

A Bayesian approach to atmospheric circulation regime assignment

Article

Published Version

Creative Commons: Attribution 4.0 (CC-BY)

Open Access

Falkena, S. K. J., de Wiljes, J., Weisheimer, A. and Shepherd, T. G. ORCID: <https://orcid.org/0000-0002-6631-9968> (2023) A Bayesian approach to atmospheric circulation regime assignment. *Journal of Climate*, 36. pp. 8619-8636. ISSN 1520-0442 doi: 10.1175/JCLI-D-22-0419.1 Available at <https://centaur.reading.ac.uk/110936/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1175/JCLI-D-22-0419.1>

Publisher: American Meteorological Society

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

A Bayesian Approach to Atmospheric Circulation Regime Assignment

SWINDA K. J. FALKENA^a,[✉] JANA DE WILJES,^{a,b} ANTJE WEISHEIMER,^{c,d} AND THEODORE G. SHEPHERD^e

^a Department of Mathematics and Statistics, University of Reading, Reading, United Kingdom

^b Institute for Mathematics, University of Potsdam, Potsdam, Germany

^c European Centre for Medium-Range Weather Forecasts, Reading, United Kingdom

^d National Centre for Atmospheric Science, Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, Oxford, United Kingdom

^e Department of Meteorology, University of Reading, Reading, United Kingdom

(Manuscript received 2 June 2022, in final form 2 February 2023, accepted 24 February 2023)

ABSTRACT: The standard approach when studying atmospheric circulation regimes and their dynamics is to use a hard regime assignment, where each atmospheric state is assigned to the regime it is closest to in distance. However, this may not always be the most appropriate approach as the regime assignment may be affected by small deviations in the distance to the regimes due to noise. To mitigate this we develop a sequential probabilistic regime assignment using Bayes's theorem, which can be applied to previously defined regimes and implemented in real time as new data become available. Bayes's theorem tells us that the probability of being in a regime given the data can be determined by combining climatological likelihood with prior information. The regime probabilities at time t can be used to inform the prior probabilities at time $t + 1$, which are then used to sequentially update the regime probabilities. We apply this approach to both reanalysis data and a seasonal hindcast ensemble incorporating knowledge of the transition probabilities between regimes. Furthermore, making use of the signal present within the ensemble to better inform the prior probabilities allows for identifying more pronounced interannual variability. The signal within the interannual variability of wintertime North Atlantic circulation regimes is assessed using both a categorical and regression approach, with the strongest signals found during very strong El Niño years.

SIGNIFICANCE STATEMENT: Atmospheric circulation regimes are recurrent and persistent patterns that characterize the atmospheric circulation on time scales of 1–3 weeks. They are relevant for predictability on these time scales as mediators of weather. In this study we propose a novel approach to assigning atmospheric states to six predefined wintertime circulation regimes over the North Atlantic and Europe, which can be applied in real time. This approach introduces a probabilistic, instead of deterministic, regime assignment and uses prior knowledge on the regime dynamics. It allows us to better identify the regime persistence and indicates when a state does not clearly belong to one regime. Making use of an ensemble of model simulations, we can identify more pronounced interannual variability by using the full ensemble to inform prior knowledge on the regimes.

KEYWORDS: Atmosphere; Blocking, North Atlantic Oscillation; Bayesian methods; Classification; Clustering

1. Introduction

A thorough understanding of extratropical circulation variability on subseasonal time scales is important for improving

 Denotes content that is immediately available upon publication as open access.

 Supplemental information related to this paper is available at the Journals Online website: <https://doi.org/10.1175/JCLI-D-22-0419.s1>.

Falkena's current affiliation: Institute for Marine and Atmospheric Research Utrecht (IMAU), Utrecht University, Utrecht, The Netherlands.

de Wiljes's current affiliations: Institute for Mathematics, Ilmenau University of Technology, Ilmenau, Germany, and Department of Computational Engineering, School of Engineering Sciences, LUT University, Lappeenranta, Finland.

Corresponding author: Swinda K. J. Falkena, s.k.j.falkena@uu.nl

DOI: 10.1175/JCLI-D-22-0419.1

© 2023 American Meteorological Society. This published article is licensed under the terms of a Creative Commons Attribution 4.0 International (CC BY 4.0) License

predictability on these time scales. Improvement of this predictability is of great societal relevance for sectors such as renewable energy. Atmospheric circulation, or weather, regimes can describe this variability by dividing the circulation into a small number of states or patterns (Hannachi et al. 2017). These regimes are recurrent patterns that represent the low-frequency variability in the atmospheric circulation. They have been studied for a long time, starting with papers focusing on their identification (e.g., Mo and Ghil 1988; Molteni et al. 1990; Vautard 1990; Michelangeli et al. 1995), with later research discussing their links with other processes and surface impacts (e.g., Straus and Molteni 2004; Cassou et al. 2005; Charlton-Perez et al. 2018; van der Wiel et al. 2019).

The most commonly used technique for identifying circulation regimes is k -means clustering (e.g., Michelangeli et al. 1995; Straus et al. 2007; Matsueda and Palmer 2018). This method separates the phase space into k clusters, where the data within each cluster are similar, but dissimilar between the different clusters. The number of clusters k has to be set a priori, for which several approaches such as a classifiability index (Michelangeli et al. 1995) or information criteria (O'Kane et al. 2013) are used. One of the drawbacks of this






FIG. 1. A conceptual example of the difficulty k -means clustering has when noise affects the data, showing what a probabilistic approach can bring. (a) An example trajectory of the data as a function of the distances to two regimes A (orange) and B (red). The 1-1 line is shown as a black dashed line and indicates that the region above is closer to regime A and the region below is closer to regime B. Numbers indicate the day corresponding to that point in the trajectory. The likelihood functions shown along the top and right of the main plot give the climatological probability of those distances given hard assignment to regime A (orange; top plot) or B (red; right plot). The dotted gray line indicates a slice through the probability space along which the pdfs in (b) are considered. (b) A slice of the likelihood functions, weighted by the prior probabilities following Bayes's theorem, for each of the regimes (solid lines: regime A is orange and regime B is red) along the gray dotted line in (a), perpendicular to the 1-1 line, for the 7th, 8th, and 9th days. The location of the data on each day is indicated by the vertical black lines, and the bars at the edge of the plots show the prior (left bar of the pair) and posterior (right hatched bar) probabilities for each of the regimes (regime A in orange at the left edge; regime B in red at the right edge). The climatological likelihood functions are shown as dashed lines in all plots in (b), and the vertical gray dotted line indicates the location of the 1-1 line. The inset plots in each main plot show an enlargement of the region around the 1-1 line.

clustering approach is that it yields a hard, deterministic, assignment of the data to each of the regimes. This means that it is difficult to quantify the uncertainty of the regime assignment, as data close to the regime center are treated the same as data that are only just (by distance) assigned to that regime.

The hard regime assignment of k -means clustering means that the result is susceptible to noise. Consider Fig. 1a, which shows the distance of the data to two regimes in time for a real case (discussed later in detail), over a period of 12 days. Initially, the data clearly are categorized to belong to regime A, being significantly closer in distance to regime A than to regime B. However, from day 7 to 9 the data make a brief excursion into a part of the phase diagram that is closer to regime B, after which they move back to being closest to regime A. The question is whether this is a real signal or simply the effect of noise. Since the regime dynamics is

quite persistent in time it is likely to be the latter, but this possibility is not picked up by the hard assignment of a standard k -means clustering approach. Often a low-pass filter is applied to remove this high-frequency variability (e.g., Straus et al. 2007; Grams et al. 2017), but in Falkena et al. (2020) it was shown that low-pass filtering can lead to a bias in the observed regime frequencies.

Another solution is to use a regularized clustering algorithm that constrains, or bounds, the number of transitions between the regimes so that it is in line with the natural metastability of the underlying dynamics. Such an approach, first introduced in the context of clustering methods by Horenko (2010), has for example been applied to discrete jump processes (Horenko 2011a) with applications in computational sociology (Horenko 2011b) and for efficient classification in the context of sparse data settings (Vecchi et al. 2022). In the context of atmospheric dynamics, time regularization has been used to study the Southern

Hemispheric circulation (O’Kane et al. 2013), the dynamics of the North Atlantic Oscillation (Quinn et al. 2021), and how to identify persistent circulation regimes (Falkena et al. 2020). A regularized clustering method allows us to better identify the signal within the noise but does require selecting a constraint parameter. This introduces a parameter selection, where, for example, an information criterion is used to decide on a suitable constraint value.

An alternative approach is to make the regime assignment probabilistic rather than deterministic, allowing for a more nuanced and informative regime assignment in the presence of noise. Methods such as mixture modeling provide such a probabilistic regime assignment (e.g., Hannachi and O’Neill 2001; Smyth et al. 1999; Baldo and Locatelli 2023) but are not widely used. Hidden Markov models (HMMs) extend the mixture modeling approach by also taking into account the dynamics of the system and not just the statistics (Majda et al. 2006; Franzke et al. 2008), but they are hard to fit for relatively short time series when the data are high dimensional. Another approach is to approximate the regime model using local Markov distance functionals with corresponding time dependent probabilities (Horenko 2011a). In studies that look into forecasting of regimes on subseasonal time scales, the probability of being in a regime is often considered by looking at the empirical distribution of the (hard) regime assignment across an ensemble (Vigaud et al. 2018; Cortesi et al. 2021; Büeler et al. 2021; Falkena et al. 2022). Such an approach is already used in an operational setting by, for example, ECMWF (Ferranti et al. 2015). A limitation of this method is that it requires availability of ensemble data, where typically the ensemble size is small, and verification is done against a hard regime assignment from reanalysis.

