Defining Southern Ocean fronts using unsupervised classification

Simon D. A. Thomas[1,2], Daniel C. Jones[1], Anita Faul[1], Erik Mackie[3,1], and Etienne Pauthenet[4]

[1]British Antarctic Survey, NERC, UKRI, UK
[2]Department of Physics, University of Cambridge, UK
[3]Cambridge Zero, University of Cambridge, UK
[4]Sorbonne Universités, UPMC Université, LOCEAN-IPSL, Paris, France

Correspondence: Simon Thomas (sithom@bas.ac.uk)

Abstract. Oceanographic fronts are transitions between thermohaline structures with different characteristics. Such transitions are ubiquitous, and their locations and properties affect how the ocean operates as part of the global climate system. In the Southern Ocean, fronts have classically been defined using a small number of continuous, circumpolar features in sea surface height or dynamic height. Modern observational and theoretical developments are challenging and expanding this traditional framework to accommodate a more complex view of fronts. Here we present a complementary new approach for calculating fronts using an unsupervised classification method called Gaussian mixture modelling (GMM) and a novel inter-class parameter called the *I*-metric. The *I*-metric approach produces a probabilistic view of front location, emphasising the fact that the boundaries between water masses are not uniformly sharp across the entire Southern Ocean. The *I*-metric approach uses thermohaline information from a range of depth levels, making it more general than approaches that only use near-surface properties. We train the GMM using an observationally-constrained state estimate in order to have more uniform spatial and temporal data coverage. The probabilistic boundaries defined by the *I*-metric roughly coincide with several classically-defined fronts, offering a novel view of this structure. The *I*-metric fronts appear to be relatively sharp in the open ocean and somewhat diffuse near large topographic features, possibly highlighting the importance of topographically-induced mixing. For comparison with a more localised method, we also use an edge detection approach for identifying fronts. We find a strong correlation between the edge field of the leading principal component and the zonal velocity; the edge detection method highlights the presence of jets, which are supported by thermal wind balance. This more localised method highlights the complex, multi-scale structure of Southern Ocean fronts, complementing and contrasting with the more domain-wide view offered by the *I*-metric. The Sobel edge detection method may be useful for defining and tracking smaller-scale fronts and jets in model or reanalysis data. The *I*-metric approach may prove to be a useful method for inter-model comparison, as it uses the thermohaline structure of those models instead of tracking somewhat ad-hoc values of sea surface height and/or dynamic height, which can vary considerably between models. In addition, the general *I*-metric approach allows front definitions to shift with changing temperature and salinity structures, which may be useful for characterising fronts in a changing climate.

1 Introduction

25 The Southern Ocean (SO) is at the centre of the global thermohaline circulation, joining the Indian, Pacific, and Atlantic oceans into a single planetary-scale heat and carbon transport system (Marshall and Speer, 2012; Talley, 2013). In the SO, upwelling and downwelling branches of the overturning circulation transport water and tracers (e.g. heat, carbon) between the surface and subsurface oceans (Sallee et al., 2010, 2012). The steeply tilted isopycnals associated with the overturning circulation also support the powerful Antarctic Circumpolar Current (ACC), with a mean combined barotropic and baroclinic volume

30 transport of roughly 173.3 ± 10.7 Sv, driven by a combination of the westerly winds and air-sea buoyancy forcing (Rintoul et al., 2001; Morrison et al., 2015; Donohue et al., 2016). In part because of its unique structure, the SO is a critical regulator of global climate, having thus far absorbed more than 75% of the excess energy and 50% of the excess carbon added to the climate system from anthropogenic emissions (Mikaloff-Fletcher et al., 2006; Frolicher et al., 2015). As such, the thermohaline structure of the Southern Ocean may be considered an important climate system parameter, as it affects how heat and carbon

35 are partitioned between the atmosphere and ocean.

Through decades of observational and theoretical effort, the global oceanographic community has curated a detailed theoretical understanding of the structure of the Southern Ocean. One of the hallmarks of this view is the presence of fronts, i.e. transitions in temperature, salinity, and/or biogeochemical properties (Deacon, 1937; Orsi et al., 1995). Although fronts are not identical to the sharp jets found in the SO, fronts and jets at the mesoscale share a close relationship partly due to thermal

40 wind balance (Sokolov and Rintoul, 2002, 2009). Traditionally, oceanographers have defined SO fronts using a small number of continuous, circumpolar features that follow contours of sea surface height or dynamic height (Kim and Orsi, 2014). However, satellite altimetry shows that the ACC features a braided and meandering structure that is not necessarily reflected in the traditional, time-averaged view of fronts as continuous property contours (Chapman, 2017; Mackie, 2018). Using individual property contours to define fronts, for example, contours of temperature or sea surface height, is somewhat limited by the

45 fact that such contours do not always line up with the locations of strong gradients (Thompson et al., 2010; Thompson and Sallée, 2012; Graham et al., 2012; Chapman, 2017). In response to more detailed SO observations, the global oceanographic community has been developing a variety of new approaches for defining and tracking fronts in more application-specific ways (Chapman et al., 2020). For example, coastal applications and open ocean applications may benefit from conceptually different treatments of ocean fronts, which are characterised by different spatial and temporal scales. For a historical view and summary

50 of advances in the area of front definition and detection, see the recent review article by Chapman et al. (2020).

In order to help us broaden our view of Southern Ocean fronts, we look to a branch of machine learning called unsupervised classification (also known as clustering). Broadly speaking, unsupervised classification attempts to identify sub-populations in data distributions that have not already been labelled or sorted. Although such methods have existed for decades, the amount of SO data has only in recent years become large enough for clustering approaches to be suitable; the application of unsupervised

55 classification to oceanographic data is in its infancy. Several recent studies have used unsupervised classification to identify coherent regimes of thermohaline structure and the transitions between them, specifically in the North Atlantic (Maze et al., 2017), Southern Ocean (Jones et al., 2019), and Indian Sector of the Southern Ocean (Rosso et al., 2020). These methods have

also been used to define coherent dynamical and biogeochemical regimes from depth-averaged ocean structure (Sonnewald et al., 2019; Le Bras et al., 2019; Jones and Ito, 2019). Recently, unsupervised classification has been used to define coherent ecological regimes from physical and biogeochemical data (Sonnewald et al., 2020). Researchers are also exploring potential connections between changes in class properties and large-scale climate phenomena. For example, a recent study tied evolution in the longitudinal extent of an algorithmically-defined class to the onset of El Niño, suggesting that unsupervised classification methods could complement existing index-based assessments of large-scale climate modes (Houghton and Wilson, 2020).

Unsupervised classification does not use specific property contours to define boundaries between thermohaline structures, so it avoids one of the fundamental limitations of many traditional front definition approaches. Given the required information, unsupervised classification methods can use more detailed thermohaline data from throughout the water column to define classes and their boundaries. Across a given front, one might expect to find not only a transition in surface values but also a change in the thermohaline structure, as indicated by a change of profile class with latitude and/or longitude. In this work, we use an unsupervised classification technique called Gaussian mixture modelling (GMM), which attempts to represent sub-populations in the data distribution using multi-dimensional Gaussian functions. Because GMM is a probabilistic method, in addition to automatically clustering the thermohaline profiles into classes, it returns for each data point a set of weights across the different classes. That is, it returns a probability distribution that can be exploited to define boundaries between coherent regimes in a novel way. In this paper, we propose that GMM can be used to represent the boundaries as "fuzzy" regions, which reflects the fact that not all transitions in the SO are uniformly sharp.

In Sect. 2, we introduce the observationally-constrained state estimate from which we draw our temperature and salinity data (Sect. 2.1), discuss principal component analysis (PCA) for dimensionality reduction (Sect. 2.2), and cover our application of GMM (Sect. 2.3). We then define the inter-class comparison metric (i.e. the *I*-metric) that we use to quantify water mass boundaries (Sect. 3.1). Next, we apply the *I*-metric to the reduced-dimension state estimate data (Sect. 3.2). For comparison, we contrast this method with a more local front-detection approach (Sect. 3.4). Finally, we discuss some caveats (Sect. 4) and offer our summary and conclusions (Sect. 5).

2   State estimate data, PCA, and unsupervised classification

Our front identification method uses a combination of principal component analysis, unsupervised classification, and a new probabilistic metric to quantify the boundaries between coherent thermohaline structures. First, we describe the dataset that we used for developing and training our method.

2.1   The Southern Ocean State Estimate

We developed our method using the Biogeochemical Southern Ocean State Estimate (B-SOSE) (Verdy and Mazloff, 2017). B-SOSE is an observationally-constrained numerical simulation created using MITgcm [mitgcm.org] (Marshall et al., 1997a, b) and a suite of Southern Ocean observations, including Argo float data, ship track data, and satellite data. B-SOSE is part of the ECCO suite of state estimates [ecco-group.org/], which includes a variety of global and regional products covering

a range of multi-year to multi-decadal time periods. Examples of other state estimates include the physics-only Southern Ocean State Estimate (SOSE) and the global ECCOv4 state estimate (Mazloff et al., 2010; Forget et al., 2015). We chose to develop our method using a state estimate because such products offer (1) uniform coverage in latitude, longitude, and time and (2) relatively high fidelity with respect to observations. We chose B-SOSE, in particular, because it represents the Southern Ocean using a spatial resolution of $1/6°$, which is eddy-permitting in the latitude range of the ACC, enabling a realistic representation of mesoscale eddy structure (Hallberg, 2013a). We expect that training our model on the physics-only SOSE would produce similar results, although we did not attempt that here. In principle, our methods can be readily applied to any gridded temperature and salinity profile dataset. It may be possible to apply these methods to in-situ data as well, if the user addresses the problem of non-uniform spatial and temporal sampling. In this paper, we focus only on applications to gridded datasets.

To construct a state estimate, researchers bring a numerical simulation into better consistency with an observational dataset using the 4DVAR method. This method uses adjoint sensitivities to calculate the required changes in the "controls" (e.g. initial conditions, mixing parameters, boundary conditions) needed to improve the agreement between the simulation and the observational dataset (Stammer et al., 2002; Wunsch and Heimbach, 2007).

The B-SOSE domain extends from the equator to $78°S$, but we only use data south of $30°S$ to focus on the Southern Ocean and to avoid the model boundary. It uses bathymetry and coastline based on Amante and Eakins (2009). B-SOSE solves the heat, salt, and momentum equations using a third-order direct space and time advection scheme with a 1-hour timestep. The time-evolving atmospheric boundary conditions use bulk formulae to solve for fluxes of heat, freshwater, and momentum, with six-hourly atmospheric state variables as inputs (Large and Yeager, 2009; Dee et al., 2011). The state estimation process iteratively adjusts the atmospheric state variables and oceanic initial conditions to improve model-data agreement. B-SOSE uses dynamic sea ice (Losch et al., 2010; Fenty and Heimbach, 2013). For vertical mixing, it uses the GLL90 mixed layer parameterization (Gaspar et al., 1990). It also uses horizontal and vertical viscosity and diffusivity. River runoff comes from the product of Dai and Trenberth (2002) augmented with an estimate of Antarctic freshwater input from iceberg and ice sheet melting (Hammond and Jones, 2016). It does not include mesoscale eddy parameterization, as this particular configuration falls into the horizontal resolution range wherein mesoscale parameterization may actually worsen the representation of the mesoscale (Hallberg, 2013b). Because we are interested in quantifying physical, large-scale fronts, we only used monthly mean temperature and salinity data. Also, because we are not interested in the surface seasonal cycle at present, we only used temperature and salinity data between 300-2000 m, following Rosso et al. (2020). We used the whole period of iteration 106 of this state estimate, which covers January 2008 to December 2012. Some key properties of B-SOSE iteration 106 are listed in Table 1. For further details, see Verdy and Mazloff (2017).

2.2 Principal component analysis

Each vertical profile in the full B-SOSE dataset is comprised of temperature and salinity values at multiple depth levels, at every grid cell and every output month. Values close to each other in the water column are correlated to some degree. Therefore, we do not necessarily need values of T and S at every depth level to capture most of the variability, and reducing the dimensionality

Property	Value
State estimate iteration number	106
Horizontal resolution	$1/6°$
Vertical resolution (variable)	4.2 m to 400 m
Number of vertical levels	52
Output frequency	Monthly averaged
Horizontal viscosity	$10 \, \mathrm{m^2 s^{-1}}$
Vertical viscosity	$10^{-3} \, \mathrm{m^2 s^{-1}}$
Horizontal diffusivity	$10 \, \mathrm{m^2 s^{-1}}$
Vertical diffusivity	$10^{-4} \, \mathrm{m^2 s^{-1}}$

Table 1. Selected properties of B-SOSE iteration 106. Output frequency refers to the output selected for this study.

of the data can improve the convergence of the training process. One specific dimension reduction technique is Principal component analysis (PCA), which identifies the functions that capture most of the variability with depth in the dataset. The result is a representation of the dataset as a linear combination of eigenfunctions (i.e. principal components), sometimes called a principal component expansion or principal component decomposition. Using this procedure, we can describe each profile using a small set of eigenvalues (i.e. coefficients of the principal component expansion) instead of a full set of temperature and salinity values. In addition to improving the speed and efficiency of the GMM algorithm, PCA reveals potential physical structures that may be useful for understanding the stratification of the SO (Pauthenet et al., 2017). We choose the number of principal components such that the percentage of variability explained (in a statistical sense) by the PCA expansion is sufficiently high for our purposes.

