was successfully able to reproduce the major features of the wave tank experiments by G16. I saw the same pattern of reflection and transmission as well as the same magnitude of the vertical velocity field. In these numerical experiments, the values of vertical energy flux through the top surface due to pressure and viscosity I calculated match theoretical predictions. The values of the vertical energy flux due to advection I calculated were magnitudes smaller than predicted by theory. However, the total values of vertical energy flux I found were almost entirely due to pressure, in agreement with theory. Overall, I believe these results show that numerical experiments of this nature can be a reliable tool for estimating the impact of staircase-like stratification structures on the propagation of internal waves and the associated vertical energy flux.

There are many avenues to build upon this work. The codes used in this work could be improved to reduce aliasing in order to handle waves forced at higher amplitudes. This could be done using high-performance computing resources to increase the spatial and temporal resolution, using an adaptive time stepping procedure, or modifying the implementation of viscosity to increase dissipation of unwanted waves. Many parameters could be adjusted to create numerical experiments that more closely resemble the physical realities of the Arctic Ocean. This includes considerations such as the stratification structure, the inclusion of the Coriolis force, and the specific properties of the waves which interact with the stratification.

The stratification structure used in these experiments was highly idealized. In the Arctic Ocean, thermohaline staircases generally consist of dozens of well-mixed layers that are several meters thick separated by interfaces ranging 0.3–0.7 m in thickness (Shibley et al., 2020) with a buoyancy frequency on the order of $N = O(10^{-2})$ rad/s while for the staircase region overall, N is on the order of $O(10^{-3})$ (Padman & Dillon, 1987; Shibley & Timmermans, 2022). Future experiments could better approximate this by adding more layers, using a more representative ratio of layer to interface thicknesses R_i , and scaling the magnitude of N within interfaces appropriately. One could also directly use a measured stratification profile from the Arctic Ocean, as did G16. However, care must be taken that the stratification profile has high enough resolution to accurately represent the thin interfaces, which is not the case for ITPs (Shibley et al., 2020). The profile G16 used was from the 2005 Beringia expedition which had a vertical resolution magnitudes finer than that of ITPs (Rainville & Winsor, 2008).

The effect of the Coriolis force increases with latitude, and, in the Arctic Ocean, is certainly not negligible. Both in the numerical experiments G16 ran with the profile from the 2005 Beringia expedition and in the study by Sutherland (2016), the authors used an inviscid yet rotating version of the streamfunction equation (4.3). As far as I am aware, no study has yet investigated the interactions of internal waves and staircase-like stratification structures while including both the effects of rotation and viscosity. While the effects of viscosity are generally considered to be negligible

at the scales which rotation becomes important, the flexibility of Dedalus makes it straight-forward to represent a system with both effects included.

As noted in both G16 and Sutherland (2016), when $N > f$ as it generally is in the Arctic Ocean, the inclusion of the Coriolis force sets a lower bound on the frequencies of waves that can be transmitted; they are evanescent outside the range $f < \omega < N$. In the Canada Basin, $f \approx 1.42 \times 10^{-4} \text{ s}^{-1}$. Near-inertial waves with frequencies close to f dominate the internal wave field in the Arctic and typically have vertical wavelengths λ_z on the order of 10's of meters and horizontal wavelengths λ_x on the order of kilometers (Dosser et al., 2014; Ghaemsaidi et al., 2016). This results in angles of propagation θ much shallower than I used in my experiments. In all my experiments, I set both N and ω to be $O(1)$. The value of λ_z I used in the one-dimensional experiments was 1 m while, for the two-dimensional experiments, I set λ_z on the order of 3 cm to match the wave tank experiments.

According to the analytical solution for one non-rotating, well-mixed layer (4.26) (Sutherland & Yewchuk, 2004), transmission may well remain comparable as both the ratio of the layer thickness to horizontal wavelength $k_x L = 2\pi L/\lambda_z$ and the angle of propagation θ decrease from values used in wave tank experiments towards those more realistic for the Arctic. When adding more layers and considering rotation, the situation is more complicated. However, numerical experiments in a flexible framework such as the one I've used in this chapter offer a way to investigate the transmission of internal waves in more realistic scenarios.

In the numerical experiments I presented here as well in those of Sutherland & Yewchuk (2004), G16, and Sutherland (2016), only weakly-nonlinear, monochromatic waves (containing only one frequency) were considered. The internal wave field of the Arctic contains many different frequencies. However, even an initially monochromatic wave can produce waves of different frequencies through the process of harmonic generation. Using a weakly nonlinear theory that includes rotation but neglects viscosity, Wunsch (2018) investigated harmonic generation from internal waves incident upon density staircases, finding analytical expressions for the amplitude of the harmonic at the top and bottom of a staircase with an arbitrary number of identical stair-steps. Other factors that would emulate more realistic waves and have been shown to affect internal wave interactions with stratification are having waves be localized in time (Supekar & Peacock, 2019) and axisymmetric (Boury et al., 2019).

Chapter 5

Conclusions

5.1 Summary of results

Thermohaline staircases have been observed in the Arctic Ocean for over 50 years. Despite this, very little is known about how these remarkably coherent structures change over long periods of time. With Arctic Sea ice in decline, there is concern that the resulting increase in wind-driven internal waves could affect vertical mixing, allowing energy to transfer from the surface to the interior. However, the stratification of thermohaline staircases make it unclear as to what extent internal waves may be able to propagate downwards. This thesis examines thermohaline staircases in the Canada Basin of the Arctic Ocean, their evolution on a decadal scale, and their interactions with internal waves.

The layers in thermohaline staircases have been shown to extend horizontally for hundreds of kilometers (Schmitt et al., 1987; Timmermans et al., 2003, 2008; Lu et al., 2022). Several studies have developed algorithmic approaches to detect staircase steps (Timmermans et al., 2008; Shibley et al., 2017; van der Boog, Otto Koetsier, et al., 2021; Lu et al., 2022). However, as all these approaches treat each profile independently, no analysis can be done on the horizontal properties of any particular layer. A few studies have developed techniques to connect layers across different profiles, but they have either relied on manual inspection, setting precise gradient thresholds, or both (Padman & Dillon, 1988; Lu et al., 2022).

In Chapter 2, I presented a novel method based on the Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) algorithm, which automatically detects and connects thermohaline staircase layers across many profiles. This method requires only a selection of a salinity range in which to search and an estimate of the typical layer thickness. High precision is not required for either and estimates could be made from a brief inspection of a dataset or from previous studies. This is in comparison to methods used previously to identify staircases, which require a choice of one or multiple thresholds on gradients in temperature, salinity, or density (Timmermans et al., 2008; Shibley et al., 2017; van der Boog, Otto Koetsier, et al., 2021; Lu et al., 2022). Choosing appropriate values for these thresholds requires precise knowledge of certain properties in the staircases which are to be detected. I provided an algorithmic approach to selecting the one required hyperparameter for HDBSCAN, m_{pts} , and developed a method of evaluating the subsequent clusters based on physical properties, marking those which do not meet expectations as outliers. Using this clustering method,

I was able to automatically track individual staircase layers for Ice-Tethered Profiler (ITP) 2 which ran for 40 days and for ITP3 which traced a 2541 km-long track over 382 days with the longest span between any two, not necessarily consecutive, profiles being 441 km.

I reproduced several key results from previous studies. By introducing the normalized inter-cluster range IR_v , I quantitatively confirmed a pattern noted in previous studies that the variations in pressure and temperature within layers is larger than with salinity (Timmermans et al., 2008; Lu et al., 2022). The salinity ranges for each layer had very little overlap, with the difference in salinity between neighboring clusters approximately seven times larger than salinity variations within any particular layer. This is in agreement with Lu et al. (2022) and supports their conclusion that salinity is the most appropriate variable by which to identify individual layers. For the lateral density ratio R_L , I found that the values matched the magnitude reported by Timmermans et al. (2008). However, I noted a possible dependence of the lateral density ratio R_L on pressure, the reason for which is unclear and worthy of further study. Along with this result, I found instances where outliers in IR_{Sp} revealed features called remnant intrusions and outliers in R_L revealed occurrences of layers splitting or merging. As the results of this study were in line with those of previous studies, this suggests that this method could be applied to datasets with larger numbers of points.

The study presented in Chapter 3 was such an extension. Using the method based upon the HDBSCAN clustering algorithm, I identified thermohaline staircase layers across 17 years of data within the Beaufort Gyre Region (BGR) including 60 different ITPs with 15,202 individual profiles for a total of 12,744,846 data points. This study filled a gap in the literature for thermohaline staircases in the Arctic Ocean. Many authors have noted the incredible persistence of these structures, particularly in the Canada Basin. Since they were first measured at Ice Island T-3 (Neal et al., 1969; Neshyba et al., 1971; Neal & Neshyba, 1973), they have since been notable features in the profiles from such major efforts as the Arctic Internal Wave Experiment (AIWEX) (1985) (Padman & Dillon, 1987, 1988, 1989), the Surface Heat Budget of the Arctic (SHEBA) experiment (1997–1998) (Shaw & Stanton, 2014), and the ITP program (2004–present) (Krishfield et al., 2008; Toole et al., 2011). Studies by Polyakov, Pnyushkov, Rembe, et al. (2012), Ménesguen et al. (2022), and Lu et al. (2022) tracked individual staircase layers, however none for a period longer than four years. In Chapter 3, I established that individual layers can indeed stay coherent on a decadal time scale.

As noted by authors in previous studies, I found the layers to have remarkable lateral coherence, on the order of 1000 km, with an aspect ratio on the order of 10^6 (Timmermans et al., 2008; Toole et al., 2011; Bebieva & Timmermans, 2019; Lu et al., 2022). I also found layers to have notable geographic distributions in pressure and temperature. To characterize these spatial variations, I fit a two-dimensional polynomial up to third order over latitude and longitude. From these, I found the same bowl-shaped distribution of pressure as noted by Lu et al. (2022) where each layer is deeper in the center of the BGR, growing shallower towards the edges. I noted that the spatial pattern of temperature changed with depth, speculating as Lu et al. (2022) did that this could be due to the differing influence of the water masses above and below, the Lower Halocline Water (LHW) and the Atlantic Water (AW).

I found those bounding water masses to be sinking at rates of 2.91 and 5.32 dbar per year for the LHW and AW, respectively. On average, the layers I detected were sinking at a rate of 2.24 dbar per year and all of these rates match the reported range of cumulative downwelling in the BGR of 2–5 meters per year (Meneghello et al., 2018; Proshutinsky et al., 2019; Timmermans & Toole, 2023),

implying the staircase layers are passively following the water column as a whole. I found layers near the top to be sinking faster than those near the bottom, implying an overall compression of the staircase. The average layer thickness I found was 1.55 dbar and thicknesses generally increase with depth, in line with previous studies (Neal et al., 1969; Padman & Dillon, 1987; Guo et al., 2016; Timmermans et al., 2008; Shibley et al., 2017; Shibley & Timmermans, 2019; Lu et al., 2022). While not conclusive, I found that layers are either thinning or staying the same thickness with time.

The AW has been warming over the past few decades and, at the same time, so too has a layer of water from the Pacific, which sits above the LHW in the water column (Timmermans et al., 2018; Dosser et al., 2021). While these water masses above and below the pycnocline have been warming, I found that the trend in time for each layer’s temperature changes signs from top to bottom. That is, layers near the top of the staircase are warming over time while layers near the bottom are cooling, with a compound trend of $(-2.29 \pm 0.23) \times 10^{-4} \text{ }^\circ\text{C/yr/dbar}$. Using these trends in temperature, I found the average net heat flux for each layer, finding the upper layers to have a positive heat flux and lower layers with negative heat flux. The trend was $(-2.82 \pm 0.39) \times 10^{-5} \text{ W/m}^2\text{/dbar}$ with a total of 0.054 W/m^2 across the staircase, matching estimates from previous studies (Guthrie et al., 2015; Shibley & Timmermans, 2019).

The development of this method using HDBSCAN presented here contributes to the wider context of clustering algorithms being applied in oceanographic studies (e.g., Koszalka & LaCasce, 2010; Espejo et al., 2014; Maze et al., 2017; C. Ma et al., 2019; Jones et al., 2019; Houghton & Wilson, 2020; Rosso et al., 2020; Desbruyères et al., 2021; Sonnewald et al., 2021). As the size of oceanographic datasets continue to grow larger and more complex, the continued development and usage of clustering algorithms could offer efficient ways to discover and analyze a wide variety of features and phenomena.

While the stratification of thermohaline staircases contain layers where internal waves become evanescent, acting as a series of barriers to their propagation, the phenomenon of internal gravity wave tunnelling allows certain waves to transmit through particular stratification configurations (Eckart, 1961; Sutherland & Yewchuk, 2004; Ghaemsaïdi et al., 2016; Sutherland, 2016). In Chapter 4, I performed numerical experiments using the Dedalus framework to investigate this phenomenon in both one-dimensional and two-dimensional configurations.

For the experiments in one spatial dimension, I quantified the transmission coefficient for waves propagating at an angle of 45° through stratification structures with up to five well-mixed layers separated by finite-width interfaces. These mixed layers had a stratification frequency $N = 0$, meaning the waves become evanescent, and a thickness based upon $k_x L$, the ratio of the layer thickness to horizontal wavelength. For the case of one layer, the transmission coefficient monotonically decreased from perfect transmission at $k_x L = 0$ to negligible transmission at $k_x L = 5$, matching analytical predictions. In stratification structures with multiple layers, I found significant peaks in transmission at values of $k_x L$ between 0 and 5 and the value of $k_x L$ for the largest of these peaks decreased as more layers were added. Ghaemsaïdi et al. (2016) numerically solved for the values of the transmission coefficient in the cases of one and two mixed layers. As both their setup and mine used finite-width interfaces, I was able to make quantitative comparisons and showed the results agreed. While Sutherland (2016) presented results for cases with more layers, I was only able to note the qualitative agreement in transmission patterns as infinitely thin interfaces were used in that study.

My numerical experiments in two spatial dimensions were focused on replicating a wave tank experiment performed by [Ghaemsaidi et al. \(2016\)](#). I successfully reproduced the pattern of reflection and transmission, finding the same magnitude of the vertical velocity field. I extended this work by calculating the vertical energy flux. In agreement with theory, I found that the total vertical energy flux through the top surface was almost entirely due to pressure. As these numerical experiments were able to reproduce the results from theory and previous studies, I believe I have shown they can be a valuable tool to analyze the impact of stratification structures that resemble staircases on the propagation of internal waves with the potential to extend to scales that would be impractical for laboratory experiments.

5.2 Future work

In this thesis, I detailed studies of thermohaline staircases, their detection in sets of hydrographic profile data, their evolution over decadal time scales, and their interactions with internal waves. For each, I presented results that both replicated those from previous studies and those that are novel. All offer avenues for further study and improvement of the methods used.

In Chapter 2, I demonstrated how my method based on the HDBSCAN clustering algorithm could automatically identify staircase layers in datasets of up to 1,532 profiles spanning up to 382 days. I extended this work in Chapter 3, applying the clustering algorithm to a dataset which has ten times the number of profiles and spans 17 years. These chapters show the potential of such a method of layer identification which could be extended in many new directions.

While, in this thesis, I have been solely concerned with thermohaline staircases in the upper Arctic Ocean, the clustering method could be used in different oceanic regimes. Staircase structures exist in the deep Arctic Ocean, at depths of 2000–3000 m ([Timmermans et al., 2003](#)), and the Argo network of autonomous profiling floats, which has been active for over 20 years, has measured staircases in different regions all across the world’s oceans ([van der Boog, Dijkstra, et al., 2021](#)). I have focused on capturing the patterns that thermohaline staircases make in temperature–salinity (T – S) space. However, the patterns which different types of layers or water masses make in T – S space, or even T – O_2 space ([Rosso et al., 2020](#)), could also be captured with this method.

As mentioned in Chapter 2, I found the clustering results to be robust even when subsampling profiles to every second, third, or fourth point, effectively reducing the vertical resolution. This follows from the fact that the clustering method does not treat profiles individually, only requiring that a sufficient number of points be available to detect clusters. This suggests that this method may be able to identify staircase layers in a set of profiles that each have too coarse resolution to resolve individual steps, or that a mix of different resolution profiles could be combined into one dataset when identifying layers.

This clustering method could be used in conjunction with other staircase detection methods, either as a first pass to identify where a more detailed process should be focused, or as a way to connect layer points identified by another method across different profiles. The results from Chapter 2 revealed instances of layer splitting or merging, as well as remnant intrusions. There is a possibility that this method could be modified to specifically identify such occurrences.

There are several technical aspects in which the clustering method could be improved. One such area is the process of selecting m_{pts} , the one required hyperparameter for HDBSCAN. The process I

presented involved comparing the Density-Based Clustering Validation (DBCV) score for clusterings of the same data across a range of different values of m_{pts} . The DBCV score gives an estimate as to the difference between the least dense regions within clusters and the most dense regions of noise, with a higher value indicating the clusters are more distinct from the noise. However, as mentioned in Chapter 3, a higher DBCV score does not always correspond to a clustering that better captures the layers. A better algorithm for selecting the most appropriate m_{pts} value could use a different clustering validation score, or perhaps a different method of validation altogether. It might also incorporate the rough correlation I found where datasets with more points generally require a larger value of m_{pts} to reasonably capture layers. Additionally, finding an appropriate value of m_{pts} involves increasingly more computations for larger datasets as larger ranges of m_{pts} values might be considered. As only one of these many clusterings is ultimately chosen as the result, an ideal method of selecting m_{pts} would limit the total number of clusterings which need be compared to reduce the computational resources required.

