CAMBRIDGE
UNIVERSITY PRESS

RESEARCH ARTICLE

A Novel Heuristic Method for Detecting Overfit in Unsupervised Classification of Climate Model Data

Emma J. D. Boland[1] ⓘ, Erin Atkinson[2] ⓘ and Dani C. Jones[1] ⓘ

[1]British Antarctic Survey, Cambridge, CB3 0ET, United Kingdom
[2]Department of Physics, University of Toronto, Toronto, M5S 1A7, Canada

Abstract

Unsupervised classification is becoming an increasingly common method to objectively identify coherent structures within both observed and modelled climate data. However, in most applications using this method, the user must choose the number of classes into which the data are to be sorted in advance. Typically, a combination of statistical methods and expertise is used to choose the appropriate number of classes for a given study, however it may not be possible to identify a single 'optimal' number of classes. In this work, we present a heuristic method, the Ensemble Difference Criterion, for unambiguously determining the maximum number of classes supported by model data ensembles. This method requires robustness in the class definition between simulated ensembles of the system of interest. For demonstration, we apply this to the clustering of Southern Ocean potential temperatures in a CMIP6 climate model, and show that the data supports between four and seven classes of a Gaussian Mixture Model.

> **Impact Statement**
>
> One method that climate scientists use to find structures in environmental data is unsupervised classification, where a machine learning model sorts data, such as temperatures from different parts of the ocean, into groups or classes. This paper proposes a new way of determining if an unsupervised classification model has 'over-fit' the data available, which means it is using too many classes. This method will be useful for climate scientists using unsupervised classification for studying climate model data, as current criteria commonly used often don't provide an upper limit on the number of classes.

1. Introduction

Studies using observational data of the climate system often seek to find structures within that data. Specifically, classifying complex data into coherent groups with similar properties is often the first step in a study that may seek to, for example, find borders between groups, describe the behaviour of each group, explain the differences between the groups, or track changes in groups over time.

Unsupervised classification methods, i.e. methods that do not know a-priori what the properties of these groups might be, have proven adept at identifying coherent spatial structures within climate data, even when no spatial information is supplied to the algorithm. In studies of ocean and atmospheric data, two commonly used unsupervised classification methods are k-means (Liu et al., 2021; Houghton

and Wilson, 2020; Yuchechen et al., 2020; Hjelmervik et al., 2015; Hjelmervik and Hjelmervik, 2014, 2013; Solidoro et al., 2007; Sonnewald et al., 2019) and Gaussian mixture modelling (GMM) (Fahrin et al., 2022; Sugiura, 2021; Zhao et al., 2021; Crawford, 2020; Jones et al., 2019; Maze et al., 2017a; Tandeo et al., 2014; Hannachi, 2007; Hannachi and O'Neill, 2001). K-means attempts to find coherent groups by "cutting" the abstract feature space using hyperplanes, whereas GMM attempts to represent the underlying covariance structure in abstract feature space using a linear combination of multi-dimensional Gaussian functions.

However, when dealing with highly-correlated data such as climate model output, the number of classes into which the data should be divided is not necessarily obvious; it remains a relatively unconstrained free parameter, and the classification methods themselves do not typically provide a way of choosing the 'best' or most appropriate number of classes to describe a given dataset. In practice, most studies repeat the training of the model multiple times for different numbers of classes before applying a statistical criterion to compare the models. Statistical criteria alone do not always show a clear single 'best' model, and so domain expertise is often additionally used to determine the appropriate number of classes for a given study.

The most commonly used statistical criterion is the minimum in the Bayesian Information Criterion (BIC, Schwarz, 1978), used in Fahrin et al. (2022); Sugiura (2021); Zhao et al. (2021); Jones et al. (2019); Maze et al. (2017b); Hjelmervik et al. (2015); Hjelmervik and Hjelmervik (2014, 2013); Sonnewald et al. (2019). The BIC is comprised of two terms: a term that rewards the statistical likelihood of the model and a term that penalises overfitting. One looks for a minimum or plateau in BIC in order to find the number of classes that maximises statistical likelihood without overfitting.

Other common statistical methods used to compare clustering results by class number include the following (a non-exhaustive list):

- Akaike Information Criterion (AIC, Akaike, 1973), used in Tandeo et al. (2014); Sonnewald et al. (2019),
- Caliński–Harabasz criterion (CHC, Caliński and Harabasz, 1974), used in Yuchechen et al. (2020),
- Davies-Bouldin index (DBI, Davies and Bouldin, 1979), used in Solidoro et al. (2007),
- Silhouette coefficient (Rousseeuw, 1987), used in Houghton and Wilson (2020),
- criteria based on confidence intervals, used in Hannachi (2007).
- other scores based on log-likelihoods, used in Hannachi and O'Neill (2001); Crawford (2020),

Studies such as Liu et al. (2021) use physical justifications only.

Most of the studies cited above use unsupervised classification on observational data. One natural extension of a study based on observational data is to carry out a similar study based on data from numerical simulations. Possibilities include assessing the simulation's fidelity, as well as greatly expanding the scope of the questions that can be asked, such as how robust clustering results are between different ensemble members, and how cluster properties vary in different simulated scenarios.

Given the rapid growth of clustering studies in recent years, and the many possibilities outlined above for motivating the use of simulated data, we believe that there will be a growing number of studies applying clustering techniques to simulated climate data. The use of multiple ensemble members is often required to reduce model uncertainty due to internal variability (Deser et al., 2020), especially at regional or decadal scales where there is generally a lower signal to noise ratio (see, for example, Dosio et al., 2019; Yeager et al., 2018). We present here a novel heuristic method for determining the maximum number of classes when classifying multi-member ensembles of climate model data, the Ensemble Difference Criterion (EDC).

We expect the clusters to be robust between ensemble members if the structures they represent are a forced response rather than dependent on internal variability. As mentioned above, multiple ensemble members may be required for a variety of reasons. One of these reasons is that the forced response is relatively weak compared with internal variability. Whilst in some cases it may be possible to fully

capture this internal variability by sub-sampling the same simulation at different points in time, if that system is in quasi-equilibrium, this may not be sufficient for a system in transition due to the effects of climate change.

