

IMPACT OF SAMPLING STRATEGIES ON THE MONITORING OF CLIMATE REGIME SHIFTS WITH A LEARNING DATA ASSIMILATION METHOD

Perrine Bauchot*

Angélique Drémeau*

Florian Sévellec^{†‡}

Ronan Fablet^{*‡}

* Lab-STICC, CNRS ENSTA Bretagne IMT Atlantique Univ Brest ENIB UBS, Brest, France

[†] Laboratoire d'Océanographie Physique et Spatiale, Univ Brest CNRS IRD Ifremer, Brest, France

[‡] Odyssey team-project, INRIA CNRS IMT Atlantique, Brest, France

ABSTRACT

In oceanography, the acquisition and processing of observations are crucial to improve the understanding of complex oceanic processes. Considering an idealized model of the North Atlantic ocean circulation, we propose to implement a variational data assimilation method optimized by deep learning to reconstruct abrupt changes in ocean circulation, representing Dansgaard-Oeschger climate events. We show that this assimilation method leads to improved reconstruction performances, particularly at low sampling frequencies. Focusing on this difficult latter case, four sampling strategies are studied more specifically. Our experiments highlight that clusters of three consecutive observations regularly sampled leads to a better monitoring of the ocean circulation regime shifts. These results pave the way for further research in optimal ocean observation.

Index Terms— Observation Sampling, Data Assimilation, Deep Learning, Ocean Monitoring, Ocean Circulation Variability

1. INTRODUCTION

Ocean observation is usually considered as a challenge because of the inhospitality, the high spatio-temporal variability and the vastness of the ocean [1]. Due to its heat capacity compared to the atmosphere and by covering up to 70% of the Earth, the ocean plays a key role in climate regulation [2]. One process involved in climate regulation is the **Atlantic Meridional Overturning Circulation** (AMOC). AMOC corresponds to the transport of warm surface water masses to the North Pole where they are transformed as cold deep water masses, which then flow at depth toward the Southern Ocean [3]. This circulation induces a heat transport in the North Atlantic and captures heat and carbon of the atmosphere in the deep ocean. Several studies highlighted its chaotic behaviour and the consequences on the climate of its possible shutdown [4, 5, 6]. In fact, the AMOC abrupt slow-downs have been hypothesized to cause abrupt climate changes during the last glacial era [7]. Displayed in $\delta^{18}O$ isotope records, these episodes are called the Dansgaard-Oeschger (D-O) events

[8]. The monitoring of oceanic processes, like the AMOC, is made possible by multiple observation platforms. While satellites provide data at a fixed frequency and increasingly fine resolution of the global ocean surface, *in situ* data reach unobserved areas, like coastal regions, the water column and the seabed with smaller spatio-temporal scales [9, 10]. However, these data are often sparse and require innovative acquisition and processing methods to compensate for the missing observations. **Data assimilation** (DA) is widely used in geosciences to tackle these sampling gaps and address forecasting and reconstruction issues. To solve these inverse problems, two main approaches have arose: sequential methods and variational methods [11, 12]. Among the latter, 4D-Var DA is widely employed and will therefore be considered as a reference in this study [13]. It relies on the minimisation of the discrepancies between the data, the model containing the physical knowledge, and the estimated reconstructed state [14]. **Deep learning** recently opened new horizons in DA. The advent of Neural ODEs and the development of physics-informed neural networks translated into new DA schemes [15, 16, 17]. Especially, 4DVarNet framework [18] provides an end-to-end trainable neural DA model to reconstruct a physical system and jointly optimise a variational representation of the inverse problem. In this paper, we investigate further whether neural inversion schemes open new avenues for the monitoring of climate dynamics, like abrupt regime shifts. We explore the interplay between the observation sampling strategy and the monitoring performance of 4DVarNet. The paper is organized as follows. First, we introduce the 4DVarNet method to solve DA problems by using deep learning. Sec. 3 details our case-study based on an idealised model of the D-O events [4]. Finally, we report numerical experiments for different observation sampling strategies in Sec. 4.

