Ocean Modelling 184 (2023) 102208

Contents lists available at ScienceDirect

Ocean Modelling

journal homepage: www.elsevier.com/locate/ocemod

Short communication

A hyper-parameterization method for comprehensive ocean models: Advection of the image point

Igor Shevchenko [a,*], P. Berloff [a,b]

[a] Department of Mathematics, Imperial College London, Huxley Building, 180 Queen's Gate, London, SW7 2AZ, UK
[b] Institute of Numerical Mathematics, Russian Academy of Sciences, Gubkina 8, 119333, Moscow, Russia

ARTICLE INFO

Keywords:
Hyper-parameterization
Build-up effect
Parameterization
Comprehensive ocean model
Primitive equations
Large scale
Small scale
Eddy parameterization problem
Nudging
Gulf Stream flow

ABSTRACT

Idealized and comprehensive ocean models at low resolutions cannot reproduce nominally-resolved flow structures similar to those presented in the high-resolution solution. Although there are various underlying physical reasons for this, from the dynamical system point of view all these reasons manifest themselves as a low-resolution trajectory avoiding the phase space occupied by the reference solution (the high-resolution solution projected onto the coarse grid). In order to solve this problem, a set of hyper-parameterization methods has recently been proposed and successfully tested on idealized ocean models. In this work, for the first time we apply one of hyper-parameterization methods (Advection of the image point) to a comprehensive, rather than idealized, general circulation model of the North Atlantic.

The results show that the hyper-parameterization method significantly outperforms the coarse-grid ocean model by reproducing both the large- and small-scale features of the Gulf Stream flow. The proposed method is much faster than even a single run of the coarse-grid ocean model, requires no modification of the model, and is easy to implement. Moreover, the method can take not only the reference solution as input data but also real measurements from different sources (drifters, weather stations, etc.), or combination of both. All this offers a great flexibility to ocean modellers working with mathematical models and/or measurements.

1. Introduction

The eddy parameterization problem (how to account for the effect of unresolved small scales onto the resolved large scales) is one of the most challenging problems in the ocean modelling, counting decades of active research. Despite a number of parameterizations (computationally affordable and physically justified mathematical models of unresolved processes) have been proposed to solve the problem (e.g., Gent and Mcwilliams, 1990; Duan and Nadiga, 2007; Frederiksen et al., 2012; Mana and Zanna, 2014; Cooper and Zanna, 2015; Grooms et al., 2015; Berloff, 2015, 2016, 2018; Ryzhov et al., 2019; Cotter et al., 2019; Ryzhov et al., 2020; Cotter et al., 2020a,b,c), it remains largely unresolved. The main general point is that most of the parameterization approaches are physics-based rather than data-driven, and the former has obvious advantage of being valid in situation when the underlying physics changes. However, in the situation when physics-based parameterizations remain in their infancy, there is a great practicality in considering data-driven parameterizations, which can reproduce nominally-resolved flow structures within their obvious limitations. For example, for many research questions a computationally-cheap data-driven solution can replace a computationally-expensive dynamical ocean simulation in climate-type models and predictions.

Our approach is both novel and orthogonal to all existing parameterization approaches. We refer to it as "hyper-parameterization" and develop it in the context of oceanic eddies, although its broader impact will be across most of the turbulence research areas. Hyper-parameterization is a new strategy, which capitalizes on the use of quasi time-invariant structures in phase space, and within this approach the main focus is shifted from representing essential local physics of cross-scale interactions towards simulating nominally resolved, most important flow patterns directly, including anti-diffusive jet-sharpening and upscale feedbacks.