Following Rosso et al. (2020), we only keep values between 300 m - 2000 m to exclude most of the surface seasonal variability from the dataset. Because the data is spaced on an irregular grid in the vertical direction, we first interpolate the temperature and salinity profiles onto a regular grid with 10 m cells in the vertical. Following Pauthenet et al. (2017), at each grid cell and time we concatenate the temperature and salinity profiles into a single vector. We normalise each depth level for both temperature and salinity separately; subtracting the mean and dividing by the standard deviation calculated for all time periods on that particular depth level and variable. That is, we standardise the temperature values at each level using the distribution of temperatures at that same depth level, and we standardise the salinity values using the distribution of salinities at that same depth level. This is a slightly different approach from Pauthenet et al. (2017), in which the authors standardise across the entire dataset. We found that for the work shown in this paper, the choice of normalisation approach does not make a large difference in the results (not shown). After normalisation, we carry out PCA expansion. We keep the first three principal components (PCs), which together statistically explain 98% of the variability across the thermohaline dataset. For completeness, we show the structure of the principal components in appendix B.

The coefficients associated with the PC1 indicate a broad division between polar, high-latitude Southern Ocean waters and the subtropics (Fig. 1(a)). The most negative PC1 coefficients are found in the Weddell Gyre, and we also see the imprints
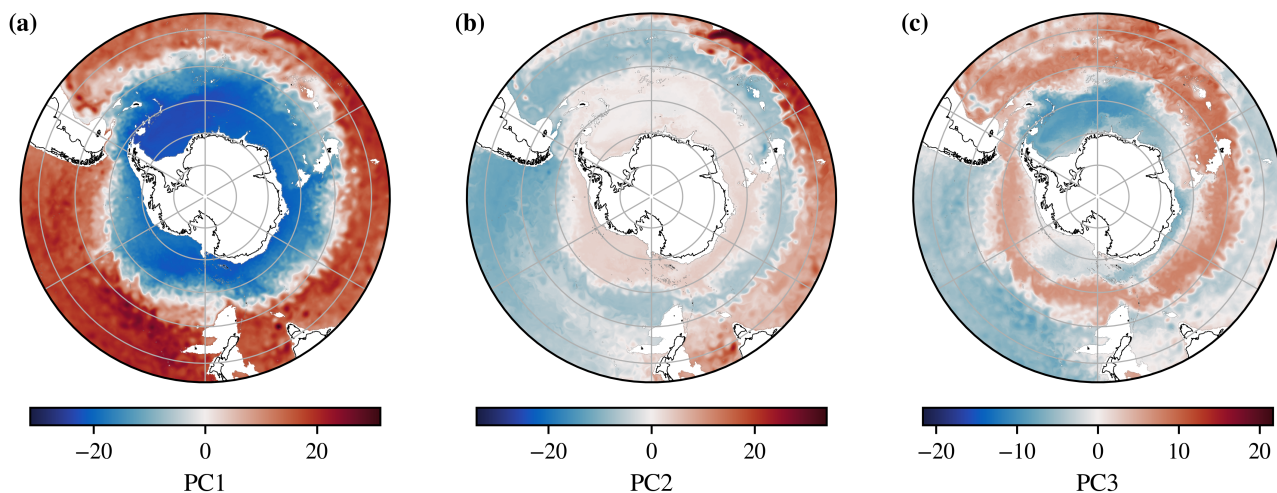
Figure 1. Each combined temperature and salinity profile can be approximated using a three-term PC expansion. Above are monthly mean coefficients of the PC expansion from June 2011. In order to limit the influence of seasonal variability, we use temperature and salinity profiles between 300m-2000 m. The first three PCs explain (a) 75%, (b) 16%, and (c) 7% of the variance respectively, together explaining a total of 98% of the variability. The white space represents bathymetry shallower than 2000 m, and its boundary is marked by a grey line.

of the South Pacific Gyre and the ACC (Vernet et al., 2019). The coefficients of PC2 bear the imprint of the ACC and of its northward flow along the East Pacific basin (Fig. 1(b)). This northward flow is associated with the formation and export of Subantarctic Mode Water and Antarctic Intermediate Water (Iudicone et al., 2007; Sallee et al., 2010; Jones et al., 2016). PC2
150   also has the imprint of the Agulhas Current around South Africa. Finally, PC3 has strong negative values in the Weddell Gyre and over most of the Pacific, with a band of circumpolar positive values that somewhat mirrors the southward drift of the ACC when considered from west to east. The spatial structure of PC1 and PC2 are largely consistent with those of Pauthenet et al. (2017), but the structure of PC3 is somewhat different from theirs, particularly in the subtropics. These differences are possibly a result of our choice of a different depth range. Given that PC3 explains a small fraction of the variability (7% of the variance
155   explained), we do not expect these differences to impact our results.

After we perform dimensionality reduction, each monthly output at each model grid cell in latitude and longitude is represented using the first three coefficients of the PC expansion. The three PC values contain combined information about both temperature and salinity, simplifying our analysis. This approach defines an abstract three-dimensional space in which we can perform unsupervised classification. In typical machine learning terminology, this abstract three-dimensional space can be
160   called the "feature space", in which each PC axis is a "feature". To be explicit, we can say: each combined temperature-salinity profile in latitude, longitude, and time is represented by a three-dimensional vector of PC values. Each three-dimensional PC vector derived from B-SOSE is an "observation". In the next section, we use unsupervised classification to identify subpopulations in the three-dimensional distribution of PC values.

2.3 Gaussian mixture modelling

Unsupervised classification attempts to identify subpopulations within a data distribution, without the assistance of any predefined labels. In our application, we attempt to identify data sub-populations in the abstract three-dimensional space defined by the PC coefficients (i.e. the "feature space"). Here we use GMM, an algorithm that attempts to fit a set of multi-dimensional Gaussian functions to the data by iteratively adjusting the means and covariances of the Gaussians (McLachlan and Basford (1988), see appendix C for more detail). This method has recently been used to classify Argo temperature profiles in the top 2 km of the North Atlantic Ocean and the Southern Ocean (Maze et al., 2017; Jones et al., 2019). GMM is well-suited to ocean applications because it offers a probabilistic measure of classification in the form of posterior probabilities, which is useful when working with a highly correlated dataset. Because GMM-derived clusters will likely feature some overlap due to the highly correlated nature of ocean data, such posterior probabilities offer an important complement to the GMM-derived class labels. In this application, we use the posterior probabilities to define coherent thermohaline regimes and their boundaries.

The GMM method attempts to represent the underlying data distribution using a set of K Gaussian functions in D dimensions (in our case $D = 3$):

$$\mathcal{N}\left(\boldsymbol{x}; \boldsymbol{\mu}_k, \Sigma_k\right) = \frac{\exp\left[-\frac{1}{2}\left(\boldsymbol{x} - \boldsymbol{\mu}_k\right)^T \left(\Sigma_k{}^{-1}\right)\left(\boldsymbol{x} - \boldsymbol{\mu}_k\right)\right]}{\sqrt{(2\pi)^D \|\Sigma_k\|}}, \tag{1}$$

where $\boldsymbol{x} \in \mathbb{R}^{D \times 1}$ is a vector in the PC space, $\boldsymbol{\mu} \in \mathbb{R}^{D \times 1}$ is the center of the Gaussian distribution expressed in vector form, $\Sigma_k \in \mathbb{R}^{D \times D}$ is the covariance matrix, and $|\Sigma_k|$ is its determinant. The covariance matrix determines the orientation of the Gaussian ellipsoids in PC space. We model the dataset, in the statistical sense of representing the dataset using a probability distribution, as a weighted sum of Gaussians:

$$\mathbb{P}(\boldsymbol{x}) \approx \sum_{k=1}^{K} \lambda_k \, \mathcal{N}\left(\boldsymbol{x} \, ; \, \boldsymbol{\mu}_k \, , \, \Sigma_k\right), \tag{2}$$

where λ_k is the weight associated with the k-th Gaussian. The process of fitting the GMM uses expectation maximisation (EM), which consists of iteratively adjusting λ_k, $\boldsymbol{\mu}_k$, and Σ_k to decrease the model-data misfit. For additional details, see Appendix C.

Once the weights, means, and covariances are fitted, each data vector $\boldsymbol{x}$ is associated with a posterior probability distribution across all of the K classes. Although we kept the random seed used in the initial guess fixed for this paper, our results are robust to the choice of random seed (not shown). This distribution is the set of likelihoods that the data vector belongs to any particular class, and the probabilities sum to one. GMM assigns each data vector to the class with the maximum posterior probability. We will now use this distribution to define an inter-class metric, which gives us a novel perspective on fronts as transitions in thermohaline structures.

3 The inter-class comparison metric (the *I*-metric)

First, we examine the structure of our profile data in PC space and introduce the *I*-metric for identifying boundaries between coherent hydrographic regimes (Sect. 3.1). Next, we examine the *I*-metric in both a monthly averaged and multi-year averaged view in latitude-longitude space, and we explore the class structure in more detail by examining the associated coherent regions and vertical profile types (Sect. 3.2). Following that, we compare our results with a local edge detection method (Sect. 3.4).

3.1 Defining the *I*-metric

For each combined temperature-salinity profile, GMM returns a probability distribution across all of the K classes. This distribution is called the *posterior probability* distribution, and it quantifies the probability that a particular profile is in a particular class. If the posterior probability is close to 1.0 for class k and very small for the other classes, then within the context of the Gaussian statistical model (i.e. GMM), the classification of the profile into class k is unambiguous and clear. However, if the posterior probability is close in value for the two classes with highest probabilities, then the classification is ambiguous and less clear. With this in mind, we can use the difference between the highest probability and the second-highest probability to quantify how clearly the profile has been classified. If the classification is unambiguous, then the profile is less likely to be associated with a boundary between coherent thermohaline regimes. If the classification is ambiguous, then the profile is more likely to be associated with a boundary. With this in mind, we propose a probabilistic inter-class comparison metric of the form:

$$I(\mathbf{x_n}) = 1 - \Big[\mathbb{P}(c = c_k)_{\text{highest}} - \mathbb{P}(c = c_l)_{\text{runner-up}}\Big], \tag{3}$$

where $\boldsymbol{x}_n$ is the n^{th} profile's PC values and $\mathbb{P}(c = c_k)_{\text{highest}}$ is the highest posterior probability that GMM has assigned the n^{th} profile as belonging to class k. The term $\mathbb{P}(c = c_l)_{\text{runner-up}}$ is the second-highest posterior probability belonging to class l. If the difference between the highest and runner-up posterior probabilities is close to one, then *I* is small. This would indicate that the profile is not likely to be associated with a boundary between thermohaline regimes. If the difference between the highest and runner-up posterior probabilities is small, then *I* is close to one, indicating that the profile is likely to be associated with a boundary between different thermohaline regimes. The *I*-metric offers an alternative method for defining boundaries as fuzzy transitions between coherent regimes. In general, some regions will feature sharp transitions across boundaries, whereas other regions will feature more gradual transitions. The relative sharpness of a transition is influenced by the processes that form, mix, and destroy water masses. In contrast with approaches that define fronts as sharp transitions located along property contours or local gradients, the *I*-metric approach allows for a wider variety of transition types between regimes.