2. A LEARNING DATA ASSIMILATION METHOD

In geosciences, DA [11, 12] relates to the solving of inverse problems for state-space models given by:

$$\begin{cases} \dot{\mathbf{X}}(t) = \mathbf{M}(\mathbf{X}(t)) \\ \mathbf{Y}(t) = \mathbf{H}_t \mathbf{X}(t) + \mathbf{q}_{\text{obs}}(t) \end{cases} \quad (1)$$

where $\mathbf{X}(t)$ is the state of the system at time t and $\mathbf{Y}(t)$ the observations acquired with a gaussian noise such as $\mathbf{q}_{\text{obs}} \sim \mathcal{N}(0, \sigma_{\text{obs}}^2)$. $\mathbf{M}$ stands for the dynamical model – the physical equations modeling the phenomenon of interest –, while $\mathbf{H}_t$ is the observation matrix and implements the sampling strategy.

2.1. Variational Data Assimilation: 4DVar classic

Variational DA relies on the minimisation of a variational cost function J over an assimilation time window ΔT to reach an optimal state sequence of the system $\hat{\mathbf{X}}$. In particular, the 4DVar DA method solves a minimization problem $\hat{\mathbf{X}} = \arg \min_{\mathbf{X}} J(\mathbf{X})$ where J is defined as:

$$J(\mathbf{X}) = \frac{1}{\Delta T} \left[\alpha_{\text{obs}} \int_{t_0}^{t_0+\Delta T} \|\mathbf{Y}(t) - \mathbf{H}_t \mathbf{X}(t)\|_2^2 dt + \alpha_{\text{dyn}} \int_{t_0}^{t_0+\Delta T} \|\phi(\mathbf{X})(t) - \mathbf{X}(t)\|_2^2 dt \right] \quad (2)$$

with $\phi(\mathbf{X}(t)) = \mathbf{X}(t-1) + \int_{t-1}^t \mathbf{M}(\mathbf{X}(t)) dt$. The first term of the variational cost J evaluates the difference between the observations and the state of the system, while the second term targets the deviation of the state sequence from the dynamical prior given by the model $\mathbf{M}$. Usually, a background regularisation [13] is added to J to constrain the stability of the assimilation at $t = t_0$. Weights α_{obs} and α_{dyn} balance the relative importance of the observation and prior terms.

4DVar DA classically uses gradient descent to minimize (2) based on adjoint methods [11]. It is considered as a state-of-the-art method, called 4DVar classic hereafter. This approach is computationally expensive and may suffer from robustness issues when dealing with strongly non-linear dynamics [19]. In this respect, 4DVarNet enables to learn an optimal parameterization and boosts the resolution of the variational problem.

2.2. 4DVarNet architecture

The 4DVarNet architecture inspires from the unrolling of an iterative gradient descent algorithm [18] for minimization problem (2). Exploiting a meta-learning approach [20], it results in a residual architecture (Fig. 1). Each residual unit applies a Long-Short Term Memory cell (LSTM) to the variational cost gradient, as, for iteration k , $\mathbf{X}_k = \mathbf{X}_{k-1} - \delta_k$ LSTM $[\nabla J(\mathbf{X}_{k-1})]$ where δ_k weights the updating of $\mathbf{X}_k$. As the LSTM cell can capture long-term dependencies, the overall architecture acts as a momentum-based gradient descent [21] to speed up a robust convergence towards an optimal state.

The model ϕ can be set a priori according to the physical model (4DVarNet-ode) or trained jointly with the solver of the minimization problem (4DVarNet-unet). In the latter case, a U-Net architecture is chosen to represent ϕ for its ability to capture the different scales of the phenomenon [22]. This also




Fig. 1. 4DVarNet scheme

inspires from recent advances in data-driven model discovery [23, 24]. In geophysics, these techniques offer new possibilities to deal with phenomena, which cannot always be equated due to their non-linearities and chaotic natures.