Advancing the hyper-parameterization approach, within broader context of data-driven parameterizations, provides the main motivation for our study, in which we continue the series of precursor works (Shevchenko and Berloff, 2021a, 2022b,a). The main novelty is to adapt the methodology for fully comprehensive and realistic ocean models, and demonstrate its success with full confidence. In turn this

* Corresponding author.
E-mail address: i.shevchenko@imperial.ac.uk (I. Shevchenko).

https://doi.org/10.1016/j.ocemod.2023.102208
Received 13 September 2022; Received in revised form 8 February 2023; Accepted 29 April 2023
Available online 6 May 2023

paves the way for broad use of hyper-parameterizations across the ocean modelling community.

2. The method

In this work, we consider the method called "Advection of the image point". The image point (also called the representative point) corresponds to the tip of the state vector of the dynamical system in the phase space. Advection means that the image point is transported by the vector flow mapped in the phase space (here, in terms of discrete observations of tendencies). The method falls into the category of data-driven methods, and has been successfully tested on a multi-layer quasi-geostrophic model and showed promising results (Shevchenko and Berloff, 2021a). The method is based on the fact that the first-order ordinary differential equation

$$\mathbf{x}'(t) = \mathbf{F}(\mathbf{x}), \quad \mathbf{x} \in \mathbb{R}^n \tag{1}$$

can be interpreted as a vector field $\mathbf{F}(\mathbf{x})$ in the phase space of Eq. (1), and the prime denotes a derivative in time. If $\mathbf{F}(\mathbf{x})$ is known, it can be used to advect the image point $\mathbf{x}$ in the phase space. Evolution of an image point can be described by the equation:

$$\mathbf{y}'(t) = \frac{1}{N}\sum_{i\in\mathcal{U}(\mathbf{y}(t))}\mathbf{F}(\mathbf{x}(t_i)) + \eta\left(\frac{1}{M}\sum_{i\in\mathcal{U}(\mathbf{y}(t))}\mathbf{x}(t_i) - \mathbf{y}(t)\right), \quad \mathbf{y}(t_0) = \mathbf{x}(t_0), \tag{2}$$

where $\mathcal{U}(\mathbf{y}(t))$ is a neighbourhood of solution $\mathbf{y}(t)$, and i is the timestep of the reference solution $\mathbf{x}(t_i)$; the integration step for (2) equals to the step with which the reference solution is sampled (6 h in our case).

In a nutshell, in Eq. (2) we nudge toward the observed reference states $\mathbf{x}(t_i)$ in the phase space neighbourhood in order to prevent runaway from the attractor. The actual dynamics is given by reconstruction of the observed references tendencies — this is the key right hand side term.

The neighbourhood is computed as the average over N (and M for the nudging term) nearest, in l_2 norm, to the solution $\mathbf{y}(t)$ points, and η is the nudging strength having units of 1/s; we will return to the choice of N, M, and η in the next section; we refer to these parameters as hyper-parameters. Note that in order to calculate the l_2-norm in the phase space of variables having different units (say, relative vorticity ζ and temperature T as in this study), it is necessary that all the variables are represented in the same units or nondimensional (as in our case). For doing so, we calculate the Z-score for these variables as follows:

$$Z_i = (X_i - \overline{X}_i)/\sigma(X_i), \quad i = \{\zeta, T\},$$

where Z_i is the new variable, X_i is the old variable, $\overline{X}_i$ is the time-mean of X_i, and $\sigma(X_i)$ is the standard deviation of X_i; X_i is a vector of all the values of the ith variable.

The hyper-parameters can be set based on the chosen metric and available data. Getting a bit ahead, we report that in our case the neighbourhood consists of $N = 15$ and $M = 5$ points, and the nudging strength is $\eta = 0.001$. The term $\mathbf{F}(\mathbf{x})$ is computed with the central finite difference in time. Originally, the method was supposed to have the same size of the neighbourhood for both the vector field $\mathbf{F}(\mathbf{x})$ and the nudging term (i.e., $N = M$). However, our experiments showed that using neighbourhoods of different sizes can prevent the so-called build-up effect which we discuss later. We refer the reader to Shevchenko and Berloff (2021a) for a more detailed discussion of the method. It is worth noting that the length of the sampling interval that is enough to properly represent the dynamics cannot be defined a priori as it strongly depends on the reference data. We recommend to use all available reference data to better approximate the structure of the reference phase space.