The rest of the manuscript is structured as follows: in section 2 we first describe the dataset used for demonstration (Southern Ocean temperature profiles from a CMIP6 climate model) and classification model (a Gaussian Mixture Model or GMM). In section 2.3 we define a number of criteria for determining the optimal number of classes, including the EDC. In section 3 we briefly present the results of an eight class model and for reference, compare them with the results of a similar study based on observations. We compare the results of the selected criteria for choosing the number of classes in section 4, discuss how to interpret the results in section 5 and finish in section 6 with a summary.

2. Data and methods

This section is arranged as follows: firstly, we describe the dataset of temperature profiles from the UK-ESM CMIP6 simulations (section 2.1), then we describe the form of the GMMs we fit to those data (section 2.2). Finally, in section 2.3, we describe the different methods we will compare to select the number of classes: the Bayesian Information Criterion (BIC), the silhouette score, and the newly defined Ensemble Difference Criterion.

2.1. The dataset

The simulation dataset is the ocean potential temperature data from the historical (2001-2014) and SSP3-7.0 (2015-2017) UK-ESM1.0 contribution to CMIP6 (Sellar et al., 2019). The ocean component is the N96ORCA1 configuration of NEMO (Kuhlbrodt et al., 2018), with a horizontal resolution of 1° and 75 vertical levels with thicknesses from 1m (at the surface) to 200m (at depth). We don't expect clustering results to necessarily be the same between different models as the results are strongly state dependant and, in general, CMIP6 models vary quite widely in their representation of the Southern Ocean (see, e.g. Beadling et al., 2020; Williams et al., 2023). We use UK-ESM as we have connections to the UK climate modelling community, but the choice of model is irrelevant for this application, which is meant illustrate a methodology - we could have used any other multi-ensemble model.

We take data from 2001-2017 inclusive (17 years) in order to allow comparison with Jones et al. (2019), who use the same time period. All future scenarios are described in detail in Meinshausen et al. (2020) – we choose the future SSP3-7.0 scenario as it has the most overlapping ensemble members with the historical period for UK-ESM1.0, and there is little to no divergence in the various scenarios in the three years used (Walton, 2020).

Seven ensembles were included, those with realisation number r1, r2, r3, r4, r8, r16 and r19. A description of the UK-ESM variant-IDs can be found in https://ukesm.ac.uk/cmip6/variant-id/. Each ensemble represents a different run or *realisation* of the Earth's climate system, forced identically from 1850 onwards, but initialised from different points in the pre-industrial control run, chosen to sample the range of multi-decadal variability of the ocean, see Sellar et al. (2019) for details.

As in Jones et al. (2019), we use ocean potential temperature-depth profiles south of 30°S that extend to at least 2000m. The dataset is thus a set of N vectors $\{\mathbf{T}_i\}$, one for each surface grid cell that fits the above criteria. This results in a total of approximately 4.5 million profiles. To keep computation costs reasonable, we use a subset of 3,000 randomly chosen profiles per month, with each month's profiles un-correlated with any other's. This results in a total of 612 thousand profiles (13.5% of the full dataset) used to train the classification model, which is then classify all profiles. The results do not change if we increase the number of profiles per month, which suggests the training set is well chosen to represent the full dataset.
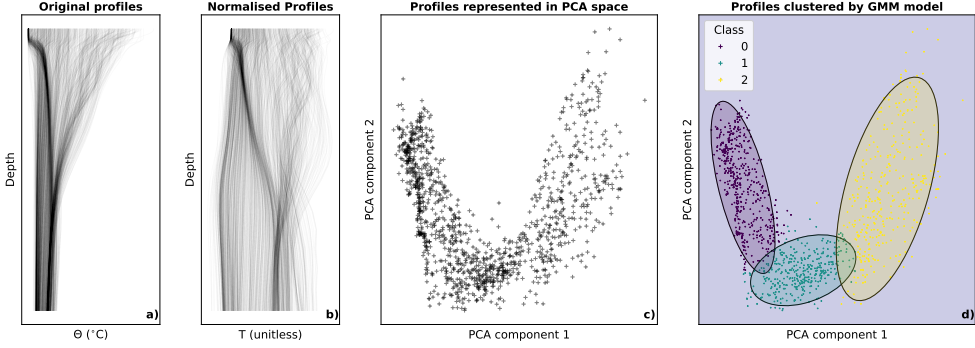
Figure 1. *An overview of the classification process using 1200 random Southern Ocean profiles from the year 2000. Raw profiles (a), are normalised (b) so each depth level has a mean of zero and a standard deviation of one. c) The profiles are transformed into PCA space. Two PCs are shown here; each point corresponds to one temperature profile. d) The data are classified according to a GMM. Three components are chosen in this figure, with an equi-probability surface shown for each, separately.*

2.2. Classification scheme

Figure 1 gives an overview of the classification process, which is carried out independently for each ensemble member described in section 2.1. As an initial step, the raw dataset is normalised so that each depth level has mean zero and standard deviation one.

Each profile is $L = 30$-dimensional (determined by the number of depth levels between 100m and 2km in the model). The large number of dimensions makes computation expensive and hinders the usefulness of distance metrics on the normalised dataset. Principal component analysis (PCA) (Abdi and Williams, 2010) is employed as a dimensionality-reduction step. The dataset in PC space is given by:

$$\mathbf{x}_i = \mathbf{P}\mathbf{T}_i, \tag{1}$$

where $\mathbf{T}_i$ is the ith temperature profile in depth space, $\mathbf{x}_i$ is the same profile in PC space, and P is the projection matrix. In the following analysis, $M = 3$ principal components are used - the first three PCs are sufficient to explain 99.8% of the variance in the training dataset.

