RMM-VAE: a machine learning method for identifying probabilistic weather regimes targeted to a local-scale impact variable

Fiona Raphaela Spuler
University of Reading
Reading, UK
f.r.spuler@pgr.readi
ng.ac.uk

Marlene Kretschmer
Universität Leipzig
Leipzig, Germany
marlene.kretschmer@un
i-leipzig.de

Yevgeniya Kovalchuck
University College London
London, UK
y.kovalchuk@ucl.ac.uk

Magdalena Alonso Balmaseda
ECMWF
Reading, UK
magdalena.balmaseda@ecmwf.int

Theodore G. Shepherd
University of Reading
Reading, UK
theodore.shepherd@reading.ac.uk

Abstract

Identifying large-scale atmospheric patterns that modulate extremes in local-scale variables such as precipitation has the potential to improve long-term climate projections as well as extended-range forecasting skill. This paper proposes a novel probabilistic machine learning method, RMM-VAE, based on a variational autoencoder architecture for identifying weather regimes targeted to a local-scale impact variable. The new method is compared to three existing methods in the task of identifying robust weather regimes that are predictive of precipitation over Morocco while capturing the full phase space of atmospheric dynamics over the Mediterranean. RMM-VAE performs well across these different objectives, outperforming linear methods in reconstructing the full phase space and predicting the target variable, highlighting the potential benefit of applying the method to various climate applications such as downscaling and extended-range forecasting.

1 Introduction

Large-scale atmospheric circulation modulates the occurrence of regional extremes such as heavy precipitation and heatwaves [1]. These extreme events cause devastating impacts to people and livelihoods across the planet as recently witnessed in the tragic flooding events in September 2023 in Libya [2]. In a changing climate, preparing for these extreme events becomes even more important which requires, on the one hand, an understanding of their projected long-term changes to make well-informed and robust adaptation decisions, and on the other hand, an improvement of their near-term prediction skill to give people and other stakeholders such as governments and humanitarian organizations more time to prepare ahead of an upcoming event.

To achieve this, investigating the low-frequency variability and patterns of the large-scale atmospheric flow and their link to regional extremes is important, as they can act as mediators between global patterns of change in the climate system and regional impacts. One common approach to analyzing this low-frequency variability is the identification of recurrent and persistent atmospheric states, so-called weather regimes. Typically, weather regimes are identified based on statistical dimensionality reduction and clustering techniques informed by a physical understanding of the underlying dynamics [3]. Such atmospheric circulation patterns have been used to disentangle

dynamic and thermodynamic components of climate change for extreme event attribution, to downscale or bias-adjust climate models [4]–[6], and to quantify the role of atmospheric internal variability in observed trends [7], [8]. Moreover, weather regimes have been shown to improve the usability and skill of forecasts at sub-seasonal-to-seasonal timescales [9]–[12].

To identify weather regimes, a combination of dimensionality reduction and clustering is commonly applied to gridded geopotential height or sea level pressure data of the region of interest. Following [13], principal component analysis (PCA) and k-means have established themselves as common choices for dimensionality reduction and clustering respectively, although various other linear statistical methods have been studied in the field [3]. Recent publications have also explored non-linear [14] and probabilistic [15] methods for identifying weather regimes.

While the relationship between weather regimes and extremes in local impact variables (e.g. extreme precipitation) is a key motivation for their investigation, there are no comprehensive methods available for identifying probabilistic weather regimes targeted to a specific impact variable, without compromising regime completeness or persistence [16]–[18]. By 'targeting' weather regimes, we mean identifying patterns in the high-dimensional space of atmospheric dynamics that are particularly predictive of the local-scale variable in question while still capturing the full dynamics of the atmospheric phase space. As this last point is particularly relevant for extending the forecast range of the impact variable through the extended range predictability of the patterns, building a purely predictive ML model of the impact variable is not useful in this context.

This paper presents a novel probabilistic machine learning method for identifying weather regimes that captures the full phase space of atmospheric dynamics in a reduced space while providing enhanced predictability of a local scale impact variable. The proposed method, which we call RMM-VAE (Regression Mixture Model Variational Autoencoder), integrates both targeted dimensionality reduction using a variational autoencoder (VAE), and probabilistic clustering using a mixture model (MM) into one coherent statistical model, extending previous machine learning architectures reported in [19], abbreviated R-VAE (Regression - VAE) here, that incorporate a co-variate into the dimensionality reduction step of a VAE model, without combining it with a clustering.