It might seem that a state $\mathbf{x}(t)$ cannot be expressed as a solution of a certain differential equation with the right hand side depending only on $\mathbf{x}(t)$, because $\mathbf{x}(t)$ can depend on many other prognostic (and/or diagnostic) variables (as in our case), because all primitive equations in MITgcm are coupled. However, it is not necessarily true, as one can reconstruct a dynamical system (and we did it in Shevchenko and Berloff, 2022b) the solution of which is $\mathbf{x}(t)$, while the right hand side depends only on $\mathbf{x}(t)$. Moreover, the idea of representing the evolution of a process as a solution of a certain differential equation with the right hand side depending only on this process is at the core of dynamical system reconstruction theory.

3. Model configuration and numerical results

For the purpose of this study, we consider the Massachusetts Institute of Technology general circulation model (MITgcm) (Marshall et al., 1997) in the North Atlantic configuration (Shevchenko and Berloff, 2021b; Jamet et al., 2019). It is a 46-layer oceanic model coupled with an atmospheric boundary model (Marshall et al., 1997; Deremble et al., 2013). The coupled model is initially spun up for 5 years and then integrated for another 2 years. Although, for this work it is not essential that the initial state of the ocean circulation is in the statistically equilibrated regime. The model is integrated at two different horizontal resolutions (1/12° and 1/3°); the oceanic and atmospheric models are implemented with the same horizontal resolution. We refer to the solution (and/or a set of diagnostic variables) computed on the 1/12°-grid and projected onto the 1/3°-grid (Fig. 1a) as the reference solution, $\mathbf{x}(t)$ in Eqs. (1)–(2), and to the solution (and/or a set of diagnostic variables) computed on the 1/3°-grid as the modelled solution, $\mathbf{y}(t)$ in Eq. (2), see Fig. 1c. The hyper-parameterized solution computed with Eq. (2) is presented in Fig. 1b. Note that the reference solution consists only of the coarse-grained variables that are modelled with the hyper-parameterization method, it is surface relative vorticity and temperature in our case. It provides a great flexibility for modelling and significant acceleration of computations. It is also worth reiterating again that we only solve Eq. (2) to model the North Atlantic, and the underlying primitive equations are not used, as it would defeat the purpose of our approach. The whole complexity of the flow dynamics enters Eq. (2) via the right hand side.

As it follows from the results in Fig. 1, the hyper-parameterized solution computed on the coarse grid (Fig. 1b) reproduces both the large scales (the Gulf Stream flow) and small scales (vortices) of the reference solution (Fig. 1a) in instantaneous and time-averaged fields, while the solution computed on the coarse grid without the hyper-parameterization (Fig. 1c) leads to no Gulf Stream or small-scale flow features. We remark that the hyper-parameterized solution is a part of a 2-year simulation from which we use only the first year of the reference solution; over the second year model (2) runs on its own.

The hyper-parameterized solution with a bad choice of hyper-parameters N, M, η (Fig. 2) still reproduces both large- and small-scales flow features but only over the period in which the reference solution is available (it is one year in our case). After that the solution almost stops evolving and eventually settles to a constant in time field like the one in the middle plot of Fig. 2. It can be seen from the time-mean over the second year (the right subplot in Fig. 2), which is very similar to the snapshot taken at year 2 (middle subplot in Fig. 2), thus showing that the solution evolution is stalled over the second year.