A probabilistic regime assignment that does not require this availability of ensemble data would help in better assessing the skill in predicting regimes, as it could be applied to reanalysis data, which are also subject to noise. Such a regime assignment would allow us to identify the instances in which the observations cannot be clearly assigned to one regime or in which a wrong hard assignment is potentially due to noise. This approach allows for a fairer verification of the model by taking some degree of observational uncertainty into account. Here it is desirable for the approach to be sequential, which allows for the regime assignment to be done in real time, making it suitable for operational applications. Most probabilistic regime assignment methods, such as mixture models or HMMs, require the availability of the full dataset when computing the regime probabilities, which would mean one has to rerun the clustering algorithm whenever a new data point is added. A method that, after training on an initial dataset, can easily be applied to data as it becomes available is more suitable for an operational setting. Such a method can also be applied to predefined regimes, to provide traceability with previous work.

The standard hard regime assignment can be considered as a random process that takes a value in the set of possible regimes at each time. The associated probability can be computed on the basis of metastability frequencies computed from previous or currently available batch data. The aim is to determine the corresponding conditional probability of being

in a regime given the data, i.e., $P(\text{Regime}|\text{Data})$. Following Bayes’s theorem, this is given by

$$P(\text{Regime}|\text{Data}) = \frac{P(\text{Data}|\text{Regime})P(\text{Regime})}{P(\text{Data})}, \quad (1)$$

combining prior knowledge of the probability of being in a regime $P(\text{Regime})$ with an observed likelihood given a regime $P(\text{Data}|\text{Regime})$. The latter can sometimes be computed from the climatological data. In Fig. 1a the observed (climatological) likelihood functions for both regimes are shown next to the trajectory. The working of Bayes’s theorem for such a trajectory is shown in Fig. 1b, which shows how the inclusion of prior information $P(\text{Regime})$ following Bayes’s theorem given in (1) affects the posterior $P(\text{Regime}|\text{Data})$ for the trajectory at days 7, 8, and 9, following a section along the dotted line in Fig. 1a. The climatological likelihood functions of the two regimes A and B, indicated by the dashed lines, are weighted (solid lines) using the prior regime probabilities, shown by the nonhatched bars at the edge of the panels. The posterior probabilities are then computed as the values of the weighted likelihood functions at the data point (vertical black line). The obtained Bayesian probabilities are indicated by the hatched bars and used to inform the prior probabilities for the next time step, using climatological information about transition probabilities.

At day 7 the prior information indicates a very high probability of being in regime A as all previous days belonged clearly to that regime. This increases the probability of $t = 7$ belonging to regime A and decreases that of belonging to regime B with respect to the climatological likelihood, which would otherwise be evenly balanced between the two regimes. Thus, there is a high probability that the data at day 7 belong to regime A. Given the known persistence of regimes, the prior information for day 8 again then indicates a high probability of being in this regime, albeit slightly smaller than at $t = 7$, which weights the likelihood functions accordingly. Although the data are closer to regime B, the prior information means that there is an approximately equal probability of being in either of the two regimes. The prior for $t = 9$ thus does not weight the likelihood functions as much as for $t = 7$ and 8, and thus the data at day 9 being equally close to both regimes means that again the probability of being in either of the regimes is close to a half. This discussion shows how the inclusion of prior information can be used to compute the probability of a regime given the data, and thereby soften the effects of noise, following the fundamental principles of probability as encoded in Bayes’s theorem [in (1)]. As noted above, the approach as discussed here is sequential and can be applied to individual realizations, making it suitable for operational applications. An initial training dataset can be used to obtain the climatological likelihood functions, after which the regime assignment can be applied to data as they become available. The latter regime assignment step is similar to finding the most probable sequence once an HMM is known (Viterbi 1967; Rabiner 1989).

Other aspects than persistence can affect the prior regime likelihood as well. It is likely that nonstationary external

factors, such as El Niño–Southern Oscillation (ENSO) or sudden stratospheric warmings (SSWs), have an influence on the prior regime probabilities (e.g., [Toniziano and Scaife 2006](#); [Ayarzagüena et al. 2018](#); [Domeisen et al. 2020](#)). The Bayesian approach allows us to incorporate such information, either by looking at, for example, an ENSO index or by making use of the availability of ensemble data. In a previous study a regularized clustering method helped to identify a more pronounced interannual regime signal by making use of the information available in an ensemble ([Falkena et al. 2022](#)). Similarly, having a more informative prior for Bayes's theorem from (1), incorporating information from external processes, can help in identifying a stronger nonstationary regime signal. The Bayesian approach discussed here is not the only method in which information on external forcing can be incorporated in the regime assignment (e.g., [Franzke et al. 2015](#)), but it is (to our knowledge) the first that allows to do this in a sequential manner.

In this paper we formalize the intuition of [Fig. 1](#) and study how to use Bayes's theorem to obtain a probabilistic regime assignment based on predefined regimes for the wintertime Euro-Atlantic sector. The use of predefined regimes respects the scientific value that has already been established for those regimes (e.g., in the relationship with particular climate impacts). In [section 2](#) we discuss the data that are used and the use of standard k -means clustering to obtain the circulation regimes that we consider for this study. The two sections that follow explain the way in which Bayes's theorem can be used for the regime assignment, where an important aim of our work is to link our method to existing work on clustering of circulation regimes. We start with the most intuitive sequential form (as discussed above) in [section 3](#), and in [section 4](#) we consider how the use of ensemble data, which picks up some external forcing signals, can help to update the prior regime probabilities to study interannual regime variability, which is discussed in [section 5](#). A discussion and conclusions are given in [section 6](#).

2. Data and clustering

For the identification of the circulation regimes the 500-hPa geopotential height fields (Z500) from two datasets are used: the ECMWF SEAS5 hindcast ensemble dataset ([Johnson et al. 2019](#)) and the ERA-Interim reanalysis dataset ([Dee et al. 2011](#)). For both datasets, daily (0000 UTC) gridpoint ($2.5^\circ \times 2.5^\circ$ resolution) Z500 data over the Euro-Atlantic sector (20° to 80°N , 90°W to 30°E) are considered for all winters [December–March (DJFM)] for which the SEAS5 ensemble data are available (1981–2016). The regimes are computed using gridpoint anomaly data, where the anomalies are computed with respect to the average DJFM climatology [see [Falkena et al. \(2020\)](#) for the rationale for this choice]. Here the climatologies of ERA-Interim and SEAS5 are used as a reference for the computation of their respective anomalies. The SEAS5 hindcast ensemble has 51 members and is initialized each year on 1 November, which means that by considering data only from December onward the effect of the atmospheric initial conditions has been effectively lost. This allows us to treat each ensemble member as an alternative, physically plausible yet not observed realization of the atmosphere

([Thompson et al. 2017](#)), subject to the nonstationary influences for that year (notably ENSO).

A standard k -means clustering algorithm ([Jain 2010](#)), with a Euclidian distance to compute the distance between the data and regimes, is used to identify six circulation regimes over the Euro-Atlantic sector for both ERA-Interim and the SEAS5 hindcast ensemble. In k -means clustering the data are sorted in k clusters that are close together within one cluster, but far from data in the other clusters based on some distance measure. These clusters are represented by their mean, which corresponds to the circulation regimes, where the number of clusters k has to be set a priori. Six was identified as a suitable number of regimes for such unfiltered data in a previous study ([Falkena et al. 2020](#)). The regimes for the SEAS5 hindcast ensemble are shown in [Fig. 2](#) and are the two phases of the North Atlantic Oscillation (NAO), the Atlantic Ridge (AR), Scandinavian blocking (SB) and both their counterparts. Note that these regimes are slightly different in their patterns from those of ERA-Interim [see [Falkena et al. \(2022\)](#) for details on this], thereby providing an inherent bias correction between the model and reanalysis. These hard regime assignments are used to compute the likelihood functions that are used in the Bayesian approach, for which a detailed discussion is given in [section 3a](#). In addition we consider the (hard) regime assignments obtained using the time-regularized clustering algorithm from [Falkena et al. \(2020\)](#). This allows for a comparison of different approaches to identify the persistent regime signal.

3. Sequential Bayesian regime assignment

In this section the Bayesian approach to regime assignment is discussed, which can be applied to ERA-Interim data as well as single ensemble realizations. We start with the details of the method itself in [section 3a](#), followed by a comparison with the results of both a standard and time-regularized k -means clustering method in [section 3b](#).

a. Method

The starting point for our sequential Bayesian regime assignment is the six regimes obtained using k -means clustering discussed in [section 2](#) and shown in [Fig. 2](#). The likelihood functions in Bayes's theorem given in (1) are computed based on the distance to these regimes, and remain fixed throughout the sequential Bayesian regime assignment. The discussion of the method as phrased below is general and can be applied to all types of regime dynamics as long as the regimes themselves and the likelihood functions are specified a priori.