To demonstrate the performance of the RMM-VAE method, it is compared to existing methods in the task of identifying weather regimes targeted to precipitation over Morocco, which serves as our target variable here. To evaluate the performance of RMM-VAE, we compare our method to two established approaches (PCA + k-means, and CCA + k-means), as well as the R-VAE method combined with k-means clustering (R-VAE + k-means). The performance of all methods is assessed on three possible objectives of targeted clustering: reconstruction of the input space, persistent and well-separated clusters, and predictive skill with respect to the target variable.

2   Proposed RMM-VAE method

Figure 1 (left side) shows a schematic of the proposed method RMM-VAE. We introduce a scalar co-variate t into the inference model, extending the architecture developed by [19] by conditioning the dimensionality reduction on a probabilistic cluster assignment [20], [21]. Using the graphical model shown in Figure 1 (right side), the joint probability distribution of the model is obtained and the loss function derived using Bayesian Variational Inference [22] (details in appendix A).
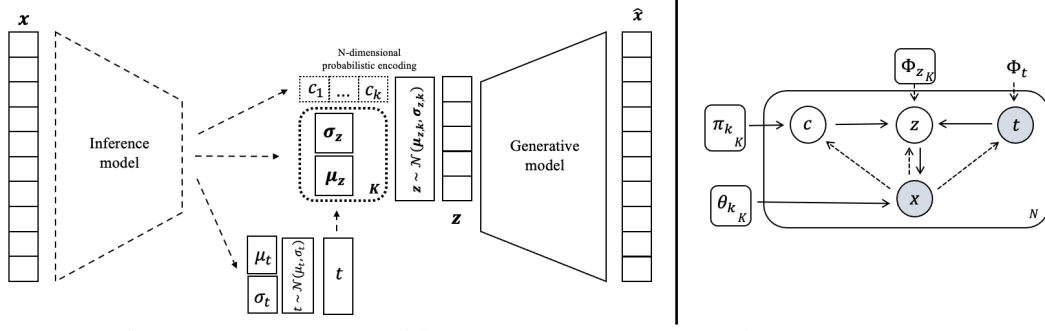
Figure 1: Left side: Schematic diagram of the RMM-VAE method which combines a regression VAE (R-VAE, based on [19]) with probabilistic clustering using mixture models (MM). Right side: RMM-VAE model written as a probabilistic graphical model in plate notation to facilitate the derivation of the loss function using Bayesian variational inference. x represents the high-dimensional input data, z the latent space with prior ϕ_z depending on the cluster assignment, as well as parameters θ depending on the cluster assignment, t the scalar target variable, with probabilistic cluster assignment c with prior π_k. In both panels, dashed lines indicate the inference model and solid lines the generative model.

3   Experiments

The performance of RMM-VAE is compared to two more established linear statistical methods, namely PCA and canonical correlation analysis (CCA), respectively combined with k-means clustering. The performance of the integrated, probabilistic clustering implemented in RMM-VAE is also compared with a version that separates targeted dimensionality reduction in an architecture based on [19] with a separate k-means clustering step (R-VAE + k-means). A description of PCA, CCA, k-means and a standard VAE architecture can be found in [23], and a derivation of the R-VAE method in [19].

Table 1 provides an overview of the compared methods along with relevant hyperparameters. A 10-dimensional latent space was implemented for all the methods, and cluster numbers between 4 and 10 were investigated. For both VAE methods, the inclusion of a hyperparameter β [24] which changes the weight of the reconstruction term in the loss function was investigated (v1 and v2).