2.3. Learning setting

We train the parameters of the LSTM cell and U-Net prior ϕ in a supervised manner. The training loss involves a trade-off of the mean square error (MSE) between the ground truth and the reconstructed state $\hat{\mathbf{X}}$, the MSE between their time derivatives and the MSE between the state $\mathbf{X}$ and its projection according to prior ϕ . The training stage is performed on batches of size 128 with an Adam optimiser [25] and a dropout of 20% to avoid over-fitting [26]. We run 450 training epochs varying the learning rate from 1×10^{-3} to 1×10^{-7} .

3. CONSIDERED CASE-STUDY: THE AMOC

The AMOC dynamics is modeled by a reduced-order representation proposed in [4]. Inspired by the Howard-Malkus loop [27] and validated against $\delta^{18}\text{O}$ paleorecords, this model reproduces the centennial to millennial variability of the D-O events. It can be written as an homogeneous three-variables ODE system $\dot{\mathbf{X}}(t) \triangleq [x_1(t), x_2(t), x_3(t)]^T$ such as:

$$\dot{\mathbf{X}}(t) = \mathbf{M}(\mathbf{X}(t)) = \begin{pmatrix} -\lambda(x_1(t) + x_3(t)) \\ (\omega_0 + x_1(t))x_3(t) - \kappa x_2(t) + f \\ -(\omega_0 + x_1(t))x_2(t) - \kappa x_3(t) \end{pmatrix}$$

where λ is a momentum friction coefficient, ω_0 a constant of the AMOC such that $\omega = \omega_0 + x_1$ is the total AMOC intensity, κ the haline damping coefficient and f the freshwater flux intensity.

As expected, the system exhibits chaotic dynamics (Fig. 2). By computing the energy of the system, three categories can be inferred: *i*) the ON Phase refers to an intense circulation; *ii*) the OFF phase indicates an AMOC shutdown; *iii*) in between, the regime shift corresponds to quick transitions between these two phases. We notice a double pseudo-periodicity in this chaotic non-linear system: a first centennial pseudo-period of around 250 years, especially on the ON




Fig. 2. Idealized AMOC system evolution and D-O events simulated over 15 000 years. Highlighted dynamical patterns: ON Phase (indigo), Regime Shift (pink) and OFF Phase (magenta).

Phase ($\omega \approx -0.05 \text{ yr}^{-1}$); a second millennial pseudo-period of 1470 ± 500 years associated with the timescale of the D-O events, and the cycle of the AMOC shutdowns (i.e. OFF Phase, $\omega \approx 0 \text{ yr}^{-1}$). More details about the parametrisation and the model dynamics can be found in [4].

4. EXPERIMENTS

4.1. Data generation

To build our dataset, we simulated $S=100$ times the AMOC system over 100 000 years from which we could draw one training set and one test set. Time series of $\Delta T=2500$ years were extracted from the first 60 000 years (resp. the last 35 000 years) of simulation for the training dataset (resp. the test dataset). As the training is supervised, the training dataset contains the ground truth, the initial state of $\mathbf{X}$, and the observations $\mathbf{Y}$.

The observations $\mathbf{Y}$ are also constructed from this simulation. To meet Shannon's criterion for the frequency of D-O events, and set up a realistic experiment with sparse observations, the latter are acquired with a maximum sampling period $T_s=120$ years. We add a Gaussian observation noise corresponding to 10% of the variance of the AMOC centennial oscillations. This percentage applies the same relative noise to each observed variable. The assimilation problem is initialised as a linear interpolation of these observations. In this study, the three variables of the state $\mathbf{X}$ are observed.