Let r be a discrete random variable indicating a regime (i.e., taking values in $\{1, \dots, k\}$ for k regimes) and let $\mathbf{d} \in \mathbb{R}^k$ be a vector containing the distances to each of the regimes (here the Euclidian distance is used, which is also the standard cost function in the k -means setting). Specifically, \mathbf{d} are the data we consider in our Bayesian approach. The use of the regime distance as data is not the only option. When one considers only a limited number of principal components (PCs) for the regime representation the PC values can be directly used. However, for the spatial fields considered here [see [Falkena et al. \(2020\)](#) for the arguments in favor of using gridpoint




FIG. 2. The six circulation regimes obtained for the SEAS5 ensemble using k -means clustering: 1) NAO+, 2) NAO-, 3) Atlantic ridge (AR+), 4) Scandinavian blocking (SB+), 5) AR-, and 6) SB-.

data] this is unfeasible as the high dimensionality means the phase space is sparsely sampled, leading to large uncertainty in the resulting distributions. Therefore, a means of dimension reduction is required for which we consider the distances to the different regimes since this is the metric used in most clustering approaches. At a given time we are interested in the probability to be in a regime r given the data [i.e., $P(r|\mathbf{d})$]. Bayes's theorem tells us that

$$P(r|\mathbf{d}) = \frac{P(\mathbf{d}|r)P(r)}{P(\mathbf{d})}. \quad (2)$$

Here, $P(r)$ is the prior probability of regime r and $P(\mathbf{d})$ is the probability of the data. Since we only consider a discrete number of regimes that are mutually exclusive and exhaustive, the latter can be computed by

$$P(\mathbf{d}) = \sum_{r=1}^k P(\mathbf{d}|r)P(r), \quad (3)$$

making it a normalization factor.

Last, $P(\mathbf{d}|r)$ is the likelihood of the data given a regime r . The likelihood of the data can be determined from the distance to each of the regimes by considering how the data fall within the conditional distance distributions (i.e., the distributions conditioned on data belonging to one of the regimes). For each data point in either the SEAS5 or ERA-Interim time series we have this distance to each of the k regimes, which

has been computed in the k -means clustering procedure to determine the hard regime assignment (section 2). This gives the distributions of the distances to each of the regimes conditional on regime r , which for SEAS5 are shown in Fig. 3.

There are a few things to note concerning these distributions. First, the distance to the regime the data is assigned to is smallest, but can still be larger than the distance to other regimes for a different data point belonging to that regime. Second, for data assigned to AR+, SB+, AR-, and SB- the distances to the other regimes are roughly equally distributed with the means being relatively close to each other. However, for data assigned to either NAO+ or NAO- the distance to the other phase is larger than that to the other four regimes. Thus these two regimes are farther away from each other than the rest of the regimes, and information on the proximity to one regime is providing information on the proximity to the other.

Also, we see that these distributions are approximately normal, justifying us to approximate the corresponding k -dimensional conditional probability density functions (pdf) by a multivariate normal. The likelihood $P(\mathbf{d}|r)$ is then given by the value of the conditional pdf, that is

$$P(\mathbf{d}|r) = \frac{\exp\left[-\frac{1}{2}(\mathbf{d} - \boldsymbol{\mu}_r)^T \boldsymbol{\Sigma}_r^{-1}(\mathbf{d} - \boldsymbol{\mu}_r)\right]}{\sqrt{(2\pi)^k |\boldsymbol{\Sigma}_r|}}, \quad (4)$$

where $||$ represents the determinant. The mean $\boldsymbol{\mu}_r$ and covariance $\boldsymbol{\Sigma}_r$, representing the variability around the cluster center,




FIG. 3. The distributions of the distances (normalized; gpm per number of grid points) to each of the regimes (colors) conditional on the SEAS5 hindcast data being assigned to the regime given in the title, based on a hard assignment. The means of each distribution are indicated by the vertical dotted lines.

are estimated from the conditional distance distributions obtained from the k -means clustering results for each regime. These estimates are done separately for ERA-Interim and SEAS5 to avoid biases due to the regimes being slightly different. The estimates of the mean and covariance are surprisingly similar between both datasets, indicating that, apart from the slight difference in regimes, the model does a reasonable job in representing the variability of the regime dynamics. A further discussion on this, including a robustness analysis of the distance distributions, is given in the online supplemental material.

To obtain the prior probability $P(r)$ there is a natural choice from propagating the probabilities of the previous time step forward. From k -means clustering an estimate of the regime dynamics is known, which is characterized by the climatological regime frequencies \mathbf{P}^c and transition probabilities T_{ij}^c between the regimes. For SEAS5 these are given by (Falkena et al. 2022; for the regimes ordered as in Fig. 2)

$$\mathbf{P}^c = \begin{pmatrix} 0.176 \\ 0.158 \\ 0.160 \\ 0.163 \\ 0.175 \\ 0.168 \end{pmatrix},$$

$$\mathbf{T}^c = \begin{pmatrix} 0.728 & 0.000 & 0.039 & 0.062 & 0.060 & 0.112 \\ 0.000 & 0.822 & 0.050 & 0.046 & 0.053 & 0.029 \\ 0.079 & 0.054 & 0.702 & 0.075 & 0.021 & 0.069 \\ 0.069 & 0.058 & 0.065 & 0.739 & 0.037 & 0.031 \\ 0.072 & 0.032 & 0.035 & 0.045 & 0.771 & 0.045 \\ 0.065 & 0.033 & 0.095 & 0.029 & 0.070 & 0.708 \end{pmatrix}. \quad (5)$$

Starting from the regime probabilities at time $t - 1$, a best estimate of the prior probabilities for the next time step is

$$\mathbf{P}(t) = \mathbf{T}^c \mathbf{P}(t - 1|\mathbf{d}), \quad (6)$$

where $\mathbf{P}(t)$ is the vector of prior probabilities $\{P(r)\}_{r=1,\dots,k}$ at time t and $\mathbf{P}(t - 1|\mathbf{d})$ is the vector of posterior probabilities $\{P(r|\mathbf{d})\}_{r=1,\dots,k}$ at time $t - 1$. Note that in the transition matrix \mathbf{T}^c the diagonal elements—corresponding to persistence of the current regime—dominate. At the start of each winter, on 1 December, there is no previous probability to use, and thus little prior information on the probability of being in any of the regimes. For that reason, the climatological regime frequencies \mathbf{P}^c are used as a prior. Note that this is nearly as uninformative as using a uniform distribution. Here the hard regime assignment is used to obtain both the initial prior for each winter and the transition probabilities to obtain subsequent priors. This is by no means the only option (e.g., one could also use a uniform prior at the start of winter). The choice made here is closest to existing methods and therefore least biased when comparing the results.

Using the prior probabilities $P(r)$ and likelihood of the data $P(\mathbf{d}|r)$ following the conditional distance distributions we can compute the posterior Bayesian probability of a regime given the data $P(r|\mathbf{d})$ using Bayes's theorem as presented in (2) in every time step. This yields a sequential probabilistic regime assignment, where the regime probabilities of one day are used to obtain a prior for the next day. Applying this method to ERA-Interim data and the ensemble members of the SEAS5 ensemble yields a probability of being in each of the six regimes at every day in winter. From here on we refer to this posterior Bayesian probability simply as the Bayesian probability. This Bayesian




FIG. 4. The (top) prior probability, (middle) conditional regime likelihood, and (bottom) Bayesian regime probability for the 23rd ensemble member in the sequential Bayesian regime assignment procedure for the winter of 1992/93. The bar at the bottom indicates the hard regime assignment following k -means clustering.

approach can be related to an HMM approach, where the regime patterns and their transition probabilities are given a priori, leaving only the hidden regime assignment to be discovered. Here the used likelihood differs from that commonly used in the standard expectation-maximization (EM) algorithm (e.g., Dempster et al. 1977; Rabiner 1989). In case the transition matrix \mathbf{T} cannot be obtained directly, as is done here through observation of the hard regime assignment, one could employ techniques to find \mathbf{T} via algorithms designed in the context of HMMs.

The above-described sequential Bayesian regime assignment is simple and allows for a straightforward comparison with the commonly used hard regime assignment, as well as with the regularized clustering results (without the need of selecting a constraint parameter). However, there are other options to model the uncertainty and to update the corresponding model parameters sequentially. For instance, one can model each regime individually and associate its center estimates with the mean of a Gaussian. The updating procedure for such a model is called the Kalman filter (Kalman 1960) or the corresponding Monte Carlo approximation the ensemble Kalman filter (Evensen and van Leeuwen 2000), and of course various other methods for more general distributions as well as iterative assimilation of incoming information exist (e.g., Kantas et al. 2014; Hu and van Leeuwen 2021; Acevedo et al. 2017). The method used here is closer to a particle filter (Del Moral 1997; Doucet et al. 2001) as our ensemble members are weighted with importance weights stemming from the likelihood rather than using an analytic formula such as is used in the Kalman filter. However, in this paper we specifically aim to stay close to existing methods and model the process of hard regime assignments as random variables in each time step. This allows for a straightforward implementation that can be readily applied in an operational setting.

Furthermore, using this method we can investigate whether the results are comparable to those found using regularized clustering methods, which have been used to improve the regime persistence in the identification procedure, without the need to select a constraint parameter.

b. Evaluation

The first question to answer is what the effect is of this Bayesian approach in practice, and whether this matches the intuition behind the method. How does the prior affect the Bayesian probabilities? A next step is to compare the probabilistic approach with results obtained using a hard regime assignment, as given by k -means clustering. Is the average regime frequency affected? What is the effect on the regime persistence? In this section we start by discussing the first question by looking at some examples to get a sense for how the method is working in practice, after which we look at the statistics of the results in comparison with a k -means clustering approach to answer the other questions.