Improvements could also be made to better capture the trends in time for layers. In Chapter 3, I found all trends using the ordinary least squares regression. This was chosen as a simple, well-known method, but other types of regression may be more appropriate for particular circumstances. Different clusters spanned different amounts of time within the 17-year period analyzed. This was mainly due to the method I used to connect clusters across each year-long time period, assigning different IDs to clusters which erroneously spanned more than one layer in certain periods. While the need to manually connect clusters between time periods could be avoided by instead clustering the entire 17 years at once, there still would be the possibility of particular clusters spanning only short amounts of time. Clusters which span shorter periods are more susceptible to extreme values affecting their calculated trend in time, something which I did not account for. More robust statistics of the trends could be found by accounting for this source of uncertainty. Another improvement could be made to the estimates of net heat flux by combining the component values at every point in time and then finding the trend rather than calculating the value for each cluster using already time-averaged components.

In Chapter 4, I presented the results of numerical experiments developed in the Dedalus framework. By reproducing the results of previous studies, I demonstrated how such numerical experiments could be reliable tools to investigate the interactions between internal waves and stratification structures. Numerical experiments have the potential to provide valuable insight into processes on scales which are too large for laboratory experiments and too small to be resolved in current global climate models.

An obvious extension to these experiments would be to use parameters which match more closely with the Arctic Ocean’s physical reality. Generally, the angles of propagation θ for internal waves in the Arctic are much shallower than those I used (Dosser et al., 2014; Ghaemsaïdi et al., 2016). In my experiments, I set the thicknesses of the layers and the intervening interfaces to be on the same order of magnitude, matching the experiments of Ghaemsaïdi et al. (2016). However, observed thermohaline staircases generally have layers which are one to two orders of magnitude thicker than the interfaces (Shibley et al., 2020). A future study could set this ratio of layer to interface thickness, as well as the overall magnitude of the buoyancy frequency N , to better match those found in the Arctic Ocean. Both could be met by directly using observed profiles, as Ghaemsaïdi et al. (2016) did with measurements from the 2005 Beringia expedition. However, the profiles would

need fine enough vertical resolution to accurately capture the details of the interfaces, which is not possible with ITP data (Shibley et al., 2020). My experiments focused on reproducing laboratory results and thus neglected the Coriolis force. However, in the Arctic, this effect is not negligible and both in the study by Sutherland (2016) and the experiment run by Ghaemsaidi et al. (2016) on the Beringia profile, an inviscid yet rotating system was considered. The effects of viscosity are usually considered negligible at the scales where rotation is important, but the flexibility of the Dedalus framework provides a straight-forward way to include both. To my knowledge, no study has considered internal wave transmission in a system with the effects of both viscosity and rotation.

The codes used for the numerical experiments could also be improved. By using high-performance computing, the spatial and temporal resolutions could be increased. An adaptive time stepping procedure could be utilized. The manner in which viscosity is implemented could be modified to increase dissipation of unwanted waves. Such stabilization measures may allow for the forcing of waves at higher amplitudes. The codes could also be modified to allow for polychromatic waves (containing more than one frequency) or non-linear interactions such as harmonic generation (Bourget et al., 2013; Wunsch, 2018). Additionally, the forced waves could be modified to more accurately represent reality by being localized in time or axisymmetric, both factors known to affect the interactions between internal waves and stratification structures (Supekar & Peacock, 2019; Boury et al., 2019).

Appendix A

Supplementary Materials for: Unsupervised clustering identifies thermohaline staircases in the Canada Basin of the Arctic Ocean

A.1 Practical salinity S_P vs. absolute salinity S_A

TEOS-10 recommends storing values of practical salinity S_P in databases in order to maintain continuity with older records and also because S_P is practically a measured quantity (McDougall & Barker, 2011). In contrast, absolute salinity S_A is a calculated variable defined as “the mass fraction of dissolved non-H₂O material in a seawater sample at its temperature and pressure,” and is recommended for use in publications (IOC, 2010). For this study, we ultimately chose to work with S_P because the two studies to which we make direct comparisons, T08 and L22, used S_P . Specifically, when we define the ranges of S_P to filter to before running the clustering algorithm, we were able to directly take the values of S_P from T08 and L22 without having to first convert to S_A .

To investigate the difference between using S_P vs. S_A on our results, we recreated Figure 2.2, but replacing S_P with S_A on all the axes, as shown in Figure A.1. We still filtered the ITP2 data to be in the practical salinity range S_P 34.05–34.75 g/kg to make sure we used the same data points. However, then we plotted those data points in Θ' – S_A space in panel (c) to run the clustering algorithm using the same $m_{pts} = 170$. This resulted in 38 clusters (as opposed to 36) with a DBCV score of 0.2177 (as opposed to 0.3034). In panel (d), the values of R_L displayed made minor changes of (from left to right): +0.39, -0.04, -0.03, -0.03, -0.01, +0.14, and +0.07.

We also reproduced Figure 2.5(a,b) in Figure A.2 using the clustering shown in Figure A.1. The two outliers in IR_{S_A} in Figure A.2(a) do not correspond to outliers in Figure 2.5(a) because these have very high IR_{S_A} , over 20, which skews the mean as used when calculating the z-score. The dark green star and the red triangle in Figure A.2(a) correspond to the two outliers in Figure 2.5(a) where they are marked by an orange 4-pointed star and a green “×”. These pairs of points have very similar values of IR_{S_A} and IR_{S_P} , respectively.

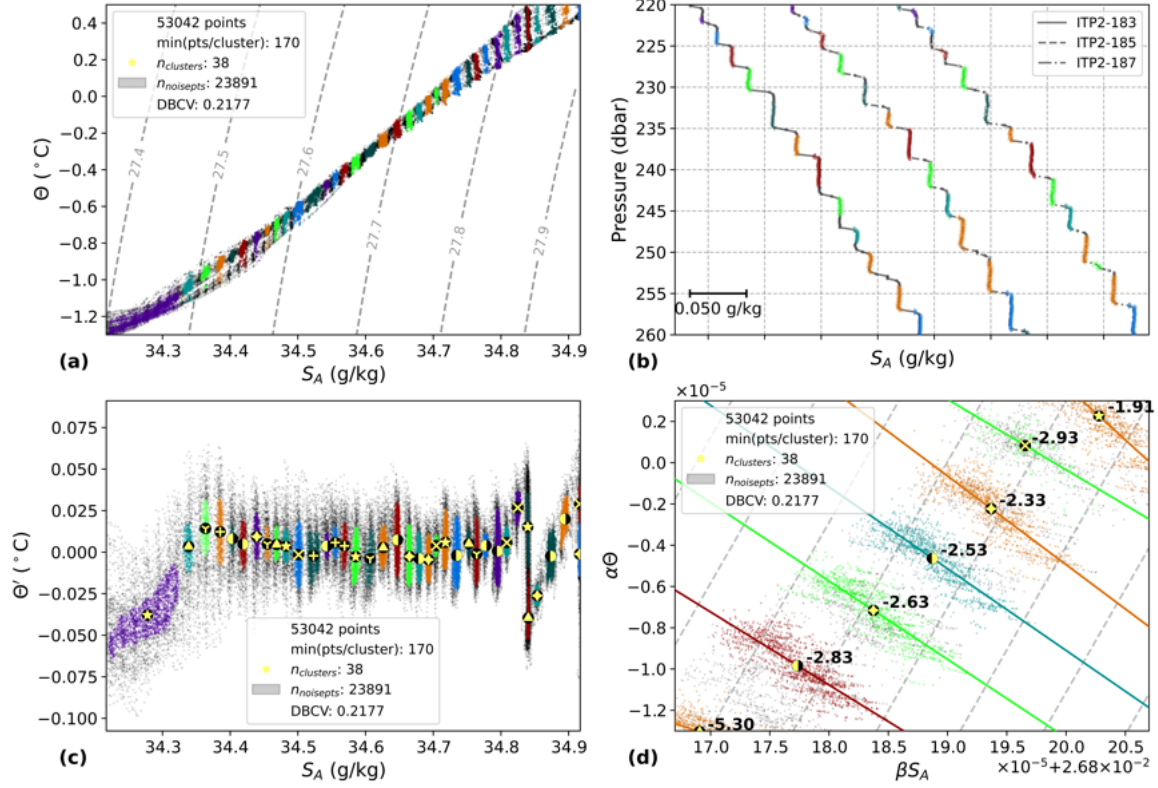


Figure A.1: Results from the clustering algorithm with $m_{pts} = 170$ and $\ell = 25$ dbar run on 53042 data points in the S_P range 34.05–34.75 g/kg from all up-going ITP2 profiles. This is the same as Figure 2.2, but the clustering algorithm was run on absolute salinity S_A instead of practical salinity S_P . (a) The data in Θ – S_A space with dashed lines of constant potential density anomaly (kg m^{-3}) referenced to the surface. (b) Profiles 183, 185, and 187 from ITP2 in a limited pressure range to show detail. Each profile is offset in S_A for clarity. (c) The spatial arrangement used as input for the algorithm where the gray points are noise and each color-marker combination indicates a cluster. The same color-marker combinations are used in each panel and the markers in panels (c) and (d) are at the cluster average for each axis. (d) A subset of the data in $\alpha\Theta$ – βS_A space with the linear regression line and inverse slope (R_L) noted for each individual cluster and with dashed lines of slope $\alpha\Theta/\beta S_A = 1$

In Figure A.2(b), we see that there are two outliers in R_L circled, both of which correspond to outliers found when using S_P . The purple star in Figure A.2(b) corresponds to the orange left half circle in Figure 2.5(b), while the dark green “Y” in Figure A.2(b) corresponds to the teal “x” in Figure 2.5(b). A notable difference is that the purple star outlier in Figure 2.5(b) (which spans $S_P = 34.233$ – 34.261 in Figure 2.2(c) encompassing what should be two distinct clusters) is no longer present in Figure A.2(b). Comparing Figures 2.2(c) and A.1(c), we see that it is now replaced by two clusters, the dark green left half circle and the red right half circle, neither of which are outliers in R_L .

We also see that the 2nd-degree polynomial fit in Figure A.2(b) is similar to that in Figure 2.5(b). We conclude that the difference between using S_A vs. S_P does not significantly affect our results.




Figure A.2: The value of each cluster’s normalized inter-cluster range for salinity IR_{S_A} in (a) and the lateral density ratio R_L in (b) as a function of the cluster’s average pressure for ITP2. This is the same as Figure 2.5(a,b), but the clustering algorithm was run on absolute salinity S_A instead of practical salinity S_P . The colors and markers are the same as the clustering shown in Figure A.1. Markers circled in red indicate outliers with a z-score greater than 2.

A.2 Varying the value of ℓ

We define the local anomaly of a temperature profile Θ' as the original profile minus a version of that profile smoothed with a moving average. Subtracting a smoothed profile from the original profile in this way gives temperature differences that are centered around zero and leads to more continuous clusters (compare Figures 2.2(a) and 2.2(c)), allowing HDBSCAN to group these points more accurately.

When calculating Θ' with Equation 3.1, we use a moving average with window size ℓ , so the resulting profile has dead regions $\ell/2$ in size with no data at the top and bottom. Therefore, ℓ cannot be larger than twice the distance available in the profile, either above or below the pressure range to be analyzed, whichever is smaller. As ℓ increases, the moving average profile gets flatter, meaning Θ' becomes closer to the original profile, just shifted in temperature space. Choosing a very small ℓ will give a moving average temperature profile that will closely match the original temperature profile, meaning that the Θ' profile will be very close to zero. This eliminates any effect temperature might have on the clustering results.

In Figure A.3, we show four pairs of profiles, restricted to the 200-400 dbar pressure range to show detail. All profiles are number 185 from ITP2, the example profile featured in Figure 2.2(b). Each panel shows both the original profile and the moving average profile with the values of ℓ : 2.5, 12.5, 25, and 37.5 dbar. The original and moving average profiles are slightly offset for clarity. When $\ell = 2.5$ dbar, the moving average profile still clearly contains some of the larger stair steps. When $\ell = 37.5$ dbar, almost all of the features in the original profile are gone.

We noted a feature in 2(c) where the cluster average Θ' increases starting at $S_P \approx 34.63$ g/kg until it jumps sharply to negative Θ' values around $S_P \approx 34.67$ g/kg, then increases again until




Figure A.3: The original profile 185 from ITP2 and the smoothed version of that profile for 4 different values of ℓ : (a) 2.5 dbar, (b) 12.5 dbar, (c) 25 dbar, and (d) 37.5 dbar. The Θ axes on each panel are offset slightly to show both lines.

$S_P \approx 34.72$ g/kg. This zig-zag pattern is due to the choice of ℓ and the presence of the AW subsurface temperature maximum in Θ profiles. In Figure A.4, we show the ITP2 data used in the study in Θ' - S_P space for the same 4 values of ℓ as in Figure A.3. This clearly illustrates that the zig-zag pattern is not present for low ℓ and becomes more pronounced as ℓ increases.

For each panel in Figure A.4, we ran a parameter sweep across m_{pts} similar to that shown in Figure 2.3(b), selecting the value of m_{pts} which gave the highest DBCV, as shown in the legends. If we always choose ℓ to be approximately five times the typical layer thickness, then the panels in Figure A.4 represent clusterings for thickness estimates of 0.5 dbar, 2.5 dbar, 5 dbar, and 7.5 dbar. We conclude that the results are not significantly sensitive to this typical layer thickness estimate as the number of clusters and their positions along the S_P axis are generally consistent across the four values of ℓ , especially between $S_P \approx 34.18$ g/kg and $S_P \approx 34.62$ g/kg. We find the most reasonable clusterings occur when ℓ is chosen to be small enough that the features outside the pressure range we analyze do not significantly affect the moving average, yet large enough that the stair steps are completely smoothed out.

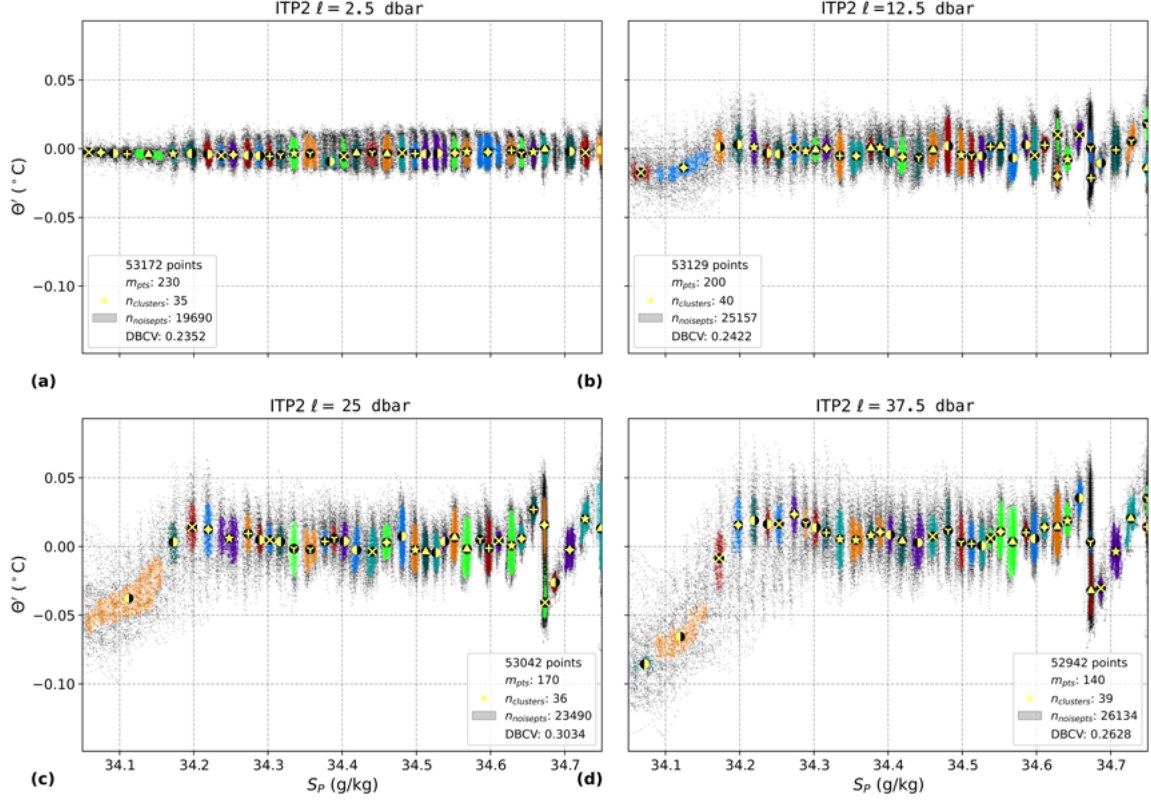


Figure A.4: Clustered data from ITP2 in the salinity range 34.05–34.75 g/kg plotted in Θ' – S_P space for 4 different values of ℓ : (a) 2.5 dbar, (b) 12.5 dbar, (c) 25 dbar (the same as in Figure 2.2(c)), and (d) 37.5 dbar, where the gray points are noise and each color-marker combination indicates a cluster.

A.3 Optional parameters for HDBSCAN

While m_{pts} has the most significant effect on the clustering results, there are several other parameters which could be adjusted when using the “hdbscan” Python package. In each case, we chose to follow the recommendation of the authors of the algorithm (Campello et al., 2013; Moulavi et al., 2014). The default distance metric is Euclidean and, as we cluster in a two dimensional space of continuous numerical variables, we maintain this selection. For the cluster selection method, we chose Leaf over Excess of Mass to avoid the results collapsing into fewer, larger clusters. HDBSCAN can be combined with DBSCAN by setting a threshold of cluster selection $\varepsilon_{clSelect}$ to explicitly group clusters a certain distance apart. However, we chose the default behavior of HDBSCAN which automatically identifies the values of ε which give the most stable clusters. HDBSCAN considers a neighborhood to be dense when m_{pts} points are within a certain distance ε . While m_{pts} , known as the minimum samples, can be independently set higher to more liberally include fringe points in clusters or lower to more conservatively classify them as noise, we follow the default recommendation to set this equal to the minimum points per cluster m_{clSize} .