Table 1: Overview of the methods applied to the identification of weather regimes and associated parameter choices

Abbr.	Method	β	Precipitation data
PCA	PCA + k-means	-	None
CCA	CCA + k-means	-	Full daily precipitation field at 0.25° resolution
R-VAE v1	R-VAE + k-means	1	Spatially averaged daily precipitation (scalar)
R-VAE v2	R-VAE + k-means	0.1	" - "
RMM-VAE v1	RMM-VAE	1	" - "
RMM-VAE v2	RMM-VAE	0.5	" - "

To analyze weather regimes over the Mediterranean region (latitude: 25°N — 50°N; longitude: 20°W — 45°E) applying these methods, geopotential height at 500hPa (z500) from ERA5 reanalysis data, 1940 — 2022 was using as input data x. The data was standardized by subtracting the daily mean and dividing by the standard deviation over the considered time period. ERA5 reanalysis data of daily total precipitation data in a region over Morocco (latitude: 30°N — 36°N; longitude: 11°W — 0°E) over the same time period were used as the target variable. Precipitation data was normalized by applying a Box-Cox transformation at each grid cell. A three-day average of daily

3

108 total precipitation around each day was taken to mirror methodologies previously applied in the
109 study of precipitation extremes over the Mediterranean [25]. The same was applied to z500 data to
110 match the two datasets.

4   Results

112 Analyzing the latent space of the four methods shown in Figure 2 (top row), we find that both of the
113 targeted VAE methods, RMM-VAE and R-VAE + k-means, disentangle the dimension in the latent
114 space associated with the target variable, which is in line with the findings presented by [19] for the
115 R-VAE method. However, what is interesting is that the cluster assignment (Figure 2, bottom row)
116 differs between the two methods: R-VAE + k-means, which carries out the cluster assignment in a
117 separate step, identifies the clusters in bands along the dimension associated with the target variable.
118 In contrast, RMM-VAE, which integrates a probabilistic cluster assignment with the targeted
119 dimensionality reduction, shows less organization of the clusters with respect to the target variable.
120
121 The performance of the dimensionality reduction component was analyzed by computing the root
122 mean square error between the reconstructed space and the original input data. We find that both
123 VAE methods outperform the two linear statistical methods, and that decreasing the weight of the
124 reconstruction loss term ($\beta < 1$) in the v2 versions increases the reconstruction loss. When
125 investigating the reconstruction of individual data points, we find that decreasing the β-parameter
126 focuses the dimensionality reduction on the region immediately surrounding Morocco. The
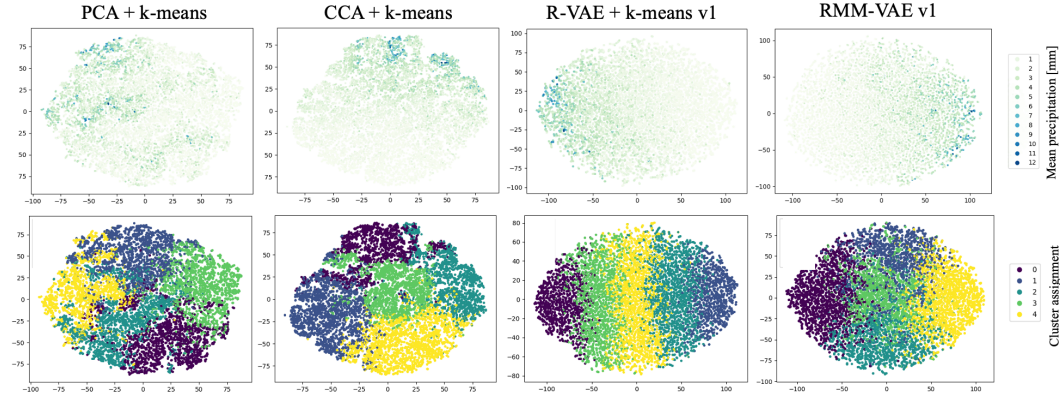127 figures to support these results are shown in the supplementary appendix B.
128

129
130 *Figure 2: Visualization of the 10-dimensional latent space reduced to two dimensions using t-distributed*
131 *stochastic nearest-neighbor embedding (t-SNE). Different values of perplexity were tested and perplexity=10*
132 *was chosen as it shows representative results. Embedded data points are colored first according to the target*
133 *variable, total mean precipitation (top row) with darker colors referring to stronger precipitation, and then*
134 *according to the cluster they are subsequently assigned to in the corresponding clustering method (bottom*
135 *row).*