To start, we consider the Bayesian regime probabilities for a single randomly chosen ensemble member. As the sequential Bayesian regime assignment works on a single-member basis this is the best way to gain insight into the workings of the Bayesian method. In Fig. 4 the prior and Bayesian regime probabilities for the 23rd ensemble member are shown together with the climatological likelihood corresponding to the observed data point. A first aspect to note is that most of the time the regime likelihood $P(\mathbf{d}|r)$ gives a clear indication of the regime the data belong to. Second, we see that the prior very closely follows the Bayesian probabilities with a delay of 1 day, corresponding to the high persistence in the transition matrix given in (5). The initial prior, given by the climatological




FIG. 5. The (top) observed regime likelihood and (bottom) Bayesian regime probability for ERA-Interim, with the hard assignment using a standard or time-regularized (persistent) k -means algorithm shown by the bars for the winter of 1993/94.

values, is uninformative and in that case the regime likelihood nearly fully determines the Bayesian probabilities. Subsequently, the prior is much more informative but in most cases the regime likelihood still strongly determines the final probability. However, when the likelihood does not clearly point toward one regime (e.g., around days 8–12), the prior information shifts the probabilities toward stronger persistence, in this case of the AR+ regime. This can also be seen around days 99–101, corresponding to days 7–9 in the example shown in Fig. 1 in section 1, where the inclusion of prior information favors persistence over a short excursion away from the most likely regime. In this way the Bayesian regime assignment allows for identifying stronger persistence (i.e., high probability of the dominant regime) without losing the signal of other regimes entering the dynamics as they still have some nonzero probability. The effect of this approach for ERA-Interim data is similar.

The Bayesian probabilistic regime assignment allows us to understand some of the subtleties of the regime dynamics; for example, regime transitions occur in the form of a decrease or increase of the regime probabilities. How does such an approach compare to the commonly used hard regime assignment obtained using k -means clustering? The bar at the bottom of Fig. 4 shows the hard regime assignment corresponding to this time series. The Bayesian regime probabilities vary more smoothly, and show fewer short back-and-forth transitions between regimes that occur several times for the hard regime assignment (e.g., around days 9 and 20). In Falkena et al. (2020) a constraint on the number of transitions between regimes was introduced to reduce the number of short back-and-forth transitions between regimes, based on the regularized clustering method introduced by Horenko (2010). This was shown to increase the regime persistence without affecting the regime occurrence rates, provided the constraint parameter was chosen appropriately. The optimal constraint parameter corresponded to an average regime duration of 6.3 days. It was selected by considering the Bayesian information criterion and falls within

the region where the regime occurrence rates are not affected by the regularization.

In Fig. 5 a comparison of the regime likelihood, Bayesian regime probabilities, and a hard regime assignment obtained using either a standard or regularized k -means approach is shown for ERA-Interim for the winter of 1993/94. The regularization does reduce the number of regime transitions by, for example, removing the NAO+ regime between two occurrences of SB- around day 18. At the same time the Bayesian probabilities show a small increase in the NAO+ likelihood, with SB- still having the highest probability. Here the regularization and Bayesian approach thus yield similar results. On the other hand, around days 84 and 107, for example, the regularization eliminates some regime transitions where the Bayesian probabilities still show some signal of the corresponding regimes. The probabilistic approach thus allows us to identify the data where the regime assignment is less clear, showing an increase in probability instead of a hard regime change. It also retains some regime transitions that the regularized clustering eliminates due to it being difficult to select the “correct” constraint value. In the probabilistic approach these show as increases in the corresponding regime probability. This analysis confirms that the Bayesian approach seems to be doing something sensible, without having to tune any parameters. When the data clearly belong to one of the six regimes, there is little benefit to the Bayesian approach. The main times where it makes a difference are the periods when one regime transitions into another, or when a regime loses some of its strength in favor of another regime but then gains strength again. Such a reduction in the regime probabilities could be an indication of increased flow instability, being close to transitioning into another of the six canonical states.

The impact of the sequential Bayesian approach on the regime frequencies, computed as the average Bayesian regime probability for this method, and (1-day) autocorrelation is shown in Fig. 6. Here the autocorrelation for the hard regime




FIG. 6. The (a) regime frequencies and (b) 1-day autocorrelation as obtained using either standard k -means clustering (circles) or a sequential Bayesian regime assignment (stars) for the SEAS5 hindcast ensemble (symbols with error bars) and ERA-Interim (symbols only), for which also the values obtained with the time-regularized k -means clustering method are shown (squares). Error bounds are determined using bootstrapping with one member per year (with replacement, 500 times), where the thick bars indicate the ± 1 standard error range, with thin bars extending showing the 95% confidence interval.

assignment is computed using a time series that is 1 when data are assigned to the corresponding regime and 0 otherwise. The average frequencies of the regimes do not change when using the Bayesian regime assignment, as can be seen in Fig. 6a. This holds both for the SEAS5 hindcast ensemble data and for ERA-Interim, where also the results of the regularized k -means clustering algorithm are shown for comparison. On the other hand, the autocorrelation, being an indication of the persistence of the regimes, is strongly affected (Fig. 6b). For ERA-Interim we see that the sequential Bayesian approach increases the autocorrelation even beyond that obtained using a regularized clustering algorithm that contains a persistence constraint. Also, for SEAS5 a strong increase in autocorrelation is found using the sequential Bayesian regime assignment relative to a standard hard assignment. For most regimes the ERA-Interim values lie at the top of the SEAS5 autocorrelation range, both for the standard and Bayesian approach. Thus we find that the Bayesian approach does not alter the regime frequencies, but does lead to more persistent regime dynamics, as we might

hope. This suggests that the transition probabilities in (5), which are used to obtain the prior regime probabilities, likely are an underestimation of the true persistence, which is improved by the use of Bayes's theorem.

4. Ensemble Bayesian regime assignment

The implicit assumption made in the sequential Bayesian approach as discussed in the previous section is that the regime dynamics is statistically stationary in time. That is, the climatological likelihood functions and transition probabilities do not change in time. This is a reasonable and minimal first assumption yielding good results, but it is likely that external factors such as ENSO affect some aspects of the regime dynamics as discussed in section 1. There are two obvious ways in which to include the effect of external forcing in the Bayesian approach. The first is to update the regime likelihood functions in time. The second is to update the prior probabilities. These two aspects are by no means the only aspects of the regime

dynamics that can be affected by external forcing. For example, one can imagine that the regimes themselves change as a consequence of external factors causing changes in the climate system. However, this is nearly impossible to quantify with the limited available data and no robust evidence for this has been found so far (e.g., Corti et al. 1999; Dorington et al. 2022). Therefore, we only discuss the above-mentioned two approaches.

In the following analysis we focus on the latter of these two approaches. The main reason for this is the lack of data availability. Even though the SEAS5 hindcast ensemble has 51 members for each year, this still is insufficient to allow for, say, weekly updating of the likelihood functions. An option for which sufficient data are available would be to compute the likelihood function during, for instance, strong El Niño years, and use those to change the likelihood functions each year. However, this relies on the hypothesis that the regions in phase space belonging to each of the regimes shift as a consequence of ENSO forcing, while it may simply be the case that some regions are visited more often than others. As there are only 36 years of data available it is impossible to test this hypothesis and thus we refrain from pursuing this approach further. On the other hand, there are sufficient data to update the prior probabilities in time. There are several ways in which this can be done. For example, one can use information on ENSO to shift the prior probabilities, or one can make use of the ensemble information by allowing the transition probabilities to change in time. We pursue the latter approach, as it makes use of the information within the SEAS5 ensemble and does not require any external information. It is explained and evaluated in the next two sections followed by an analysis of the resulting interannual variability in section 5.

a. Updating the transition probabilities

To obtain more informative prior regime probabilities, the transition probabilities T_{ij} from regime i to j are updated following the ensemble behavior. This allows not only for (fixed) persistence to inform the prior, but also nonstationary external factors, such as ENSO, through the ensemble statistics. Although there are not sufficient data to robustly estimate the transition probabilities directly, they can be inferred from the occurrence rates. The main assumption we make when updating the transition matrix \mathbf{T} in time is that the regime probabilities are approximately stationary with respect to the current best estimate of the transition matrix. That is, we look for a transition matrix $\mathbf{T}(t)$ for which the regime probabilities averaged over the ensemble at time t , $\bar{\mathbf{P}}(t)$, are approximately stationary:

$$\mathbf{T}(t)\bar{\mathbf{P}}(t) = \bar{\mathbf{P}}(t) + \boldsymbol{\epsilon}' \quad (7)$$

Here $\boldsymbol{\epsilon}'$ is a noise term. Note that the climatological transition probabilities \mathbf{P}^c are (nearly) stationary with respect to the transition matrix \mathbf{T}^c . The aim thus is to find a transition matrix $\mathbf{T}(t)$ for which (7) holds. In addition we have that a transition matrix is normalized, meaning its columns each sum to unity:

$$\sum_{i=1}^k T_{ij} = 1, \quad \forall j \in 1, \dots, k. \quad (8)$$

This gives two equations that are used to update $\mathbf{T}(t)$ at each time step t . The problem of finding the values of the transition matrix $\mathbf{T}(t)$ is ill posed as there are not sufficient constraints, which means some choices need to be made in determining its values. The approach we propose in the following paragraph is one that follows the regime dynamics closely and is least biased in the sense that the deviations from \mathbf{T}^c are equally distributed over all six regimes.