A.4 Clusters from ITP3

In Figure A.5, we create a figure similar to that of Figure 2.2, but for ITP3. In Figure A.6, we create a figure similar to that of Figure 2.3(a), but for ITP3. Note that the run with $m_{pts} = 710$ had a DBCV score of 0.3885, which is slightly higher than the run at $m_{pts} = 580$ with a DBCV score of 0.3862. However, because the differences in DBCV were not great, we decided to use $m_{pts} = 580$ to reduce computation time as the clustering algorithm takes longer to run with higher values of m_{pts} .

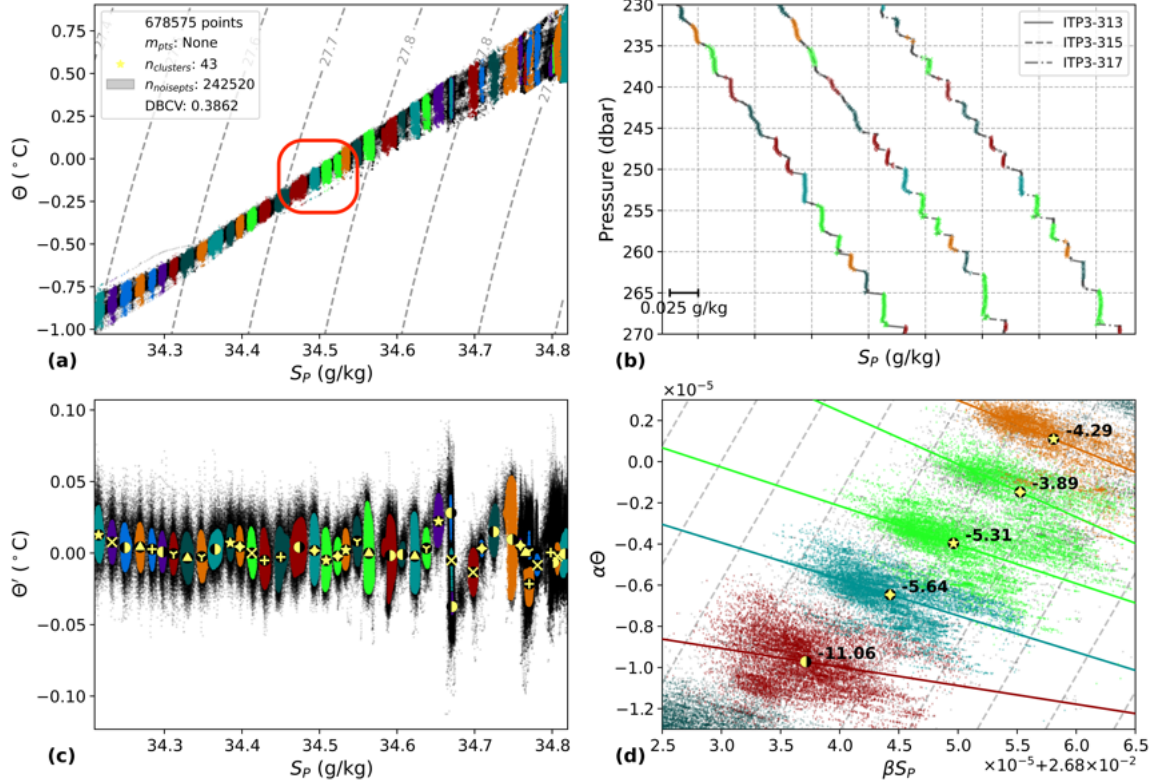


Figure A.5: Results from the clustering algorithm with $m_{pts} = 580$ and $\ell = 25$ dbar run on 678575 data points in the S_P range 34.21–34.82 g/kg from all up-going ITP3 profiles. (a) The data in θ – S_P space with dashed lines of constant potential density anomaly (kg m^{-3}) referenced to the surface. (b) Profiles 313, 315, and 317 from ITP3 in a limited pressure range to show detail. Each profile is offset in S_P for clarity. (c) The spatial arrangement used as input for the algorithm where the gray points are noise and each color-marker combination indicates a cluster. The same color-marker combinations are used in each panel and the markers in panels (c) and (d) are at the cluster average for each axis. (d) A subset of the data in $\alpha\theta$ – βS_P space with the linear regression line and inverse slope (R_L) noted for each individual cluster and with dashed lines of slope $\alpha\theta/\beta S_A = 1$

A.5 Clusters across time for ITP2

In Figure A.7, we create a figure similar to that of Figure 2.4, but for ITP2.




Figure A.6: A parameter sweep showing the number of clusters found (solid lines) and DBCV (dashed lines) in ITP3 as a function of 72 different values of m_{pts} with $\ell = 25$ dbar.

A.6 Differences from T08

Figure 6 from [Timmermans et al. \(2008\)](#) shows clusters in $\alpha\Theta - \beta S_P$ space with panel (a) showing five values of R_L for ITP2, which range from -3.5 to -3.0, and panel (b) showing 12 values of R_L for ITP5, which range from -3.9 to -2.4. It is unclear whether their reported $R_L = -3.7 \pm 0.9$ applies to data from ITP2, ITP5, both, or all six ITPs they included in their study. It is not explicitly stated whether that value applies to just the clusters shown in that figure or it includes the dozens of clusters they identified earlier in their study. T08 never explicitly state that they did not find a trend in R_L with respect to depth. However, given that T08 reported one overall value of R_L and that [Bebieva & Timmermans \(2019\)](#) (which extended the work of T08) show a constant value of R_L in the depth range we consider, we assume T08 also found a constant value of R_L .

There are several factors that could explain why our results for R_L differ from those of T08. The fact that they used potential temperature θ while we worked with conservative temperature Θ is not relevant, since when we repeat our calculations using θ instead of Θ , the values of R_L we find are negligibly different. It is unclear whether T08 calculated the values of α and β for each point or for each layer. We calculated α and β for each point in each profile individually using functions from the Gibbs Seawater (GSW) Oceanographic Toolbox, the Python implementation of TEOS-10 ([McDougall & Barker, 2011](#)), while T08 used the [McDougall \(1987\)](#) algorithm coded by [Morgan \(1994\)](#) as GSW was not yet released. This difference in the calculation of α and β may conceivably explain why the clusters we show in Figure 2.2(d) are tilted several degrees compared to those shown in Figure 6(a) from T08, and why the values of R_L found here are all less negative than corresponding values in Figure 6(a) of T08. The difference in calculating α and β may also account for the shift in the ranges of βS_P between the two figures. Both plots show a span in βS_P of 4.0×10^{-5} , but our plot ranges from 2.6838×10^{-2} to 2.6878×10^{-2} , while theirs ranges from 2.7002×10^{-2} to 2.7042×10^{-2} . Because the clustering algorithm only considers Θ' and S_P , any difference in calculating α and β would not affect the clusters we found, only the values of R_L .

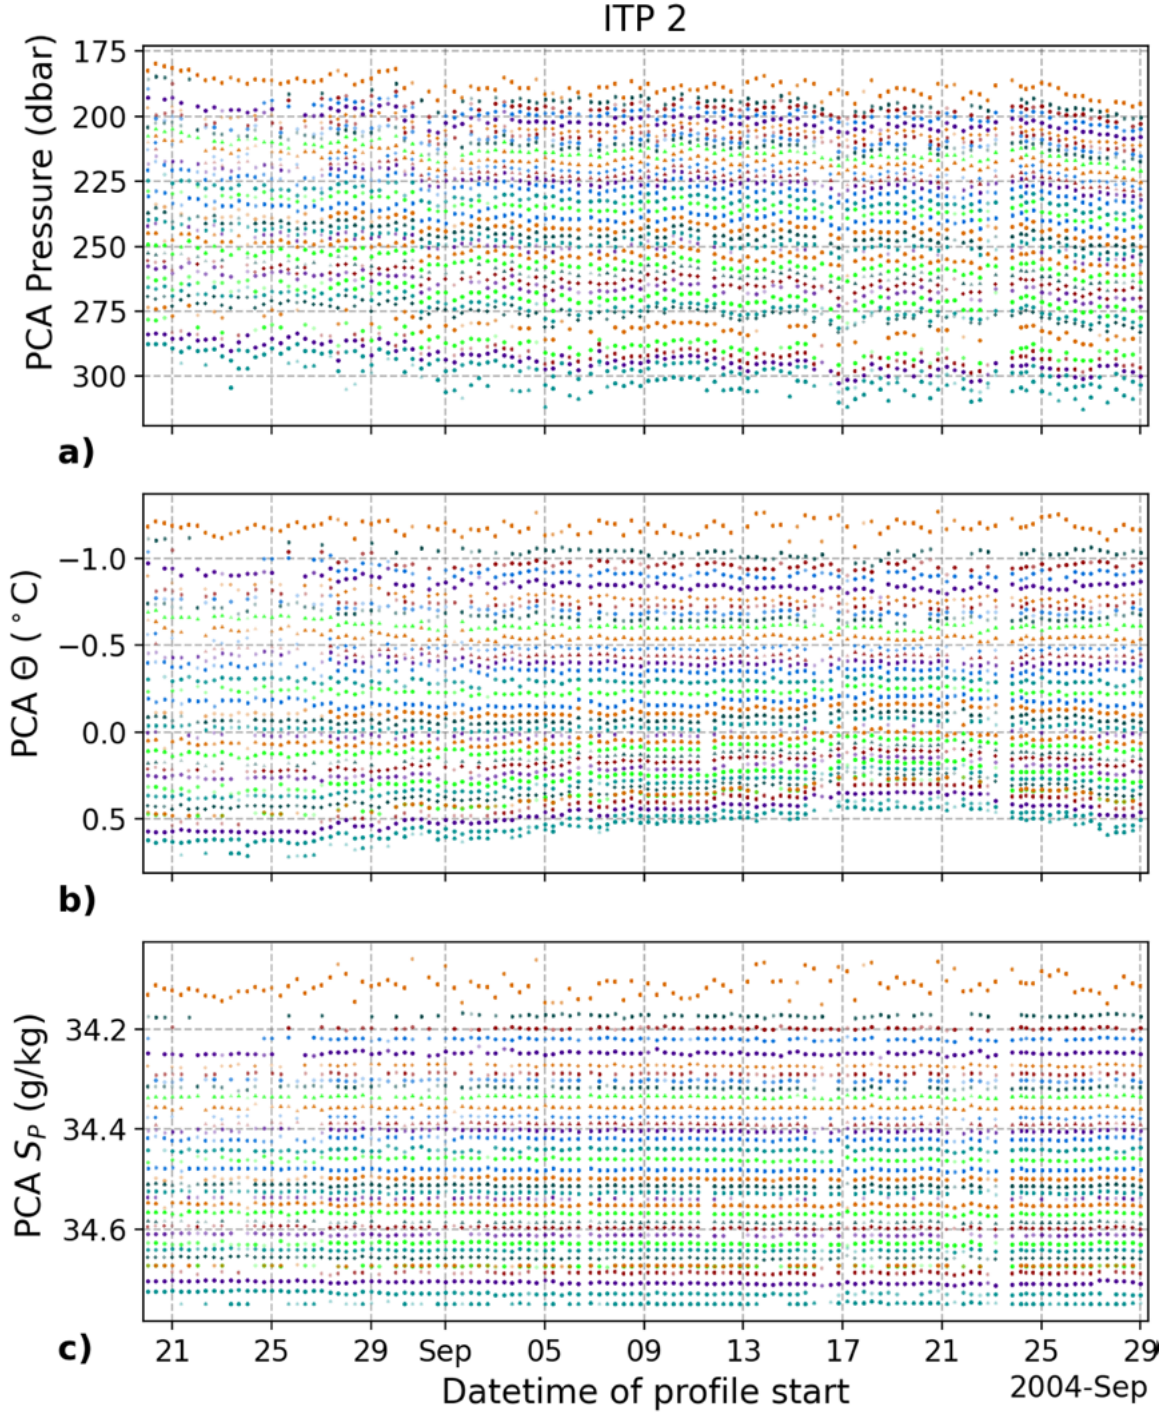


Figure A.7: The average (a) pressure, (b) Θ , and (c) S_P for the points within each cluster for each profile (profile cluster average, PCA) across time. The clustering algorithm was run with $m_{pts} = 170$ and $\ell = 25$ dbar on 53042 data points in the salinity range 34.05–34.75 g/kg from all up-going ITP2 profiles.

A.7 Differences from L22

L22 used 758 profiles from ITP3 while we used all 766 available up-going profiles. The reasons 8 profiles were not used was not stated, nor could we determine exactly which profiles they are. Based on the gaps in Figure 3 of L22, we believe that the 8 missing up-going profiles are from July 2006. We do not, however, believe that this difference of 8 profiles accounts in any significant way for the differences between our results and those of L22.

Instead, we believe the reason our results differ from those of L22 has to do with the difference in the methods used to identify layers. Because the method used by L22 to delineate different layers only considered salinity, they would not separate layers with temperature inversions. Remnant intrusions are homogeneous in salinity but, compared to stair steps, are inverted in temperature and have different patterns of heat and salt flux. They tend to be present near the bottom of the staircase, which is indeed where we find the largest differences between our results and those of L22. A remnant intrusion’s pattern of a section of warmer water above a section of colder water is distinct enough for the clustering algorithm to find two clusters. However, if only salinity is considered, as L22 did, it would appear to be one regular staircase layer. Whether a remnant intrusion represents one or multiple layers is subjective, and we do not claim that one method is superior to the other. Here, in fact, it is the disagreement between these two methods is arguably the most interesting point, because it indicates that a specific physical process is operating.

While we agree with L22 that salinity is the best measured quantity by which to identify the layers of a staircase, it is still important to consider temperature. Another topic of future study would be to adapt this method to automatically distinguish between well-mixed layers and intrusions.

A.8 Subsampling profiles

As the clustering algorithm does not consider profiles individually, only requiring sufficient numbers of observations to be made within the same region and time period, it could potentially detect layers in profiles that have too low resolution to resolve individual steps. While we did not conduct an in-depth analysis, Figure A.8 shows clusterings of the ITP3 data where the profiles were subsampled. We subsampled every profile to only every second, third, or fourth point and found similar clusters to using all available data points in Figure A.5(c). Subsampling effectively reduces the vertical resolution and even though, especially in the case of using only every fourth point, the individual steps are no longer well-resolved, the staircase layers are still found because across all the profiles, a sufficient number of observations occurred within each layer. Note that we did not perform a parameter sweep to find the ideal value of m_{pts} for these clusterings, simply choosing a reasonable estimation based upon the total number of data points left after subsampling.

We had originally investigated the effects of subsampling in anticipation of applying HDBSCAN to data from the Arctic Ice Dynamics Joint Experiment (AIDJEX) during which hundreds of hydrographic profiles were taken in the Canada Basin in 1975-1976. It is known that staircase structures were observed in profiles from this experiment as shown in Figure A.9, reproduced from Bauer et al. (1980). The publicly available versions of these profiles can be found at <https://doi.org/10.34992/4xak-8r05>. However, these have much lower vertical resolutions than the example figure shows, varying around 1 measurement every 1–1.5 m.

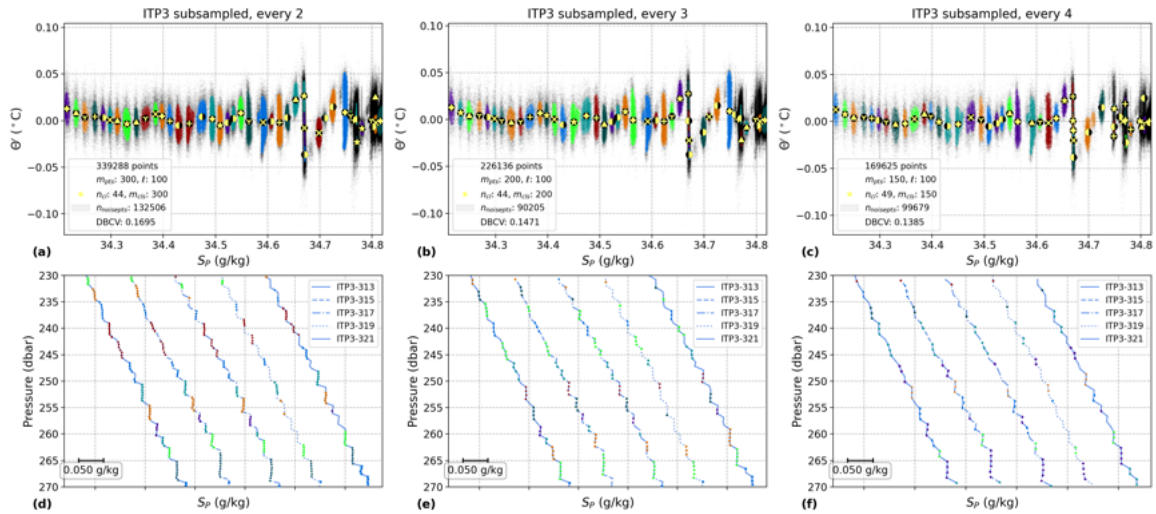


Figure A.8: Examples of clusterings when subsampling the ITP3 data to every (a,d) second, (b,e) third, or (c,f) fourth point. The top row shows the resulting clusterings and the bottom row shows a set of five example profiles to show the effect of subsampling on individual profiles.

We found that the clustering algorithm still identified the majority of the same clusters when subsampled down to even 1 in every 12 points (see Figure A.10) which gives an effective resolution much lower than 1.5 m. Therefore, we expected that the clustering algorithm should be able to identify staircase layers in the AIDJEX data, despite the lower vertical resolution. However, the available versions of the profiles from AIDJEX were not simply subsampled from the high-resolution versions, an example of which is in Figure A.9. A smoothing process was applied to each profile and therefore the signal of the staircase structures is not detectable.




Figure A.9: A section of the temperature and salinity profile taken at Camp Snowbird, Station 1, on 16 May 1975. The step structure can be observed, although distorted by an eddy. Reproduced from Figure 20 of [Bauer et al. \(1980\)](#).