4.2. Regular sampling

In a first experiment, we aim to evaluate the advantages of 4DVarNet-unet compared to the state-of-the-art 4DVar classic method and the 4DVarNet-ode scheme where ϕ comes from a Runge-Kutta 4 integration of the physical model $\mathbf{M}$. For this purpose, a regular sampling strategy is considered, with a sampling period (T_s) varying from 10 to 120 years. The Normalised Mean Square Error (NMSE) is chosen as performance metric defined for each variable $i \in \{1, 2, 3\}$ as:

$$\text{NMSE}(i) = \frac{1}{S \Delta T \sigma_{\text{GT}}^2} \sum_{s=0}^S \sum_{t=0}^{\Delta T} (\hat{x}_{i,s}(t) - x_{i,s}^{\text{GT}}(t))^2$$




Fig. 3. NMSE in function of T_s for the 3 DA methods


where $\mathbf{X}_s^{\text{GT}}$ is the ground truth and σ_{GT}^2 its variance. As expected, the NMSE increases with the sampling period (Fig. 3). In particular, the 4DVar classic method is very quickly affected by the data scarcity, with an error above 10% on the reconstruction of x_2 and x_3 as soon as the sampling period exceeds 40 years. Conversely, we note the excellent performances of the 4DVarNet-ode/unet methods. This first result proves that the LSTM cell optimises the solving of a variational DA problem. While 4DVarNet-ode method is constrained physically, the trajectory adopted by 4DVarNet-unet is more likely to reduce the NMSE since it is fully learned by optimising a learning cost aimed at minimising this error. This result is in line with the growing field of data-driven models discovery, where physical models are not necessarily considered as optimal models to perform the specific task of reconstructing (non-linear chaotic) dynamics. We also notice that x_1 is overall better estimated than (x_2, x_3) regardless the method. It can be explained by the scale differences (apparent on Fig. 2). The gap between the estimates is however wider with the 4DVar classic method, possibly revealing its difficulty to deal with multi-scale processes.

Even with a sampling period greater than 100 years, 4DVarNet-unet provides very satisfactory results until the maximum sampling period of $T_s=120$ years. It appears as a powerful tool to provide high-quality reconstructions with fewer observations. In a second set of experiments, we propose to evaluate the robustness of 4DVarNet-unet in face of more difficult sampling strategies, where *i*) the sampling period is set to $T_s=100$ years (large enough to discredit the other two DA methods) and *ii*) the sampling is not necessarily regular.

4.3. Varying sampling

Four different sampling strategies are tested out:

1. **Regular sampling** : Observations are regularly acquired, corresponding to $H_t = I_3 \forall t \in [1 : 100 : \Delta T]$;
2. **Regular cluster sampling** : We consider observa-



Sampling strategy	Regular sampling	Regular cluster sampling	Random sampling	Random cluster sampling
NMSE ON Phase	3.19×10^{-3}	8.47×10^{-4}	1.10×10^{-2}	8.07×10^{-3}
NMSE Regime Shift	1.05×10^{-2}	2.26×10^{-3}	6.64×10^{-2}	4.51×10^{-2}
NMSE OFF Phase	5.40×10^{-2}	8.47×10^{-3}	2.93×10^{-1}	2.02×10^{-1}

Table 1. NMSE of the reconstructed signal of each variable for each DA method. Bold values exhibit the best performance.

tions clusters composed of three finer consecutive time steps ($\delta_t=10$ years), acquired every T_s years, namely $H_{t-\delta_t} = H_t = H_{t+\delta_t} = I_3 \forall t \in [1 : 100 : \Delta T]$;

3. **Random sampling** : Observations are sampled randomly in time. Budget of observations is fixed according to the observations number in Strategy 1. We thus have $H_t = I_3$ for t randomly chosen in the ΔT -time window;

4. **Random cluster sampling** : We consider observation clusters of three finer consecutive time steps ($\delta_t=10$ years), separated by intervals of random length. Number of clusters is fixed according to the observations number in Strategy 1. We thus have $H_{t-\delta_t} = H_t = H_{t+\delta_t} = I_3$ for t randomly chosen in the ΔT -time window.

Note that strategies 2 and 4 lead to a larger number of samples than strategies 1 and 3. This choice is guided by an operational context in which it is costly to deploy acquisition resources but which can be left in place over local time windows of a few years.