In our *I*-metric application, GMM clusters the profiles in feature space (Fig. 2(a)). The structure of the data shown in PC space is broadly consistent with that found in other studies (e.g. Pauthenet et al. (2017, 2018, 2019)). The data distribution is reasonably well represented by a linear combination of multi-dimensional Gaussian functions (Fig. 2). The *I*-metric values indicate transition regions between classes, where the class labelling is relatively ambiguous (Fig. 2(b)). We choose $K = 5$ to represent the general, large-scale pattern of the data; we explore the sensitivity of our results to K in Sect. 4.4. In the next section, we examine the *I*-metric and class structure in physical space.
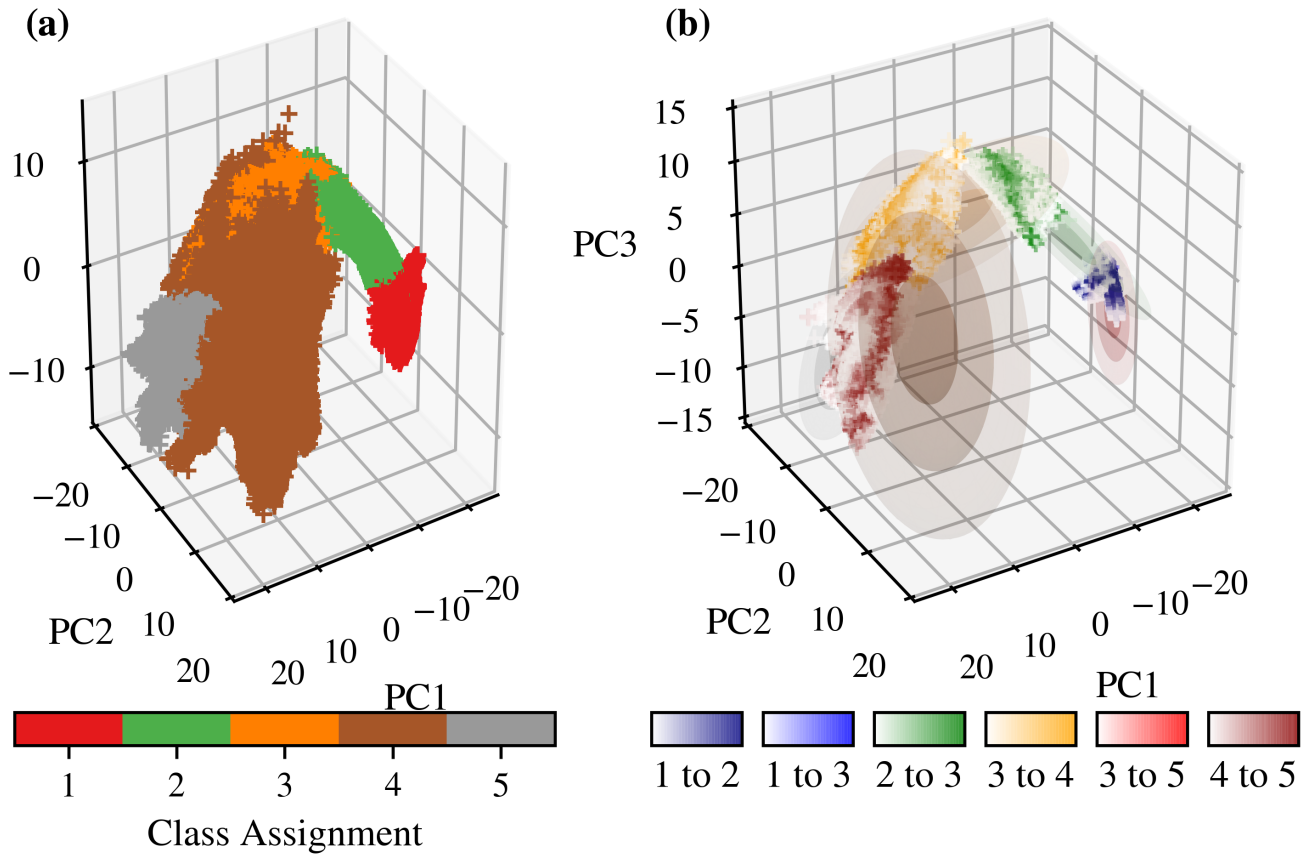
Figure 2. (a) The classification analysis takes place in the abstract PC space. Each point represents a 3D vector of principal component values that describe a single combined temperature and salinity profile. The three axes are the three principal components. Class assignments are indicated by the colours. (b) The *I*-metric highlights transitions between classes in the abstract PC space. The Gaussian ellipsoids of the GMM are shown in red, and the *I*-metric values associated with each point are shown using six different colour scales. Each colour scale corresponds to a particular transition between classes. Points with low *I*-metric values are not shown. The above is a subset of data taken from 12 months of monthly averaged B-SOSE data, inclusively between August 2011 and July 2012.

3.2 Geographic view of the *I*-metric

The *I*-metric viewed in latitude-longitude space illustrates the rich variety of transition types found in the Southern Ocean (Fig. 3). In all sectors of the SO, we see sharp transitions where the regions of high *I* values are narrow and more gradual transitions where the regions of high *I* values are more spread out. Some features are circumpolar, in consistency with the view of SO fronts as continuous lines that encircle Antarctica. However, we also see regions where the continuity and circumpolar nature of the fronts is not as clear, suggesting that a broader view may be appropriate (Chapman et al., 2020). The fronts are not

uniformly sharp across all longitudes; for example, the northernmost transition is broad and gradual in the Atlantic sector, sharp in the Indian sector, and relatively broad in the Pacific sector. The southernmost band of high I-metric values is relatively sharp in the Atlantic sector, becoming increasingly broad as we follow it into the Indian and Pacific sectors. In the Pacific sector, it extends into an especially broad region in the Amundsen Sea, in consistency with the intersection of the classically-defined

235   southern boundary (SBdy) with the Antarctic continent (Kim and Orsi, 2014). Upstream of Kerguelen Plateau, there is a region where the I-metric is spread out and diffuse between classes 2 and 3; this region also features a standing meander associated with enhanced eddy kinetic energy (Frenger et al., 2015; Siegelman et al., 2019). The enhanced mesoscale eddy kinetic energy associated with the meander is consistent with increased lateral mixing and the spread out pattern in the I-metric found in the same region. Closer to the Antarctic continent, we also see the imprints of both the Weddell Gyre and the Ross Gyre, in regions

240   of coherent structures with low I-metric values, in part enforced by the gyre circulation.

The monthly mean I-metric (Fig. 3(a)) also highlights individual ring-like eddies; although these features are not typically considered fronts, they are small-scale transition regions between different hydrographic structures. We do expect the I-metric to be non-zero across these features. The monthly view also features mesoscale meanders, highlighting the detailed structure of the SO, which is partly a result of the energetic mesoscale eddy field. The I-metric does feature some month-to-month

245   variability; in some locations the fronts meander in their north-south extent, and in others they are relatively stationary, likely due to bathymetric constraints (see animations in Thomas (2021), in the "gifs" directory).

By averaging the four years worth of monthly means, we obtain a map of the climatological I-metric, which is averaged over many eddy lifetimes (Fig. 3(b)). Comparing an example monthly field with the climatological field, we can examine the imprint of eddy spatial variability and the meandering of the fronts on the I-metric pattern. Most of our observations about the

250   metric are unchanged by this averaging; we identify three roughly circumpolar bands of high I-metric values, with significant spatial variability and some overlap. The three bands are fairly distinct in the Atlantic sector, with the northernmost transition being the broadest. Upstream of Kerguelen plateau, the two northernmost bands become somewhat hard to distinguish. This is possibly a consequence of the eddy mixing and upwelling hotspot in that region, which tends to spread out hydrographic features in latitude-longitude space, increasing the degree of spatial correlation found there. Upstream of Kerguelen plateau,

255   the Polar Front features strong seasonal variability (Pauthenet et al., 2018). Note that the I-metric band aligned roughly with the Polar Front only passes south of the plateau (e.g. south of Heard Island), in consistency with other studies of the subsurface component of the Polar Front (e.g. Pauthenet et al. (2018)).

The three bands of higher I-metric values are distinct downstream of the Kerguelen plateau in the Indian sector; notably, the southernmost band features especially high I-metric values in this sector. This pattern is associated with the transition

260   between the Antarctic Circumpolar Current and the Antarctic Slope Current (ASC), which tend to flow in opposite directions (Thompson et al., 2018; Pauthenet et al., 2021). In the Pacific sector, we see the southernmost band turn into the Amundsen Sea and intersect with the Antarctic continental slope, spreading out in a diffuse region that is consistent with the behaviour of the southernmost extent of the ACC, the eastern boundary of the Ross Gyre, and the eastward shelf circulation along the West Antarctic Peninsula (Nakayama et al., 2018). In this same sector, two large regions of low I-metric values spatially correspond

265   to export pathways of Subantarcitc Mode Water and Antarctic Intermediate Water (Iudicone et al., 2007; Sallee et al., 2010,
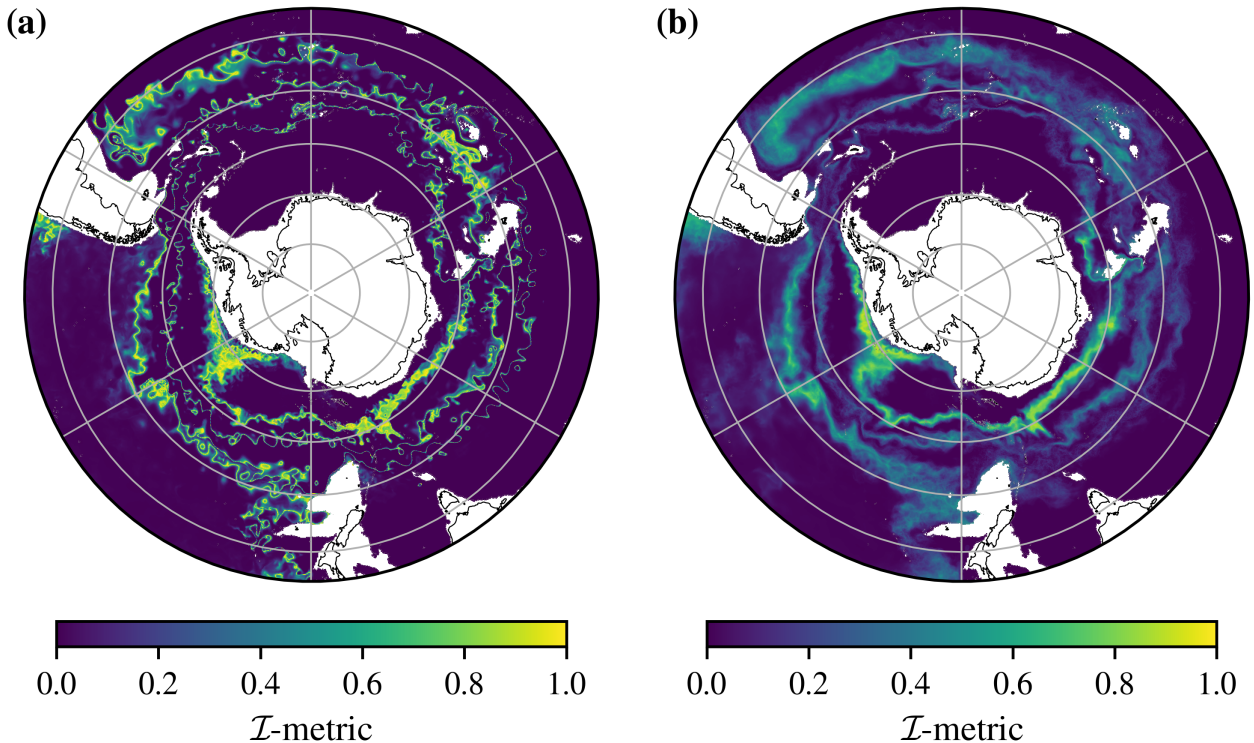
Figure 3. The magnitude of the *I*-metric highlights transitions between coherent thermohaline regimes. Panel (a) is the *I*-metric for a single month (April 2012) and (b) for the time average of the B-SOSE iteration 106 dataset (60 months). Latitude lines are shown between 80°S and 40°S every 10°, and longitude lines are shown every 60°. Animations showing month-to-month and interannual variability are available in the software release (Thomas, 2021).

2012; Jones et al., 2016). The higher I-metric values delimit the edges of these more coherent thermohaline regimes (Fig. 3(b)), which are influenced by basin-scale stratification and the structure of the South Pacific Gyre.

3.3 Properties of the thermohaline regimes

In order to better understand the coherent thermohaline regimes underlying our *I*-metric results, we examine their lateral extents
270 and their vertical properties. Despite not being given any latitude or longitude information, the underlying GMM captures several coherent, large-scale features of Southern Ocean thermohaline structure (Fig. 4(a)). Class 1 contains the coldest waters in the SO, covering both the Weddell and Ross gyres near Antarctica. The mean profile in this class features cold temperatures that are nearly uniform with depth; in general, they are salt stratified in that the near-surface waters are fresher than the subsurface waters, ensuring that the density profile is stable overall (Fig. 5). The boundary between class 1 and class 2 broadly
275 lies between the classically-defined Southern ACC front (SACCF) and the Southern Boundary (SBDY) (Kim and Orsi, 2014),