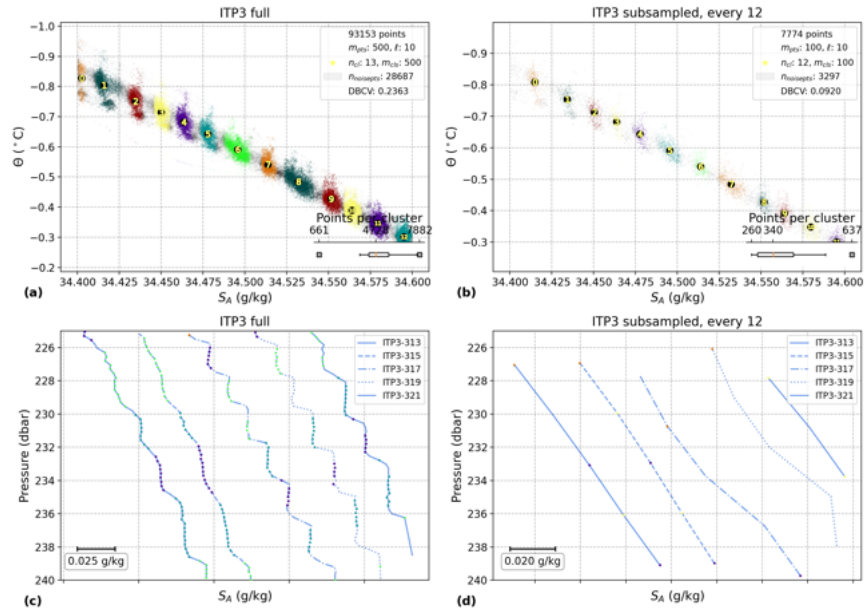


Figure A.10: Examples of clusterings when subsampling the ITP3 data to every (a,d) second, (b,e) third, or (c,f) fourth point. The top row shows the resulting clusterings and the bottom row shows a set of five example profiles to show the effect of subsampling on individual profiles.

Appendix B

Supplementary Materials for: Tracking Beaufort Gyre staircase layers across 17 years

This appendix contains supplemental figures and tables for Chapter [3](#).

Compound Trend	Value	Units	R^2
$d_p(d_t p)$	$(1.62 \pm 27.56) \times 10^{-3}$	dbar/yr/dbar	0.007
$d_p(d_t p)$ of residuals	$(-2.23 \pm 0.60) \times 10^{-2}$	dbar/yr/dbar	-0.395
$d_p(d_t S_A)$	$(2.48 \pm 3.67) \times 10^{-6}$	g/kg/yr/dbar	0.078
$d_p(d_t S_A)$ of residuals	$(1.74 \pm 2.37) \times 10^{-6}$	g/kg/yr/dbar	0.085
$d_p(d_t \Theta)$	$(-2.52 \pm 0.56) \times 10^{-4}$	$^{\circ}\text{C}/\text{yr}/\text{dbar}$	-0.463
$d_p(d_t \Theta)$ of residuals	$(-2.29 \pm 0.23) \times 10^{-4}$	$^{\circ}\text{C}/\text{yr}/\text{dbar}$	-0.759
$d_p(d_t \sigma_1)$	$(1.93 \pm 0.47) \times 10^{-5}$	$\text{kg}/\text{m}^3/\text{yr}/\text{dbar}$	0.431
$d_p(d_t \sigma_1)$ of residuals	$(1.78 \pm 0.24) \times 10^{-5}$	$\text{kg}/\text{m}^3/\text{yr}/\text{dbar}$	0.654
$d_{S_A}(d_t p)$	-4.55 ± 4627	dbar/yr/(g/kg)	-1.14×10^{-4}
$d_{S_A}(d_t p)$ of residuals	-3.10 ± 1.04	dbar/yr/(g/kg)	-0.328
$d_{S_A}(d_t S_A)$	$(4.29 \pm 6.16) \times 10^{-4}$	g/kg/yr/(g/kg)	0.081
$d_{S_A}(d_t S_A)$ of residuals	$(2.44 \pm 3.99) \times 10^{-4}$	g/kg/yr/(g/kg)	0.071
$d_{S_A}(d_t \Theta)$	$(-5.23 \pm 0.87) \times 10^{-2}$	$^{\circ}\text{C}/\text{yr}/(\text{g}/\text{kg})$	-0.572
$d_{S_A}(d_t \Theta)$ of residuals	$(-3.99 \pm 0.36) \times 10^{-2}$	$^{\circ}\text{C}/\text{yr}/(\text{g}/\text{kg})$	-0.789
$d_{S_A}(d_t \sigma_1)$	$(3.92 \pm 0.75) \times 10^{-3}$	$\text{kg}/\text{m}^3/\text{yr}/(\text{g}/\text{kg})$	0.521
$d_{S_A}(d_t \sigma_1)$ of residuals	$(3.02 \pm 0.40) \times 10^{-3}$	$\text{kg}/\text{m}^3/\text{yr}/(\text{g}/\text{kg})$	0.660
$d_{\Theta}(d_t p)$	-0.23 ± 2.11	dbar/yr/ $^{\circ}\text{C}$	-0.013
$d_{\Theta}(d_t p)$ of residuals	-1.61 ± 0.47	dbar/yr/ $^{\circ}\text{C}$	-0.373
$d_{\Theta}(d_t S_A)$	$(2.08 \pm 2.81) \times 10^{-4}$	g/kg/yr/ $^{\circ}\text{C}$	0.086
$d_{\Theta}(d_t S_A)$ of residuals	$(1.23 \pm 1.82) \times 10^{-4}$	g/kg/yr/ $^{\circ}\text{C}$	0.078
$d_{\Theta}(d_t \Theta)$	$(-2.34 \pm 0.40) \times 10^{-2}$	$^{\circ}\text{C}/\text{yr}/^{\circ}\text{C}$	-0.561
$d_{\Theta}(d_t \Theta)$ of residuals	$(-1.86 \pm 0.16) \times 10^{-2}$	$^{\circ}\text{C}/\text{yr}/^{\circ}\text{C}$	-0.805
$d_{\Theta}(d_t \sigma_1)$	$(1.78 \pm 0.34) \times 10^{-3}$	$\text{kg}/\text{m}^3/\text{yr}/^{\circ}\text{C}$	0.517
$d_{\Theta}(d_t \sigma_1)$ of residuals	$(1.42 \pm 0.18) \times 10^{-3}$	$\text{kg}/\text{m}^3/\text{yr}/^{\circ}\text{C}$	0.682
$d_{\sigma_1}(d_t p)$	0.16 ± 7.34	dbar/yr/(kg/m^3)	0.003
$d_{\sigma_1}(d_t p)$ of residuals	-4.72 ± 1.66	dbar/yr/(kg/m^3)	-0.314
$d_{\sigma_1}(d_t S_A)$	$(6.68 \pm 9.77) \times 10^{-4}$	g/kg/yr/(kg/m^3)	0.079
$d_{\sigma_1}(d_t S_A)$ of residuals	$(3.77 \pm 6.33) \times 10^{-4}$	g/kg/yr/(kg/m^3)	0.069
$d_{\sigma_1}(d_t \Theta)$	$(-8.30 \pm 1.38) \times 10^{-2}$	$^{\circ}\text{C}/\text{yr}/(\text{kg}/\text{m}^3)$	-0.572
$d_{\sigma_1}(d_t \Theta)$ of residuals	$(-6.27 \pm 0.58) \times 10^{-2}$	$^{\circ}\text{C}/\text{yr}/(\text{kg}/\text{m}^3)$	-0.782
$d_{\sigma_1}(d_t \sigma_1)$	$(6.20 \pm 1.18) \times 10^{-3}$	$\text{kg}/\text{m}^3/\text{yr}/(\text{kg}/\text{m}^3)$	0.520
$d_{\sigma_1}(d_t \sigma_1)$ of residuals	$(4.72 \pm 0.64) \times 10^{-3}$	$\text{kg}/\text{m}^3/\text{yr}/(\text{kg}/\text{m}^3)$	0.652

Table B.1: The compound trends for each variable v across the cluster average of each variable $d_v(d_t v)$ for both the original variables and the residuals with the corresponding `polyfit2d`.




Figure B.1: The dates on which profiles were taken within the BGR by each instrument in the time period analyzed in this study. This shows only the up-going profiles which extend below 400 dbar. Note that, while the same colors are used here as in the clustering plots, the colors for each ITP are only to make them easier to distinguish and do not correspond to certain clusters with the same color.

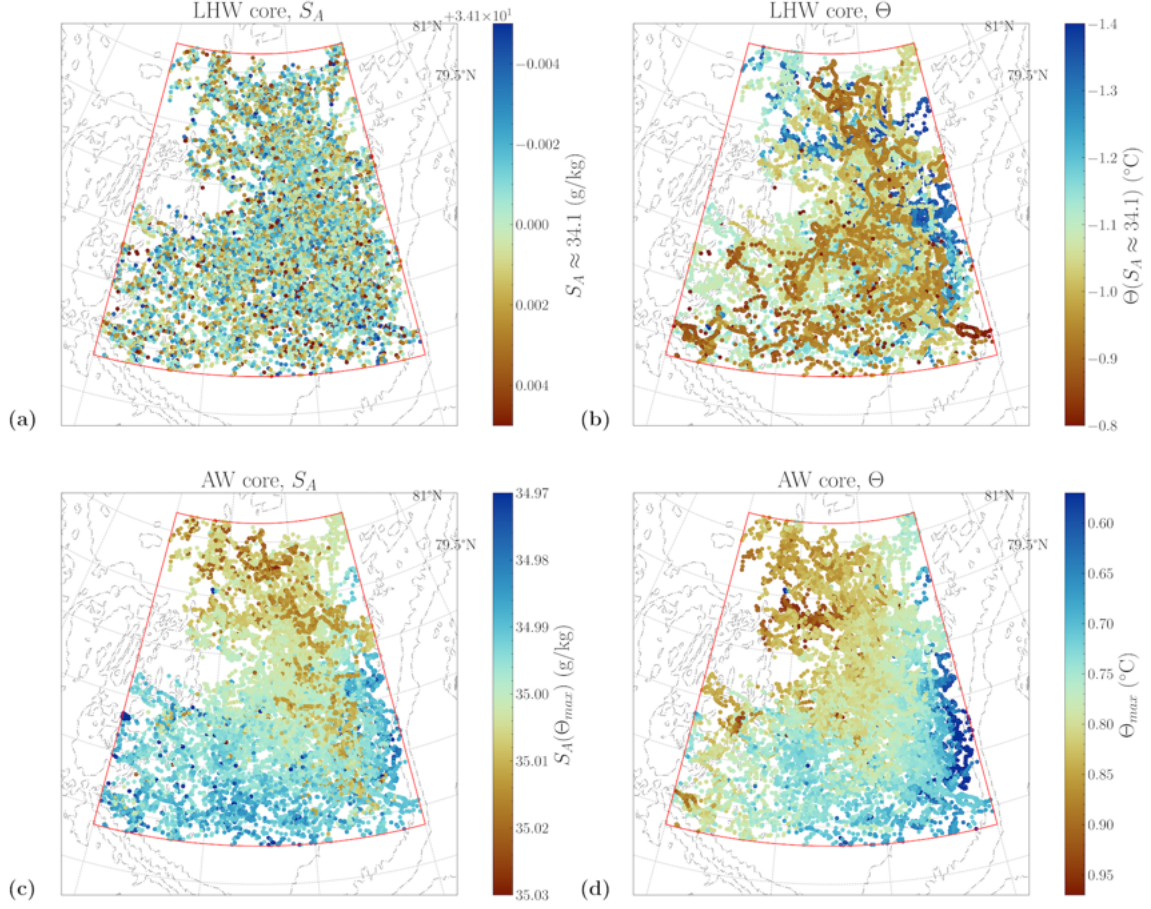


Figure B.2: Maps showing, for each of the 15,247 profiles in the BGR, the (a) salinity and (b) temperature of the LHW core as well as the (c) salinity and (d) temperature of the AW core. The color bars are oriented such that values physically higher in the water column are closer to the top of the plot. Similar plots for the pressure of the LHW and AW cores are in Figure 3.2.




Figure B.3: The salinity across all time periods, colored by cluster, before relabeling to create Figure 3.6. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.




Figure B.4: The pressure across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.




Figure B.5: The temperature across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.




Figure B.6: The density anomaly, referenced to 1000 m, across all time periods, colored by cluster. Noise points are marked in black. The blue and pink triangle markers denote the values of the LHW and AW, respectively, for each profile. The vertical red dashed lines appear on every August 15th to denote the boundary between time periods. The IDs for clusters in BGR0506 (BGR2122) are shown by the colored numbers on the left (right) of the plot.




Figure B.7: Maps of temperature for clusters 6, 23, 35, 52, 63, 69, 79, 90, and 96 over the corresponding `polyfit2d`. The cluster number increases with depth and none of the clusters here are outliers in IR_{S_A} or R_L .

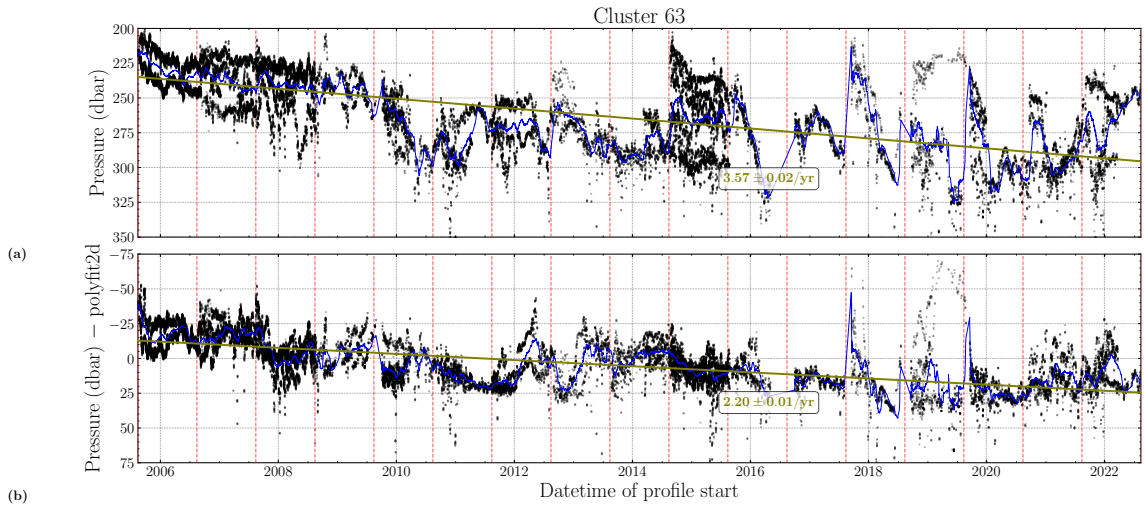


Figure B.8: Trends in time of pressure for Cluster 63 both (a) before ($R^2 = 0.638$) and (b) after correcting with the `polyfit2d` ($R^2 = 0.632$). The blue lines are the 30-day moving average. This includes data from the 66,395 points within Cluster 63.

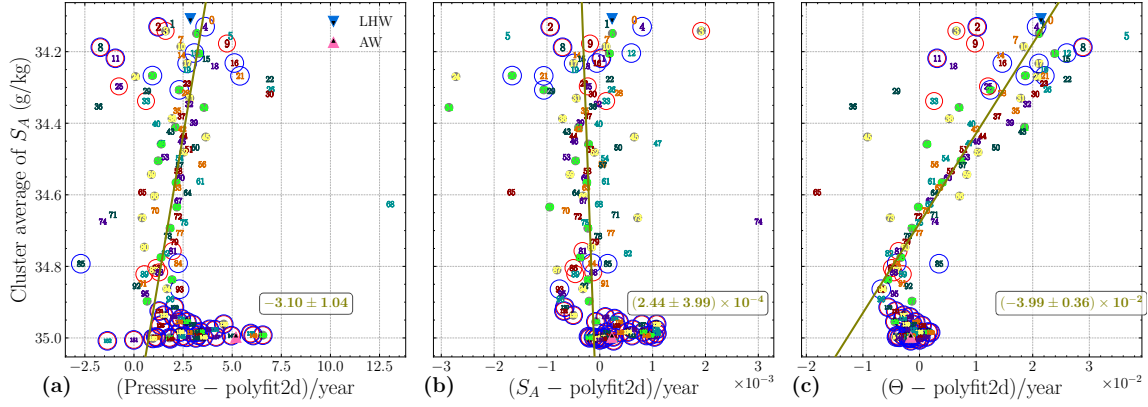


Figure B.9: For all clusters, the trends in time of the polyfit2d residual for (a) pressure, (b) salinity, and (c) temperature with respect to the cluster average of salinity. This is similar to Figure 3.9, but with each cluster marked with its ID and color. Blue and red circles denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have (from left to right) R^2 values of -0.328, 0.071, and -0.789. The blue and pink triangle markers denote the values for the LHW and AW, respectively.




Figure B.10: The thicknesses (in dbar) of the detected layers. (a) The non-zero thickness values against salinity of each cluster that is neither an outlier in IR_{S_A} nor R_L for every profile in the BGR, colored by cluster. For all clusters, the (a) cluster average of layer thickness and (b) trends of percent change in layer thickness over time with respect to the cluster average of salinity. This is similar to Figure 3.10, but with each cluster marked with its ID and color. The values in (c) are calculated by taking the trend in time for the cluster's non-zero thicknesses and dividing by that cluster's average thickness shown in (b), however are highly uncertain, see text for explanation. Blue and red circles denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have R^2 values of 0.609 for (b) and -0.041 for (c).




Figure B.11: (a) The non-zero thicknesses (in dbar) of the detected layers for just ITP13 against pressure for each cluster that is neither an outlier in IR_{SA} nor R_L . The linear trend has an R^2 value of 0.693. The red box denotes the domain bounds of the inset. (b) Reproduced from [Shibley & Timmermans \(2019\)](#). The thickness of layers they detected for ITP13. The blue dots represent the average depth and thickness for the data binned to 0.5-m increments and the linear fit line through them has a slope of 0.12 ± 0.01 m/m.