The regime dynamics is dominated by persistence (i.e., the probability of a regime to transition to itself, corresponding to the diagonal elements of the transition matrix), as can be seen in (5). Therefore we focus on these diagonal elements $T_{ii}(t)$ for updating the matrix $\mathbf{T}(t)$ in time. Writing out (7) element-wise while separating the diagonal and off-diagonal elements yields

$$T_{ii}(t)\bar{P}_i(t) + \sum_{j \neq i}^k T_{ij}(t)\bar{P}_j(t) = \bar{P}_i(t) + \boldsymbol{\epsilon}'_i, \quad \forall i \in 1, \dots, k. \quad (9)$$

Because the diagonal terms dominate, we assume the off-diagonal elements do not differ much from the climatological values, that is $T_{ij}(t) \approx T_{ij}^c$ for all $i \neq j$. This yields an approximate equation for the diagonal elements of $\mathbf{T}(t)$:

$$T_{ii}(t)\bar{P}_i(t) \approx \bar{P}_i(t) - \sum_{j \neq i}^k T_{ij}^c(t)\bar{P}_j(t). \quad (10)$$

When a particular regime is less populated than it is in climatology, the other regimes will conversely be more populated, implying a larger negative term on the right-hand side of (10) and thus a smaller value of the self-transition probability, which makes physical sense. Note that this approximation breaks down when $\bar{P}_i(t)$ is very small relative to the other $\bar{P}_j(t)$, in which case we set $T_{ii}(t) = 0$ to prevent negative values. Starting from the updated diagonal elements, the off-diagonal elements are computed using (8) with an equal distribution of the perturbation from the climatological value over the off-diagonal terms.

The estimation of the transition matrix \mathbf{T} in essence is the same as trying to fit an HMM to the data. The difficulty here is the limited availability of data, where we only consider data at one point in time to retain the sequential nature of the method. This makes the use of less heuristic, more sophisticated methods unreliable due to the large impact of noise on the data. If many more ensemble members were available, something like the Baum-Welch algorithm might be a worthwhile approach for estimating \mathbf{T} (Baum et al. 1970). Starting the updating of $\mathbf{T}(t)$ from the diagonal elements and adjusting the off-diagonal elements equally is not the only option. It might even be better to not adjust the off-diagonal elements equally. However, since $\bar{\mathbf{P}}(t)$ is an average over only 51 ensemble members, robustness would be an issue when making any further assumptions in updating $\mathbf{T}(t)$ and hence we stick to the simplest approach.

The above method is equivalent to considering $\mathbf{T}(t)$ as the climatological transition matrix plus a perturbation




FIG. 7. The (top) prior probability, (top middle) conditional regime likelihood, and (bottom middle) Bayesian regime probability for the 42nd ensemble member in the Bayesian regime assignment procedure for the winter of 1992/93. The solid line shows the sequential Bayesian approach, and the dashed line shows the ensemble approach discussed in section 4. (bottom) The difference between the updated self-transition probabilities in the ensemble approach and the climatological values.

[i.e., $\mathbf{T}(t) = \mathbf{T}^c + \mathbf{T}'(t)$], and subsequently assuming that the perturbations to the off-diagonal terms are small. An alternative way of looking at this is by considering it as a Markov regression model (Hamilton 1989; Krolzig 1997). That is, we write the transition matrix \mathbf{T} as

$$\mathbf{T}(t) = \mathbf{T}^c + \sum_m \alpha_m(t) \mathbf{T}_m. \quad (11)$$

Here \mathbf{T}_m are matrices that set the shape of the perturbations to the climatological transition matrix, where the sum over each of the columns is zero for every m , and $\alpha_m(t)$ gives the strength of that term at time t . For a choice of

$$\mathbf{T}_m = \begin{pmatrix} 0 & \dots & -\frac{1}{k-1} & \dots & 0 \\ \vdots & & 1 & & \vdots \\ 0 & & -\frac{1}{k-1} & & 0 \end{pmatrix}, \quad (12)$$

where the m th column is nonzero, this is exactly equivalent to the approach mentioned before. Here the α_m can be computed using the same assumptions as discussed before. This shows that there are several ways of looking at the problem

that yield the same outcome, increasing the confidence in this approach.

b. Evaluation

To get an idea of how this approach can inform the prior probabilities consider Fig. 7, which shows both the sequential and ensemble Bayesian regime assignments for the (randomly chosen) 42nd ensemble member during the winter of 1992/93. This is the same winter for which the 23rd ensemble member is shown in Fig. 4. As an example, consider the probability of AR-. Around days 5–10 the ensemble indicates this regime is less likely, as shown by a lower self-transition probability, lowering the prior probability of the regime. On the other hand, from day 25 onward AR- is more likely according to the ensemble, increasing its prior probability relative to the sequential approach. In most cases changes to the final probabilities are small. The only exceptions occur when a regime is deemed very unlikely (i.e., does not occur in any of the other ensemble members), as happens twice for the SB+ regime between day 60 and 90. In these two cases a high observed likelihood for SB+ is reduced substantially in the Bayesian probabilities in favor of the second most-likely regime according to the likelihood (e.g., a 90% likelihood is reduced to a 35% Bayesian probability). Yet importantly, the Bayesian probability of this regime is still nonzero, so it can quickly respond to new information. The overall regime




FIG. 8. The interannual variability of the occurrence rates for the ensemble Bayesian regime assignment for SEAS5 (colors, with 95% confidence interval shaded), with the sequential Bayesian approach indicated by the black dashed lines. The gray-shaded areas bounded by the gray dotted lines indicate the 10th and 90th percentile of the ensemble Bayesian assignment for each regime. The black dotted curve shows the ERA-Interim variability, and the box-and-whisker plots on the right show the average occurrence rate during very strong El Niño (indicated by the vertical red solid lines) and strong La Niña years (indicated by the vertical blue dash-dotted lines).

frequencies and autocorrelation are not affected and remain as shown in Fig. 6 for the sequential approach.

5. Interannual variability

The interannual variability as obtained using the ensemble Bayesian regime assignment is shown in Fig. 8, with the result of the sequential Bayesian approach shown for reference (the interannual variability of the sequential Bayesian approach is nearly identical to that obtained for the k -means clustering assignment). The primary signal in the variability is found during very strong El Niño years (vertical red solid lines) with SB- and NAO- showing an increase in frequency, while AR+, AR-, and NAO+ show a decrease in frequency. The signal during strong La Niña years (vertical blue dash-dotted lines) is less clear, with on average an increase in NAO+ and decrease of NAO- frequency. However, not every individual event matches this behavior. To define El Niño and La Niña years the Niño-3.4 index is used (Trenberth 1997). Strong years correspond to a threshold of ± 1.5 , and very strong years to a threshold of ± 2 . The asymmetry in the thresholds used for El Niño and La Niña years is due to there being no very strong La Niña events in the considered time period. These results, with a less-pronounced regime response to La Niña than to El Niño, reflect the well-known nonlinearity of the response to ENSO (Straus and Molteni 2004; Toniazzo and

Scaife 2006) and are in line with those obtained in Falkena et al. (2022) using a regularization on the ensemble members. The boxes on the right of each panel show the average regime frequencies during the identified El Niño and La Niña years for both the sequential and ensemble Bayesian approach, where there is an asymmetric response to ENSO for both methods. Some enhancement of the signal is found using the ensemble Bayesian regime assignment, which is most clear for the AR- and SB- regimes. The ERA-Interim variability from the sequential Bayesian approach is shown as well to give a perspective on the magnitude of the interannual variability.

To further consider the effect the updating of the transition matrix in the ensemble approach has on the interannual variability, consider Fig. 9, which shows the difference between the sequential and ensemble Bayesian regime assignment as well as the yearly average change to the self-transition probabilities, or persistence, of the regimes following the ensemble approach. Note that on average the perturbation to the self-transition probabilities is negative. The effect of the ensemble Bayesian approach on the regime frequencies is clearly visible for AR+, AR-, and SB-, where the signal in response to El Niño is enhanced. For NAO+ a strong increase in regime frequency is found for the 1988/89 La Niña, together with a weak change during El Niño years. NAO- and SB+ do not show much difference in interannual variability between the two methods, although in the latter case there is little signal to




FIG. 9. The difference in interannual variability of the occurrence rates between the standard and ensemble Bayesian regime probabilities (solid lines; left y axis), and the change in the self-transition probability for the regimes following the ensemble (dashed lines; right y axis).

enhance. The changes in the self-transition probabilities in general match those found in the regime frequencies, as expected. One aspect to note here is that for NAO+ the changes in the self-transition probability are relatively larger than those in the regime frequencies, especially when comparing with SB-.