The training dataset is fitted to a Gaussian mixture model (GMM) (Anderson and Moore, 2012) as in Jones et al. (2019). This assumes that the data points $\mathbf{x}_i$ are sampled from a joint distribution over profiles $\mathbf{x}$ and class k (a *Gaussian mixture*), with probability

$$\rho_k(\mathbf{x}) = \lambda_k \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_k, \Sigma_k), \tag{2}$$

where

$$\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_k, \Sigma_k) = \frac{1}{\sqrt{(2\pi)^M |\Sigma_k|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^T \Sigma_k^{-1} (\mathbf{x} - \boldsymbol{\mu}_k)\right] \tag{3}$$

is a multivariate Gaussian distribution with mean $\boldsymbol{\mu}_k$ and covariance matrix Σ_k, and

$$\lambda_k \in [0, 1], \sum_{k=1}^{K} \lambda_k = 1 \tag{4}$$

is any distribution over the set $\{k_i\} = [1, K]$ where K is the total number of classes. The marginal distribution of the profiles $\mathbf{x}$ is

$$\rho(\mathbf{x}) = \sum_{k=1}^{K} \rho_k(\mathbf{x}). \tag{5}$$

This work uses the GMM classification model from the `scikit-learn` Python package (Pedregosa et al., 2011). This implements an expectation-maximisation algorithm to learn the set of parameters $\{\lambda_k, \boldsymbol{\mu}_k, \Sigma_k\}$ from the training dataset. Each data point is then assigned to a single class according to the Gaussian distribution it most likely belongs to, i.e. the class with the greatest probability $\rho_k(\mathbf{x}_i)$:

$$k_i = \arg\max_k \rho_k(\mathbf{x}_i). \tag{6}$$

The final result of the classification is a function that maps from a temperature profile to a class index, which is used to classify the full dataset, including the training data, for that ensemble member:

$$f : \mathbf{T}_i \mapsto k_i, \tag{7}$$

as well as providing the probability distributions ρ_k. This allows one to define the overall certainty of the model as the sum of the log-likelihoods $\mathcal{L}(K)$ over all N profiles:

$$\mathcal{L}(K) = \sum_i^N \log(\rho(\mathbf{x}_i)), \tag{8}$$

which is a measure of how well a given GMM represents the data.

2.3. Definition of Metrics to determine the Optimal Number of Classes

We here define three criteria for determining the optimal number of classes that we will use to analyse our results: the Bayesian Information Criterion (BIC), the silhouette coefficient, and a novel method, the Ensemble Difference Criterion (EDC).

2.3.1. BIC

The BIC is a cost metric that rewards a high total likelihood of the data $\mathcal{L}(K)$, while regularising based on the estimated number of free parameters in the GMM $N_f(K)$.

$$\text{BIC}(K) = -2\mathcal{L}(K) + N_f(K)\ln(N), \tag{9}$$

$$N_f(K) = K - 1 + KM + \frac{1}{2}KM(M-1). \tag{10}$$

A low BIC implies that the model fits the data well, while also avoiding overfitting. An optimal value of K can be determined by, for example, a minimum in the BIC curve, or for a value of K above which there are diminishing returns, determined by looking at the ΔBIC curve, that is the change in score for adding an extra class. The AIC is similarly defined but weights the number of free parameters differently.

2.3.2. Silhouette coefficient

The silhouette coefficient is a measure of how well the clusters in a model are separated from each other. For each classified sample, the sample silhouette coefficient S_i is based on comparing the mean distance between point i and all other points in the same cluster (a_i), and the mean distance between point i and all points in the "nearest" cluster (b_i):

$$S_i = \frac{b_i - a_i}{max(a_i, b_i)}. \tag{11}$$

Note that in our example, the distance between samples is calculated in PCA space, the space on which the samples are classified (see figure 1d), and not in geographical space.

S_i can vary from -1 (a 'poor' score, indicating the sample is closer to samples in the neighbouring class than its own) to 1 (a 'good' score, indicating the sample is far from samples in the neighbouring class), whereas a score of 0 indicates overlapping classes. The overall silhouette coefficient for a given model is the mean of the sample silhouettes, i.e. $\text{SIL} \equiv 1/N \sum_i^N S_i$. Maximising the silhouette coefficient therefore desirable for applications where well-separated classes is a priority.

2.3.3. Ensemble Difference Criterion

If multiple realisations of the dataset are available, for example different ensemble members of UK-ESM, the simulated data can be leveraged to determine a physically-motivated and unambiguous value of the maximum number of classes.

Firstly we define a map between class indices in different ensembles

$$f^{E_1 \rightarrow E_2}(k) = \arg\min_l ||\mathbf{T}_k^{E_1} - \mathbf{T}_l^{E_2}||^2, \tag{12}$$

where $\mathbf{T}_i^{E_j}$ is the mean temperature profile of the ith class in ensemble E_j, such that $f^{E_1 \rightarrow E_2}(k)$ is the index of the class in ensemble E_2 that is most similar to the class with index k in ensemble E_1.

Here 'most similar' is defined as the profile that minimises the Euclidean distance in temperature/depth space, but the method is not dependent on this exact choice as long as it is allows one to identify correspondences between classes from different GMMs. Other maps could include, for example, the Euclidean distance between class centres in PCA space, the variance of each depth level, or the mean temperature of each class. We tested calculating the Euclidean distances in PCA space for the results reported here (in sections 3 and 4.3), and most mappings were identical to those using temperature/depth space, with occasional differences for one class in a given ensemble member (not shown).

The Ensemble Difference Criterion (EDC) requires that, having defined $f^{E_1 \rightarrow E_2}(k)$:

$$\forall_{x,y} f^{E_x \leftrightarrow E_y}, \ x, y \in E \tag{13}$$

where x and y are pairs of ensemble members, and $\leftrightarrow$ indicates the map is bijective, that is each element in ensemble x maps to one and only one member in ensemble y, and vice versa. If this is satisfied, the model passes the EDC, if it is not, the EDC is violated. In practice, to find the maximum number of classes allowed by the EDC, one successively increases K until the EDC is violated.

3. Comparison with Observations

We first compare an eight-class GMM fit on the model data, as described in section 2, with the results of Jones et al. (2019), who fit an eight class GMM on Argo data from the same time period. We present this in order to provide the context for discussing the choice of class number in section 4, and leave detailed discussion of the physical interpretation of these results to another study.

Whilst an ideal model would exactly reproduce the 'true' Southern Ocean, and would therefore have the same number of appropriate classes, we do not expect the number of classes most appropriate for the UK-ESM data is necessarily the same as the number of classes appropriate for the Argo data. This is partly because the Argo data is under-sampled in many regions (see Jones et al., 2019), whereas the UK-ESM data is not, inevitably leading to datasets with different properties. Even if the observations had better coverage, the Southern Ocean is imperfectly replicated by any climate model and therefore does not necessarily support the same number of classes. Nonetheless, it is informative to asses how well the results of observations and modelling studies match, in order to understand in which ways results from the modelled data might apply to the real world.