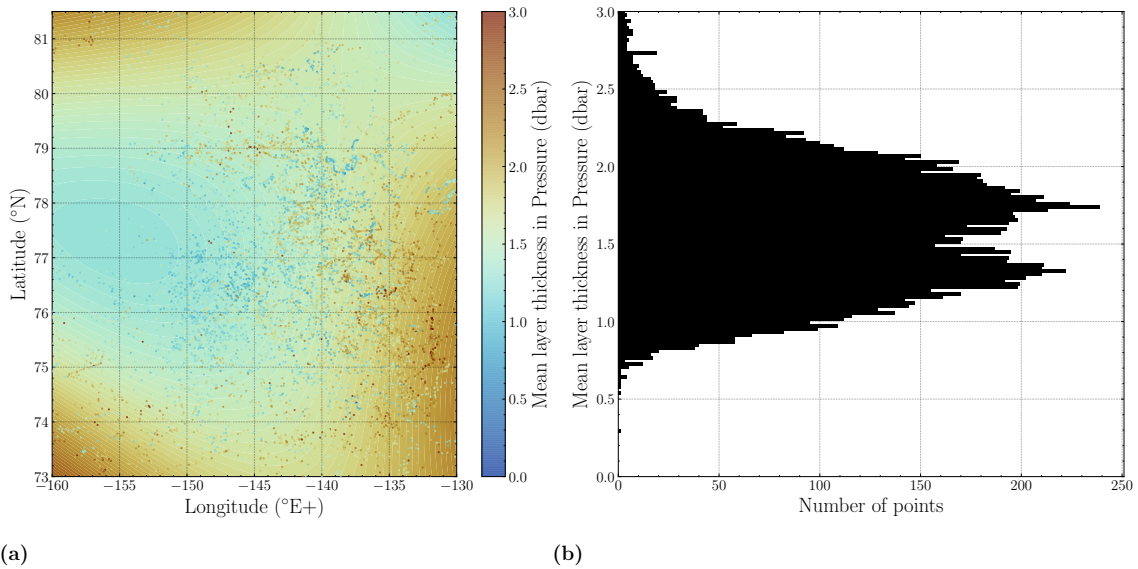


Figure B.12: The mean of the non-zero thicknesses (in dbar) per profile of the detected layers for all clusters that are neither outliers in IR_{SA} nor R_L . (a) A map of these mean layer thicknesses per profile on top of the `polyfid2d`. (b) A histogram of the mean layer thicknesses per profile.

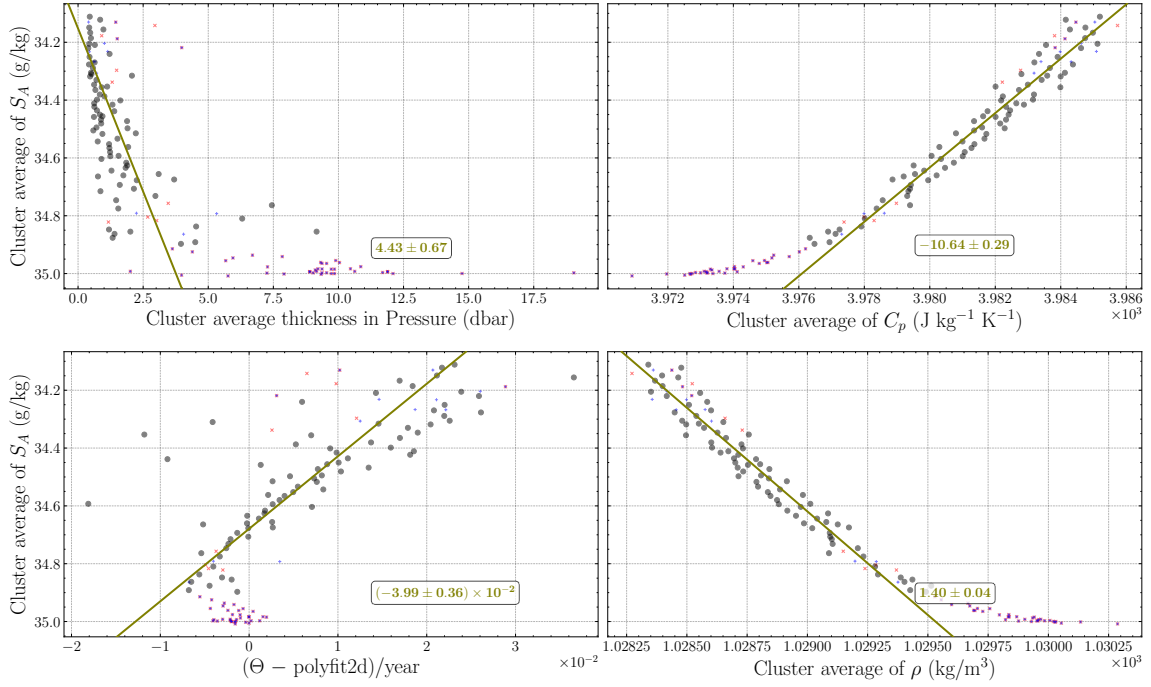


Figure B.13: The cluster average of each component of heat flux against salinity. (a) The cluster average layer thickness. (b) The cluster average isobaric heat capacity. A blue “+” and a red “x” denote clusters that are outliers in IR_{S_A} and R_L , respectively. All outliers were ignored when calculating the fit lines which have R^2 values of 0.609 for (a) and -0.107 for (b).

Appendix C

Supplementary equations and derivations

C.1 The Boussinesq Equations of Motion

The simulations in Chapter 4 are based upon the Boussinesq equations of motion which come from taking the Boussinesq approximation on the Navier Stokes equations. However, many variations of these particular equations exist. Therefore, I derive here the set of equations I use, starting from the fundamental relations and making every assumption and approximation clear along the way.

The three fundamental relations for the system in question are the conservation of mass, the conservation of momentum, and the thermodynamic equation. For each, I consider the properties of a small fluid parcel, bounded by some arbitrary control volume.

C.1.1 Conservation of Momentum

While the **conservation of momentum** is commonly known as *Newton's second law*, $a = F/m$, it can be more generally expressed as

$$\frac{D\vec{v}}{Dt} = \frac{\sum \vec{F}}{\rho\Delta V} = \sum \vec{F}^* \quad (\text{C.1})$$

where $D/Dt = \partial_t + \vec{v} \cdot \nabla$ is the material derivative (Sutherland, 2010; Cushman-Roisin & Beckers, 2011; Kundu et al., 2015; Vallis, 2017) and $\sum \vec{F}^*$ [N/kg] represents the sum of forces on the fluid per unit mass $\rho\Delta V$ [kg] (Sutherland, 2010; Cushman-Roisin & Beckers, 2011; Lautrup, 2011; Vallis, 2017). I use an asterisk superscript to denote quantities that are “per unit mass.” Note that (C.1) actually represents multiple equations, one for each component in \vec{v} , where usually $\vec{v} = (u, v, w)$. I make the assumption that there are four relevant forces to consider in this system:

$$\sum \vec{F}^* = \vec{F}_C^* + \vec{F}_\nu^* + \vec{F}_p^* + \vec{F}_g^*. \quad (\text{C.2})$$

where \vec{F}_C^* , \vec{F}_ν^* , \vec{F}_p^* , and \vec{F}_g^* are the forces per unit mass due to Coriolis, viscosity, pressure, and gravity, respectively (Sutherland, 2010).

Forces

For a parcel of geophysical fluid on a rotating Earth, there is an apparent Coriolis force which arises from taking a reference frame which is stationary with respect to the planet's surface. The Coriolis force per unit mass can be expressed as:

$$\vec{F}_C^* = -2\vec{\Omega} \times \vec{u} = (fv, -fu, 0) \quad (\text{C.3})$$

where $\vec{\Omega}$ [rad/s] is the rotation rate of the Earth and $f = 2\Omega \sin \phi$ is the Coriolis parameter at a latitude ϕ (Sutherland, 2010; Vallis, 2017). If the scales of motion cover a sufficiently small range of latitudes, approximately no larger than $\Delta\phi \approx 10^\circ$, then the Coriolis parameter can be considered constant. This is known as the 'f-plane' approximation and, at midlatitudes, $f_0 \simeq 10^{-4} \text{ s}^{-1}$ (Sutherland, 2010).

Viscosity can be thought of as the diffusion of momentum with the tendency to bring a fluid to a uniform velocity. A fast-moving fluid parcel is slowed by contact with an adjacent slow-moving parcel and, similarly, the velocity of the flow-moving parcel is increased by contact with the fast moving parcel (Sutherland, 2010). If we assume the stress of these parcels sliding past each other in the x direction establishes a proportional velocity gradient, then the momentum flux in that direction is

$$\vec{\mathcal{F}}_{\nu x} = -\nu\rho\nabla u$$

where ν [m^2/s] is the kinematic viscosity (Sutherland, 2010). We take seawater to be a homogeneous Newtonian fluid and assume that the viscosity does not change significantly over the relevant range of temperatures, i.e. ν is constant (Sutherland, 2010; Lautrup, 2011). We also assume an isotropic fluid and so the momentum flux will have the same relation in the other directions for $\vec{\mathcal{F}}_{\nu y}$ and $\vec{\mathcal{F}}_{\nu z}$ (Lautrup, 2011). Therefore, a non-zero divergence in the total momentum flux $\vec{\mathcal{F}}_\nu$ will cause an acceleration which means the viscous force per unit mass is (Sutherland, 2010; Vallis, 2017)

$$\vec{F}_\nu^* = -\frac{1}{\rho} \nabla \cdot \vec{\mathcal{F}}_\nu = \nu \nabla^2 \vec{v}. \quad (\text{C.4})$$

Pressure p [N/m^2] is the normal force per unit area of the surface of a control volume $d\vec{S}$ [m^2], so the pressure force per unit mass \vec{F}_p^* [N/kg] would be the pressure integrated over that surface

$$\vec{F}_p^* = \frac{-1}{\rho\Delta V} \oint_S p d\vec{S}$$

where ΔV [m^3] is the control volume and the negative sign orients the normal outward from the control volume, insuring the force acts from high to low pressure (Sutherland, 2010; Vallis, 2017). Applying the divergence theorem gives

$$\vec{F}_p^* = \frac{-1}{\rho\Delta V} \int_V \nabla p dV = -\frac{\nabla p}{\rho}. \quad (\text{C.5})$$

Note that this implicitly assumes that p is isotropic, having no dependence on direction. This assumption will become explicitly valid when I later assume hydrostatic balance, making Pascal's

Law applicable (Lautrup, 2011).

I define gravity as acting in the downward z -direction, so the force per unit mass due to gravity \vec{F}_g^* [N/kg] would be

$$\vec{F}_g^* = -g\hat{z}. \quad (\text{C.6})$$

where $g \simeq 9.80 \text{ m/s}^2$ (Sutherland, 2010).

The Navier-Stokes Equation

Substituting for all the force terms (C.3), (C.4), (C.5), and (C.6) into equation (C.1) yields

$$\frac{D\vec{v}}{Dt} = -(2\vec{\Omega}) \times \vec{v} + \nu \nabla^2 \vec{v} - \frac{\nabla p}{\rho} - g\hat{z}. \quad (\text{C.7})$$

This is one form of the *Navier-Stokes equation* (Sutherland, 2010; Lautrup, 2011; Vallis, 2017). We can greatly simplify (C.7) by invoking hydrostatic balance and the Boussinesq approximation.

The Boussinesq Approximation

The ocean varies very little in density, only changing by a few percent from the surface to the ocean floor (Sutherland, 2010). Since most of the variation is in the vertical, it is convenient to subdivide ρ into a background density profile $\bar{\rho}(z)$ and the density perturbations $\rho'(x, y, z, t)$ where $\rho = \bar{\rho} + \rho'$ with $\bar{\rho} \gg \rho'$. The same is true for pressure so we define $p = \bar{p} + p'$ where $\bar{p}(z)$ is the hydrostatically balanced pressure, $p'(x, y, z, t)$ are the pressure fluctuations, and with $\bar{p} \gg p'$ (Vallis, 2017).

When a fluid is at rest, $\vec{v} = 0$, $\rho' = 0$, and $p' = 0$ so the Navier-Stokes equation (C.7) becomes

$$\frac{\nabla \bar{p}}{\bar{\rho}} = -g\hat{z}$$

and for geophysical fluids on a local scale, gravity acts in the vertical, meaning that the z direction is the only non-zero component of the derivative (Lautrup, 2011) which gives

$$\partial_z \bar{p} = -\bar{\rho}g\hat{z}, \quad (\text{C.8})$$

also known as *hydrostatic balance* (Vallis, 2017).

The Boussinesq approximation considers density variations to be negligible unless they affect buoyancy forces because g is relatively large (Vallis, 2017). This is a reasonable approximation for the ocean because density variations are small, especially in horizontal directions (Sutherland, 2010). We multiply (C.7) by ρ and, by taking the Boussinesq approximation, we neglect all terms with ρ' except for the one multiplying the buoyancy force:

$$\bar{\rho} \frac{D\vec{v}}{Dt} = -\bar{\rho}(2\vec{\Omega}) \times \vec{v} + \bar{\rho}\nu \nabla^2 \vec{v} - \nabla(\bar{p} + p') - (\bar{\rho} + \rho')g\hat{z} \quad (\text{C.9})$$

where, because of hydrostatic balance (C.8) and the fact that \bar{p} only depends on z , the terms $-\nabla \bar{p}$ and $-\bar{\rho}g\hat{z}$ cancel out.

Because of how little density varies in the ocean, we can further subdivide the background density profile $\bar{\rho}(z) = \rho_0 + \hat{\rho}(z)$, taking a constant reference density ρ_0 and a profile of background density

variations $\hat{\rho}(z)$ where $\rho_0 \gg \hat{\rho}$. Invoking the Boussinesq approximation again to neglect all terms with ρ' or $\bar{\rho}$ except those related to buoyancy and dividing by ρ_0 gives:

$$\frac{D\vec{v}}{Dt} = -(2\vec{\Omega}) \times \vec{v} + \nu \nabla^2 \vec{v} - \frac{\nabla p'}{\rho_0} + \frac{\rho'}{\rho_0} g \hat{z}. \quad (\text{C.10})$$

By defining specific pressure to be $p^* = p'/\rho_0$ [$\text{m}^2 \text{s}^{-2}$] and buoyancy to be $\vec{b} = -\frac{\rho'}{\rho_0} g \hat{z}$ [m/s^2] (Vallis, 2017), we can write the momentum equation as

$$\frac{D\vec{v}}{Dt} = -(2\vec{\Omega}) \times \vec{v} + \nu \nabla^2 \vec{v} - \nabla p^* + \vec{b} \quad (\text{C.11})$$

C.1.2 Conservation of Mass

I express the **conservation of mass** as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0, \quad (\text{C.12})$$

where ρ [kg/m^3] is the fluid density, t [s] is time, and \vec{v} [m/s] is the velocity field of the fluid. This says that, for some fixed control volume, the change in mass of fluid inside needs to be equal to the flow of fluid in or out (Sutherland, 2010). Or, as stated by Lautrup (2011), “It expresses the obvious fact that the mass you gain while you eat equals the mass of the food you pass into your mouth.” It is also known as the *local equation of mass conservation* or the *equation of continuity*.

Water is an effectively incompressible fluid (Cushman-Roisin & Beckers, 2011). Using the subdivisions of density used earlier, $\rho(t, x, y, z) = \rho_0 + \hat{\rho}(z) + \rho'(t, x, y, z)$ where $\rho_0 \gg \hat{\rho} \gg \rho'$. This “effective” incompressibility arises because $\rho \approx \rho_0$ and so can generally be assumed to have a constant density. It cannot accumulate anywhere, that is, for a given control volume, equal amounts of fluid must both enter and leave the boundaries for each unit of time (Lautrup, 2011). Therefore, the mass within the control volume cannot change with time, $\partial \rho / \partial t = 0$, and conservation of mass is now equivalent to a conservation of volume (Cushman-Roisin & Beckers, 2011). Using this in (C.12) results in the *local incompressibility condition*

$$\nabla \cdot \vec{v} = 0, \quad (\text{C.13})$$

which states that the divergence of the velocity vanishes for incompressible flow.

C.1.3 Thermodynamic Equation

The last fundamental relation, **the thermodynamic equation**, comes from the first law of thermodynamics. I express this as a conservation of internal energy

$$dE_{int} = dW + dQ + dC \quad (\text{C.14})$$

which states that the change in internal energy dE_{int} [J/s] of a parcel is equal to the sum of the change in heat of the parcel dQ , the work done on the parcel dW , and the change in energy due to its chemical composition dC (Vallis, 2017).

Work is defined as the product of the force times the displacement (Cushman-Roisin & Beckers,

2011). In this case, the work done on the parcel is due to pressure, which itself is a force per unit area, then $W = pV$ [kg m²/s²] where V [m³] is a volume. Then, the change in work per unit mass is

$$dW^* = \frac{p}{\rho^2} \frac{\partial \rho}{\partial t}. \quad (\text{C.15})$$

Because of the conservation of mass (C.12) and the incompressibility condition (C.13), dW' [m²/s³] goes to zero, meaning that we neglect work when considering dE_{int} for this system (Cushman-Roisin & Beckers, 2011).