The response of the changes in regime frequency to El Niño events found using the ensemble Bayesian approach appears to show a true signal and is very unlikely to have arisen by chance. To understand this, consider the change in regime frequency for SB-. The marginal probability of a very strong El Niño event is 3/36 (3 events in 36 years), so the chance of the first increase in SB- frequency aligning with El Niño is 3/36. Then, given that the first El Niño event has already happened, the probability of the second spike aligning is 2/35 and for the third 1/34. This gives a probability of $3/36 \times 2/35 \times 1/34 \approx 10^{-4}$ for the alignment occurring by chance. The alignment of the increase/decrease in frequency for the other regimes only further decreases the probability of this being by chance. Also note that the response of both AR+ and AR- is a decrease in regime frequency during El Niño years, indicating another aspect of nonlinearity in the circulation response to ENSO.

Some of these signals in response to ENSO can already be picked up using 10-member ensembles. In Fig. 10 the interannual variability of the regime frequency is shown for 50 random 10-member ensembles obtained from the full SEAS5 ensemble. For the full ensemble the strongest signal was found for SB- during very strong El Niño years, and this is the signal that jumps out most strongly again. To quantify this

the probability of detection (POD) and false alarm ratio (FAR) for the 10-member ensembles are considered for peaks or troughs in regime frequency aligning with El Niño and La Niña (Fig. 11). Here, peaks and troughs are considered as exceedances with respect to the n th percentile. The POD is computed as the number of peaks/troughs aligning with El Niño/La Niña years over the total number of El Niño/La Niña years, and the FAR is computed as the number of peaks/troughs outside those El Niño/La Niña years divided by the total number of peaks/troughs. As expected, for El Niño there is a high POD for peaks in the SB- regime frequency with a relatively low FAR (Fig. 11a). Also for NAO- (peaks), NAO+, AR+, and AR- (troughs) there is some signal, with the FAR being comparable to the POD. For La Niña years there is some signal for NAO+, AR+ (peaks), and NAO- (troughs), but it is not as strong as for SB- in El Niño years (Fig. 11b). This is to be expected as we cannot expect to identify strong signals using a smaller ensemble if they are not clear in the full ensemble. Nevertheless, the relatively high PODs for these three regimes are encouraging.

To see whether the found response to ENSO for some regimes also reflects a predictable signal in the observations we regress the ERA-Interim interannual variability onto the SEAS5 one, as in Falkena et al. (2022). The results for this, looking at the sequential and ensemble Bayesian approach, are shown in Table 1. In addition to the p value, we also compute the Bayes factor, which is the ratio of the probabilities of the data D given two different hypotheses H_1 and H_2 , that is, $P(D|H_1)/P(D|H_2)$ (Kass and Raftery 1995). Here the first




FIG. 10. The interannual variability of the regime frequency for the ensemble Bayesian approach when applied to (random) ensembles of 10 members. In total 50 random ensembles are shown. The solid red and dash-dotted blue lines indicate very strong El Niño and strong La Niña years, respectively.

hypothesis H_1 is that of a linear regression model, whereas the second hypothesis H_2 assumes a constant, climatological, regime frequency. For its computation we follow the Bayesian information criterion approximation from Wagenmakers (2007). Values of the Bayes factor above 1 indicate that H_1 is more likely, with values between 3 and 20 constituting positive evidence and values over 20 yielding strong evidence toward it (Kass and Raftery 1995).

Using the sequential Bayesian approach we already find some predictable signal for the NAO+ and SB- regimes, with Bayes factors of 7.6 and 5.1 respectively (Table 1). The Bayes factor for NAO- is also above 3, but here the p value is larger, reducing the confidence in this being a true signal. These results are comparable to those found in Falkena et al. (2022), with the regression coefficients being close to 1 for NAO+, NAO-, and SB-. These regression coefficients around 1 indicate that the signal in SEAS5 is of similar magnitude to that in ERA-Interim, showing no evidence of a signal-to-noise paradox for the regime frequencies, in contrast to the NAO index (Falkena et al. 2022). Using the ensemble information to update the transition probabilities increases the predictable signal for NAO+ and SB-, with smaller p values and higher Bayes factors. Also the AR- signal is enhanced with a Bayes factor over 3 although the p value is still relatively large. The enhancement of the NAO+ signal is comparable to that found using a regularized clustering approach, whereas the change for SB- is weaker (a Bayes factor of 13.2 as compared with 5.5; Falkena et al. 2022). On the other hand, the decrease in Bayes factors for NAO- and AR- using a

regularized approach is not found using the ensemble Bayesian method, which shows small increases of the Bayes factors. In Falkena et al. (2022) a significant signal was found using multiple linear regression of ERA-Interim NAO- onto the SEAS5 NAO+ and SB-, which we find here as well with Bayes factors of 21.1 for the sequential method increasing to 26.6 using the ensemble approach. In comparing the two methods, we find that the ensemble Bayesian regime assignment allows us to identify more pronounced interannual variability signals for some regimes while still accounting for the signal of the other regimes.

6. Conclusions and discussion

A new approach exploiting Bayes's theorem given in (1) is proposed to obtain a probabilistic regime assignment of the atmospheric state on a given day, based on preexisting definitions of the regimes. The approach combines climatological likelihood functions with prior information from the previous day, using climatological estimates of regime persistence, to obtain a Bayesian regime probability. This sequential probabilistic regime assignment allows for smoother transitions between the regimes and indicates whenever data do not clearly belong to one regime. In contrast to previously studied methods that used a regularized k -means clustering algorithm (Falkena et al. 2020, 2022) there is no parameter, other than the number of regimes k , that has to be selected. Also, the method can be applied in real time as new data come in. Applying the approach to six wintertime circulation regimes over the Euro-Atlantic sector yields




FIG. 11. The probability of detection (solid) and false alarm ratio (dashed) for a (left) peak or (right) trough in regime frequency in 10-member subsamples of the SEAS5 ensembles occurring in the same year as a (a) very strong El Niño or (b) strong La Niña, as a function of the percentile used for the definition of the peaks and troughs. The colored lines indicate the regime values, and the gray lines show the values for peaks and troughs occurring in random years (i.e., no signal).

an increase in persistence, without affecting the average regime frequencies for both SEAS5 and ERA-Interim (Fig. 6). In addition, for ERA-Interim the 1-day autocorrelation was found to be higher than that obtained using a regularized k -means approach containing a persistence constraint (Falkena et al. 2020). The Bayesian probabilistic regime assignment can help to overcome the need for some of the heuristic devices, such as a “no regime” category, that are commonly used in circulation regime studies (e.g., Cassou et al. 2005; Grams et al. 2017). The regime probabilities indicate when data cannot be clearly assigned to one regime, whereas the incorporation of prior information ensures persistent regime dynamics. Here, the focus has been on the regime dynamics within the winter season and on interannual time scales, leaving the challenging problem of seasonality of regimes aside (e.g., Breton et al. 2022).

A yet more informative prior for the Bayesian approach can be obtained by continuously updating the prior probabilities by

taking information from the full SEAS5 ensemble into account. Starting from the assumption of approximate stationarity of the ensemble mean regime frequencies at each day, the regime transition matrix is updated. This update is started from the diagonal of the transition matrix since the persistence dominates the regime dynamics. The limited availability of data is not sufficient to reliably apply other approaches such as hidden Markov models. This updated transition matrix in turn affects the prior probabilities, leading to more pronounced interannual variability for some regimes. When considering the interannual variability, the response to three very strong El Niño events in recent decades clearly stands out (Fig. 8). During these three winters SB- and NAO- increase in frequency, while NAO+, AR+, and AR- decrease. The signals for AR+, AR-, and SB- are enhanced by the ensemble Bayesian approach relative to the sequential method. The signal during La Niña winters is less pronounced,

TABLE 1. The regression coefficient, p value, and Bayes factor for linear regression of the interannual variability in regime frequency (ERA-Interim onto SEAS5) for all six regimes. In addition, the result of multiple linear regression of the ERA-Interim NAO– frequency against the SEAS5 ensemble mean NAO+ and SB– regime frequencies is shown. Values for both the sequential and the ensemble Bayesian approach are shown.

	Regime	NAO+	NAO–	AR+	SB+	AR–	SB–	MLR	NAO–
Sequential Bayes	Reg. coef	1.170	1.094	–0.504	0.258	1.207	1.083	NAO+	–1.369
	p value	0.052	0.139	0.592	0.795	0.174	0.082	SB–	–1.838
	Bayes factor	7.579	3.251	1.167	1.037	2.696	5.054		0.047
Ensemble Bayes	Reg. coef	1.066	1.035	–0.435	0.225	1.037	0.785	NAO+	–1.429
	p value	0.044	0.133	0.527	0.782	0.136	0.075	SB–	–1.412
	Bayes factor	8.910	3.365	1.240	1.042	3.306	5.487		0.041
									26.641

with the increase in NAO+ frequency during 1988/89 standing out most clearly.

This response to ENSO in the SEAS5 ensemble can already be identified using only a 10-member ensemble. The increase in SB– occurrence during El Niño years is a particularly strong signal and is found in nearly all 10-member ensembles considered (Fig. 10). Also, for NAO+, NAO–, AR+, and AR– significant probabilities of detection for peaks or troughs coinciding with El Niño are found. However, here there also is a substantial false alarm ratio indicating that many peaks or troughs in the ensemble occur in non–El Niño years. For La Niña there also is some signal, but not as strong as for El Niño years. These results suggest that one may not need a very large ensemble to identify regime signals in response to ENSO.