including the turn of the SBDY towards almost being perpendicular with the Antarctic continent in the Pacific sector of the SO. Class 2 is circumpolar, with excursions into the Amundsen Sea and the area just south of Kerguelen Plateau. It features salt stabilisation, with a fresh layer near 300 m (Fig. 5). Class 3 is also circumpolar, with a northward excursion in the Atlantic sector. The boundary between classes 2 and 3 roughly follows the Polar Front (PF), separating the colder, fresher Antarctic waters from the warmer, saltier subtropical waters further north. Finally, there are two subtropical classes; class 4 represents the Atlantic and Indian sectors of the subtropics, and class 5 represents the large-scale South Pacific Gyre. The boundary between classes 3 and 4 roughly aligns with the Subantarctic Front (SAF), particularly over large portions of the Indian and Pacific sectors (Fig. 4). The mean of class 5 has a salinity minimum around 700 m, corresponding to the presence of the Antarctic Intermediate Water layer (Iudicone et al., 2007).
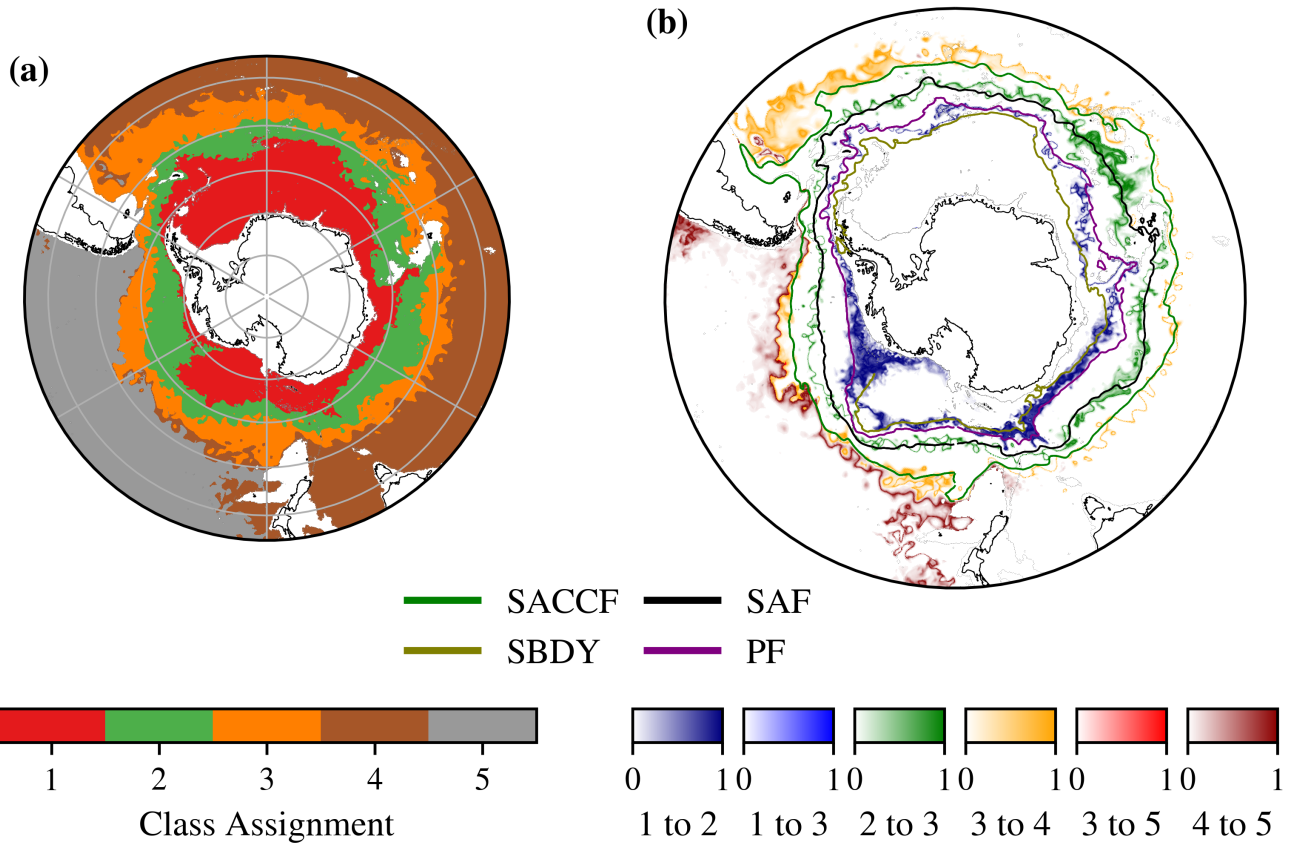
Figure 4. (a) The cluster assignments with $K = 5$ and (b) the I-metric for all present class transitions. This view highlights the transitions between specific classes. The transitions in (b) have some similarities to the altimetric fronts from Kim and Orsi (2014). These fronts are shown overlain on (b) (SBDY: Southern boundary, SACCF: Southern ACC front, PF: Polar front, SAF: Sub-Antarctic front). Data from June 2011 as a representative month.

3.4 An edge detection approach towards identifying fronts

For comparison with the GMM method, which uses properties of an entire training dataset to detect changes in thermohaline structure, we use a more local front detection method implimented by Hjelmervik and Hjelmervik (2019) in the North Atlantic. This method, called the Sobel method, directly examines spatial gradients in the principal component fields using a Sobel operator (Sobel and Feldman, 1968). To do this, the PCs of each grid point are placed onto a rectangular grid with the same spacing as the data sampling, where points without data are masked. The strength of an edge at a point is found by the two dimensional convolution (represented by *) of the gridded PCs and the following two matrices. In the x direction the Sobel
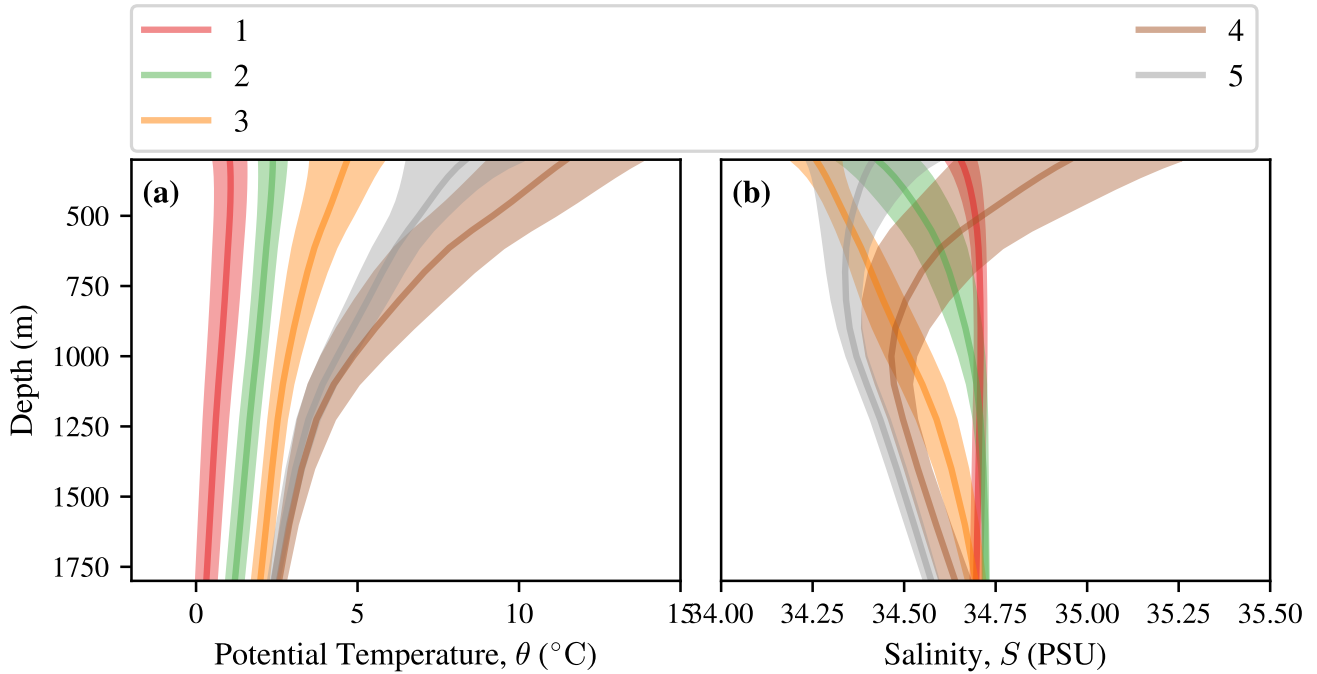
Figure 5. Profiles of the five GMM clusters between 300 m and 1800 m in (a) temperature and (b) salinity. This is calculated from the profiles classified using the statistical model fitted on training data itself. The central line is the mean and the envelope on either sides indicates one standard deviation.

operator is

$$G_x = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}, \tag{4}$$

and in the y direction the Sobel operator is

$$G_y = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}. \tag{5}$$

The effect of this operator is similar to a gradient operator with some smoothing. There is a correlation coefficient of 0.99 between $G_y * PC1$ and the y gradient, and there is a correlation coefficient of 0.999 between $G_x * PC1$ and the x gradient of PC1. The motivation for using the Sobel operator rather than the gradient operator is principally that it can reduce the noise in data, as shown by application to photographs (Vincent et al., 2009). Hjelmervik and Hjelmervik (2019) used the magnitude of the x and y Sobel operators, which approximates the magnitude of the gradient, to examine fronts in the Arctic and North Atlantic. They show that the magnitude of the Sobel gradient can be thresholded to highlight features such as the Gulf stream.
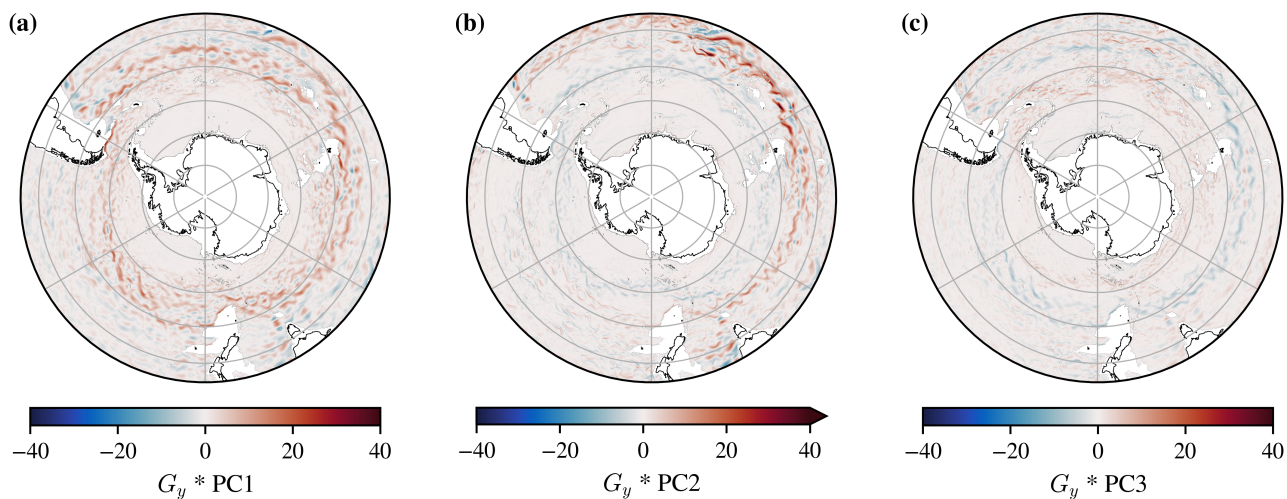
Figure 6. G_y Sobel edge operator convolved in 2D (*) with the principal component coefficient fields for the month of June 2011. The correlation coefficient between panel (a) for PC1, and the zonal velocity U for the same month in B-SOSE iteration 106 is 0.85, showing that the structure it highlights is substantially similar to the ACC. Panels (b) and (c) for PC2 and PC3 are also related to the ACC, (correlation coefficients of 0.18, and -0.18 respectively). The grey line is the 2000 m isobath.

Rather than working with the gradient magnitude, Fig. 6 shows G_y*PC1, G_y*PC2, G_y*PC3 alone. This is more interpretable, as the G_y*PC1 component is strongly correlated to the zonal velocity U ($r = 0.85$). Hjelmervik and Hjelmervik (2019) use a threshold value to define fronts, but instead we plot the gradient directly as a colormap for each PC, which is useful as it does not obscure any information about the fronts themselves. Appendix A shows that the correlation between the Sobel G_y gradient of PC1 with the meridional velocity, V, and the correlation between G_x*PC1 and zonal velocity U increase for roughly the first 2 years of B-SOSE iteration 106, suggesting that the model is still spinning up to geostrophic balance.

The GMM and Sobel methods are complementary. GMM reveals the large-scale temperature and salinity structure associated with changes in stratification, which has traditionally been used to define the fronts, whereas edge detection methods like the Sobel method used here reveals the smaller-scale structure of multiple jets, which can merge and separate. As such, both approaches may be useful ways of characterising ocean structure without making ad-hoc assumptions related to particular property values or strict requirements that the structures be circumpolar and continuous. The present proliferation of front definition and analysis methods is driven by the need to expand how the oceanographic community deals with ocean structure across a wide variety of spatial and temporal scales (Chapman et al., 2020).

4   Discussion

In this section, we discuss the sensitivity of our results to our choice of dataset (Sect. 4.1), touch on the temporal variability in our results (Sect. 4.2), discuss a possible connection with the Antarctic Slope Current (Sect. 4.3), examine the sensitivity of the results to the choice of the number of classes K (Sect. 4.4), and discuss the interpretation of posterior probabilities (Sect. 4.5).

4.1   Sensitivity to choice of dataset

We chose to use B-SOSE data for this study in order to (1) work with a dataset that features relatively uniform coverage in latitude-longitude and (2) to allow us to examine temporal variability as well as spatial variability. B-SOSE is an observationally-constrained estimate of the hydrographic structure of the Southern Ocean, so it does accurately capture many features of large-scale and mesoscale structure (Verdy and Mazloff, 2017). However, because B-SOSE is a numerical model run, it will no doubt have some biases with respect to observations, particularly on smaller scales. We expect that our results would not change dramatically on basin-wide scales across different state estimate and reanalysis products.

To examine the differences of this bias on the class structure and the structure of the inter-class comparison metric, this study could be repeated with a purely observational dataset such as Argo. One trade-off for such a study would be the fact that observational datasets are relatively sparse in terms of both spatial and temporal coverage relative to a state estimate or other numerical model run. One could attempt to use the same GMM trained on B-SOSE with Argo data, but possible biases between B-SOSE and the Argo dataset could make this challenging. It might be possible to adjust for those biases in the data cleaning and preparation step of the analysis; the standardisation process, which is already a part of the analysis presented here, is a step towards this bias removal and correction that may facilitate comparisons between models and observations. Alternatively, one could attempt to re-train the GMM using Argo data alone. This has been done in other studies, so it should be possible in principle (e.g. (Maze et al., 2017; Jones et al., 2019; Rosso et al., 2020)).