Assuming no internal sources of heat, the change in the heat of the parcel can only be due to diffusion and is given by Fourier's heat equation

$$dQ = \frac{k_T}{\rho} \nabla^2 T$$

where k_T [kg m/K s³] is the thermal conductivity and T [K] is temperature (Cushman-Roisin & Beckers, 2011). The change in the heat of the parcel can also be thought of as the flow of temperature and so $dQ = C_v(DT/Dt)$, where C_v [m²/K s²] is the heat capacity at constant volume. Therefore, I can express this relation as

$$\frac{DT}{Dt} = \kappa_T \nabla^2 T. \quad (\text{C.16})$$

by defining the thermal diffusivity $\kappa_T = k_T/(\rho C_v)$ [m²/s], (Lautrup, 2011; Kundu et al., 2015)

For seawater, I make the assumption that the change in the energy due to chemical composition dC will be only due to variations in salinity. In an analogous manner to heat, that leads to the following relation

$$\frac{DS}{Dt} = \kappa_S \nabla^2 S \quad (\text{C.17})$$

where S [g/kg] is salinity and κ_S [m²/s] is the coefficient of salt diffusion (Cushman-Roisin & Beckers, 2011).

We have already assumed water is incompressible, and therefore the only two factors that affect density are temperature and salinity. For small fluctuations in both, the density of seawater can be represented by the linear equation of state (1.1)

$$\rho = \rho_0[1 - \alpha(T - T_0) + \beta(S - S_0)] \quad (\text{C.18})$$

where ρ_0 , T_0 , and S_0 are some reference density, temperature, and salinity, respectively; α [K⁻¹] is the thermal expansion coefficient (Timmermans et al., 2008; Cushman-Roisin & Beckers, 2011), and β [unitless] is the haline contraction coefficient (Bebieva & Timmermans, 2019; van der Boog, Otto Koetsier, et al., 2021). Taking the material derivative of (C.18) gives

$$\frac{D\rho}{Dt} = -\rho_0\alpha\frac{DT}{Dt} + \rho_0\beta\frac{DS}{Dt}$$

Substituting (C.16) and (C.17) into that expression and then rearranging (C.18) to eliminate S gives

$$\frac{D\rho}{Dt} = \kappa_S \nabla^2 \rho - \rho_0 \alpha \kappa_T \nabla^2 T + \rho_0 \alpha \kappa_S \nabla^2 T$$

This expression simplifies greatly if we assume the diffusivities for heat κ_T and salinity κ_S are equal. As described in Section 1.1, the double-diffusive process which allows thermohaline staircases to form only takes place because these two diffusivities are significantly different in magnitude, $\kappa_T = 1.4 \times 10^{-7} \text{m}^2/\text{s}$, and $\kappa_S = 1.1 \times 10^{-9} \text{m}^2/\text{s}$ (Shibley & Timmermans, 2019). However, for the numerical experiments I present in Chapter 4, I assume the stratification is fixed in time and that the processes that rely on the difference between κ_T and κ_S are small in scale compared to the wave motions I am interested in (Cushman-Roisin & Beckers, 2011). Therefore, I adopt the assumption that $\kappa_T = \kappa_S = \kappa = 1.4 \times 10^{-7} [\text{m}^2/\text{s}]$, which leads to

$$\frac{D\rho}{Dt} = \kappa \nabla^2 \rho, \quad (\text{C.19})$$

the thermodynamic equation in terms of density (Winters et al., 2004; Lautrup, 2011). It may seem that (C.19) is in direct conflict with the conservation of mass (C.12). However, by taking into account the subtleties of water being nearly incompressible, this conflict is resolved by re-expressing (C.19) in terms of buoyancy.

In Terms of Buoyancy

To resolve the apparent conflict between the thermodynamic equation (C.19) and the conservation of mass (C.12), we start by expanding the material derivative in (C.19):

$$\partial_t \rho = \kappa \nabla^2 \rho - \vec{v} \cdot \nabla \rho \quad (\text{C.20})$$

Next, as before, we subdivide ρ into the constant reference density, the background density variations, and the density perturbations, $\rho(t, x, y, z) = \rho_0 + \hat{\rho}(z) + \rho'(t, x, y, z)$:

$$\partial_t (\rho') = \kappa \nabla^2 (\hat{\rho} + \rho') - \vec{v} \cdot \nabla (\hat{\rho} + \rho') \quad (\text{C.21})$$

where, since ρ_0 is constant, it disappears from every term and since $\hat{\rho}$ does not depend on t , it is absent from the left-hand side.

Then, we multiply by $-g/\rho_0$ and substitute in buoyancy $\vec{b} = -\frac{\rho'}{\rho_0} g \hat{z} [\text{m/s}^2]$ where applicable

$$\partial_t \vec{b} = -\kappa \frac{g}{\rho_0} \partial_z^2 \hat{\rho} + \kappa \nabla^2 \vec{b} + w \frac{g}{\rho_0} \partial_z \hat{\rho} - \vec{v} \cdot \nabla \vec{b}. \quad (\text{C.22})$$

For the projects in Chapter 4, I will assume linear background density variations so $\partial_z^2 \hat{\rho}$ goes to zero.

We define the background stratification frequency $N(z) [\text{s}^{-1}]$ to be

$$N^2(z) = -\frac{g}{\rho_0} \partial_z \hat{\rho}. \quad (\text{C.23})$$

which is also sometimes called the *Brunt-Väisälä frequency* (Sutherland, 2010; Cushman-Roisin &

Beckers, 2011; Kundu et al., 2015). Substituting this into (C.22) gives

$$\partial_t \vec{b} + \vec{v} \cdot \nabla \vec{b} = \kappa \nabla^2 \vec{b} - N^2(z)w, \quad (\text{C.24})$$

the thermodynamic equation in terms of buoyancy.

C.1.4 The Boussinesq Equations of Motion

Here, I bring together equations (C.11), (C.13), and (C.24) the Boussinesq equations of motion in terms of buoyancy, valid in three dimensions:

$$\frac{D\vec{v}}{Dt} = -(2\vec{\Omega}) \times \vec{v} + \nu \nabla^2 \vec{v} - \nabla p^* + \vec{b} \quad (\text{C.25a})$$

$$\nabla \cdot \vec{v} = 0 \quad (\text{C.25b})$$

$$\frac{D\vec{b}}{Dt} = \kappa \nabla^2 \vec{b} - N^2(z)w \quad (\text{C.25c})$$

where $D/Dt = \partial_t + \vec{v} \cdot \nabla$ is the material derivative, $\vec{v} = (u\hat{x}, v\hat{y}, w\hat{z})$ is the velocity in [m/s], $\vec{\Omega}$ is the rotation rate of the Earth in [rad/s], ν is the kinematic viscosity in [m²/s], $\nabla^2 = \partial_x^2 + \partial_y^2 + \partial_z^2$ is the Laplacian, $\nabla = (\partial_x, \partial_y, \partial_z)$ is the gradient, $p^* = p'/\rho_0$ is the specific pressure in [m²/s²], $\vec{b} = -\frac{\rho'}{\rho_0}g\hat{z}$ is the buoyancy in [m/s²], κ is the diffusivity in [m²/s], and $N^2(z)$ is the background stratification frequency in [s⁻²].

In terms of the streamfunction

Here, I show how I take the Boussinesq equations (C.25) and express them in terms of the streamfunction, ψ , resulting in the equation I employ for some of my numerical experiments in Chapter 4. Assuming the amplitude of the waves is small enough such that the terms proportional to amplitude squared (in this case, the nonlinear advection terms) can be neglected, the material derivatives collapse to be just time derivatives, $D/Dt \rightarrow \partial_t$. This is known as the *small amplitude assumption* (Sutherland, 2010; Cushman-Roisin & Beckers, 2011). Here, I maintain the viscous terms in order to compare to other studies in Chapter 4. By assuming an isotropic fluid and taking the f -plane assumption, there are no distinctions between horizontal directions. This allows me to always orient the x -axis along the horizontal direction of wave motion and assume a two-dimensional system (Sutherland, 2010). Additionally, I take the length scale of my system to be small enough that the Coriolis force is negligible ($f = 0$). This gives

$$\partial_x u + \partial_z w = 0 \quad (\text{C.26a})$$

$$\partial_t b = -N^2(z)w \quad (\text{C.26b})$$

$$\partial_t u = \nu(\partial_x^2 u + \partial_z^2 u) - \partial_x p^* \quad (\text{C.26c})$$

$$\partial_t w = \nu(\partial_x^2 w + \partial_z^2 w) - \partial_z p^* + b. \quad (\text{C.26d})$$

Because this flow is incompressible and restricted to two dimensions, I can define the stream-

function ψ to be $(u, w) = (-\partial_z \psi, \partial_x \psi)$ (Sutherland, 2010). Substituting this into (C.26) gives:

$$-\partial_x \partial_z \psi + \partial_x \partial_z \psi = 0 \quad (\text{C.27a})$$

$$\partial_t b = -N^2(z) \partial_x \psi \quad (\text{C.27b})$$

$$-\partial_t \partial_z \psi = \nu(-\partial_x^2 \partial_z \psi - \partial_z^3 \psi) - \partial_x p^* \quad (\text{C.27c})$$

$$\partial_t \partial_x \psi = \nu(\partial_x^3 \psi + \partial_x \partial_z^2 \psi) - \partial_z p^* + b \quad (\text{C.27d})$$

where the continuity equation (C.27a) is automatically satisfied. Next, I apply ∂_x to (C.27b), $\partial_t \partial_z$ to (C.27c), and $\partial_t \partial_x$ to (C.27d):

$$\partial_t \partial_x b = -N^2(z) \partial_x^2 \psi \quad (\text{C.28a})$$

$$-\partial_t^2 \partial_z^2 \psi = -\nu \partial_t (\partial_x^2 \partial_z^2 \psi + \partial_z^4 \psi) - \partial_t \partial_x \partial_z p^* \quad (\text{C.28b})$$

$$\partial_t^2 \partial_x^2 \psi = \nu \partial_t (\partial_x^4 \psi + \partial_x^2 \partial_z^2 \psi) - \partial_t \partial_x \partial_z p^* + \partial_t \partial_x b \quad (\text{C.28c})$$

Then, subtracting (C.28b) from (C.28c), I get

$$\partial_t^2 (\partial_x^2 \psi + \partial_z^2 \psi) = \nu \partial_t (\partial_x^4 \psi + 2\partial_x^2 \partial_z^2 \psi + \partial_z^4 \psi) + \partial_t \partial_x b \quad (\text{C.29})$$

where the pressure terms have cancelled out. Finally, to reduce the system to one equation, I substitute (C.28a) for $\partial_t \partial_x b$:

$$\partial_t^2 \nabla^2 \psi = \nu \partial_t \nabla^4 \psi - N^2(z) \partial_x^2 \psi \quad (\text{C.30})$$

where $\nabla^4 = \partial_x^4 + 2\partial_x^2 \partial_z^2 + \partial_z^4$ is the biharmonic operator, the square of the Laplacian (Ghaemsaïdi et al., 2016).

C.2 The Polarization Relation

In order to correctly force waves in my numerical experiments using the Boussinesq equations (C.25) in two dimensions, I need the polarization relation between u , w , p^* , and b . As with finding the dispersion relation (4.5), I start by taking the standard plane wave assumption (4.4) which I formulate here as

$$\eta(t, x, z) = \eta_0 \exp[i(k_x x + k_z z - \omega t)], \quad (\text{C.31})$$

where k_x and k_z are the horizontal and vertical wavenumbers, ω is the frequency of the boundary forcing, η represents any one of the variables u , w , p^* , or b , and η_0 is an arbitrary coefficient.

I then make the *small amplitude assumption*, that is, η_0 is small enough such that the nonlinear terms, which scale with the amplitude squared, can be neglected, linearizing the equations of motion

(Sutherland, 2010; Cushman-Roisin & Beckers, 2011):

$$\partial_t u = -\partial_x p^* \quad (\text{C.32a})$$

$$\partial_t w = -\partial_z p^* + b \quad (\text{C.32b})$$

$$\partial_x u + \partial_z w = 0 \quad (\text{C.32c})$$

$$\partial_t b = -N^2(z)w. \quad (\text{C.32d})$$

Similar to Section 4.2.2, the plane wave assumption causes the partial derivatives become

$$\partial_t \eta = -i\omega \eta, \quad \partial_x \eta = ik_x \eta, \quad \partial_z \eta = ik_z \eta$$

which can then be substituted in to (C.32) to give

$$-i\omega u = -ik_x p^* \quad (\text{C.33a})$$

$$-i\omega w = -ik_z p^* + b \quad (\text{C.33b})$$

$$ik_x u + ik_z w = 0 \quad (\text{C.33c})$$

$$-i\omega b = -N^2(z)w. \quad (\text{C.33d})$$

Expanding (C.33d) with the plane waves assumption,

$$-i\omega b_0 \exp[i(k_x x + k_z z - \omega t)] = -N^2(z)w_0 \exp[i(k_x x + k_z z - \omega t)]$$

and defining $b_0 = -\hat{A}g$ means that the coefficient for the vertical velocity is

$$w_0 = -\hat{A}g \frac{i\omega}{N^2}.$$

Using this in (C.33c),

$$ik_x u_0 \exp[i(k_x x + k_z z - \omega t)] - ik_z \frac{i\omega \hat{A}g}{N^2} \exp[i(k_x x + k_z z - \omega t)] = 0$$

means that the coefficient for the horizontal velocity is

$$u_0 = \hat{A}g \frac{i\omega k_z}{N^2 k_x}.$$

Using that in (C.33a),

$$i\omega \hat{A}g \frac{i\omega k_z}{N^2 k_x} \exp[i(k_x x + k_z z - \omega t)] = ik_x p_0^* \exp[i(k_x x + k_z z - \omega t)]$$

means that the coefficient for the pressure is

$$p_0^* = \hat{A}g \frac{i\omega^2 k_z}{N^2 k_x^2}.$$

Expressing the plane wave assumption for all four variables with the above coefficients and

rewriting the complex exponentials using Euler's formula gives

$$u = +\hat{A}g \frac{i\omega k_z}{N^2 k_x} [i \sin(k_x x + k_z z - \omega t) + \cos(k_x x + k_z z - \omega t)] \quad (\text{C.34a})$$

$$w = -\hat{A}g \frac{i\omega}{N^2} [i \sin(k_x x + k_z z - \omega t) + \cos(k_x x + k_z z - \omega t)] \quad (\text{C.34b})$$

$$b = -\hat{A}g [i \sin(k_x x + k_z z - \omega t) + \cos(k_x x + k_z z - \omega t)] \quad (\text{C.34c})$$

$$p^* = +\hat{A}g \frac{i\omega^2 k_z}{N^2 k_x^2} [i \sin(k_x x + k_z z - \omega t) + \cos(k_x x + k_z z - \omega t)]. \quad (\text{C.34d})$$

Taking the real part of (C.34) gives the polarization relation

$$u = -\hat{A}g \frac{\omega k_z}{N^2 k_x} \sin(k_x x + k_z z - \omega t) \quad (\text{C.35a})$$

$$w = +\hat{A}g \frac{\omega}{N^2} \sin(k_x x + k_z z - \omega t) \quad (\text{C.35b})$$

$$b = -\hat{A}g \cos(k_x x + k_z z - \omega t) \quad (\text{C.35c})$$

$$p^* = -\hat{A}g \frac{\omega^2 k_z}{N^2 k_x^2} \sin(k_x x + k_z z - \omega t). \quad (\text{C.35d})$$

References

- Aagaard, K., Coachman, L., & Carmack, E. (1981, jun). On the halocline of the Arctic Ocean. *Deep Sea Research Part A. Oceanographic Research Papers*, 28(6), 529–545. doi: 10.1016/0198-0149(81)90115-1
- Aagaard, K., & Greisman, P. (1975, sep). Toward new mass and heat budgets for the Arctic Ocean. *Journal of Geophysical Research*, 80(27), 3821–3827. doi: 10.1029/JC080i027p03821
- Aguilar, D. A., Sutherland, B. R., & Muraki, D. J. (2006). Laboratory generation of internal waves from sinusoidal topography. *Deep-Sea Research Part II: Topical Studies in Oceanography*, 53(1-2), 96–115. doi: 10.1016/j.dsr2.2005.09.015
- André, Q., Barker, A. J., & Mathis, S. (2017, sep). Layered semi-convection and tides in giant planet interiors. *Astronomy & Astrophysics*, 605, A117. doi: 10.1051/0004-6361/201730765
- Ascher, U. M., Ruuth, S. J., & Spiteri, R. J. (1997). Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations. *Applied Numerical Mathematics*, 25(2-3), 151–167. doi: 10.1016/S0168-9274(97)00056-1
- Bauer, E., Hunkins, K. L., Manley, T. O., & Tiemann, W. (1980). Salinity, Temperature and Depth Data Camp Snowbird. *Arctic Ice Dynamics Joint Experiment 1975-1976 Physical Oceanography Data Report*, 3(May). Retrieved from https://archive.org/details/DTIC_{_}ADA118204/page/n1/mode/2up
- Bebieva, Y., & Speer, K. (2019). The Regulation of Sea Ice Thickness by Double-Diffusive Processes in the Ross Gyre. *Journal of Geophysical Research: Oceans*, 124(10), 7068–7081. doi: 10.1029/2019JC015247
- Bebieva, Y., & Timmermans, M.-L. (2017). The relationship between double-diffusive intrusions and staircases in the Arctic Ocean. *Journal of Physical Oceanography*, 47(4), 867–878. doi: 10.1175/JPO-D-16-0265.1
- Bebieva, Y., & Timmermans, M. L. (2019). Double-Diffusive Layering in the Canada Basin: An Explanation of Along-Layer Temperature and Salinity Gradients. *Journal of Geophysical Research: Oceans*, 124(1), 723–735. doi: 10.1029/2018JC014368
- Bourget, B., Dauxois, T., Joubaud, S., & Odier, P. (2013). Experimental study of parametric subharmonic instability for internal plane waves. *Journal of Fluid Mechanics*, 723, 1–20. doi: 10.1017/jfm.2013.78