In order to compare the classification across ensembles, it is first necessary to align the classes between the ensembles, as the GMM labels the classes 1-8 at random each time it is fit. Thus, for
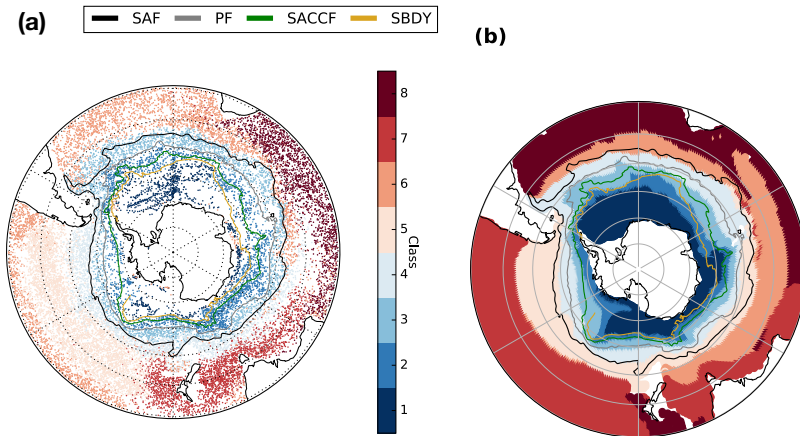
Figure 2. *Comparison of the class assignments from Argo (a, Jones et al., 2019) and the ensemble mode class from UK-ESM (b), for 2001-2017. See text for further details. The ocean fronts from Kim and Orsi (2014) are marked by the coloured lines as labelled.*

each ensemble member, the classes are first sorted from coldest to warmest by the class mean profile temperature value at the shallowest level (5m). Figure 2a) shows the ensemble mode (i.e. most common) class for each point in the domain, across the dataset time period (2001-2017) and all seven ensemble members after sorting by surface temperature (which are shown individually in the supplementary information, figures S1 and S2).

Note that, for eight classes, sorting by surface temperature does not always result in alignment between the classes that are closest to one another when considering the overall Euclidean distance between mean profiles at all depths, see figure 7 and the discussion in section 4.

Despite this, there is relatively little inter-ensemble variation in the mean temperature profiles for the classes sorted in this manner: figure 3b) shows the ensemble and time mean profiles (solid lines), with the standard deviations of the time mean profiles from each ensemble (dotted lines) smaller than the ensemble mean of the temporal standard deviations (dashed lines). For comparison, figure 3a) shows the mean and standard deviations of the temperature profiles from Jones et al. (2019), sorted by the shallowest temperature.

Comparing the mean profiles between the two observationally-derived GMM and the model-derived GMM using the mean Euclidean distance in temperature/depth space (see section 4.3 for more details), we can map from each class in one model to the class it is most similar to in the other model (blue and red arrows in figure 3).

The two GMMs show some structural similarities with the coldest four classes (1-4) largely being circumpolar, with boundaries somewhat aligned with the climatological ACC fronts from Kim and Orsi (2014), marked in figure 2. The warmer classes (5-8) show a more basin dominated structure in both models. There are, however, differences in the detailed structures: within the Southern Boundary (yellow line) the UK-ESM model has two distinct classes, 1 and 2, whereas the Argo model has only one (class 1). This could be due to the extreme sparsity of Argo data close to the continent, as well as the fact it is likely to be from Austral summer only (when the region is ice-free), whereas the UK-ESM data is year-round.

Both models then have two circumpolar classes that lie roughly between the Southern Boundary and the SAF - class 3 & 4 in UK-ESM and 2 & 3 in Argo, with mean profiles that match well (see arrows in figure 3). Both class 5 in UK-ESM and Class 4 in Argo extend in a roughly zonal band from the South West coast of South America westward to the south Indian ocean, and show similar mean profiles. This
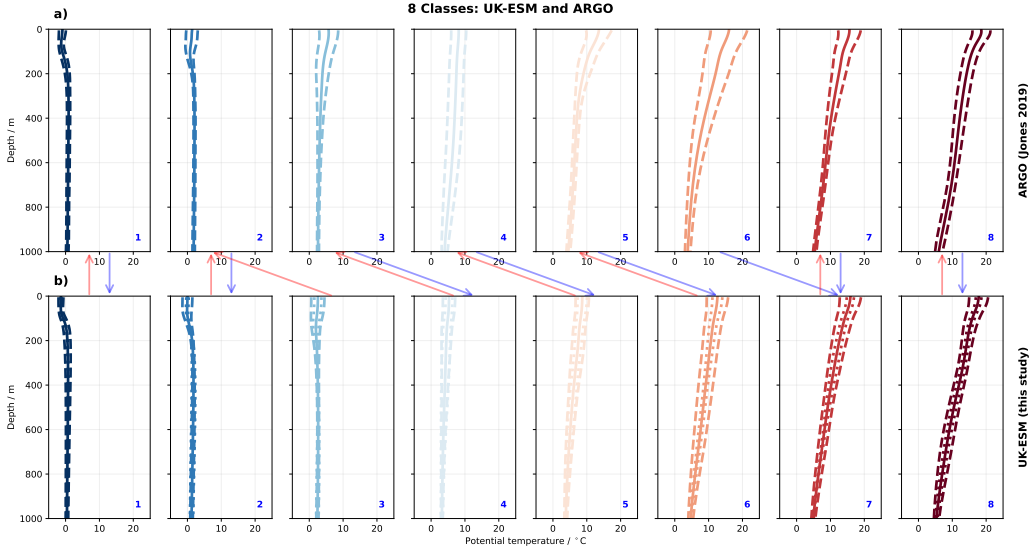
Figure 3. *The mean temperature profiles for an 8 class GMM from Jones et al. (2019) (a) compared with the ensemble mean profiles from this study (b). Dashed lines indicate the standard deviation across samples for the Argo data, and the ensemble average of the temporal standard deviations for the UK-ESM data. The dotted line indicates the ensemble standard deviation in the temporal mean profiles for the UK-ESM data. Red arrows indicate the mapping from a UK-ESM class to the closest class in the Argo data, and the blue arrows indicate the same for the Argo classes and the closest classes in the UK-ESM data.*

suggests the structure of the ACC is a strong bound on the water properties, producing similar class boundaries in physical space and similar mean profiles in both models.