We also use a linear regression analysis to identify predictable signals in the observations on interannual time scales. Here, as in Falkena et al. (2022), NAO+ and SB– were found to be predictable from the SEAS5 ensemble with regression coefficients around 1 (Table 1), suggesting no signal-to-noise deficit for these regimes. The ensemble approach leads to an increase in Bayes factor relative to the sequential method for all regimes, with the largest improvement for NAO+.

ENSO is certainly part of the reason for the predictable signal found with the regression approach, but it is likely that other processes play a role as well. Previous studies have linked the frequency of Euro-Atlantic circulation regimes to the Madden–Julian oscillation (e.g., Cassou 2008; Straus et al. 2015; Lee et al. 2019, 2020) and the stratospheric polar vortex (e.g., Charlton-Perez et al. 2018; Domeisen et al. 2020), and it would be interesting to see whether the Bayesian approach to regime assignment can aid in better understanding the links between these processes and the regime frequencies. In that respect, the clear improvement in persistence obtained from the sequential method (Fig. 5) should be useful for such S2S applications, even if the seasonal averages are not much affected. Information about other climatic processes that are known to affect the regime occurrence can be used to obtain an informative prior for the regime probabilities. For example, knowledge of the states of ENSO or the stratospheric vortex can inform the prior regime probabilities. Such priors

can be used for both model ensembles as well as reanalysis datasets and aid in better distinguishing the signal from the noise.

The use of the Bayesian regime assignment approach is not limited to atmospheric circulation regimes; it can be applied to any case in which the data can be separated into two or more regimes. For example, one can think of the two phases of the NAO or the jet latitude (Woolings et al. 2010). For the application one needs some information on the regime likelihood function and a way to obtain an informative prior. In most cases the latter will be the most challenging and requires a thorough understanding of the processes involved. For circulation regimes a prior based on climatological transition probabilities, which automatically builds in persistence, was shown to be a suitable and natural choice, and incorporating information from a full ensemble enhanced the interannual signal. Depending on the regime process considered other choices for the prior may be more suitable.

Acknowledgments. Author Falkena was supported by the Centre for Doctoral Training in Mathematics of Planet Earth, with funding from the U.K. Engineering and Physical Sciences Research Council (EPSRC) (Grant EP/L016613/1). The research of author de Wiljes has been partially funded by Deutsche Forschungsgemeinschaft (DFG)—SFB1294/1, 318763901. We thank the three reviewers for their constructive feedback.

Data availability statement. ERA-Interim data are available online (<https://www.ecmwf.int/en/forecasts/dataset/ecmwf-reanalysis-interim>). SEAS5 data are available via the Copernicus Climate Data Store (<https://cds.climate.copernicus.eu/>).

REFERENCES

Acevedo, W., J. de Wiljes, and S. Reich, 2017: Second-order accurate ensemble transform particle filters. *SIAM J. Sci. Comput.*, **39**, A1834–A1850, <https://doi.org/10.1137/16M1095184>.

Ayarzagüena, B., S. Ineson, N. J. Dunstone, M. P. Baldwin, and A. A. Scaife, 2018: Intraseasonal effects of El Niño–Southern Oscillation on North Atlantic climate. *J. Climate*, **31**, 8861–8873, <https://doi.org/10.1175/JCLI-D-18-0097.1>.

Baldo, A., and R. Locatelli, 2023: A probabilistic view on modeling weather regimes. *Int. J. Climatol.*, **43**, 1710–1730, <https://doi.org/10.1002/joc.7942>.

Baum, L. E., T. Petrie, G. Soules, and N. Weiss, 1970: A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Stat.*, **41**, 164–171, <https://doi.org/10.1214/aoms/1177697196>.

Breton, F., M. Vrac, P. Yiou, P. Vaittinada Ayar, and A. Jézéquel, 2022: Seasonal circulation regimes in the North Atlantic: Towards a new seasonality. *Int. J. Climatol.*, **42**, 5848–5870, <https://doi.org/10.1002/joc.7565>.

Büeler, D., L. Ferranti, L. Magnusson, J. F. Quinting, and C. M. Grams, 2021: Year-round sub-seasonal forecast skill for Atlantic–European weather regimes. *Quart. J. Roy. Meteor. Soc.*, **147**, 4283–4309, <https://doi.org/10.1002/qj.4178>.

Cassou, C., 2008: Intraseasonal interaction between the Madden–Julian Oscillation and the North Atlantic Oscillation. *Nature*, **455**, 523–527, <https://doi.org/10.1038/nature07286>.

—, T. Laurent, and A. S. Phillips, 2005: Tropical Atlantic influence on European heat waves. *J. Climate*, **18**, 2805–2811, <https://doi.org/10.1175/JCLI3506.1>.

Charlton-Perez, A. J., L. Ferranti, and R. W. Lee, 2018: The influence of the stratospheric state on North Atlantic weather regimes. *Quart. J. Roy. Meteor. Soc.*, **144**, 1140–1151, <https://doi.org/10.1002/qj.3280>.

Cortesi, N., V. Torralba, L. Lledó, A. Manrique, S. Nube, and G. Reviriego, 2021: Yearly evolution of Euro-Atlantic weather regimes and of their sub-seasonal predictability. *Climate Dyn.*, **56**, 3933–3964, <https://doi.org/10.1007/s00382-021-05679-y>.

Corti, S., F. Molteni, and T. N. Palmer, 1999: Signature of recent climate change in frequencies of natural atmospheric circulation regimes. *Nature*, **398**, 799–802, <https://doi.org/10.1038/19745>.

Dee, D. P., and Coauthors, 2011: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quart. J. Roy. Meteor. Soc.*, **137**, 553–597, <https://doi.org/10.1002/qj.828>.

Del Moral, P., 1997: Nonlinear filtering: Interacting particle resolution. *C. R. Acad. Sci., Ser. I: Math.*, **325**, 653–658, [https://doi.org/10.1016/S0764-4442\(97\)84778-7](https://doi.org/10.1016/S0764-4442(97)84778-7).

Dempster, A. P., N. M. Laird, and D. B. Rubin, 1977: Maximum likelihood from incomplete data via the EM algorithm. *J. Roy. Stat. Soc., 39B* (1), 1–22, <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>.

Domeisen, D. I. V., C. M. Grams, and L. Papritz, 2020: The role of North Atlantic–European weather regimes in the surface impact of sudden stratospheric warming events. *Wea. Climate Dyn.*, **1**, 373–388, <https://doi.org/10.5194/wcd-1-373-2020>.

Dorrrington, J., K. Strommen, F. Fabiano, and F. Molteni, 2022: CMIP6 models trend toward less persistent European blocking regimes in a warming climate. *Geophys. Res. Lett.*, **49**, e2022GL100811, <https://doi.org/10.1029/2022GL100811>.

Doucet, A., N. de Freitas, and N. Gordon, 2001: *Sequential Monte Carlo Methods in Practice*. Vol. 1. Springer, 581 pp.

Evensen, G., and P. J. van Leeuwen, 2000: An ensemble Kalman smoother for nonlinear dynamics. *Mon. Wea. Rev.*, **128**, 1852–1867, [https://doi.org/10.1175/1520-0493\(2000\)128<1852:AEKSFN>2.0.CO;2](https://doi.org/10.1175/1520-0493(2000)128<1852:AEKSFN>2.0.CO;2).

Falkena, S. K. J., J. de Wiljes, A. Weisheimer, and T. G. Shepherd, 2020: Revisiting the identification of wintertime atmospheric circulation regimes in the Euro-Atlantic sector. *Quart. J. Roy. Meteor. Soc.*, **146**, 2801–2814, <https://doi.org/10.1002/qj.3818>.

—, —, —, and —, 2022: Detection of interannual ensemble forecast signals over the North Atlantic and Europe using atmospheric circulation regimes. *Quart. J. Roy. Meteor. Soc.*, **148**, 434–453, <https://doi.org/10.1002/qj.4213>.

Ferranti, L., S. Corti, and M. Janousek, 2015: Flow-dependent verification of the ECMWF ensemble over the Euro-Atlantic sector. *Quart. J. Roy. Meteor. Soc.*, **141**, 916–924, <https://doi.org/10.1002/qj.2411>.

Franzke, C. L. E., D. T. Crommelin, A. Fischer, and A. J. Majda, 2008: A hidden Markov model perspective on regimes and metastability in atmospheric flows. *J. Climate*, **21**, 1740–1757, <https://doi.org/10.1175/2007JCLI1751.1>.

—, T. J. O’Kane, D. P. Monselesan, J. S. Risbey, and I. Horenko, 2015: Systematic attribution of observed Southern Hemisphere circulation trends to external forcing and internal variability. *Nonlinear Processes Geophys.*, **22**, 513–525, <https://doi.org/10.5194/npg-22-513-2015>.

Grams, C. M., R. Beerli, S. Pfenninger, I. Staffell, and H. Wernli, 2017: Balancing Europe’s wind-power output through spatial deployment informed by weather regimes. *Nat. Climate Change*, **7**, 557–562, <https://doi.org/10.1038/nclimate3338>.

Hamilton, J. D., 1989: A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, **57**, 357–384, <https://doi.org/10.2307/1912559>.

Hannachi, A., and A. O’Neill, 2001: Atmospheric multiple equilibria and non-Gaussian behaviour in model simulations. *Quart. J. Roy. Meteor. Soc.*, **127**, 939–958, <https://doi.org/10.1002/qj.49712757312>.