- Boury, S., Odier, P., & Peacock, T. (2019). Axisymmetric internal wave transmission and resonance in non-linear stratifications. *arXiv*, 1906.07844, 1–21.
- Boury, S., Supekar, R., Fine, E. C., Musgrave, R., Mickett, J. B., Voet, G., ... Alford, M. H. (2022, nov). Observations of Double Diffusive Staircase Edges in the Arctic Ocean. *Journal of Geophysical Research: Oceans*, 127(11), 1–14. doi: 10.1029/2022JC018906
- Boyd, J. P. (2001). *Chebyshev and Fourier Spectral Methods* (2nd ed.). Courier Corporation.
- Briggs, W. L., & Henson, V. E. (1995). *The DFT: An Owner's Manual for the Discrete Fourier Transform*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
- Brown, J. M., & Radko, T. (2022, oct). Disruption of Arctic Staircases by Shear. *Geophysical Research Letters*, 49(20). doi: 10.1029/2022GL100605
- Burns, K. J., Vasil, G. M., Oishi, J. S., Lecoanet, D., & Brown, B. P. (2020, apr). Dedalus: A flexible framework for numerical simulations with spectral methods. *Physical Review Research*, 2(2), 023068. doi: 10.1103/PhysRevResearch.2.023068
- Callies, J. (2018, sep). Restratification of Abyssal Mixing Layers by Submesoscale Baroclinic Eddies. *Journal of Physical Oceanography*, 48(9), 1995–2010. doi: 10.1175/JPO-D-18-0082.1
- Campello, R. J. G. B., Moulavi, D., & Sander, J. (2013). Density-Based Clustering Based on Hierarchical Density Estimates. *Lecture Notes in Computer Science*, 7819, 160–172. doi: 10.1007/978-3-642-37456-2_14
- Carmack, E. C. (2007). The alpha/beta ocean distinction: A perspective on freshwater fluxes, convection, nutrients and productivity in high-latitude seas. *Deep-Sea Research Part II: Topical Studies in Oceanography*, 54(23–26), 2578–2598. doi: 10.1016/j.dsr2.2007.08.018
- Carmack, E. C., Aagaard, K., Swift, J. H., Macdonald, R. W., McLaughlin, F. A., Jones, E. P., ... Killiush, L. R. (1997). Changes in temperature and tracer distributions within the Arctic Ocean: Results from the 1994 Arctic Ocean section. *Deep-Sea Research Part II: Topical Studies in Oceanography*, 44(8), 1487–1493. doi: 10.1016/S0967-0645(97)00056-8
- Chen, L. G. (1995, apr). Mixed Layer Density Ratio from the Levitus Data. *Journal of Physical Oceanography*, 25(4), 691–701. doi: 10.1175/1520-0485(1995)025<0691:MLDRFT>2.0.CO;2
- Cohen, J., Screen, J. A., Furtado, J. C., Barlow, M., Whittleston, D., Coumou, D., ... Jones, J. (2014). Recent Arctic amplification and extreme mid-latitude weather. *Nature Geoscience*, 7(9), 627–637. doi: 10.1038/ngeo2234
- Comiso, J. C., Parkinson, C. L., Gersten, R., & Stock, L. (2008). Accelerated decline in the Arctic sea ice cover. *Geophysical Research Letters*, 35(1), 1–6. doi: 10.1029/2007GL031972
- Croquette, V., & Williams, H. (1989, jul). Nonlinear waves of the oscillatory instability on finite convective rolls. *Physica D: Nonlinear Phenomena*, 37(1–3), 300–314. doi: 10.1016/0167-2789(89)90138-3
- Cushman-Roisin, B., & Beckers, J.-M. (2011). *Introduction to Geophysical Fluid Dynamics. Physical and Numerical Aspects*. Elsevier Ltd. doi: 10.1016/B978-0-12-088759-0.00022-5

- De Abreu, S., Cormier, R. M., Schee, M. G., Zemskova, V. E., Rosenblum, E., & Grisouard, N. (2024). Two-dimensional Numerical Simulations of Mixing under Ice Keels. *The Cryosphere*, 18(7), 3159–3176. doi: 10.5194/tc-18-3159-2024
- Desbruyères, D., Chafik, L., & Maze, G. (2021, feb). A shift in the ocean circulation has warmed the subpolar North Atlantic Ocean since 2016. *Communications Earth & Environment*, 2(1). doi: 10.1038/s43247-021-00120-y
- Dosser, H. V., Chanona, M., Waterman, S., Shibley, N. C., & Timmermans, M. (2021). Changes in internal wave-driven mixing across the Arctic Ocean: Finescale estimates from an 18-year pan-Arctic record. *Geophysical Research Letters*, 48(8), 1–10. doi: 10.1029/2020gl091747
- Dosser, H. V., & Rainville, L. (2016, feb). Dynamics of the Changing Near-Inertial Internal Wave Field in the Arctic Ocean. *Journal of Physical Oceanography*, 46(2), 395–415. doi: 10.1175/JPO-D-15-0056.1
- Dosser, H. V., Rainville, L., & Toole, J. M. (2014, feb). Near-Inertial Internal Wave Field in the Canada Basin from Ice-Tethered Profilers. *Journal of Physical Oceanography*, 44(2), 413–426. doi: 10.1175/JPO-D-13-0117.1
- Echeverri, P. (2009). *Internal tide generation by tall ocean ridges* (Doctoral dissertation, Woods Hole, MA). doi: 10.1575/1912/3144
- Eckart, C. (1961). Internal waves in the ocean. *Physics of Fluids*, 4(7), 791–799. doi: 10.1063/1.1706408
- Ekwurzel, B., Schlosser, P., Mortlock, R. A., Fairbanks, R. G., & Swift, J. H. (2001). River runoff, sea ice meltwater, and Pacific water distribution and mean residence times in the Arctic Ocean. *Journal of Geophysical Research: Oceans*, 106(C5), 9075–9092. doi: 10.1029/1999jc000024
- Espejo, A., Camus, P., Losada, I. J., & Méndez, F. J. (2014, aug). Spectral Ocean Wave Climate Variability Based on Atmospheric Circulation Patterns. *Journal of Physical Oceanography*, 44(8), 2139–2152. doi: 10.1175/JPO-D-13-0276.1
- Ester, M., Kriegel, H.-P., Sander, J., & Xu, X. (1996). A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Proceedings of the second international conference on knowledge discovery and data mining* (pp. 226–231). Elsevier.
- Fer, I. (2014). Near-Inertial Mixing in the Central Arctic Ocean. *Journal of Physical Oceanography*, 44(8), 2031–2049. doi: 10.1175/JPO-D-13-0133.1
- Fernando, H. J. (1989). Oceanographic implications of Laboratory Experiments on Diffusive Interfaces. *Journal of Physical*, 19, 1707–1715. doi: 10.1175/1520-0485(1989)019<1707:OIOLEO>2.0.CO;2
- Fokas, A. S., & Smitheman, S. A. (2012, nov). The Fourier Transforms of the Chebyshev and Legendre Polynomials. *arXiv e-prints*, 1211.4943v, 1–17.
- Foran, K. (2017). *Modelling Internal Wave Propagation through Layered Fluid Using Spectrally-Based DNS* (MSc Thesis). University of Toronto.

- Garcia-Soto, C., Cheng, L., Caesar, L., Schmidtko, S., Jewett, E. B., Cheripka, A., ... Abraham, J. P. (2021, sep). An Overview of Ocean Climate Change Indicators: Sea Surface Temperature, Ocean Heat Content, Ocean pH, Dissolved Oxygen Concentration, Arctic Sea Ice Extent, Thickness and Volume, Sea Level and Strength of the AMOC (Atlantic Meridional Overturning Circulation). *Frontiers in Marine Science*, 8(September). doi: 10.3389/fmars.2021.642372
- Ghaemsaidi, S. J., Dosser, H. V., Rainville, L., & Peacock, T. (2016, jan). The impact of multiple layering on internal wave transmission. *Journal of Fluid Mechanics*, 789, 617–629. doi: 10.1017/jfm.2015.682
- Gjelstrup, C. V., & Stedmon, C. A. (2024, may). A switch in thermal and haline contributions to stratification in the Greenland Sea during the last four decades. *Progress in Oceanography*. doi: 10.1016/j.pocean.2024.103283
- Grisouard, N., & Thomas, L. N. (2015). Critical and near-critical reflections of near-inertial waves off the sea surface at ocean fronts. *Journal of Fluid Mechanics*, 765, 273–302. doi: 10.1017/jfm.2014.725
- Gulev, S., Thorne, P., Ahn, J., Dentener, F., Domingues, C., Gerland, S., ... Vose, R. (2021). The Changing State of the Climate. In *Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change* (pp. 287–422). doi: 10.1017/9781009157896.004
- Guo, S., Zhou, S., Qu, L., & Lu, Y. (2016, oct). Laboratory experiments on diffusive convection layer thickness and its oceanographic implications. *Journal of Geophysical Research: Oceans*, 121(10), 7517–7529. doi: 10.1002/2016JC012172
- Guthrie, J. D., Fer, I., & Morison, J. (2015, dec). Observational validation of the diffusive convection flux laws in the Amundsen Basin, Arctic Ocean. *Journal of Geophysical Research: Oceans*, 120(12), 7880–7896. doi: 10.1002/2015JC010884
- Guthrie, J. D., Morison, J. H., & Fer, I. (2013). Revisiting internal waves and mixing in the Arctic Ocean. *Journal of Geophysical Research: Oceans*, 118(8), 3966–3977. doi: 10.1002/jgrc.20294
- Hassol, S. J. (2004). *Impacts of a Warming Arctic: Arctic Climate Impact Assessment* (Tech. Rep.).
- Hinneburg, A., & Keim, D. A. (2003, nov). A General Approach to Clustering in Large Databases with Noise. *Knowledge and Information Systems*, 5(4), 387–415. doi: 10.1007/s10115-003-0086-9
- Holmes, R. M., de Lavergne, C., & McDougall, T. J. (2019, oct). Tracer Transport within Abyssal Mixing Layers. *Journal of Physical Oceanography*, 49(10), 2669–2695. doi: 10.1175/JPO-D-19-0006.1
- Houghton, I. A., & Wilson, J. D. (2020). El Niño Detection Via Unsupervised Clustering of Argo Temperature Profiles. *Journal of Geophysical Research: Oceans*, 125(9), 1–12. doi: 10.1029/2019JC015947
- Huppert, H. E., & Linden, P. F. (1979). On heating a stable salinity gradient from below. *Journal of Fluid Mechanics*, 95(3), 431–464. doi: 10.1017/S0022112079001543

- Iglewicz, B., & Hoaglin, D. C. (1993). *How to Detect and Handle Outliers*. ASQ Quality Press. doi: 10.2307/1269377
- IOC. (2010). *The international thermodynamic equation of seawater – 2010 : Calculation and use of thermodynamic properties* (Vol. 56). UNESCO.
- Jablonowski, C., & Williamson, D. L. (2011). The Pros and Cons of Diffusion, Filters and Fixers in Atmospheric General Circulation Models. In T. D. Ringler (Ed.), *Numerical techniques for global atmospheric models* (Vol. 80, pp. 381–493). doi: 10.1007/978-3-642-11640-7_13
- Jones, D. C., Holt, H. J., Meijers, A. J., & Shuckburgh, E. (2019). Unsupervised Clustering of Southern Ocean Argo Float Temperature Profiles. *Journal of Geophysical Research: Oceans*, 124(1), 390–402. doi: 10.1029/2018JC014629
- Kelley, D. E. (1990). Fluxes through diffusive staircases: A new formulation. *Journal of Geophysical Research*, 95(C3), 3365. doi: 10.1029/jc095ic03p03365
- Kimura, S., Nicholls, K. W., & Venables, E. (2015, jan). Estimation of Ice Shelf Melt Rate in the Presence of a Thermohaline Staircase. *Journal of Physical Oceanography*, 45(1), 133–148. doi: 10.1175/JPO-D-14-0106.1
- Klymak, J. M., & Legg, S. M. (2010). A simple mixing scheme for models that resolve breaking internal waves. *Ocean Modelling*, 33(3-4), 224–234. doi: 10.1016/j.ocemod.2010.02.005
- Koszalka, I. M., & LaCasce, J. H. (2010, aug). Lagrangian analysis by clustering. *Ocean Dynamics*, 60(4), 957–972. doi: 10.1007/s10236-010-0306-2
- Krishfield, R., Toole, J., & Timmermans, M. (2008). ITP Data Processing Procedures. *Woods Hole Oceanographic Institution*(March), 1–24. Retrieved from <http://www.whoi.edu/filesserver.do?id=35803&pt=2&p=41486>
- Kundu, P. K., Cohen, I. M., & Dowling, D. R. (2015). *Fluid Mechanics* (6th ed.). Academic Press.
- Lautrup, B. (2011). *Physics of Continuous Matter: Exotic and Everyday Phenomena in the Macroscopic World* (2nd ed.). CRC Press.
- Li, X., Wang, Q., Danilov, S., Koldunov, N., Liu, C., Müller, V., ... Jung, T. (2024, feb). Eddy activity in the Arctic Ocean projected to surge in a warming world. *Nature Climate Change*, 14(2), 156–162. doi: 10.1038/s41558-023-01908-w
- Lincoln, B. J., Rippeth, T. P., Lenn, Y.-D., Timmermans, M. L., Williams, W. J., & Bacon, S. (2016, sep). Wind-driven mixing at intermediate depths in an ice-free Arctic Ocean. *Geophysical Research Letters*, 43(18), 9749–9756. doi: 10.1002/2016GL070454
- Linden, P. F., & Shirtcliffe, T. G. L. (1978, aug). The diffusive interface in double-diffusive convection. *Journal of Fluid Mechanics*, 87(03), 417–432. doi: 10.1017/S002211207800169X
- Logan, C. H. A., & Fotopoulou, S. (2020, jan). Unsupervised star, galaxy, QSO classification. *Astronomy & Astrophysics*, 633, A154. doi: 10.1051/0004-6361/201936648

- Lu, Y.-Z., Guo, S.-X., Zhou, S.-Q., Song, X.-L., & Huang, P.-Q. (2022, nov). Identification of Thermohaline Sheet and Its Spatial Structure in the Canada Basin. *Journal of Physical Oceanography*, 52(11), 2773–2787. doi: 10.1175/JPO-D-22-0012.1
- Ma, C., Li, S., Yang, Y., Yang, J., & Chen, G. (2019, sep). Extraction of Revolving Channels of Drifters around Mesoscale Eddy Centers Based on Spatiotemporal Trajectory Clustering. *Journal of Atmospheric and Oceanic Technology*, 36(9), 1903–1916. doi: 10.1175/JTECH-D-19-0007.1
- Ma, Y., & Peltier, W. (2022). Thermohaline staircase formation in the diffusive convection regime: a theory based upon stratified turbulence asymptotics. *Journal of Fluid Mechanics*, 931, 1–13. doi: 10.1017/jfm.2021.945
- MacKinnon, J. A., Simmons, H. L., Hargrove, J., Thomson, J., Peacock, T., Alford, M. H., ... Wood, K. R. (2021, dec). A warm jet in a cold ocean. *Nature Communications*, 12(1), 2418. doi: 10.1038/s41467-021-22505-5
- Maykut, G. A., & Untersteiner, N. (1971). Some results from a time-dependent thermodynamic model of sea ice. *J Geophys Res*, 76(6), 1550–1575. doi: 10.1029/jc076i006p01550
- Maze, G., Mercier, H., Fablet, R., Tandeo, P., Lopez Radcenco, M., Lenca, P., ... Le Goff, C. (2017). Coherent heat patterns revealed by unsupervised classification of Argo temperature profiles in the North Atlantic Ocean. *Progress in Oceanography*, 151, 275–292. doi: 10.1016/j.pocean.2016.12.008
- McDougall, T. J. (1987). Neutral Surfaces. *Journal of Physical Oceanography*, 17(11), 1950–1964. doi: 10.1175/1520-0485(1987)017<1950:ns>2.0.co;2
- McDougall, T. J., & Barker, P. M. (2011). *Getting started with TEOS-10 and the Gibbs Seawater (GSW) Oceanographic Toolbox* (3.06.12 ed.). Retrieved from www.TEOS-10.org
- McInnes, L., Healy, J., & Astels, S. (2017). hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, 2(11), 205. doi: 10.21105/joss.00205
- Mcphee, M. G. (2012). Advances in understanding ice – ocean stress during and since AIDJEX. *Cold Regions Science and Technology*, 76-77, 24–36. doi: 10.1016/j.coldregions.2011.05.001
- Meneghello, G., Marshall, J., Timmermans, M.-L., & Scott, J. (2018, apr). Observations of Seasonal Upwelling and Downwelling in the Beaufort Sea Mediated by Sea Ice. *Journal of Physical Oceanography*, 48(4), 795–805. doi: 10.1175/JPO-D-17-0188.1
- Ménèsguen, C., Lique, C., & Caspar-Cohen, Z. (2022, nov). Density Staircases Are Disappearing in the Canada Basin of the Arctic Ocean. *Journal of Geophysical Research: Oceans*, 127(11). doi: 10.1029/2022JC018877
- Mercier, M. J., Garnier, N. B., & Dauxois, T. (2008, aug). Reflection and diffraction of internal waves analyzed with the Hilbert transform. *Physics of Fluids*, 20(8). doi: 10.1063/1.2963136
- Millero, F. J. (2010). History of the equation of state of seawater. *Oceanography*, 23(3), 18–33. doi: 10.5670/oceanog.2010.21