The geographic distributions of the warmest classes is most different between the two models, with classes 5 to 8 in the Argo model mostly confined to a single basin or sub-region. The three sub-tropical class in the UK-ESM model show a more cross-basin structure. This could be because there are fewer sub-tropical classes in the UK-ESM model than the Argo model, (three vs four). This can be tested by increasing the number of classes in the UK-ESM model to nine. This does lead to an extra sub-tropical class in the model, but the structure is no more similar to the Argo data than for the eight class model (see figures S3 and S4 in the supplementary info).

The difference in structures of the sub-tropical classes between the two models could also be because the temperature profiles in this region are too similar to cluster well. This is indicated by the very similar structure of the warmest mean profiles in figure 3, and may indicate that the datasets do not robustly support eight unique classes. A comparison of various methods to find a robust choice of classes is carried out in the next section.

4. Choosing the number of classes

The key free parameter in the fitting of a GMM is the total number of classes K. Choosing a value of K presents a trade-off: low values will exhibit fewer structures and may fail to capture some of the underlying covariance structure whereas high values risk overfitting, losing physical interpretability and predictability. We here compare the results of three criteria, defined in section 2.3: the Bayesian Information Criterion (BIC), the silhouette coefficient, and the Ensemble Difference Criterion (EDC).
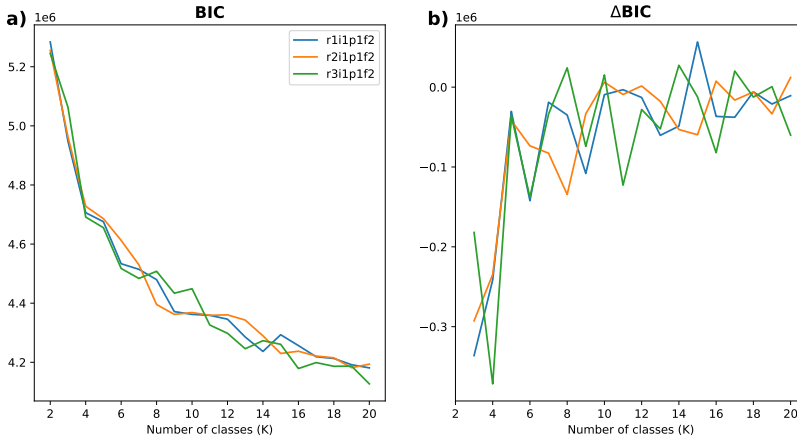
4.1. Bayesian information criterion

Figure 4. *The a) BIC curves and b) ΔBIC curves for three of the ensembles.*

The BIC and ΔBIC curves for three ensemble members, clustered into a range of class numbers, as described in section 2.2, is shown in figure 4. We only calculated the curve for the first three members due to the computational expense of fitting multiple models. There is no clear minimum in the BIC curve (figure 4a) for up to 20 classes, so no clearly optimal value of K. The ΔBIC curve (figure 4b) shows that adding an extra class leads to the largest improvements for 3 and 4 classes, then the gains begin to diminish for > 5 classes. However, there is large variability between and within ensemble members, and some class values lead to an *increase* in the BIC for a given ensemble member.

For this study, the BIC and AIC were almost identical for all clustering models studied (not shown), indicating the dominance of the log-likelihood term over the free-parameter terms in equations 9.
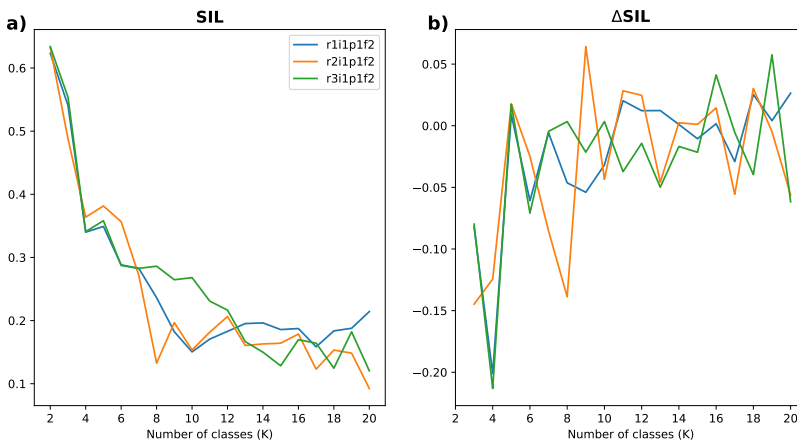
4.2. Silhouette coefficient

Figure 5. *The a) SIL curves and b) ΔSIL curves for three of the ensembles.*

The SIL and ΔSIL curves for three ensemble members, clustered into a range of class numbers, as described in section 2.2, are shown in figure 5. For each ensemble member, we calculate the SIL for a random subset of 10,000 samples (approx 1.6%) from the training set, due to the computational intensity of the calculation. Conversely to the BIC (section 4.1), it is desirable to *maximise* the SIL score.