—, D. M. Straus, C. L. E. Franzke, S. Corti, and T. Woollings, 2017: Low-frequency nonlinearity and regime behavior in the Northern Hemisphere extratropical atmosphere. *Rev. Geophys.*, **55**, 199–234, <https://doi.org/10.1002/2015RG000509>.

Horenko, I., 2010: On clustering of non-stationary meteorological time series. *Dyn. Atmos. Oceans*, **49**, 164–187, <https://doi.org/10.1016/j.dynatmoce.2009.04.003>.

—, 2011a: Nonstationarity in multifactor models of discrete jump processes, memory, and application to cloud modeling. *J. Atmos. Sci.*, **68**, 1493–1506, <https://doi.org/10.1175/2011JAS3692.1>.

—, 2011b: On analysis of nonstationary categorical data time series: Dynamical dimension reduction, model selection, and applications to computational sociology. *Multiscale Model. Simul.*, **9**, 1700–1726, <https://doi.org/10.1137/100790549>.

Hu, C.-C., and P. J. van Leeuwen, 2021: A particle flow filter for high-dimensional system applications. *Quart. J. Roy. Meteor. Soc.*, **147**, 2352–2374, <https://doi.org/10.1002/qj.4028>.

Jain, A. K., 2010: Data clustering: 50 years beyond K-means. *Pattern Recognit. Lett.*, **31**, 651–666, <https://doi.org/10.1016/j.patrec.2009.09.011>.

Johnson, S. J., and Coauthors, 2019: SEAS5: The new ECMWF seasonal forecast system. *Geosci. Model Dev.*, **12**, 1087–1117, <https://doi.org/10.5194/gmd-12-1087-2019>.

Kalman, R. E., 1960: A new approach to linear filtering and prediction problems. *J. Basic Eng.*, **82**, 35–45, <https://doi.org/10.1115/1.3662552>.

Kantas, N., A. Beskos, and A. Jasra, 2014: Sequential Monte Carlo methods for high-dimensional inverse problems: A case study for the Navier–Stokes equations. *SIAM/ASA J. Uncertainty Quantif.*, **2**, 464–489, <https://doi.org/10.1137/130930364>.

Kass, R. E., and A. E. Raftery, 1995: Bayes factors. *J. Amer. Stat. Assoc.*, **90**, 773–795, <https://doi.org/10.1080/01621459.1995.10476572>.

Krolzig, H.-M., 1997: *Markov-Switching Vector Autoregressions*. Springer-Verlag, 357 pp.

Lee, J. C. K., R. W. Lee, S. J. Woolnough, and L. J. Boxall, 2020: The links between the Madden-Julian Oscillation and European weather regimes. *Theor. Appl. Climatol.*, **141**, 567–586, <https://doi.org/10.1007/s00704-020-03223-2>.

Lee, R. W., S. J. Woolnough, A. J. Charlton-Perez, and F. Vitart, 2019: ENSO modulation of MJO teleconnections to the North Atlantic and Europe. *Geophys. Res. Lett.*, **46**, 13 535–13 545, <https://doi.org/10.1029/2019GL084683>.

Majda, A. J., C. L. Franzke, A. Fischer, and D. T. Crommelin, 2006: Distinct metastable atmospheric regimes despite nearly Gaussian statistics: A paradigm model. *Proc. Natl. Acad. Sci. USA*, **103**, 8309–8314, <https://doi.org/10.1073/pnas.0602641103>.

Matsueda, M., and T. N. Palmer, 2018: Estimates of flow-dependent predictability of wintertime Euro-Atlantic weather regimes in medium-range forecasts. *Quart. J. Roy. Meteor. Soc.*, **144**, 1012–1027, <https://doi.org/10.1002/qj.3265>.

Michelangeli, P.-A., R. Vautard, and B. Legras, 1995: Weather regimes: Recurrence and quasi stationarity. *J. Atmos. Sci.*, **52**, 1237–1256, [https://doi.org/10.1175/1520-0469\(1995\)052<1237:WRRAQS>2.0.CO;2](https://doi.org/10.1175/1520-0469(1995)052<1237:WRRAQS>2.0.CO;2).

Mo, K., and M. Ghil, 1988: Cluster analysis of multiple planetary flow regimes. *J. Geophys. Res.*, **93**, 10 927–10 952, <https://doi.org/10.1029/JD093iD09p10927>.

Molteni, F., S. Tibaldi, and T. N. Palmer, 1990: Regimes in the wintertime circulation over northern extratropics. I: Observational evidence. *Quart. J. Roy. Meteor. Soc.*, **116**, 31–67, <https://doi.org/10.1002/qj.49711649103>.

O’Kane, T. J., J. S. Risbey, C. Franzke, I. Horenko, and D. P. Monselesan, 2013: Changes in the metastability of the midlatitude Southern Hemisphere circulation and the utility of non-stationary cluster analysis and split-flow blocking indices as diagnostic tools. *J. Atmos. Sci.*, **70**, 824–842, <https://doi.org/10.1175/JAS-D-12-028.1>.

Quinn, C., D. Harries, and T. J. O’Kane, 2021: Dynamical analysis of a reduced model for the North Atlantic Oscillation. *J. Atmos. Sci.*, **78**, 1647–1671, <https://doi.org/10.1175/JAS-D-20-0282.1>.

Rabiner, L., 1989: A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE*, **77**, 257–286, <https://doi.org/10.1109/5.18626>.

Smyth, P., K. Ide, and M. Ghil, 1999: Multiple regimes in Northern Hemisphere height fields via mixture model clustering. *J. Atmos. Sci.*, **56**, 3704–3723, [https://doi.org/10.1175/1520-0469\(1999\)056<3704:MRINHH>2.0.CO;2](https://doi.org/10.1175/1520-0469(1999)056<3704:MRINHH>2.0.CO;2).

Straus, D. M., and F. Molteni, 2004: Circulation regimes and SST forcing: Results from large GCM ensembles. *J. Climate*, **17**, 1641–1656, [https://doi.org/10.1175/1520-0442\(2004\)017<1641:CRASFR>2.0.CO;2](https://doi.org/10.1175/1520-0442(2004)017<1641:CRASFR>2.0.CO;2).

—, S. Corti, and F. Molteni, 2007: Circulation regimes: Chaotic variability versus SST-forced predictability. *J. Climate*, **20**, 2251–2272, <https://doi.org/10.1175/JCLI4070.1>.

—, E. Swenson, and C.-L. Lappen, 2015: The MJO cycle forcing of the North Atlantic circulation: Intervention experiments with the Community Earth System Model. *J. Atmos. Sci.*, **72**, 660–681, <https://doi.org/10.1175/JAS-D-14-0145.1>.

Thompson, V., N. J. Dunstone, A. A. Scaife, D. M. Smith, J. M. Slingo, S. Brown, and S. E. Belcher, 2017: High risk of unprecedented UK rainfall in the current climate. *Nat. Commun.*, **8**, 107, <https://doi.org/10.1038/s41467-017-00275-3>.

Tonizzzo, T., and A. A. Scaife, 2006: The influence of ENSO on winter North Atlantic climate. *Geophys. Res. Lett.*, **33**, L24704, <https://doi.org/10.1029/2006GL027881>.

Trenberth, K. E., 1997: The definition of El Niño. *Bull. Amer. Meteor. Soc.*, **78**, 2771–2778, [https://doi.org/10.1175/1520-0477\(1997\)078<2771:TDOENO>2.0.CO;2](https://doi.org/10.1175/1520-0477(1997)078<2771:TDOENO>2.0.CO;2).

van der Wiel, K., H. C. Bloomfield, R. W. Lee, L. P. Stoop, R. Blackport, J. A. Screen, and F. M. Selten, 2019: The influence of weather regimes on European renewable energy production and demand. *Environ. Res. Lett.*, **14**, 094010, <https://doi.org/10.1088/1748-9326/ab38d3>.

Vautard, R., 1990: Multiple weather regimes over the North Atlantic: Analysis of precursors and successors. *Mon. Wea. Rev.*, **118**, 2056–2081, [https://doi.org/10.1175/1520-0493\(1990\)118<2056:MWROTN>2.0.CO;2](https://doi.org/10.1175/1520-0493(1990)118<2056:MWROTN>2.0.CO;2).

Vecchi, E., L. Pospíšil, S. Albrecht, T. J. O’Kane, and I. Horenko, 2022: eSPA+: Scalable entropy-optimal machine learning classification for small data problems. *Neural Comput.*, **34**, 1220–1255, https://doi.org/10.1162/neco_a_01490.

Vigaud, N., A. W. Robertson, and M. Tippett, 2018: Predictability of recurrent weather regimes over North America during winter from submonthly reforecasts. *Mon. Wea. Rev.*, **146**, 2559–2577, <https://doi.org/10.1175/MWR-D-18-0058.1>.

Viterbi, A., 1967: Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans. Inf. Theory*, **13**, 260–269, <https://doi.org/10.1109/TIT.1967.1054010>.

Wagenmakers, E.-J., 2007: A practical solution to the pervasive problems of p values. *Psychon. Bull. Rev.*, **14**, 779–804, <https://doi.org/10.3758/BF03194105>.

Woollings, T., A. Hannachi, and B. Hoskins, 2010: Variability of the North Atlantic eddy-driven jet stream. *Quart. J. Roy. Meteor. Soc.*, **136**, 856–868, <https://doi.org/10.1002/qj.625>.