- Morgan, P. P. (1994). *SEAWATER: A Library of MATLAB® Computational Routines for the Properties of Sea Water: Version 1.2* (Tech. Rep.). Hobart, Tasmania: CSIRO Marine Laboratories. Retrieved from <http://hdl.handle.net/102.100.100/239771?index=1>
- Morozov, E. G., & Paka, V. T. (2010). Internal waves in a high-latitude region. *Oceanology*, 50(5), 668–674. doi: 10.1134/S0001437010050048
- Moulavi, D., Jaskowiak, P. A., Campello, R. J. G. B., Zimek, A., & Sander, J. (2014, apr). Density-Based Clustering Validation. In *Proceedings of the 2014 siam international conference on data mining* (Vol. 2, pp. 839–847). Philadelphia, PA: Society for Industrial and Applied Mathematics. doi: 10.1137/1.9781611973440.96
- Nash, J. D., Shroyer, E. L., Kelly, S. M., Inall, M. E., Duda, T. F., Levine, M. D., ... Musgrave, R. C. (2012). Are any coastal internal tides predictable? *Oceanography*, 25(2), 80–95. doi: 10.5670.oceanog.2012.44
- Neal, V. T., & Neshyba, S. (1973, may). Microstructure Anomalies in the Arctic Ocean. *Journal of Geophysical Research*, 78(15), 2695–2701. doi: 10.1029/JC078i015p02695
- Neal, V. T., Neshyba, S., & Denner, W. (1969, oct). Thermal Stratification in the Arctic Ocean. *Science*, 166(3903), 373–374. doi: 10.1126/science.166.3903.373
- Neshyba, S., Neal, V. T., & Denner, W. (1971, nov). Temperature and Conductivity Measurements under Ice Island T-3. *Journal of Geophysical Research*, 76(33), 8107–8120. doi: 10.1029/JC076i033p08107
- Neshyba, S., Neal, V. T., & Denner, W. W. (1972, jan). Spectra of Internal Waves: In-Situ Measurements in a Multiple-Layered Structure. *Journal of Physical Oceanography*, 2(1), 91–95. doi: 10.1175/1520-0485(1972)002<0091:SOIWSM>2.0.CO;2
- Newman, F. C. (1976). Temperature Steps in Lake Kivu: A Bottom Heated Saline Lake. *Journal of Physical Oceanography*, 6(2), 157–163. doi: 10.1175/1520-0485(1976)006<0157:TSILKA>2.0.CO;2
- NOAA. (2022). *Monthly Global Climate Report for Annual 2021*.
- Padman, L., & Dillon, T. M. (1987). Vertical Heat Fluxes Through the Beaufort Sea Thermohaline Staircase. *Journal of Geophysical Research*, 92(C10), 10799. doi: 10.1029/JC092iC10p10799
- Padman, L., & Dillon, T. M. (1988, oct). On the Horizontal Extent of the Canada Basin Thermohaline Steps. *Journal of Physical Oceanography*, 18(10), 1458–1462. doi: 10.1175/1520-0485(1988)018<1458:OTHEOT>2.0.CO;2
- Padman, L., & Dillon, T. M. (1989). Thermal microstructure and internal waves in the Canada Basin diffusive staircase. *Deep Sea Research Part A, Oceanographic Research Papers*, 36(4), 531–542. doi: 10.1016/0198-0149(89)90004-6
- Peralta-Ferriz, C., & Woodgate, R. A. (2015). Seasonal and interannual variability of pan-Arctic surface mixed layer properties from 1979 to 2012 from hydrographic data, and the dominance of stratification for multiyear mixed layer depth shoaling. *Progress in Oceanography*, 134, 19–53. doi: 10.1016/j.pocean.2014.12.005

- Polyakov, I. V., Pnyushkov, A. V., Alkire, M. B., Ashik, I. M., Baumann, T. M., Carmack, E. C., ... Yulin, A. (2017). Greater role for Atlantic inflows on sea-ice loss in the Eurasian Basin of the Arctic Ocean. *Science*, 356(6335), 285–291. doi: 10.1126/science.aai8204
- Polyakov, I. V., Pnyushkov, A. V., Rembe, R., Ivanov, V. V., Lenn, Y. D., Padman, L., & Carmack, E. C. (2012). Mooring-based observations of double-diffusive staircases over the laptev sea slope. *Journal of Physical Oceanography*, 42(1), 95–109. doi: 10.1175/2011JPO4606.1
- Polyakov, I. V., Pnyushkov, A. V., & Timokhov, L. A. (2012, dec). Warming of the Intermediate Atlantic Water of the Arctic Ocean in the 2000s. *Journal of Climate*, 25(23), 8362–8370. doi: 10.1175/JCLI-D-12-00266.1
- Polyakov, I. V., Walsh, J. E., & Kwok, R. (2012). Recent changes of Arctic multiyear sea ice coverage and the likely causes. *Bulletin of the American Meteorological Society*, 93(2), 145–151. doi: 10.1175/BAMS-D-11-00070.1
- Pontin, C. M., Barker, A. J., Hollerbach, R., André, Q., & Mathis, S. (2021). Wave propagation in semiconvective regions of giant planets. *Monthly Notices of the Royal Astronomical Society*, 493(4), 5788–5806. doi: 10.1093/mnras/staa664
- Proshutinsky, A., Krishfield, R., Toole, J. M., Timmermans, M. L., Williams, W., Zimmermann, S., ... Zhao, J. (2019). Analysis of the Beaufort Gyre Freshwater Content in 2003–2018. *Journal of Geophysical Research: Oceans*, 124(12), 9658–9689. doi: 10.1029/2019JC015281
- Radko, T. (2013). *Double-Diffusive Convection*. Cambridge University Press. Retrieved from <https://doi.org/10.1017/CB09781139034173>
- Radko, T., Bulters, A., Flanagan, J. D., & Campin, J. M. (2014). Double-diffusive recipes. Part I: Large-scale dynamics of thermohaline staircases. *Journal of Physical Oceanography*, 44(5), 1269–1284. doi: 10.1175/JPO-D-13-0155.1
- Rainville, L., Lee, C., & Woodgate, R. (2011, sep). Impact of Wind-Driven Mixing in the Arctic Ocean. *Oceanography*, 24(3), 136–145. doi: 10.5670/oceanog.2011.65
- Rainville, L., & Winsor, P. (2008). Mixing across the Arctic Ocean: Microstructure observations during the Beringia 2005 Expedition. *Geophysical Research Letters*, 35(8), 2–6. doi: 10.1029/2008GL033532
- Rantanen, M., Karpechko, A. Y., Lipponen, A., Nordling, K., Hyvärinen, O., Ruosteenoja, K., ... Laaksonen, A. (2022). The Arctic has warmed nearly four times faster than the globe since 1979. *Communications Earth and Environment*, 3(1), 1–10. doi: 10.1038/s43247-022-00498-3
- Richards, A. E., Johnson, H. L., & Lique, C. (2022, sep). Spatial and Temporal Variability of Atlantic Water in the Arctic From 40 Years of Observations. *Journal of Geophysical Research: Oceans*, 127(9). doi: 10.1029/2021JC018358
- Rosso, I., Mazloff, M. R., Talley, L. D., Purkey, S. G., Freeman, N. M., & Maze, G. (2020). Water Mass and Biogeochemical Variability in the Kerguelen Sector of the Southern Ocean: A Machine Learning Approach for a Mixing Hot Spot. *Journal of Geophysical Research: Oceans*, 125(3), 1–23. doi: 10.1029/2019JC015877

- Rousseeuw, P. J., & Croux, C. (1993, dec). Alternatives to the Median Absolute Deviation. *Journal of the American Statistical Association*, 88(424), 1273–1283. doi: 10.1080/01621459.1993.10476408
- Rousseeuw, P. J., & Leroy, A. M. (1987). *Robust Regression and Outlier Detection*. (Vol. 152) (No. 1). John Wiley and Sons Inc. doi: 10.2307/2982847
- Schee, M. G., Rosenblum, E., Lilly, J. M., & Grisouard, N. (2024, may). Unsupervised clustering identifies thermohaline staircases in the Canada Basin of the Arctic Ocean. *Environmental Data Science*, 3, e13. doi: 10.1017/eds.2024.13
- Schmitt, R., Perkins, H., Boyd, J., & Stalcup, M. (1987, oct). C-SALT: an investigation of the thermohaline staircase in the western tropical North Atlantic. *Deep Sea Research Part A. Oceanographic Research Papers*, 34(10), 1655–1665. doi: 10.1016/0198-0149(87)90014-8
- Shaw, W. J., & Stanton, T. P. (2014, aug). Vertical diffusivity of the Western Arctic Ocean halocline. *Journal of Geophysical Research: Oceans*, 119(8), 5017–5038. doi: 10.1002/2013JC009598
- Shaw, W. J., Stanton, T. P., McPhee, M. G., Morison, J. H., & Martinson, D. G. (2009). Role of the upper ocean in the energy budget of Arctic sea ice during SHEBA. *Journal of Geophysical Research: Oceans*, 114(6), 1–21. doi: 10.1029/2008JC004991
- Shibley, N. C., & Timmermans, M. (2022, jul). The Beaufort Gyre’s Diffusive Staircase: Finescale Signatures of Gyre-Scale Transport. *Geophysical Research Letters*, 49(13). doi: 10.1029/2022GL098621
- Shibley, N. C., & Timmermans, M.-L. (2019, mar). The Formation of Double-Diffusive Layers in a Weakly Turbulent Environment. *Journal of Geophysical Research: Oceans*. doi: 10.1029/2018JC014625
- Shibley, N. C., Timmermans, M.-L., Carpenter, J. R., & Toole, J. M. (2017, feb). Spatial variability of the Arctic Ocean’s double-diffusive staircase. *Journal of Geophysical Research: Oceans*, 122(2), 980–994. doi: 10.1002/2016JC012419
- Shibley, N. C., Timmermans, M.-L., & Stranne, C. (2020). Analysis of Acoustic Observations of Double-Diffusive Finestructure in the Arctic Ocean. *Geophysical Research Letters*, 47(18), 1–11. doi: 10.1029/2020GL089845
- Shu, Q., Wang, Q., Årthun, M., Wang, S., Song, Z., Zhang, M., & Qiao, F. (2022, jul). Arctic Ocean Amplification in a warming climate in CMIP6 models. *Science Advances*, 8(30), 1–10. doi: 10.1126/sciadv.abn9755
- Simmons, H. L., Hallberg, R. W., & Arbic, B. K. (2004, dec). Internal wave generation in a global baroclinic tide model. *Deep Sea Research Part II: Topical Studies in Oceanography*, 51(25-26), 3043–3068. doi: 10.1016/j.dsr2.2004.09.015
- Sledd, A., & L’Ecuyer, T. S. (2021). Emerging Trends in Arctic Solar Absorption. *Geophysical Research Letters*, 48(24), 1–9. doi: 10.1029/2021gl095813

- Sonneveld, M., Lguensat, R., Jones, D. C., Dueben, P. D., Brajard, J., & Balaji, V. (2021, jul). Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environmental Research Letters*, *16*(7), 073008. doi: 10.1088/1748-9326/ac0eb0
- Steele, M., & Boyd, T. (1998, may). Retreat of the cold halocline layer in the Arctic Ocean. *Journal of Geophysical Research: Oceans*, *103*(C5), 10419–10435. doi: 10.1029/98JC00580
- Stewart, K. D., & Haine, T. W. (2016). Thermobaricity in the transition zones between alpha and beta oceans. *Journal of Physical Oceanography*, *46*(6), 1805–1821. doi: 10.1175/JPO-D-16-0017.1
- Stranne, C., Mayer, L., Weber, T. C., Ruddick, B. R., Jakobsson, M., Jerram, K., ... Gårdfeldt, K. (2017, nov). Acoustic Mapping of Thermohaline Staircases in the Arctic Ocean. *Scientific Reports*, *7*(1), 15192. doi: 10.1038/s41598-017-15486-3
- Su, H., Jiang, J., Wang, A., Zhuang, W., & Yan, X.-H. (2022, jul). Subsurface Temperature Reconstruction for the Global Ocean from 1993 to 2020 Using Satellite Observations and Deep Learning. *Remote Sensing*, *14*(13), 3198. doi: 10.3390/rs14133198
- Supekar, R., & Peacock, T. (2019). Interference and transmission of spatiotemporally locally forced internal waves in non-uniform stratifications. *Journal of Fluid Mechanics*, *866*, 350–368. doi: 10.1017/jfm.2019.106
- Sutherland, B. R. (2010). *Internal gravity waves*. Cambridge University Press.
- Sutherland, B. R. (2016). Internal wave transmission through a thermohaline staircase. *Physical Review Fluids*, *1*(1), 013701. doi: 10.1103/PhysRevFluids.1.013701
- Sutherland, B. R., & Yewchuk, K. (2004, jul). Internal wave tunnelling. *Journal of Fluid Mechanics*, *511*, 125–134. doi: 10.1017/S0022112004009863
- Timmermans, M.-L., Garrett, C., & Carmack, E. (2003). The thermohaline structure and evolution of the deep waters in the Canada Basin, Arctic Ocean. *Deep Sea Research Part I: Oceanographic Research Papers*, *50*, 1305–1321. doi: 10.1016/S0967-0637(03)00125-0
- Timmermans, M.-L., & Marshall, J. (2020). Understanding Arctic Ocean Circulation: A Review of Ocean Dynamics in a Changing Climate. *Journal of Geophysical Research: Oceans*, *125*(4), 1–35. doi: 10.1029/2018JC014378
- Timmermans, M.-L., & Pickart, R. S. (2023, jul). The Arctic Ocean’s Changing Beaufort Gyre System: An Assessment of Current Understanding, Open Questions and Future Research Directions. *Bulletin of the American Meteorological Society*, *104*(7), E1282–E1289. doi: 10.1175/BAMS-D-23-0129.1
- Timmermans, M.-L., Toole, J., & Krishfield, R. (2018). Warming of the interior Arctic Ocean linked to sea ice losses at the basin margins. *Science Advances*, *4*(8), 1–7. doi: 10.1126/sciadv.aat6773
- Timmermans, M.-L., Toole, J., Krishfield, R., & Winsor, P. (2008). Ice-Tethered Profiler observations of the double-diffusive staircase in the Canada Basin thermocline. *Journal of Geophysical Research*, *113*, 1–10. doi: 10.1029/2008jc004829

- Timmermans, M.-L., & Toole, J. M. (2023, jan). The Arctic Ocean’s Beaufort Gyre. *Annual Review of Marine Science*, 15(1), 223–248. doi: 10.1146/annurev-marine-032122-012034
- Toole, J. M., Krishfield, R. A., O’Brien, J. K., & Cole, S. T. (2016). *Ice-Tethered Profiler observations: Vertical profiles of temperature, salinity, oxygen, and ocean velocity from an Ice-Tethered buoy system*. NOAA National Centers for Environmental Information. doi: 10.7289/v5mw2f7x
- Toole, J. M., Krishfield, R. A., Timmermans, M. L., & Proshutinsky, A. (2011). The Ice-Tethered profiler: Argo of the Arctic. *Oceanography*, 24(3), 126–135. doi: 10.5670/oceanog.2011.64
- Turner, J. S. (1965). The coupled turbulent transports of salt and and heat across a sharp density interface. *International Journal of Heat and Mass Transfer*, 8(5), 759–767. doi: 10.1016/0017-9310(65)90022-0
- Turner, J. S. (2010, jan). The Melting of Ice in the Arctic Ocean: The Influence of Double-Diffusive Transport of Heat from Below. *Journal of Physical Oceanography*, 40(1), 249–256. doi: 10.1175/2009JPO4279.1
- Vallis, G. K. (2017). *Atmospheric and oceanic fluid dynamics: Fundamentals and large-scale circulation, second edition* (2nd ed.). Cambridge University Press. doi: 10.1017/9781107588417
- van der Boog, C. G., Dijkstra, H. A., Pietrzak, J. D., & Katsman, C. A. (2021, feb). Double-diffusive mixing makes a small contribution to the global ocean circulation. *Communications Earth & Environment*, 2(1), 46. doi: 10.1038/s43247-021-00113-x
- van der Boog, C. G., Otto Koetsier, J., Dijkstra, H. A., Pietrzak, J. D., & Katsman, C. A. (2021). Global dataset of thermohaline staircases obtained from Argo floats and Ice-Tethered Profilers. *Earth System Science Data*, 13(1), 43–61. doi: 10.5194/essd-13-43-2021
- Wells, M. G., & Wettlaufer, J. S. (2007). The long-term circulation driven by density currents in a two-layer stratified basin. *Journal of Fluid Mechanics*, 572, 37–58. doi: 10.1017/S0022112006003478
- Wenegrat, J. O., Callies, J., & Thomas, L. N. (2018, nov). Submesoscale Baroclinic Instability in the Bottom Boundary Layer. *Journal of Physical Oceanography*, 48(11), 2571–2592. doi: 10.1175/JPO-D-17-0264.1
- Winters, K. B., MacKinnon, J. A., & Mills, B. (2004, jan). A Spectral Model for Process Studies of Rotating, Density-Stratified Flows. *Journal of Atmospheric and Oceanic Technology*, 21(1), 69–94. doi: 10.1175/1520-0426(2004)021<0069:ASMFPS>2.0.CO;2
- Winton, M. (2011, aug). Do Climate Models Underestimate the Sensitivity of Northern Hemisphere Sea Ice Cover? *Journal of Climate*, 24(15), 3924–3934. doi: 10.1175/2011JCLI4146.1
- Wunsch, S. (2018, nov). Nonlinear harmonic generation by internal waves in a density staircase. *Physical Review Fluids*, 3(11), 114803. doi: 10.1103/PhysRevFluids.3.114803
- Xue, N., Khodaparast, S., Zhu, L., Nunes, J. K., Kim, H., & Stone, H. A. (2017). Laboratory layered latte. *Nature Communications*, 8(1), 1–6. doi: 10.1038/s41467-017-01852-2

- Yang, Y., Verzicco, R., Lohse, D., & Caulfield, C. (2022, feb). Layering and vertical transport in sheared double-diffusive convection in the diffusive regime. *Journal of Fluid Mechanics*, 933, A30. doi: 10.1017/jfm.2021.1091
- Yu, L. S., Bosse, A., Fer, I., Orvik, K. A., Bruvik, E. M., Hessequik, I., & Kvalsund, K. (2017). The Lofoten Basin eddy: Three years of evolution as observed by Seagliders. *Journal of Geophysical Research: Oceans*, 122(8), 6814–6834. doi: 10.1002/2017JC012982