The overall maximum SIL is for two classes, and there appears to be a local maximum at five classes for the three ensembles shown. As with the BIC, there is large variability both between and within ensemble members, especially at higher numbers of classes.
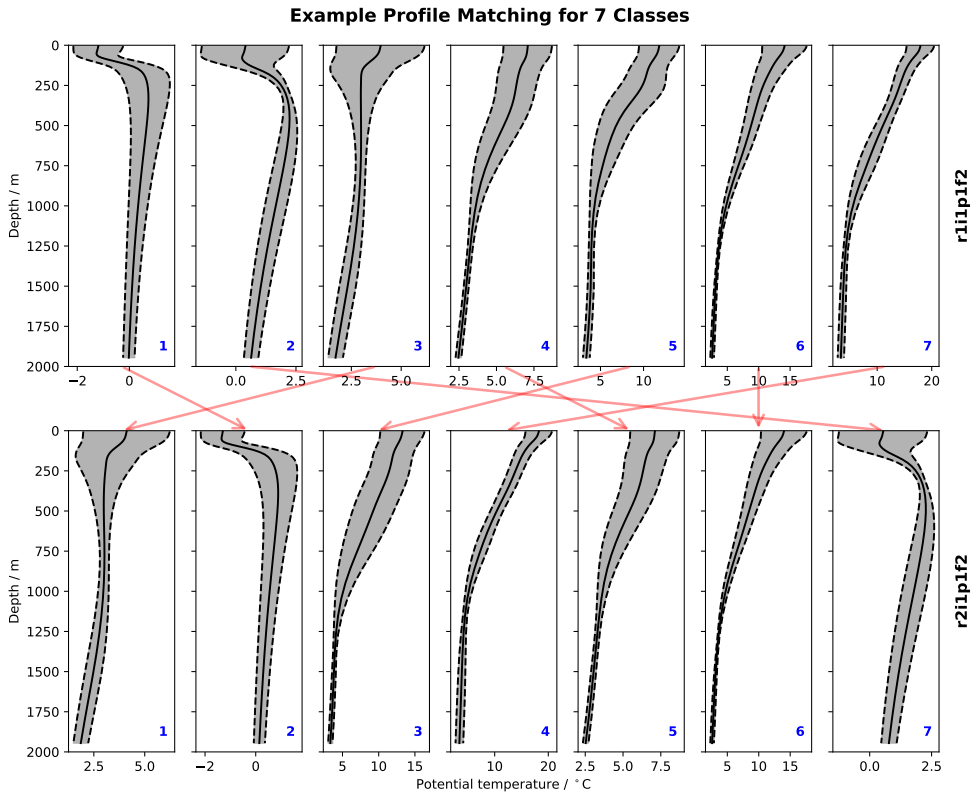
4.3. Ensemble Difference Criterion

Figure 6. *An example of using the Ensemble Difference Criterion (EDC) to match classes resulting from a 7-class GMM of two different ensemble members (r1i1p1f2 and r2i1p1f2).*

An example of using the EDC for the UK-ESM data is shown in figures 6 and 7. Figure 6 shows an example of matching between 7-class GMMs fit to two ensemble members. The matching is successful because the map is bijective for all ensemble pairs (only one is shown here): each class in one model is matched by only one class in any other model.

Figure 7 shows an example of matching between 8-class GMMs to the same two ensemble members. The matching is not bijective in this case. For example, classes 6 and 7 in the r1i1p1f2 model are both closer to class 1 in the r2i1p1f2 model than to any other. The mapping fails for all pairs between the seven ensembles (only one is shown here), with some mappings finding three classes match to the same referent. The EDC thus determines 7 classes as the maximum supported by this dataset.
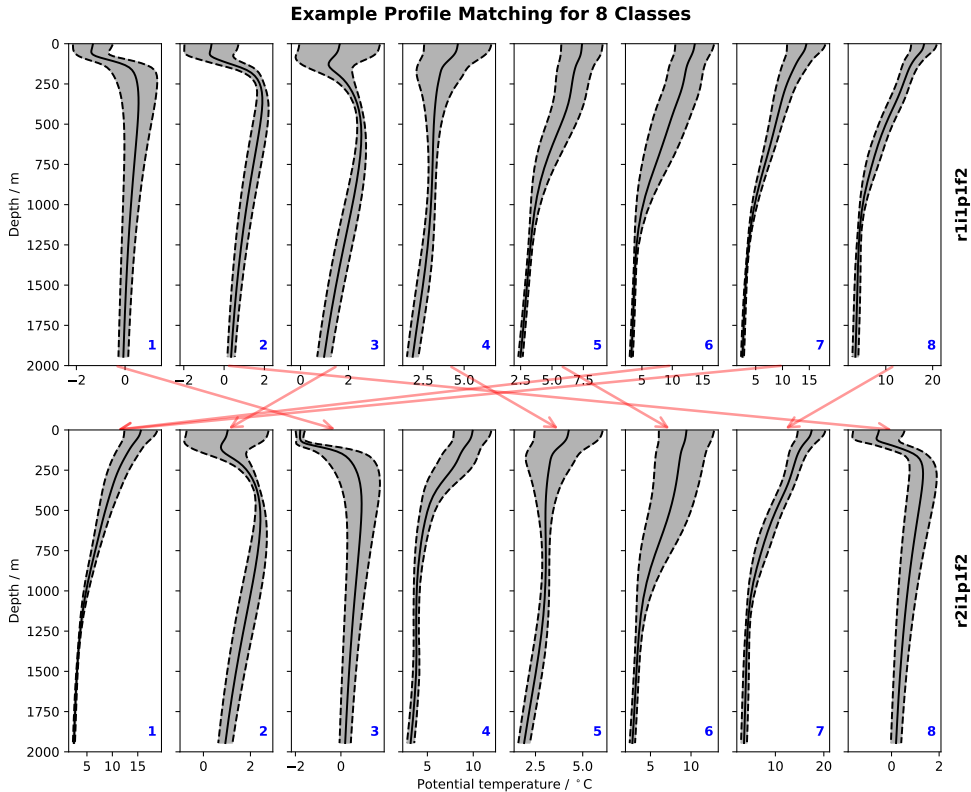
Figure 7. *An example of using the Ensemble Difference Criterion (EDC) to match classes resulting from a 8-class GMM of two different ensemble members (r1i1p1f2 and r2i1p1f2). Note that the map between classes of each ensemble is not bijective - 6 and 7 in the first ensemble are both matched to class 1 in the second.*

5. Discussion

The lack of a clear optimal number of classes from any of the methods tested is perhaps expected, the ocean is a complex system and the GMM representation of the covariance structure in abstract feature space will always be an approximation. Rather, the number of classes will be chosen based on what produces a physically reasonable, and practically useful, distribution. This cannot be determined by statistical tests alone and indeed, in other clustering work, tests such as BIC are accompanied by simply plotting the class distribution for different values of *K* and determining which is most reasonable, as well as investigating the dependence of the results on the number of classes (see, for example, Jones et al., 2019; Houghton and Wilson, 2020). This is typically sufficient; however it may be antithetic to the motivation for automatic clustering in this application, where the identification of different classes is to be performed without the influence of previous clustering studies.