The NMSE is again chosen as performance metric, as:

$$\text{NMSE} = \frac{1}{MS\Delta T\sigma_{\text{GT}}^2} \sum_{m=1}^M \sum_{s=1}^S \sum_{t=0}^{\Delta T} \|\hat{\mathbf{X}}_{m,s}(t) - \mathbf{X}_{m,s}^{\text{GT}}(t)\|_2^2$$

with $M=10$ the number of models trained for each strategy.

Tab. 1 summarizes the performance of 4DVarNet-unet for the four sampling strategies. We first observe that random sampling strategies lead to a NMSE overall 10 times as big as the regular sampling strategies on the ON Phase reconstruction. Even if randomness could be expected to sweep up more different states of the system and statistically lead to a better reconstruction, the pseudo-periodicities of the AMOC system might have been missed by these random sampling strategies, making them less reliable. Adding more trajectories to the learning process could probably improve the performance of such strategies, but priority is given here to methods which require less data to be trained, since the ocean is usually sparsely sampled. It is however interesting to see that the clustering of observations is more efficient on the regular sampling strategy than on the random one. The observation clusters improve the reconstruction of the ON Phase by

more than 26% for the random sampling and even by 73% for the regular sampling. We can suppose that acquiring three consecutive observations regularly helps to detect the system variations, especially when a regime shift occurs. Therefore, acquiring regular observations over a sufficiently (local) long time window could help monitor regime changes similar to those caused by D-O events, especially when data is rare due to low sampling frequency. The discrepancy between regular and random strategies is even bigger on the reconstruction of the extremes, associated with abrupt climate changes. Due to their chaotic dynamics, OFF phases are the most difficult to reconstruct by our trained models. However, the regular cluster sampling strategy provides a precise reconstruction of the D-O events with an error below 1%. In addition, the gain in accuracy in the reconstruction of regime shifts could capture the premise of an abrupt change. We could take advantage of this to improve the reconstruction of OFF phases by adapting the sampling strategy on-line, e.g. by increasing the sampling frequency when a regime shift is detected.

5. CONCLUSION

In this paper, we solved a DA problem to monitor abrupt climate regime changes – D-O events – using an idealized deterministic 3-variable chaotic model of the North Atlantic Ocean circulation variability. Deep learning allowed us to improve the solving of this inverse problem by implementing a ResNet composed of convolutional LSTM cells, validated against the 4DVar DA method. A U-Net was also exploited to model our nonlinear chaotic phenomenon. We focused our effort on the impact of sampling on the reconstruction performance of such an approach. These preliminary results obtained pave the way for more in-depth work on the optimisation of these sampling strategies with regard to the physical processes considered, in particular non-linear chaotic ones such as the AMOC system.