4.2   Temporal variability of the fronts

We found that the class structure and boundary positions did not feature large temporal variability with respect to the mean state, but much more work could be done to examine this variability and its connection to the processes that determine thermohaline structure (e.g. surface forcing, subsurface mixing, advection). This is outside the scope of our present study, which is focused on proposing a new metric for identifying and tracking boundaries in Southern Ocean structure.

4.3   The Antarctic Slope Current

The Antarctic Slope Current (ASC) that separates warmer open ocean waters from the colder waters on the Antarctic continental shelf is an important component of heat transport in the Southern Ocean. It acts to control the flow of warm water onto the continental shelf and eventually under the floating ice shelves. In a recent paper, Thompson et al. (2020) suggest that if the source of the Antarctic Slope Current (ASC) intersects with the ACC in the Bellingshausen Sea, then the ASC source would be considered a major component of the overturning circulation. In our study, we found a diffuse boundary between classes in

the Bellingshausen Sea region, which may be relevant for the physical context of the ASC, which is still under investigation (Fig. 4(b)).

4.4 Sensitivity to the maximum number of classes

In this study, we chose $K = 5$ as the number of classes based on sensitivity tests and also based on a priori knowledge. Specifically, previous studies used a front structure with five broad regions, delineated by four fronts, so we might expect a value around $K = 5$ based on this (e.g. Orsi et al. (1995); Pollard et al. (2002); Kim and Orsi (2014)).

Generally, the choice of the maximum number of classes K can be thought of as a way to select models of varying degrees of complexity. Statistical models with lower K values are potentially easier to interpret, only capturing the most dominant structures in the dataset. For example, the probabilistic boundary between the two classes in a $K = 2$ statistical model roughly separates colder, fresher Antarctic waters from the warmer, saltier subtropical waters (Fig. 7(a)). Notably, in this case, the magnitude of the *I*-metric appears to largely decrease as we follow it from the Atlantic and Indian basins and into the Pacific basin, indicating that the boundary becomes less sharp with longitude. This possibly reflects the fact that the Pacific basin hosts some of the dominant northward export pathways of Subantarctic Mode Water and Antarctic Intermediate Water, consistent with a less sharp transition between polar and subtropical waters (Iudicone et al., 2007; Herraiz-Borreguero and Rintoul, 2011; Jones et al., 2016). A statistical model with $K = 4$ retains most of the features of our analysis with $K = 5$, but the transition region closest to Antarctica in $K = 5$ is no longer present.

The $K = 5$ statistical model we used in this work captures near-Antarctic and circumpolar structure, as well as some subtropical structure. A more complex statistical model with higher K would capture more of the subtropical structure (not shown). This is consistent with sensitivity studies using temperature-only Argo data, where increasing K added details to the subtropical class structure while leaving the circumpolar class structure largely unchanged (Jones et al., 2019). Statistical models with much higher K values may capture more structure in the data, but increasing K also risks overfitting. That is, if we tune the GMM statistical model to match an increasing number of structures in PC space, we risk losing generality; the goal is to represent the dominant structures of the dataset without overfitting every small variation, some of which could represent noise in the data. This has a direct analogue with overfitting in terms of simple statistical models; it is unwise to use a 10th-order polynomial when a quadratic captures the dominant features of the dataset, because the higher-order polynomial is less likely to generalise to other similar datasets. In addition, statistical models with very high K values are increasingly difficult to interpret in terms of our current physical and biogeochemical understanding. Note that regional studies, e.g. those carried out in specific sectors of the SO may find it useful to increase K based on local structure (e.g. Rosso et al. (2020)). This is consistent with the suggestion by Chapman et al. (2020) that front definitions may need to be more flexible and region-specific, as opposed to expecting a particular definition to apply globally (or even across a single ocean basin).

4.5 Interpreting posterior probabilities

The posterior probabilities returned by a Gaussian mixture model are affected by our choice of K. We should be careful not to over-interpret the posterior probabilities as confidences in the correctness of the assigned labels. Notably, GMM does not
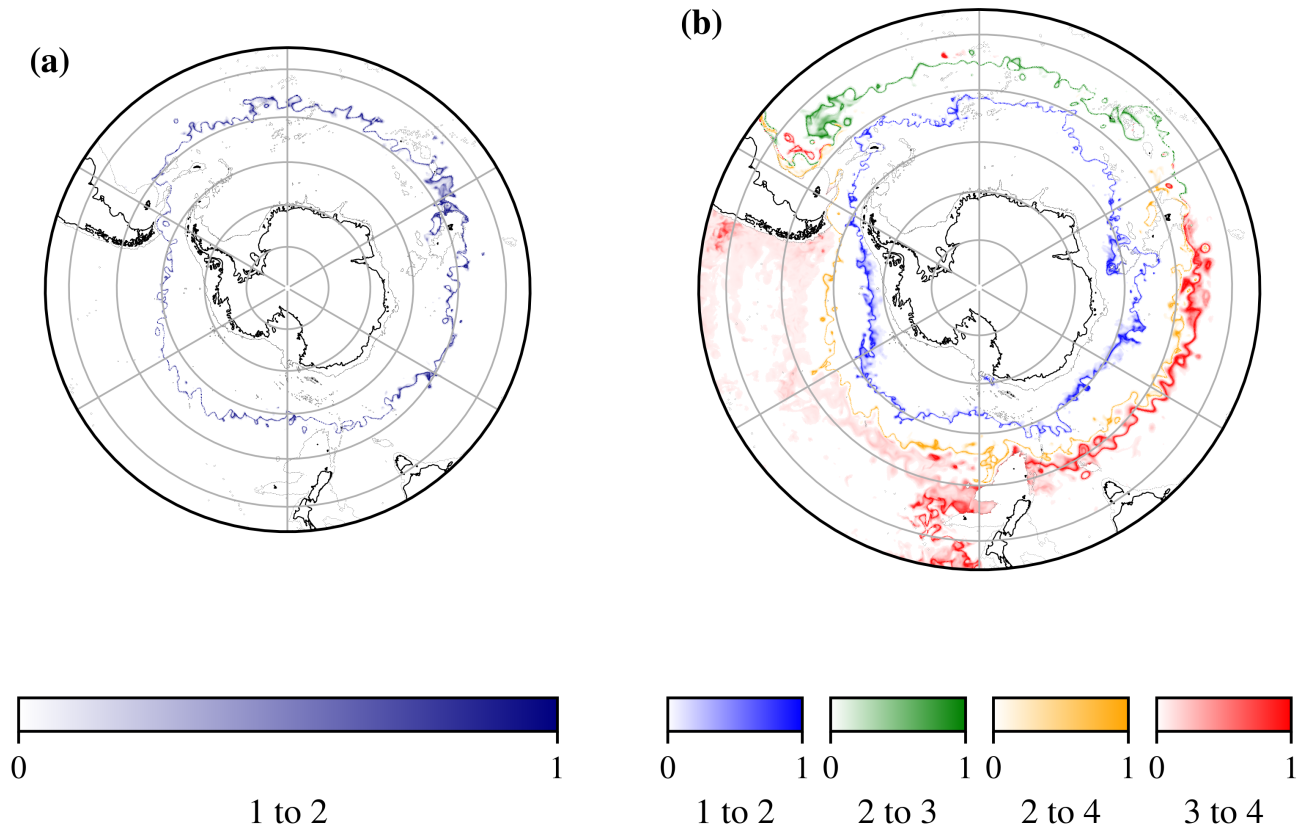
Figure 7. Decreasing K removes details from the statistical description of Southern Ocean thermohaline structure. Shown is the GMM-derived I-metric, using (a) $K = 2$ and (b) $K = 4$, for a monthly average over June 2011. Grey line is 2000 m isobath.

indicate the probability that a given profile belongs to *none* of the classes in a given statistical model. With that in mind, we can interpret the posterior probability as a measure of unambiguity *in the context of* a given statistical model. When one probability is larger than all others with some margin, the profile is unambiguously classified, while probabilities of similar magnitude indicate that the profile cannot be unambiguously classified in the current statistical model with the specified number of classes. In this study, we used the posterior probability distribution to identify boundaries between coherent thermohaline regimes, taking advantage of this property of GMM.

5 Conclusions

In this study, we proposed a new metric for defining and identifying boundaries between coherent regimes of temperature and salinity structure. Our method uses Gaussian mixture modelling, a type of unsupervised machine learning, to establish a

statistical model of thermohaline structure that is intended to capture the large-scale features of the dataset in both PC space and in geographic space. We developed our method in the Southern Ocean due to the presence of circumpolar structures and

390   relatively clear fronts, but our approach could be applied to other regions or even to the global ocean as a whole. The *I*-metric provides a flexible, probabilistic method to define and identify boundaries in an oceanographic dataset without using ad hoc property contours; the boundaries are derived in a generalised method that reflects the structure of the dataset. The *I*-metric has potential as a method for comparing different observational and numerical modelling datasets in a robust, algorithmic way that is not heavily affected by biases in the mean state between datasets. It features a parameter K that allows users to

395   increase and decrease the level of complexity of the statistical model; the optimal value of K will vary between applications. The Sobel edge detection method may be useful for defining and tracking smaller-scale fronts and jets in model or reanalysis data. As discussed in Chapman et al. (2020), the field of oceanography needs to consider fronts and boundaries in a more general, application-specific way, due in part to the richness of ocean structure on different spatial scales. The *I*-metric was designed with this problem in mind; it is intended to be a complementary addition to the oceanographic toolbox as opposed to

400   a replacement for any particular method.

Code and data availability. B-SOSE iteration 106 state estimate data is available from the Scripps Institution of Oceanography (http://sose. ucsd.edu/B-SOSE6_iter106_solution.html). The MITgcm source code that was used to create B-SOSE is available on GitHub (https://github. com/MITgcm/MITgcm). Original climatological front positions from Kim and Orsi (2014) are available on Researchgate (https://www. researchgate.net/publication/338420242_ACC_fronts). The code used to carry out the analysis and figure creation for this paper is available

405   via Zenodo (Thomas, 2021) (Up-to-date repository: https://github.com/so-wise/so-fronts). This software uses scikit-learn (Pedregosa et al., 2011) and pyxpcm (Maze, 2020) as foundations. We used Cartopy for mapping (Met Office, 2010 - 2015).
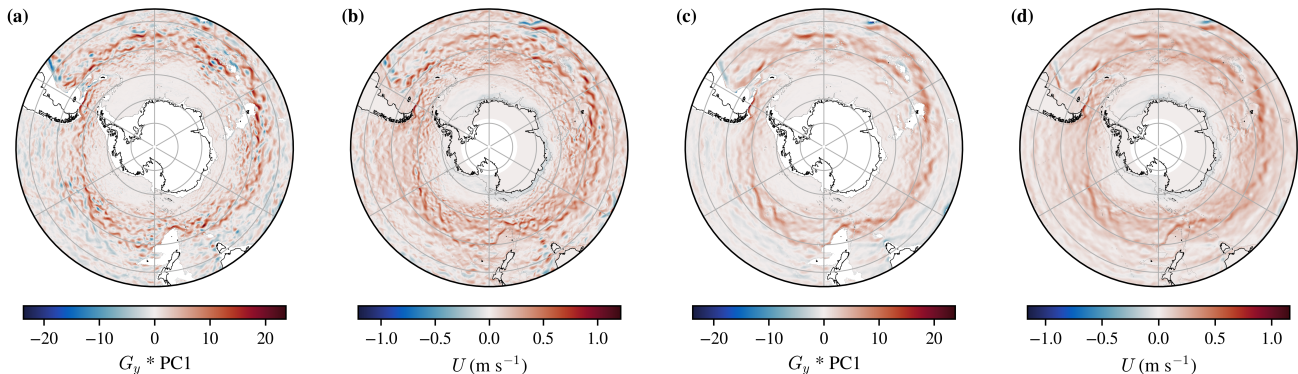
Figure A1. A comparison between the G_x*PC1 Sobel edge detection field and the zonal velocity, U, at 2m. Panels (a) and (b) are from the monthly mean over June 2011, whereas (c) and (d) are the mean over all of the monthly means in the dataset.

Appendix A: The relation between edge detection and the velocity field

Fig. A1 and Fig. A3 illustrate the spatial resemblance between G_y*PC1 and U, G_x*PC1 and V, respectively, compared over June 2011 or as an average over the full B-SOSE period. The domain-averaged correlation is shown quantitatively in Fig. A2 and Fig. A4, where in the last couple of years of the reanalysis product there is an especially high correlation between the two. That Fig. A2 and Fig. A4 show opposite signs in the correlation is equivalent to the reversal in sign between that we would expect (Cushman-Roisin and Beckers, 2011, chapters 15 and 18). Those figures also show that the magnitude of the correlation between the respective variables increases during the first two years of the dataset before flattening off. This is suggestive of the model spinning up towards geostrophic balance. As the first principal component statistically explains the first-order structure in the ocean, it primarily represents the density contrast produced by the thermohaline structure from the tropics to the poles.

Appendix B: Principal component structure

In this work, we use principal component expansion for dimension reduction and to examine the structure of the Southern Ocean. In this appendix we display the principal components (i.e. eigenvectors) used in this expansion. First, we examine the means of the temperature and salinity structure across the entire dataset (Fig. B1). The temperature decreases with depth, whilst the salinity has a minimum around 750 m, in part associated with the presence of Antarctic Bottom Water (AAIW).