In our particular use case, we are attempting to classify bodies of water that are geographically and physically linked across a range of spatial and temporal scales, and so we do not expect exact separation between classes, but rather expect a degree of overlap between classes. Indeed, this is why the GMM is chosen, which explicitly allows for such overlap. The nature of the GMM also allows for classes that may be very different shapes and sizes in PCA space. This means that metrics such as the SIL may not be ideal for assessing a GMM, as having a tightly shaped class neighbouring a broader class will give a relatively low SIL for the broader class, when in fact the model is a 'good' fit by other measures. Thus

we believe the SIL is a less fitting metric than scores based on the overall likelihood (BIC/AIC) for our application, but it nonetheless helps by providing complementary information.

Comparing the SIL and BIC, we can see a trade-off between worsening class separation and increased likelihood, as indicated by decreasing SIL and BIC with increasing class number K. The BIC or AIC would seem to be better metrics for this application as they are designed to balance goodness of fit with a penalty for over-fitting. However, we did not see any difference between the BIC and AIC scores for the models we generated, implying that the over-fitting penalty is very small at these number of classes. This essentially means the BIC and AIC do not provide an upper limit on the number of classes in this case, as the over-fitting penalty is much smaller than the likelihood factor for the models studied here. This plateauing of the BIC and AIC is also seen in, for example, Sonnewald et al. (2019).

The physical motivation of the EDC scheme is advantageous over purely statistical tests such as BIC. However, there are two caveats to the EDC that should be highlighted. Firstly, using it assumes that, for the time period sampled, we should expect the Southern Ocean in each ensemble to be in comparable conditions. This need not be true, if, for example, they are in different phases of the SAM or ENSO. Secondly, in effect all this method has done is move the ambiguity in the process from the specific choice of cost metric used, and subsequent method of interpreting that metric, to the choice of the how to determine the correspondence between classes in each ensemble.

Given the limitations of all the criteria discussed, the best estimate of the number of classes may come from trying to combine the information from all three (the BIC, the SIL and the EDC). For the example shown here, the SIL score prefers 2 or 5 classes, depending on the absolute or local maximum, the BIC curve points to 4 or more classes, and we can use the EDC to provide an upper limit of 7 classes for our models, resulting in an overall range of 4 to 7 classes supported by our analysis.

6. Summary

Unsupervised classification studies are becoming increasingly common in the study of observed and simulated climate data. In many cases, techniques such as Gaussian Mixture Models (GMMs) and K-means clustering have been remarkably successful at identifying structures that are coherent in time and space, despite not being provided with temporal or spatial information. The property driven-approach of these methods is attractive when the aim is to objectively identify structures. That is not to say that domain knowledge is un-necessary, and in fact experience of the climate system is required to interpret the results and place them in context.

Unsupervised classification methods generally require no external input beyond the data, and the number of classes into which the data should be divided. However, there is not necessarily an obvious choice for this number, which is a relatively unconstrained free parameter. In practice, most studies use statistical criteria such as the Bayesian Information Criterion (BIC) or Silhouette score to compare the results of training the model multiple times for a range of number of classes.

In this paper, we have demonstrated a novel new method for determining over-fit, the Ensemble Difference Criterion (EDC), which can be used, for example, on climate data from ensemble simulations. In demanding robustness of the class definitions between different ensemble members, the EDC provides a definitive upper limit on the number of classes supported by the data.

In the example presented here, the GMM classification of Southern Ocean temperature profiles from historical climate simulations of the UK-ESM model, we find that the BIC does not give a strong constraint on over-fitting, and so the EDC is an extremely useful extra tool. Taking into account the information provided by the BIC, the EDC, and Silhouette scores, we find our data can support between four and seven classes.

For reference, we also compare an eight class GMM based on the UK-ESM data with a previous eight class GMM based on Argo observations. We find that there is good agreement between GMMs within the ACC region, however there are too few Argo observations South of the Southern Boundary to provide a good match with the UK-ESM data in this region. Additionally, the sub-tropical classes don't

match well between GMMs, either because the simulations don't reproduce the temperature structure of the observed sub-tropical Southern Ocean, or because the temperature profiles are not distinct enough in these regions to support robust classes for an eight class model.

References

Abdi, H. and Williams, L. J. (2010). Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459.

Akaike, H. (1973). Information Theory and an Extension of the Maximum Likelihood Principle. In Petrov, B. and Caski, F., editors, *Second International Symposium on Information Theory*, pages 267–281, Budapest. Akademiai Kiado.

Anderson, B. D. and Moore, J. B. (2012). *Optimal filtering*. Courier Corporation.

Beadling, R. L., Russell, J., Stouffer, R., Mazloff, M., Talley, L., Goodman, P., Sallée, J.-B., Hewitt, H., Hyder, P., and Pandde, A. (2020). Representation of Southern Ocean properties across coupled model intercomparison project generations: CMIP3 to CMIP6. *Journal of Climate*, 33(15):6555–6581.

Boland, E. and Atkinson, E. (2023). emmomp/OceanClusteringMethods: Second release.

Caliński, T. and Harabasz, J. (1974). A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1):1–27.

Crawford, A. (2020). The Use of Gaussian Mixture Models with Atmospheric Lagrangian Particle Dispersion Models for Density Estimation and Feature Identification. *Atmosphere*, 11(12):1369.

Davies, D. L. and Bouldin, D. W. (1979). A cluster separation measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-1(2):224–227.

Deser, C., Lehner, F., Rodgers, K. B., Ault, T., Delworth, T. L., DiNezio, P. N., Fiore, A., Frankignoul, C., Fyfe, J. C., Horton, D. E., et al. (2020). Insights from Earth system model initial-condition large ensembles and future prospects. *Nature Climate Change*, 10(4):277–286.

Dosio, A., Jones, R. G., Jack, C., Lennard, C., Nikulin, G., and Hewitson, B. (2019). What can we know about future precipitation in Africa? Robustness, significance and added value of projections from a large ensemble of regional climate models. *Climate Dynamics*, 53:5833–5858.

Fahrin, F., Jones, D. C., Wu, Y., Keeble, J., and Archibald, A. T. (2022). Unsupervised classification of ozone profiles in ukesm1. *Atmospheric Chemistry and Physics Discussions*, pages 1–27.