136 Investigating the predictive performance of the clusters using the Ranked Probability Skill
137 Score (supporting derivation and figures shown in supplementary appendix C), we find that
138 both VAE methods outperform both CCA and PCA up to a cluster number of 9, which shows
139 the potential of targeting weather regimes to a local impact variable by introducing a co-variate
140 to a VAE architecture. While the R-VAE + k-means method outperforms RMM-VAE in terms
141 of predictive skill, the R-VAE + k-means clusters are less persistent and separable compared
142 to RMM-VAE, assessed using the silhouette score. In both methods, the v2 versions which
143 deprioritize the reconstruction term in the loss function ($\beta < 1$) identify clusters with higher
144 predictive skill but lower cluster persistence and separability, compared to the respective v1
145 versions ($\beta = 1$). This finding highlights the trade-off between the predictive skill of weather
146 regimes for a given target variable, which would be maximized in a purely predictive model,
147 and identifying robust clusters, which is important for using the weather regimes for extended-
148 range prediction and for linking regional extremes to global patterns of climate change.

4

5 Conclusion

The paper introduced a novel machine learning method, RMM-VAE, for the identification of weather regimes with respect to a scalar target variable. The novelty of this method lies in developing an architecture that combines non-linear, probabilistic, and targeted dimensionality reduction with probabilistic clustering using mixture models in a coherent Bayesian probabilistic framework. The new method performs well across all the different objectives analyzed, outperforming the linear methods in reconstructing the full phase space and in predicting the target variable. Compared to the other machine learning method, R-VAE + k-means, the proposed RMM-VAE method loses in predictive skill but identifies more persistent and separable clusters, which is relevant for various climate applications such as extended-range forecasting or the downscaling of climate models, highlighting the potential benefit of applying the RMM-VAE method to these use-cases. In future work, we are planning to analyze the dynamical processes captured by the targeted clusters and extend the architecture to identify joint clusters between two high-dimensional spaces.

Acknowledgments

References

[1]   T. G. Shepherd, 'Atmospheric circulation as a source of uncertainty in climate change projections', *Nat. Geosci.*, vol. 7, no. 10, Art. no. 10, Oct. 2014, doi: 10.1038/ngeo2253.

[2]   'Libya floods: Why damage to Derna was so catastrophic', *BBC News*, Sep. 13, 2023. Accessed: Sep. 27, 2023. [Online]. Available: https://www.bbc.com/news/world-africa-66799518

[3]   Abdel. Hannachi, D. M. Straus, C. L. E. Franzke, S. Corti, and T. Woollings, 'Low-frequency nonlinearity and regime behavior in the Northern Hemisphere extratropical atmosphere', *Rev. Geophys.*, vol. 55, no. 1, pp. 199–234, 2017, doi: 10.1002/2015RG000509.

[4]   P. Ailliot, C. Thompson, and P. Thomson, 'Space–Time Modelling of Precipitation by Using a Hidden Markov Model and Censored Gaussian Distributions', *J. R. Stat. Soc. Ser. C Appl. Stat.*, vol. 58, no. 3, pp. 405–426, Jul. 2009, doi: 10.1111/j.1467-9876.2008.00654.x.

[5]   D. Maraun *et al.*, 'Precipitation downscaling under climate change: Recent developments to bridge the gap between dynamical models and the end user', *Rev. Geophys.*, vol. 48, no. 3, p. RG3003, Sep. 2010, doi: 10.1029/2009RG000314.

[6]   C. Yang, R. E. Chandler, V. S. Isham, and H. S. Wheater, 'Spatial-temporal rainfall simulation using generalized linear models', *Water Resour. Res.*, vol. 41, no. 11, 2005, doi: 10.1029/2004WR003739.

[7]   C. Heinze-Deml, S. Sippel, A. G. Pendergrass, F. Lehner, and N. Meinshausen, 'Latent Linear Adjustment Autoencoder v1.0: a novel method for estimating and emulating dynamic precipitation at high resolution', *Geosci. Model Dev.*, vol. 14, no. 8, pp. 4977–4999, Aug. 2021, doi: 10.5194/gmd-14-4977-2021.

[8]   L. Terray, 'A dynamical adjustment perspective on extreme event attribution', *Weather Clim. Dyn.*, vol. 2, no. 4, pp. 971–989, Oct. 2021, doi: 10.5194/wcd-2-971-2021.