6. REFERENCES

- [1] Eonjoo Kim and Mae L. Seto, "An Approach to Choose Observation Systems to Observe Ocean Phenomena," in *2022 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, Sept. 2022, pp. 306–313.
- [2] Angélique V. Melet et al., "Chapter 2 - The role of ocean mixing in the climate system," in *Ocean Mixing*, pp. 5–34. Elsevier, Jan. 2022.
- [3] Martha W. Buckley and John Marshall, "Observations, inferences, and mechanisms of the Atlantic Meridional Overturning Circulation: A review," *Reviews of Geophysics*, vol. 54, no. 1, pp. 5–63, 2016.
- [4] Florian Sévellec and Alexey V. Fedorov, "Millennial Variability in an Idealized Ocean Model: Predicting the AMOC Regime Shifts," *Journal of Climate*, vol. 27, no. 10, pp. 3551–3564, May 2014.
- [5] Agathe Germe et al., "Chaotic Variability of the Atlantic Meridional Overturning Circulation at Subannual Time Scales," *Journal of Physical Oceanography*, vol. 52, no. 5, pp. 929–949, May 2022.
- [6] Peter Ditlevsen and Susanne Ditlevsen, "Warning of a forthcoming collapse of the Atlantic meridional overturning circulation," *Nature Communications*, vol. 14, no. 1, pp. 4254, July 2023.
- [7] Florian Sévellec and Alexey V. Fedorov, "Unstable AMOC during glacial intervals and millennial variability: The role of mean sea ice extent," *Earth and Planetary Science Letters*, vol. 429, pp. 60–68, Nov. 2015.
- [8] W. Dansgaard et al., "Evidence for general instability of past climate from a 250-kyr ice-core record," *Nature*, vol. 364, no. 6434, pp. 218–220, July 1993.
- [9] Walter Munk, "Chapter 1 Oceanography before, and after, the advent of satellites," in *Elsevier Oceanography Series*, vol. 63 of *Satellites, oceanography and society*, pp. 1–4. Elsevier, Jan. 2000.
- [10] Bàrbara Barceló-Llull et al., "Fine-Scale Ocean Currents Derived From in situ Observations in Anticipation of the Upcoming SWOT Altimetric Mission," *Frontiers in Marine Science*, vol. 8, 2021.
- [11] Geir Evensen, *Data Assimilation*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [12] A. Carrassi et al., "Data Assimilation in the Geosciences - An overview on methods, issues and perspectives," *Wiley Interdisciplinary Reviews: Climate Change*, vol. 9, Sept. 2017.
- [13] Dusanka Zupanski, "A General Weak Constraint Applicable to Operational 4DVAR Data Assimilation Systems," *Monthly Weather Review*, vol. 125, no. 9, pp. 2274–2292, Sept. 1997.
- [14] Sibor Cheng et al., "Machine Learning With Data Assimilation and Uncertainty Quantification for Dynamical Systems: A Review," *IEEE/CAA Journal of Automatica Sinica*, vol. 10, no. 6, pp. 1361–1387, June 2023.
- [15] Ricky T. Q. Chen et al., "Neural Ordinary Differential Equations," Dec. 2019.
- [16] Emmanuel de Bezenac et al., "Deep Learning for Physical Processes: Incorporating Prior Scientific Knowledge," Jan. 2018.
- [17] Rossella Arcucci et al., "Deep Data Assimilation: Integrating Deep Learning with Data Assimilation," *Applied Sciences*, vol. 11, no. 3, pp. 1114, Jan. 2021.
- [18] Ronan Fablet et al., "Learning Variational Data Assimilation Models and Solvers," *Journal Of Advances In Modeling Earth Systems*, vol. 13, no. 10, Oct. 2021.
- [19] Max Yaremchuk and Paul Martin, "On Sensitivity Analysis within the 4DVAR Framework," *Monthly Weather Review*, vol. 142, no. 2, pp. 774–787, Feb. 2014.
- [20] Joaquin Vanschoren, "Meta-Learning: A Survey," Oct. 2018.
- [21] Qingping Zhou et al., "Deep Unrolling Networks with Recurrent Momentum Acceleration for Nonlinear Inverse Problems," Aug. 2023.
- [22] Olaf Ronneberger et al., "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 2015, pp. 234–241, Springer International Publishing.
- [23] Jens Berg and Kaj Nyström, "Data-driven discovery of PDEs in complex datasets," *Journal of Computational Physics*, vol. 384, pp. 239–252, May 2019.
- [24] Samuel H. Rudy et al., "Data-driven discovery of partial differential equations," *Science Advances*, vol. 3, no. 4, pp. e1602614, Apr. 2017.
- [25] Diederik P. Kingma and Jimmy Ba, "Adam: A Method for Stochastic Optimization," Jan. 2017.
- [26] Nitish Srivastava et al., "Dropout: A Simple Way to Prevent Neural Networks from Overfitting," *Journal of Machine Learning Research*, vol. 15, 2014.
- [27] Pierre Welander, "Thermohaline Effects in the Ocean Circulation and Related Simple Models," in *Large-Scale Transport Processes in Oceans and Atmosphere*, pp. 163–200. Springer Netherlands, Dordrecht, 1986.