The structure of the first three principal components (i.e. eigenvectors) reflects small variations on the mean structure (Fig. B2). The mean profiles are similar, but the variation associated with an increase or decrease in the principal component value changes with depth. The first principal component (PC1) explains 76% of the variability in the dataset, notably consisting of variations throughout the mid-range of the profiles in temperature and throughout nearly the entire depth range in salinity. The second principal component explains 16% of the variability and consists of larger changes in the upper part of the profile, above roughly 1000 m. The third principal component (PC3) explains 7% of the variance and exhibits variations
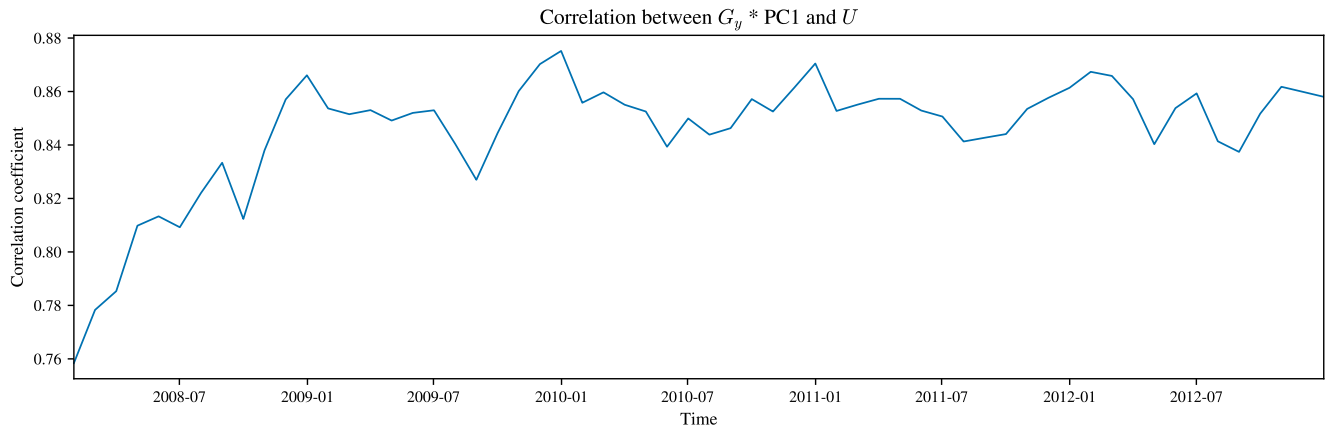
Figure A2. The correlation between G_x*PC1 and the zonal velocity, U, at 2m, for each monthly mean in the B-SOSE iteration 106 dataset. The increase in the correlation over the first two years could be interpreted as the spin-up.
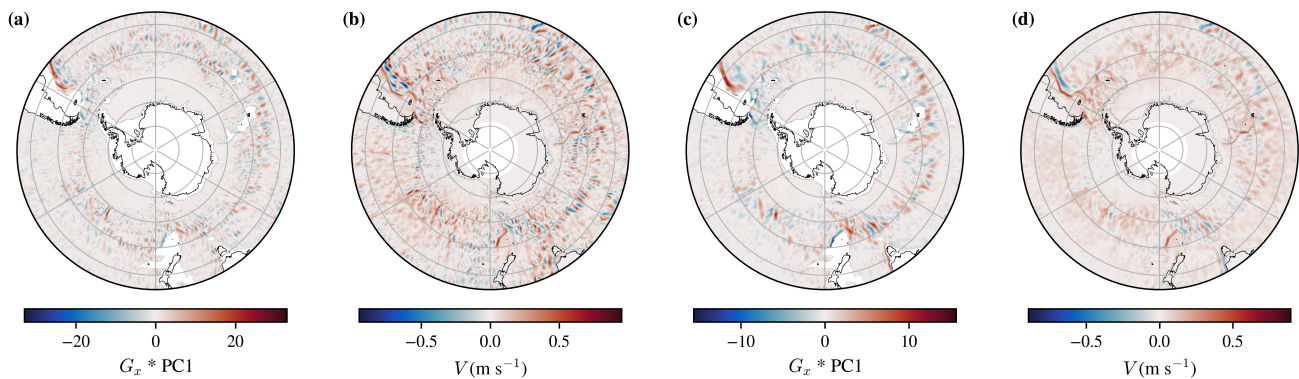
Figure A3. A comparison between the G_y*PC1 Sobel edge detection field and the meridional velocity, V, at 2m. Panels (a) and (b) are from monthly mean over June 2011, whereas (c) and (d) are the mean over all of the monthly means in the dataset.

Correlation between G_x * PC1 and V
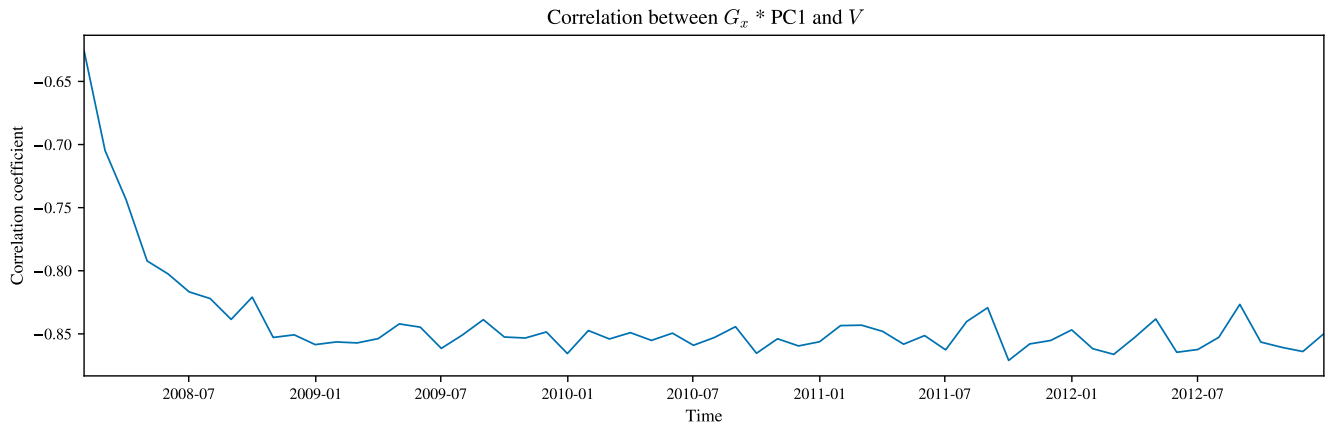
Figure A4. The correlation between the G_y*PC1 and the meridional velocity, V, at 2m for each monthly mean in the B-SOSE iteration 106 dataset. The increase in the correlation coefficient over the first two years could be interpreted as the spin up as in Fig. A2
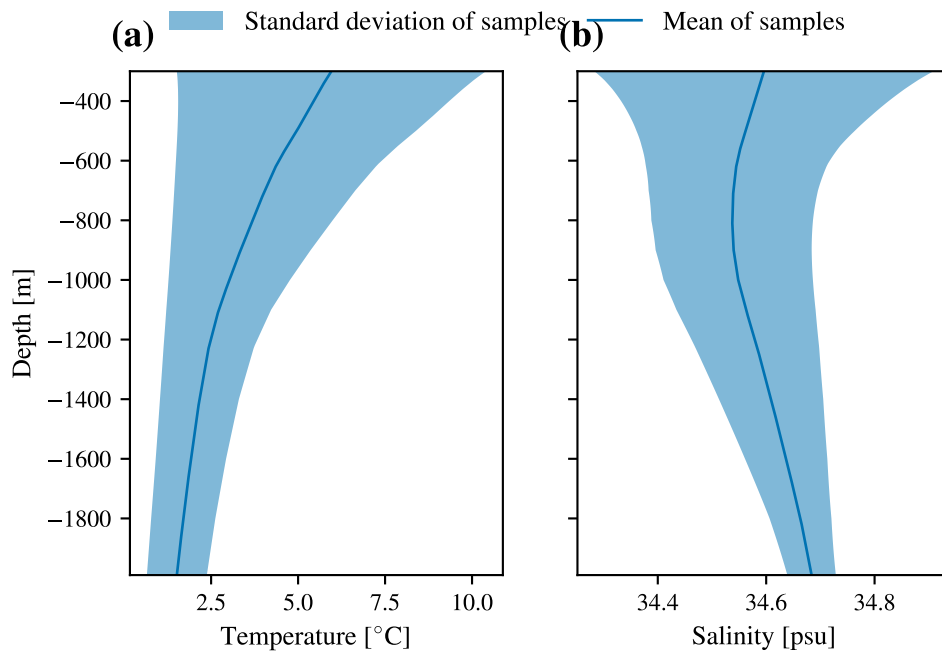
Figure B1. The means and standard deviation of samples taken from B-SOSE iteration 106.

above and below a somewhat fixed mid-point. After principal component expansion is applied, each profile is represented by just three numbers, i.e. the eigenvalues of the principal component expansion.

Appendix C: Gaussian mixture modelling

430   A Gaussian mixture model (GMM) attempts to represent a dataset using a linear combination of multi-dimensional Gaussian distributions. A multi-dimensional Gaussian (Eq. C1), is a simple generalisation of a Gaussian to D dimensions.

$$\mathcal{N}\left(\boldsymbol{x}_n; \boldsymbol{\mu}_k, \Sigma_k\right) = \frac{\exp\left[-\frac{1}{2}\left(\boldsymbol{x}_n - \boldsymbol{\mu}_k\right)^T \left(\Sigma_k^{-1}\right)\left(\boldsymbol{x}_n - \boldsymbol{\mu}_k\right)\right]}{\sqrt{(2\pi)^D \|\Sigma_k\|}}, \tag{C1}$$

where k is the index for the k-th cluster of K clusters, n is the index for the n-th data point of N data points, and D corresponds to the three principal components of the data.

435   We make the assumption that the probability distribution that generated the dataset can be approximated by a set of multivariate Gaussians (Eq. C2):

$$\widetilde{\mathbb{P}}(\boldsymbol{x}_n) \simeq \mathbb{P}(\boldsymbol{x}_n) = \sum_{k=1}^{K} \lambda_k \, \mathcal{N}\left(\boldsymbol{x}_n \, ; \, \boldsymbol{\mu}_k \, , \, \Sigma_k\right). \tag{C2}$$

Any probability distribution function (PDF) could be described by an arbitrarily large number of Gaussians (Eq. C3), but to be a good method of describing the data this should be a manageable number.

$$440 \quad \widetilde{\mathbb{P}}(\boldsymbol{x}_n) = \lim_{K \to \infty} \sum_{k=1}^{K} \lambda_k \, \mathcal{N}\left(\boldsymbol{x}_n \, ; \, \boldsymbol{\mu}_k \, , \, \Sigma_k\right) \tag{C3}$$

In this paper, we showed that our Southern Ocean thermohaline dataset can be fairly represented as a series of plateau-like regions in PC variable space, so it can be approximated by a PDF made from a set of multivariate Gaussians, where the boundaries between these Gaussians correspond to the fronts (Fig. 2).

C1   Expectation Maximisation

445   To initialize the method, the first K clusters are created randomly. Next, the set of Gaussians is iteratively adjusted (Eq.s C4, C5 and C6) until it reaches a local minimum in the cost function Maze et al. (2017). It is generally expected that reducing the number of dimensions in the pre-processing step helps improve the convergence. The following section draws heavily from Maze et al. (2017).