[9]   S. Allen, G. R. Evans, P. Buchanan, and F. Kwasniok, 'Incorporating the North Atlantic Oscillation into the post-processing of MOGREPS-G wind speed forecasts', *Q. J. R. Meteorol. Soc.*, vol. 147, no. 735, pp. 1403–1418, 2021, doi: 10.1002/qj.3983.

[10] H. C. Bloomfield, D. J. Brayshaw, P. L. M. Gonzalez, and A. Charlton-Perez, 'Pattern-based conditioning enhances sub-seasonal prediction skill of European national energy variables', *Meteorol. Appl.*, vol. 28, no. 4, Jul. 2021, doi: 10.1002/met.2018.

[11] P. L. M. Gonzalez *et al.*, 'Weather patterns in Southeast Asia: Enhancing high-impact weather subseasonal forecast skill', *Q. J. R. Meteorol. Soc.*, 2022, doi: 10.1002/qj.4378.

[12] N. Mastrantonas, L. Magnusson, F. Pappenberger, and J. Matschullat, 'What do large-scale patterns teach us about extreme precipitation over the Mediterranean at medium- and extended-range forecasts?', *Q. J. R. Meteorol. Soc.*, vol. 148, no. 743, pp. 875–890, 2022, doi: 10.1002/qj.4236.

[13] P.-A. Michelangeli, R. Vautard, and B. Legras, 'Weather Regimes: Recurrence and Quasi Stationarity', *J. Atmospheric Sci.*, vol. 52, no. 8, pp. 1237–1256, Apr. 1995, doi: 10.1175/1520-0469(1995)052<1237:WRRAQS>2.0.CO;2.

[14] A. Baldo and R. Locatelli, 'A probabilistic view on modelling weather regimes', *Int. J. Climatol.*, vol. 1, no. 21, 2022, doi: 10.1002/joc.7942.

[15] S. K. J. Falkena, J. de Wiljes, A. Weisheimer, and T. G. Shepherd, 'A Bayesian Approach to Atmospheric Circulation Regime Assignment'. arXiv, Jun. 23, 2022. Accessed: Jan. 11, 2023. [Online]. Available: http://arxiv.org/abs/2206.11576

[16] F. Gadouali, N. Semane, Á. g. Muñoz, and M. Messouli, 'On the Link Between the Madden-Julian Oscillation, Euro-Mediterranean Weather Regimes, and Morocco Winter Rainfall', *J. Geophys. Res. Atmospheres*, vol. 125, no. 8, p. e2020JD032387, 2020, doi: 10.1029/2020JD032387.

215 [17] H. C. Bloomfield, D. J. Brayshaw, and A. J. Charlton-Perez, 'Characterizing the winter
216   meteorological drivers of the European electricity system using targeted circulation types',
217   *Meteorol. Appl.*, vol. 27, no. 1, p. e1858, 2020, doi: 10.1002/met.1858.
218 [18] E. Rouges, L. Ferranti, H. Kantz, and F. Pappenberger, 'European heatwaves: Link to large-
219   scale circulation patterns and intraseasonal drivers', *Int. J. Climatol.*, vol. 1, no. 21, 2023,
220   doi: 10.1002/joc.8024.
221 [19] Q. Zhao, E. Adeli, N. Honnorat, T. Leng, and K. M. Pohl, 'Variational AutoEncoder For
222   Regression: Application to Brain Aging Analysis', *Med. Image Comput. Comput.-Assist.*
223   *Interv. MICCAI Int. Conf. Med. Image Comput. Comput.-Assist. Interv.*, vol. 11765, pp.
224   823–831, Oct. 2019, doi: 10.1007/978-3-030-32245-8_91.
225 [20] Z. Jiang, Y. Zheng, H. Tan, B. Tang, and H. Zhou, 'Variational Deep Embedding: An
226   Unsupervised and Generative Approach to Clustering'. arXiv, Jun. 27, 2017. Accessed:
227   May 19, 2023. [Online]. Available: http://arxiv.org/abs/1611.05148
228 [21] Q. Zhao, N. Honnorat, E. Adeli, and K. M. Pohl, 'Variational Autoencoder with Truncated
229   Mixture of Gaussians for Functional Connectivity Analysis', *Inf. Process. Med. Imaging*
230   *Proc. Conf.*, vol. 11492, pp. 867–879, Jun. 2019, doi: 10.1007/978-3-030-20351-1_68.
231 [22] D. P. Kingma and M. Welling, 'Auto-Encoding Variational Bayes'. arXiv, 2013. Accessed:
232   Jul. 10, 2023. [Online]. Available: http://arxiv.org/abs/1312.6114
233 [23] K. Murphy, *Probabilistic Machine Learning: An Introduction*. 2022.
234 [24] I. Higgins *et al.*, 'β-VAE: LEARNING BASIC VISUAL CONCEPTS WITH A
235   CONSTRAINED VARIATIONAL FRAMEWORK', *ICLR*, 2017.
236 [25] N. Mastrantonas, P. Herrera-Lormendez, L. Magnusson, F. Pappenberger, and J.
237   Matschullat, 'Extreme precipitation events in the Mediterranean: Spatiotemporal
238   characteristics and connection to large-scale atmospheric flow patterns', *Int. J. Climatol.*,
239   vol. 41, no. 4, pp. 2710–2728, 2020, doi: 10.1002/joc.6985.
240 [26] T. Gneiting and A. E. Raftery, 'Strictly Proper Scoring Rules, Prediction, and Estimation',
241   *J. Am. Stat. Assoc.*, vol. 102, no. 477, pp. 359–378, Mar. 2007, doi:
242   10.1198/016214506000001437.