The build-up effect and solution degradation. The lack of reference data for the hyper-parameterized model and/or bad choice of hyper-parameters can steer the trajectory to leave the reference phase space (the phase space occupied by the reference solution; note that the phase space of the primitive equations consists of all prognostic variables and its dimension is equal to the number of degrees of freedom, while the reference phase space consists of only the variables that are hyper-parameterized, which are surface relative vorticity and temperature in our case). This escape may result in a significant degradation of the hyper-parameterized solution and even lead to a numerical blow-up. We refer to this as the build-up effect, meaning that after a period of time, say T, the neighbourhood of the nearest

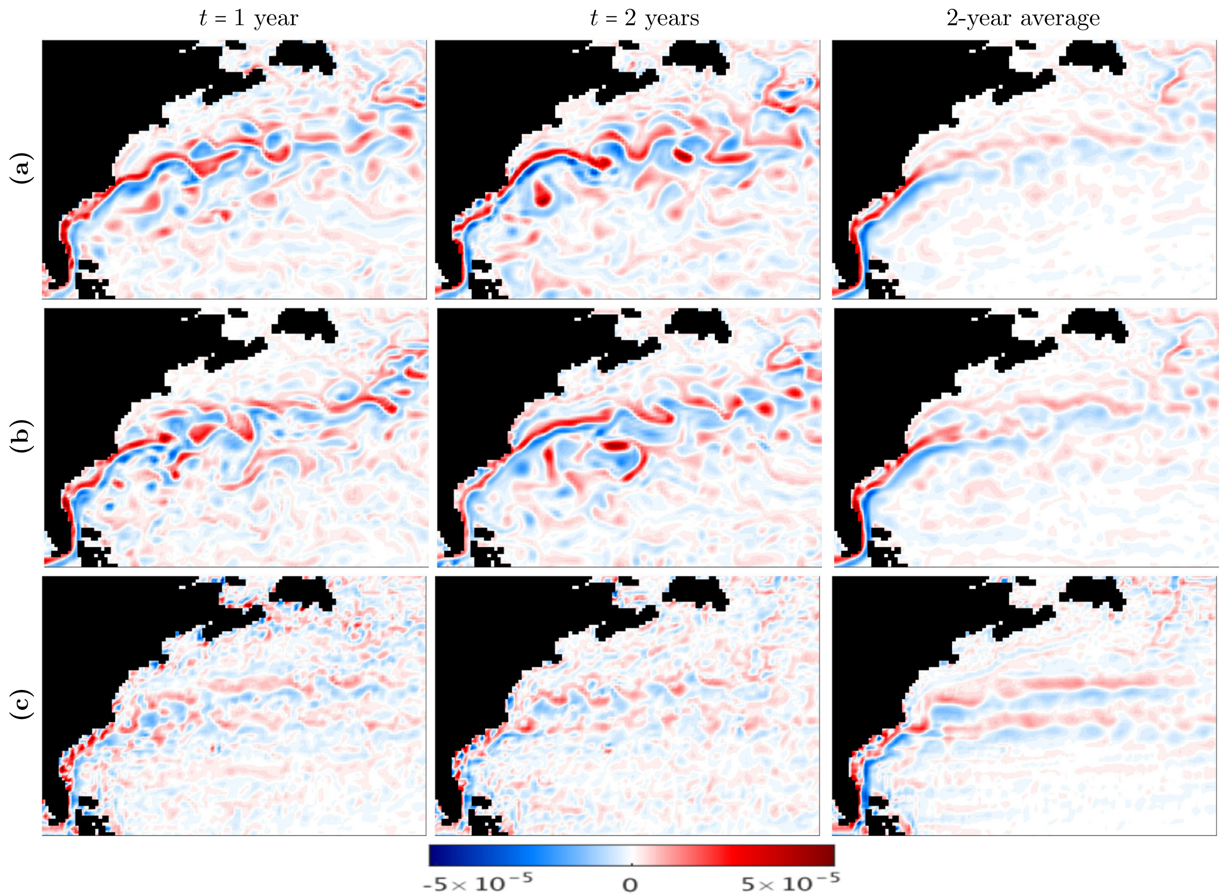

Fig. 1. Shown are snapshots of the surface relative vorticity $\zeta = v_x - u_y$ [1/s] for **(a)** the reference solution (computed at horizontal resolution 1/12° and then projected on the 1/3°-grid), **(b)** hyper-parameterized solution computed at horizontal resolution 1/3° for $N = 15$, $M = 5$, $\eta = 0.001$, **(c)** modelled solution computed at horizontal resolution 1/3°, and the 2-year time-average (last column). Snapshots are taken after 1 year (left column) and 2 years (middle column) of simulations. Note the modelled solution **(c)** fails to reproduce important large-scale (the Gulf Stream eastward jet extension) and small-scale (vortices) structures of the flow dynamics in both instantaneous and time-averaged fields.

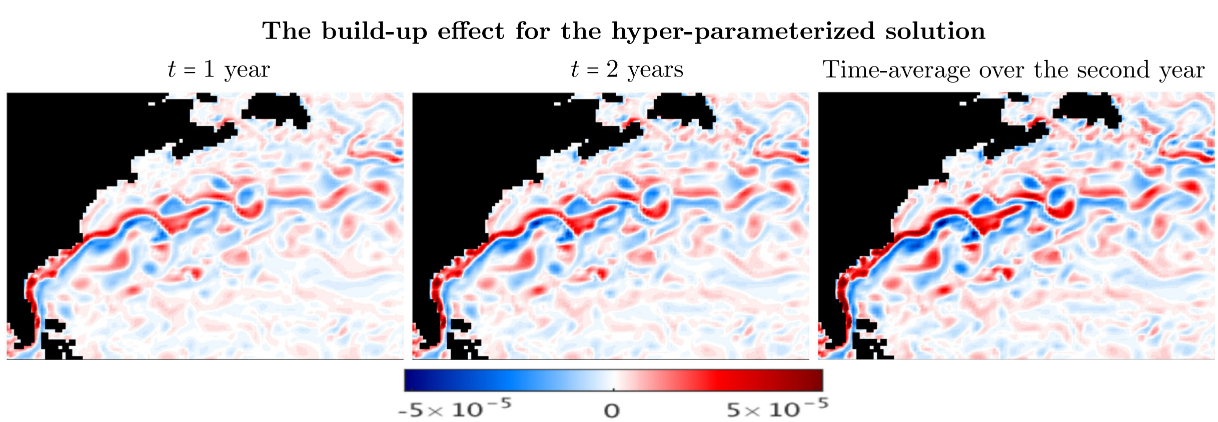