The expectation of the model given the data is increased by updating the weights λ_k, means $\boldsymbol{\mu}_k$, and covariance matrices Σ_k
450   in the following way:

$$\lambda_k^{(t+1)} = \frac{1}{N} \sum_{n=1}^{N} \mathbb{P}\left(c_n = k \mid \boldsymbol{x}_n \, ; \, \{\lambda_k \, , \, \boldsymbol{\mu}_k \, , \, \Sigma_k\}^{(t)}\right), \tag{C4}$$
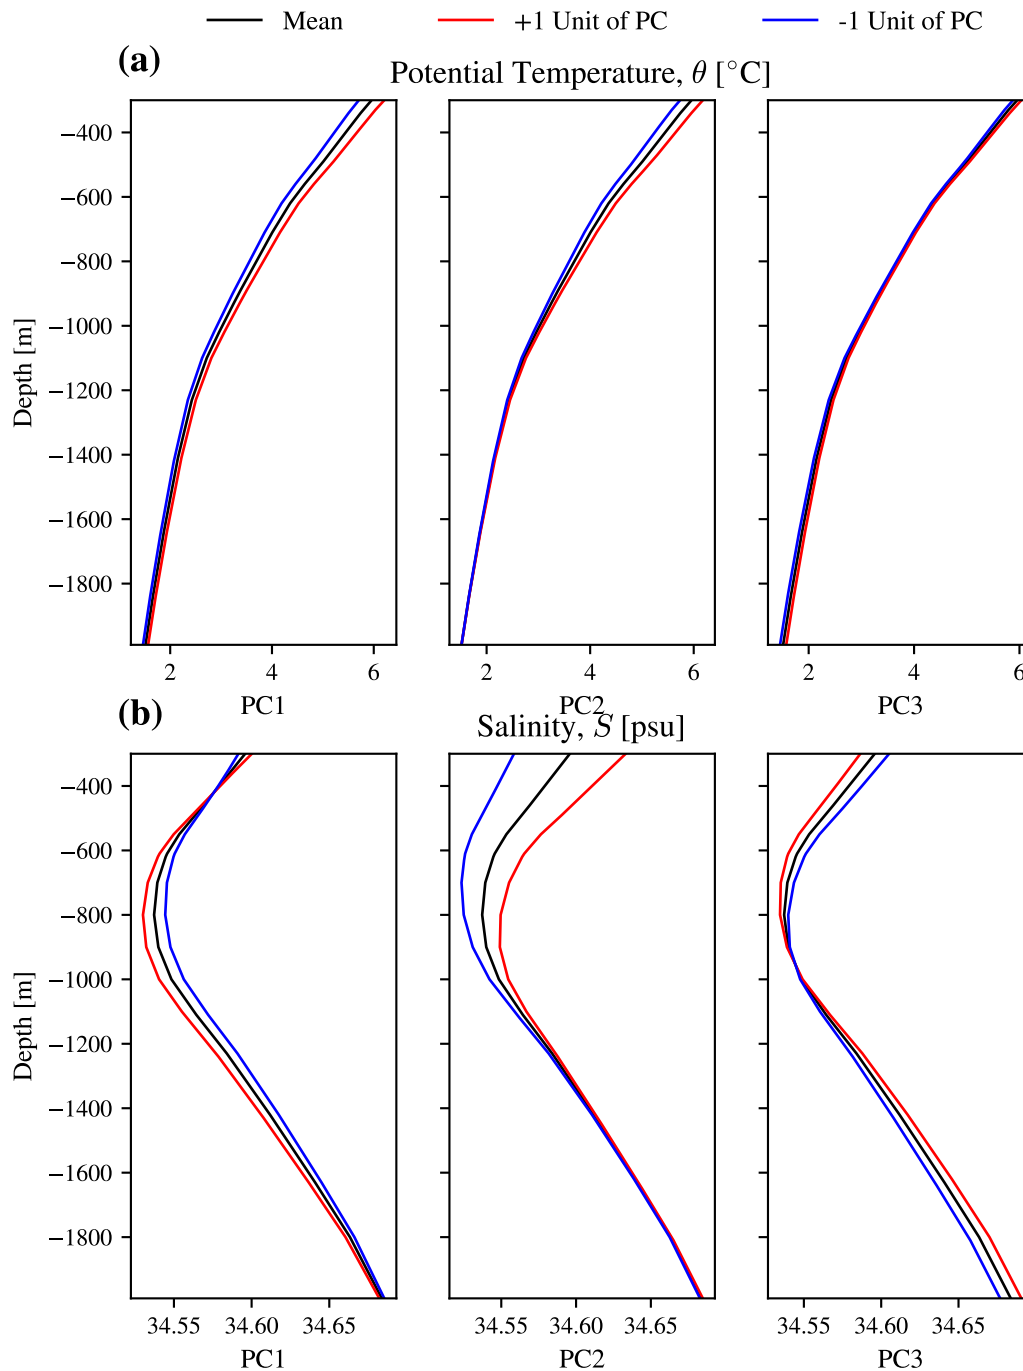
Figure B2. The temperature and salinity components of the three retained principal components that were used in this work. Specifically, they are PC1 (left coumn), PC2 (middle column), and PC3 (right column). Shown are the mean structures (black lines) with the effect of adding (red line) or removing (blue line) one unit of a principal component as a deviation from the mean profile, after Figure 4 in Pauthenet et al. (2017). When compared to Fig. 1 we can see that PC1 corresponds to the hot-cold north-south contrast.

$$\boldsymbol{\mu}_k^{(t+1)} = \frac{\sum_{n=1}^{N} \mathbb{P}\left(c_n = k \mid \boldsymbol{x}_n \, ; \, \{\lambda_k \, , \, \boldsymbol{\mu}_k \, , \, \Sigma_k\}^{(t)}\right) \boldsymbol{x}_n}{\sum_{n=1}^{N} \mathbb{P}\left(c_n = k \mid \boldsymbol{x}_n \, ; \, \{\lambda_k \, , \, \boldsymbol{\mu}_k \, , \, \Sigma_k\}^{(t)}\right)}, \tag{C5}$$

$$\Sigma_k^{(t+1)} = \frac{\sum_{n=1}^{N} \mathbb{P}\left(c^n = k \mid \boldsymbol{x}^n; \, \{\lambda^k, \underline{\mu}^k, \Sigma^k\}^{(t)}\right) \cdot \left(\boldsymbol{x}_n - \boldsymbol{\mu}_k^{(t+1)}\right)\left(\boldsymbol{x}_n - \boldsymbol{\mu}_k^{(t+1)}\right)^T}{\sum_{n=1}^{N} \mathbb{P}\left(c_n = k \mid \boldsymbol{x}_n; \, \{\lambda_k, \boldsymbol{\mu}_k, \Sigma_k\}^{(t)}\right)}, \tag{C6}$$

where c_n is the classification of the n-th cluster which could be any one of the K clusters. This is repeated until the parameters of the model have converged.

C2 Information Criterion

GMM needs an input hyperparameter K that sets the number of clusters that will be fitted to the data. GMM is relatively cheap to run, and so it is reasonable to run it with a large range of K and choose the K which best describes them. The often-used criterions are the Bayesian Information Criterion (BIC) (Eq. C7) and Akaike Information Criterion (AIC) (Eq. C8). They both essentially contain a term that measures the agreement of the model to the data, and have a penalty term for the number of parameters that have been used to achieve this (related to K). So we are looking for minima in AIC / BIC to guide our choice of K. There is no clear minimum for this dataset in K for $2 \le K \le 100$, which is typical of oceanographic applications due in part to the highly correlated nature of the data (e.g. (Sonnewald et al., 2019; Jones et al., 2019)). Because K is weakly constrained, we are able to select a lower value of K for ease of interpretation, having verified that it captures the large-scale structure of the data in PC space, which is suitable for our purposes. BIC and AIC take the forms:

$$\begin{aligned} \text{BIC} &= -2\mathcal{L}(K) + \eta_f(k)\log(N), & \text{(C7)} \\ \text{with} \quad \eta_f &= K - 1 + KD + \frac{KD(D-1)}{2}, \\ \text{AIC} &= 2K - 2\mathcal{L}, & \text{(C8)} \end{aligned}$$

where the log-likelihood is expressed as:

$$\mathcal{L} = \log[\mathbb{P}(X)] = \sum_{n=0}^{N-1} \log\left(\sum_{k=1}^{K} \lambda_k \, \mathcal{N}\left(\boldsymbol{x}_n \, ; \, \lambda_k \, , \, \boldsymbol{\mu}_k \, , \, \Sigma_k\right)\right) \tag{C9}$$

C3 Labelling the Data Set

Each data point is assigned a posterior probability distribution across the K clusters (Eq. C10). This uncertainty information is one of the useful features of GMM. The probability takes the form:

$$\mathbb{P}\left(c_n = k \mid \boldsymbol{x}_n \, ; \, \lambda_k \, , \, \boldsymbol{\mu}_k \, , \, \Sigma_k\right) = \frac{\lambda_k \, \mathcal{N}\left(\boldsymbol{x}_n \, ; \, \boldsymbol{\mu}_k \, , \, \sum_k\right)}{\sum_{k=1}^{K} \lambda_k \, \mathcal{N}\left(\boldsymbol{x}_n \, ; \, \boldsymbol{\mu}_k \, , \, \Sigma_k\right)}. \tag{C10}$$

475   To label a dataset, each data point is assigned a label from the cluster that it would be the most likely to be generated by, in
a statistical sense (Eq. C11).

$$\mathcal{C} = \arg\max_k \left(\mathbb{P}\left(c_n = k \mid \boldsymbol{x}_n \,;\, \lambda_k \,,\, \boldsymbol{\mu}_k \,,\, \Sigma_k \right), 1 : k \right) \tag{C11}$$

Author contributions. DJ designed the initial project, and both ST and DJ developed it further. ST wrote the software (Thomas, 2021), performed the analysis, and created the figures. AF proposed a significant improvement to the inter-class metric. EM and EP provided expert
480 guidance on fronts, Southern Ocean structure, and dynamics. ST and DJ wrote the initial manuscript, and all authors assisted with edits.

Competing interests. No competing interests

Acknowledgements. This work originated as a Natural Environment Research Council (NERC) Research Experience Placement (REP) project funded by the SPITFIRE Doctoral Training Partnership (reference NE/S007210/1). ST is supported by studentship 2413578 from the UKRI Centre for Doctoral Training in Application of Artificial Intelligence to the study of Environmental Risks (reference EP/S022961/1).
485 DJ is supported by a UKRI Future Leaders Fellowship (reference MR/T020822/1). DJ and ST also received funding from the NERC ACSIS project (reference NE/N018028/1). EP received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement 637770). The authors would like to thank Emma Boland, Peter Haynes, Guillaume Maze, and John Taylor for comments that improved the quality of this work.

References

490   Amante, C. and Eakins, B. W.: ETOPO1 Global Relief Model converted to PanMap layer format, PANGAEA, https://doi.org/10.1594/PANGAEA.769615, backup Publisher: NOAA-National Geophysical Data Center Type: data set, 2009.

Chapman, C. C.: New perspectives on frontal variability in the Southern Ocean, Journal of Physical Oceanography, 47, 1151–1168, https://doi.org/10.1175/JPO-D-16-0222.1, 2017.

Chapman, C. C., Lea, M.-A., Meyer, A., Sallée, J.-B., and Hindell, M.: Defining Southern Ocean fronts and their influence on biological and 495   physical processes in a changing climate, Nature Climate Change, https://doi.org/10.1038/s41558-020-0705-4, 2020.

Cushman-Roisin, B. and Beckers, J.-M.: Introduction to geophysical fluid dynamics: physical and numerical aspects, Academic press, 2011.

Dai, A. and Trenberth, K. E.: Estimates of Freshwater Discharge from Continents: Latitudinal and Seasonal Variations, Journal of Hydrometeorology, 3, 660–687, https://doi.org/10.1175/1525-7541(2002)003<0660:EOFDFC>2.0.CO;2, 2002.

Deacon, G.: The hydrology of the Southern Ocean, Cambridge University Press ; Macmillan, London; New York, 1937.

500   Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Halm, E. V., Isaksen, L., K\aallberg, P., Köhler, M., Matricardi, M., McNally, A. P., Monge Sanz, B. M., Morcrette, J. J., Park, B. K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J. N., and Vitart, F.: The ERA-Interim reanalysis: configuration and performance of the data assimilation system, Quarterly Journal of the Royal Meteorological Society, 137, 553–597, 505   https://doi.org/10.1002/qj.828, 2011.

Donohue, K. A., Tracey, K. L., Watts, D. R., Chidichimo, M. P., and Chereskin, T. K.: Mean Antarctic Circumpolar Current transport measured in Drake Passage: MEAN ACC TRANSPORT, Geophysical Research Letters, 43, 11,760–11,767, https://doi.org/10.1002/2016GL070319, 2016.

Fenty, I. and Heimbach, P.: Coupled sea ice–ocean-state estimation in the Labrador Sea and Baffin Bay, Journal of Physical Oceanography, 510   https://doi.org/10.1175/JPO-D-12-065.1, 2013.

Forget, G., Campin, J. M., Heimbach, P., Hill, C. N., Ponte, R. M., and Wunsch, C.: ECCO version 4: an integrated framework for non-linear inverse modeling and global ocean state estimation, Geoscientific Model Development, 8, 3071–3104, https://doi.org/10.5194/gmd-8-3071-2015, 2015.

Frenger, I., Munnich, M., Gruber, N., and Knutti, R.: Southern Ocean eddy phenomenology, Journal of Geophysical Research: Oceans, 120, 515   7413–7449, https://doi.org/https://doi.org/10.1002/2015JC011047, 2015.

Frolicher, T. L., Sarmiento, J. L., Paynter, D. J., Dunne, J. P., Krasting, J. P., and Winton, M.: Dominance of the Southern Ocean in Anthropogenic Carbon and Heat Uptake in CMIP5 Models, Journal of Climate, 28, 862–886, https://doi.org/10.1175/JCLI-D-14-00117.1, 2015.

Gaspar, P., Grégoris, Y., and Lefevre, J. M.: A simple eddy kinetic energy model for simulations of the oceanic vertical mixing: 520   Tests at station Papa and long-term upper ocean study site, Journal of Geophysical Research: Atmospheres, 95, 16 179–16 193, https://doi.org/10.1029/JC095iC09p16179, 1990.

Graham, R. M., Boer, A. M. d., Heywood, K. J., Chapman, M. R., and Stevens, D. P.: Southern Ocean fronts: Controlled by wind or topography?, Journal of Geophysical Research: Oceans, 117, https://doi.org/10.1029/2012JC007887, 2012.

Hallberg, R.: Using a resolution function to regulate parameterizations of oceanic mesoscale eddy effects, Ocean Modelling, 72, 92–103, 525   https://doi.org/10.1016/j.ocemod.2013.08.007, 2013a.

Hallberg, R.: Using a resolution function to regulate parameterizations of oceanic mesoscale eddy effects, Ocean Modelling, 72, 92–103, https://doi.org/10.1016/j.ocemod.2013.08.007, 2013b.