243

Supplementary Appendices

A - RMM-VAE method: loss function

The assumptions of the model can be expressed in a probabilistic graphical model and the joint probability distributions of the model can then be written as:

$$p_\theta(x, z, t, c^k) = p_\theta(x|z)p(z|t)p(t)\,p(z|c^k)p(c^k)$$

$$q_\phi(z, t, c^k|x) = q_\phi(z|x)q_\phi(t|x)\,q_\phi(c^k|x)$$

The loss function can then be derived as:

$$\mathcal{L}(x) = -D_{KL}\big(q_\phi(z, c, t|x)\big|\,p_\theta(x, z, t, c)\big)$$

$$= \sum_k q_\phi(c^k|x)\left[E_{q_\phi(z|x)}\big[\log\big(p_\theta(x|z)\big)\big] - E_{q_\phi(t|x)}\big[D_{KL}\big(q_\phi(z|x)\big|p(z|t)\big)\big]\right.$$

$$\left. - D_{KL}\big(q_\phi(t|x)\big|p(t)\big) - D_{KL}\big(q_\phi(z|x)\big|p(z|c^k)\big)\right] - D_{KL}\big(q_\phi(c^k|x)\big|p(c^k)\big).$$

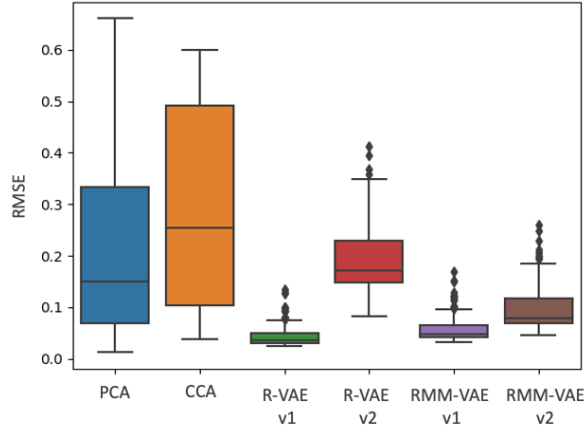### B - Additional figures - reconstruction loss