Hannachi, A. (2007). Tropospheric Planetary Wave Dynamics and Mixture Modeling: Two Preferred Regimes and a Regime Shift. *Journal of the Atmospheric Sciences*, 64(10):3521–3541. Publisher: American Meteorological Society Section: Journal of the Atmospheric Sciences.

Hannachi, A. and O'Neill, A. (2001). Atmospheric multiple equilibria and non-Gaussian behaviour in model simulations. *Quarterly Journal of the Royal Meteorological Society*, 127(573):939–958. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/qj.49712757312.

Hjelmervik, K. and Hjelmervik, K. T. (2014). Time-calibrated estimates of oceanographic profiles using empirical orthogonal functions and clustering. *Ocean Dynamics*, 64(5):655–665.

Hjelmervik, K., Hjelmervik, K. T., and Østenstad, P. (2015). Estimation of oceanographic profiles and climatological regions in the Barents Sea. In *OCEANS 2015 - Genova*, pages 1–6.

Hjelmervik, K. T. and Hjelmervik, K. (2013). Estimating temperature and salinity profiles using empirical orthogonal functions and clustering on historical measurements Topical Collection on the 16th biennial workshop of the Joint Numerical Sea Modelling Group (JONSMOD) in Brest, France 21-23 May 2012. *Ocean Dynamics*, 63(7):809–821. Publisher: Springer.

Houghton, I. A. and Wilson, J. D. (2020). El Niño Detection Via Unsupervised Clustering of Argo Temperature Profiles. *Journal of Geophysical Research: Oceans*, 125(9).

Jones, D. C., Holt, H. J., Meijers, A. J. S., and Shuckburgh, E. (2019). Unsupervised Clustering of Southern Ocean Argo Float Temperature Profiles. *Journal of Geophysical Research: Oceans*, 124(1):390–402.

Kim, Y. S. and Orsi, A. H. (2014). On the variability of antarctic circumpolar current fronts inferred from 1992–2011 altimetry. *Journal of Physical Oceanography*, 44(12):3054–3071.

Kuhlbrodt, T., Jones, C. G., Sellar, A., Storkey, D., Blockley, E., Stringer, M., Hill, R., Graham, T., Ridley, J., Blaker, A., et al. (2018). The low-resolution version of HadGEM3 GC3. 1: Development and evaluation for global climate. *Journal of Advances in Modeling Earth Systems*, 10(11):2865–2888.

Liu, Y., Chen, W., Chen, Y., Chen, W., Ma, L., and Meng, Z. (2021). Ocean Front Reconstruction Method Based on K-Means Algorithm Iterative Hierarchical Clustering Sound Speed Profile. *JMSE*, 9(11):1233.

Maze, G., Mercier, H., and Cabanes, C. (2017a). Profile classification models. *Mercator Ocean Journal*, (55):48–56.

Maze, G., Mercier, H., Fablet, R., Tandeo, P., Lopez Radcenco, M., Lenca, P., Feucher, C., and Le Goff, C. (2017b). Coherent heat patterns revealed by unsupervised classification of Argo temperature profiles in the North Atlantic Ocean. *Progress in Oceanography*, 151:275–292. Publisher: Pergamon.

Meinshausen, M., Nicholls, Z. R., Lewis, J., Gidden, M. J., Vogel, E., Freund, M., Beyerle, U., Gessner, C., Nauels, A., Bauer, N., et al. (2020). The shared socio-economic pathway (SSP) greenhouse gas concentrations and their extensions to 2500. *Geoscientific Model Development*, 13(8):3571–3605.

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. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65.

Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, pages 461–464.

Sellar, A. A., Jones, C. G., Mulcahy, J. P., Tang, Y., Yool, A., Wiltshire, A., O'connor, F. M., Stringer, M., Hill, R., Palmieri, J., et al. (2019). UKESM1: Description and evaluation of the UK Earth System Model. *Journal of Advances in Modeling Earth Systems*, 11(12):4513–4558.

Solidoro, C., Bandelj, V., Barbieri, P., Cossarini, G., and Fonda Umani, S. (2007). Understanding dynamic of biogeochemical properties in the northern Adriatic Sea by using self-organizing maps and k-means clustering. *Journal of Geophysical Research: Oceans*, 112(C7). _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2006JC003553.

Sonnewald, M., Wunsch, C., and Heimbach, P. (2019). Unsupervised Learning Reveals Geography of Global Ocean Dynamical Regions. *Earth and Space Science*, 6(5):784–794. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2018EA000519.

Sugiura, N. (2021). Clustering global ocean profiles according to temperature-salinity structure. arXiv:2103.14165 [physics].

Tandeo, P., Chapron, B., Ba, S., Autret, E., and Fablet, R. (2014). Segmentation of Mesoscale Ocean Surface Dynamics Using Satellite SST and SSH Observations. *IEEE Trans. Geosci. Remote Sensing*, 52(7):4227–4235.

Walton, J. (2020). 3D visualisation of CMIP6 data. *UKESM Newsletter*, (11).

Williams, R. G., Ceppi, P., Roussenov, V., Katavouta, A., and Meijers, A. J. (2023). The role of the Southern Ocean in the global climate response to carbon emissions. *Philosophical Transactions of the Royal Society A*, 381(2249):20220062.

Yeager, S., Danabasoglu, G., Rosenbloom, N., Strand, W., Bates, S., Meehl, G., Karspeck, A., Lindsay, K., Long, M., Teng, H., et al. (2018). Predicting near-term changes in the Earth System: A large ensemble of initialized decadal prediction simulations using the Community Earth System Model. *Bulletin of the American Meteorological Society*, 99(9):1867–1886.

Yuchechen, A. E., Lakkis, S. G., Caferri, A., Canziani, P. O., and Muszkats, J. P. (2020). A Cluster Approach to Cloud Cover Classification over South America and Adjacent Oceans Using a k-means/k-means++ Unsupervised Algorithm on GOES IR Imagery. *Remote Sensing*, 12(18):2991.

Zhao, J., Guo, Z., Guo, Y., Lin, W., and Zhu, W. (2021). A self-organizing forecast of day-ahead wind speed: Selective ensemble strategy based on numerical weather predictions. *Energy*, 218:119509.