Hammond, M. D. and Jones, D. C.: Freshwater flux from ice sheet melting and iceberg calving in the Southern Ocean, Geoscience Data Journal, 3, 60–62, https://doi.org/10.1002/gdj3.43, 2016.

530 Herraiz-Borreguero, L. and Rintoul, S. R.: Subantarctic mode water: distribution and circulation, Ocean Dynamics, 61, 103–126, https://doi.org/10.1007/s10236-010-0352-9, 2011.

Hjelmervik, K. B. and Hjelmervik, K. T.: Detection of oceanographic fronts on variable water depths using empirical orthogonal functions, IEEE Journal of Oceanic Engineering, https://doi.org/10.1109/JOE.2019.2917456, 2019.

Houghton, I. A. and Wilson, J. D.: El Niño Detection Via Unsupervised Clustering of Argo Temperature Profiles, Journal of Geophysical 535 Research: Oceans, 125, https://doi.org/10.1029/2019JC015947, 2020.

Iudicone, D., Rodgers, K., Schopp, R., and Madec, G.: An exchange window for the injection of Antarctic Intermediate Water into the South Pacific, Journal of Physical Oceanography, 37, 31–49, https://doi.org/10.1175/JPO2985.1, 2007.

Jones, D. C. and Ito, T.: Gaussian mixture modeling describes the geography of the surface ocean carbon budget, pp. 108–113, NCAR, https://doi.org/10.5065/y82j-f154, 2019.

540 Jones, D. C., Meijers, A. J. S., Shuckburgh, E., Sallée, J.-B., Haynes, P., McAufield, E. K., and Mazloff, M. R.: How does Subantarctic Mode Water ventilate the Southern Hemisphere subtropics?, Journal of Geophysical Research - Oceans, 121, 6558–6582, https://doi.org/10.1002/2016jc011680, 2016.

Jones, D. C., Holt, H. J., Meijers, A. J., and Shuckburgh, E.: Unsupervised clustering of Southern Ocean Argo float temperature profiles, Journal of Geophysical Research: Oceans, 124, https://doi.org/10.1029/2018JC014629, 2019.

545 Kim, Y. S. and Orsi, A. H.: On the variability of Antarctic Circumpolar Current fronts inferred from 1992–2011 altimetry, Journal of Physical Oceanography, https://doi.org/10.1175/JPO-D-13-0217.1, 2014.

Large, W. and Yeager, S.: The global climatology of an interannually varying air‚Äìsea flux data set, Climate Dynamics, 33, 341–364, https://doi.org/10.1007/s00382-008-0441-3, 2009.

Le Bras, I. A.-A., Sonnewald, M., and Toole, J. M.: A Barotropic Vorticity Budget for the Subtropical North Atlantic Based on Observa-550 tions, Journal of Physical Oceanography, 49, 2781–2797, https://doi.org/10.1175/JPO-D-19-0111.1, publisher: American Meteorological Society, 2019.

Losch, M., Menemenlis, D., Campin, J.-M., Heimbach, P., and Hill, C.: On the formulation of sea-ice models. Part 1: Effects of different solver implementations and parameterizations, Ocean Modelling, 33, 129–144, https://doi.org/10.1016/j.ocemod.2009.12.008, 2010.

Mackie, E. J. B.: Southern Ocean Circulation and Frontal Dynamics from Cryosat-2 Along-track Radar Altimetry, Ph.D. thesis, University 555 of Bristol, School of Geographical Sciences, 2018.

Marshall, J. and Speer, K.: Closure of the meridional overturning circulation through Southern Ocean upwelling, Nature Geoscience, 5, 171–180, https://doi.org/10.1038/ngeo1391, 2012.

Marshall, J., Adcroft, A., Hill, C., and Perelman, L.: A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers, Journal of Geophysical Research, 102, 5753–5766, https://doi.org/10.1029/96JC02775, 1997a.

560 Marshall, J., Hill, C., Perelman, L., and Adcroft, A.: Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling, Journal of Geophysical Research, 102, 5733–5752, https://doi.org/10.1029/96JC02776, 1997b.

Maze, G.: Ocean Profile Classification Model in python, https://doi.org/10.5281/zenodo.3906236, 2020.

Maze, G., Mercier, H., Fablet, R., Tandeo, P., Radcenco, M. L., Lenca, P., Feucher, C., and Le Goff, C.: Coherent heat patterns revealed by unsupervised classification of Argo temperature profiles in the North Atlantic Ocean, Progress in Oceanography, https://doi.org/10.1016/j.pocean.2016.12.008, 2017.

Mazloff, M. R., Heimbach, P., and Wunsch, C.: An Eddy-Permitting Southern Ocean State Estimate, Journal of Physical Oceanography, 40, 880–899, https://doi.org/10.1175/2009jpo4236.1, 2010.

McLachlan, G. J. and Basford, K. E.: Mixture models: Inference and applications to clustering, vol. 38, M. Dekker New York, 1988.

Met Office: Cartopy: a cartographic python library with a matplotlib interface, Exeter, Devon, http://scitools.org.uk/cartopy, 2010 - 2015.

Mikaloff-Fletcher, S., Gruber, N., Jacobson, A. R., Doney, S. C., Dutkiewicz, S., Gerber, M., Follows, M., Joos, F., Lindsay, K., Menemenlis, D., Mouchet, A., Müller, S. A., and Sarmiento, J. L.: Inverse estimates of anthropogenic CO2 uptake, transport, and storage by the ocean, Global Biogeochemical Cycles, 20, https://doi.org/10.1029/2005gb002530, 2006.

Morrison, A. K., Frölicher, T. L., and Sarmiento, J. L.: Upwelling in the Southern Ocean, Physics Today, 68, 27–32, https://doi.org/10.1063/PT.3.2654, 2015.

Nakayama, Y., Menemenlis, D., Zhang, H., Schodlok, M., and Rignot, E.: Origin of Circumpolar Deep Water intruding onto the Amundsen and Bellingshausen Sea continental shelves, Nature Communications, 9, 3403, https://doi.org/10.1038/s41467-018-05813-1, 2018.

Orsi, A., Whitworth, T., and Nowlin, W.: On the meridional extent and fronts of the Antarctic Circumpolar Current, Deep Sea Research Part I, 42, 641–673, https://doi.org/10.1016/0967-0637(95)00021-W, 1995.

Pauthenet, E., Roquet, F., Madec, G., and Nerini, D.: A linear decomposition of the Southern Ocean thermohaline structure, Journal of Physical Oceanography, 47, 29–47, https://doi.org/10.1175/JPO-D-16-0083.s1, 2017.

Pauthenet, E., Roquet, F., Madec, G., Guinet, C., Hindell, M., McMahon, C. R., Harcourt, R., and Nerini, D.: Seasonal Meandering of the Polar Front Upstream of the Kerguelen Plateau, Geophysical Research Letters, 45, 9774–9781, https://doi.org/10.1029/2018GL079614, 2018.

Pauthenet, E., Roquet, F., Madec, G., Sallee, J.-B., and Nerini, D.: The Thermohaline Modes of the Global Ocean, Journal of Physical Oceanography, 49, 2535–2552, https://doi.org/10.1175/JPO-D-19-0120.1, 2019.

Pauthenet, E., Sallee, J.-B., Schmidtko, S., and Nerini, D.: Seasonal Variation of the Antarctic Slope Front Occurrence and Position Estimated from an Interpolated Hydrographic Climatology, Journal of Physical Oceanography, 51, 1539–1557, https://doi.org/10.1175/JPO-D-20-0186.1, 2021.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, Journal of Machine Learning Research, 12, 2825–2830, 2011.

Pollard, R., Lucas, M., and J, R.: Physical controls on biogeochemical zonation in the Southern Ocean, Deep Sea Research Part II, 49, 3289–3305, https://doi.org/10.1016/S0967-0645(02)00084-X, 2002.

Rintoul, S., Hughes, C., and Olbers, D.: The Antarctic Circumpolar Current system, International Geophysics Series, pp. 271–302, 2001.

Rosso, I., Mazloff, M. R., Talley, L. D., Purkey, S. G., Freeman, N. M., and Maze, G.: 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, https://doi.org/10.1029/2019JC015877, 2020.

Sallee, J., Speer, K., Rintoul, S., and Wijffels, S.: Southern Ocean thermocline ventilation, Journal of Physical Oceanography, 40, 509–529, https://doi.org/10.1175/2009JPO4291.1, 2010.

Sallee, J.-B., Matear, R. J., Rintoul, S. R., and Lenton, A.: Localized subduction of anthropogenic carbon dioxide in the Southern Hemisphere oceans, Nature Geoscience, 5, 579–584, https://doi.org/10.1038/ngeo1523, 2012.

Siegelman, L., OToole, M., Flexas, M., Riviere, P., and Klein, P.: Submesoscale ocean fronts act as biological hotspot for southern elephant seal, Scientific Reports, 9, 5588, https://doi.org/10.1038/s41598-019-42117-w, 2019.

Sobel, I. and Feldman, G.: A 3x3 isotropic gradient operator for image processing, Pattern Classification and Scene Analysis, pp. 271–272, 1968.

Sokolov, S. and Rintoul, S. R.: Structure of Southern Ocean fronts at 140¬∞E, Journal of Marine Systems, 37, 151 – 184, https://doi.org/10.1016/S0924-7963(02)00200-2, 2002.

Sokolov, S. and Rintoul, S. R.: Circumpolar structure and distribution of the Antarctic Circumpolar Current fronts: 1. Mean circumpolar paths, Journal of Geophysical Research: Atmospheres, 114, 3675, https://doi.org/10.1029/2008JC005108, 2009.

Sonnewald, M., Wunsch, C., and Heimbach, P.: Unsupervised Learning Reveals Geography of Global Ocean Dynamical Regions, Earth and Space Science, pp. 1–32, https://doi.org/10.1029/2018EA000519, 2019.

Sonnewald, M., Dutkiewicz, S., Hill, C., and Forget, G.: Elucidating ecological complexity: Unsupervised learning determines global marine eco-provinces, Science Advances, 6, eaay4740, https://doi.org/10.1126/sciadv.aay4740, 2020.

Stammer, D., Wunsch, C., Giering, R., Eckert, C., Heimbach, P., Marotzke, J., Adcroft, A., Hill, C. N., and Marshall, J.: Global ocean circulation during 1992-1997, estimated from ocean observations and a general circulation model, Journal of Geophysical Research: Oceans, 107, https://doi.org/10.1029/2001JC000888, 2002.

Talley, L.: Closure of the Global Overturning Circulation Through the Indian, Pacific, and Southern Oceans: Schematics and Transports, Oceanography, 26, 80–97, https://doi.org/10.5670/oceanog.2013.07, 2013.

Thomas, S. D. A.: so-wise/so-fronts: Defining Southern Ocean fronts using unsupervised classification, https://doi.org/10.5281/zenodo.5500666, 2021.

Thompson, A. F. and Sallée, J.-B.: Jets and topography: Jet transitions and the impact on transport in the Antarctic Circumpolar Current, Journal of physical Oceanography, 42, 956–972, https://doi.org/10.1175/JPO-D-11-0135.1, 2012.

Thompson, A. F., Haynes, P. H., Wilson, C., and Richards, K. J.: Rapid Southern Ocean front transitions in an eddy-resolving ocean GCM, Geophysical Research Letters, 37, https://doi.org/10.1029/2010GL045386, 2010.

Thompson, A. F., Stewart, A. L., Spence, P., and Heywood, K. J.: The Antarctic Slope Current in a Changing Climate, Reviews of Geophysics, 56, 741–770, https://doi.org/10.1029/2018RG000624, 2018.

Thompson, A. F., Speer, K. G., and Schulze Chretien, L. M.: Genesis of the Antarctic Slope Current in West Antarctica, Geophysical Research Letters, 47, https://doi.org/10.1029/2020GL087802, 2020.

Verdy, A. and Mazloff, M. R.: A data assimilating model for estimating Southern Ocean biogeochemistry, Journal of Geophysical Research: Oceans, 122, 6968–6988, https://doi.org/10.1002/2016jc012650, 2017.

Vernet, M., Geibert, W., Hoppema, M., Brown, P. J., Haas, C., Hellmer, H. H., Jokat, W., Jullion, L., Mazloff, M., Bakker, D. C. E., Brearley, J. A., Croot, P., Hattermann, T., Hauck, J., Hillenbrand, C.-D., Hoppe, C. J. M., Huhn, O., Koch, B. P., Lechtenfeld, O. J., Meredith, M. P., Naveira Garabato, A. C., Nothig, E.-M., Peeken, I., Rutgers van der Loeff, M. M., Schmidtko, S., Schroder, M., Strass, V. H., Torres-Valdes, S., and Verdy, A.: The Weddell Gyre, Southern Ocean: present knowledge and future challenges, Reviews of Geophysics, https://doi.org/10.1029/2018RG000604, 2019.

Vincent, O. R., Folorunso, O., et al.: A descriptive algorithm for sobel image edge detection, in: Proceedings of informing science & IT education conference (InSITE), vol. 40, pp. 97–107, Informing Science Institute California, 2009.

Wunsch, C. and Heimbach, P.: Practical global oceanic state estimation, Physica D: Nonlinear Phenomena, 230, 197 – 208, https://doi.org/10.1016/j.physd.2006.09.040, 2007.