Figure 3: Distribution of root mean squared error (RMSE) between original input data and data reconstructed from dimensionality-reduced space over all data points, for the different methods.
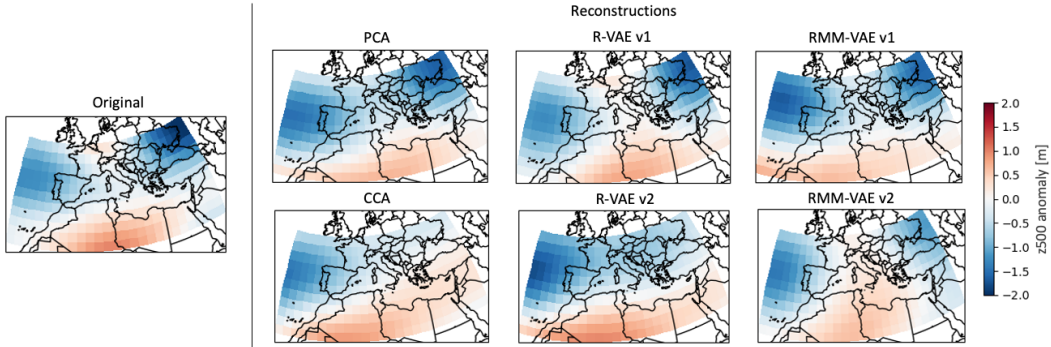

Figure 4: Gridded and normalized z500 anomalies, as detailed in section 2, on an example day 1940-01-04, showing the original data on the left and the reconstructions using different methods on the right.

C - Assessment of predictive skill

To evaluate the informativeness of the weather regimes with respect to the target variable, and in particular their utility for extended-range forecasting, we analyze the Ranked Probability Skill Score (RPSS) which is widely used in forecast evaluation. The RPS is defined as $RPS = \frac{1}{N}\sum_{n=1}^{N}\sum_{j=1}^{m}\big(\delta_{i_n j} - p_j\big)^2 = -\frac{1}{N}\sum_{n=1}^{N}(1 - 2p_{i_n} + \sum_{j=1}^{m} p_j^2)$, and the score (RPSS) calculated with respect to a reference forecast, chosen here to be the climatology over the entire period:

268      $RPSS = 1 - \frac{RPS_{forecast}}{RPS_{climatology}}$, with m forecast categories and N timesteps. $\delta_{i_n j}$ is the Kronecker delta

269   and equals 1 if the observation i at timestep n corresponds to category j, and 0 otherwise. The RPSS

270   is a strictly proper scoring rule to measure the accuracy of a probabilistic prediction of mutually

271   exclusive discrete outcomes [26]. A RPSS of 1 indicates a perfect forecast while low values indicate

272   little, or no skill compared to the reference forecast. Here, we construct a forecast of the target

273   variable given the occurrence of a weather regime and the conditional probability of the target

274   variable given that weather regime: each discrete target is forecast with a probability corresponding

275   to the probability of the weather regime at this given day, multiplied by the climatological

276   conditional probability of the target given the weather regime.

277
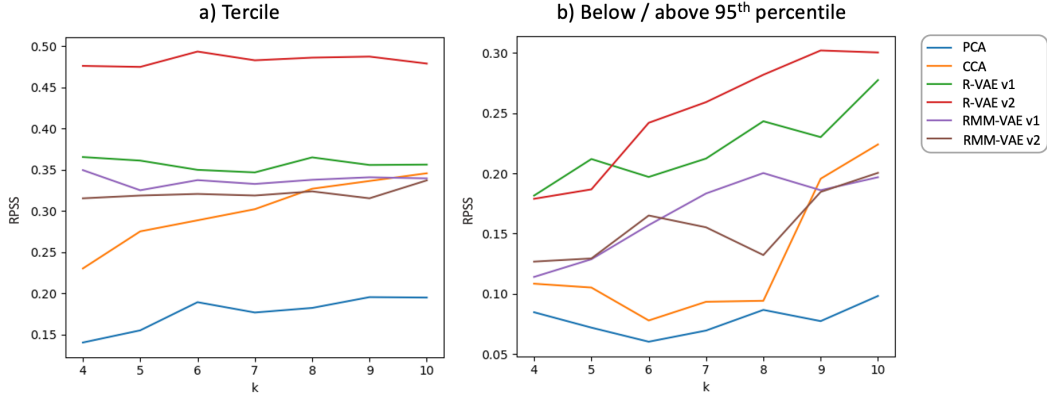
278
279 *Figure 5: Ranked Probability Skill Score for forecast of mean precipitation over Morocco given weather*
280 *regimes at zero lead time for different numbers of weather regimes. a) Skill score for binary forecast above or*
281 *below 95th percentile and c) Skill score for forecast of the tercile of the precipitation distribution. For*
282 *probabilistic clustering, this skill score is computed using the most likely cluster at the given data point.*

283