Fig. 2. Shown are snapshots of the surface relative vorticity $\zeta = v_x - u_y$ [1/s] demonstrating the build-up effect for the hyper-parameterized solution computed at horizontal resolution 1/3° for $N = M = 5$, $\eta = 0.001$. Snapshots are taken after 1 year (left column) and 2 years (middle column) of simulations.

points stalls (Fig. 2), i.e. the points in the neighbourhood become the closest ones to the image point $\mathbf{y}(t)$ for $\forall t > T$ (in the present case $T = 1$ year); therefore, the same points are used again and again during the integration of Eq. (2) thus driving the image point away from the reference phase space. In principle, building up numerical errors may terminate this runaway, and the solution can return back to the reference phase space region, but this is case dependent and should be kept in mind. However, if the return time is relatively long (longer than the characteristic time of the reference solution) then the flow dynamics can be seriously distorted over the period of the trajectory injection.

In order to prevent the build-up effect, one should properly set up the hyper-parameters N, M, and η by using weak nudging towards a slightly lager set of neighbours. After some experiments we have found that $N = 15$, $M = 5$, and $\eta = 0.001$ lead to no build-up effect (it is not necessarily the only choice, and other sets of hyper-parameters providing no build-up may exist). The optimal hyper-parameters for a reference solution $\mathbf{x}(t)$ can be computed by solving the following problem

$$\min_{N,M,\eta} \mathcal{F}(\mathbf{x}(t), \mathbf{y}(t)), \quad t \in [0, T] \tag{3}$$

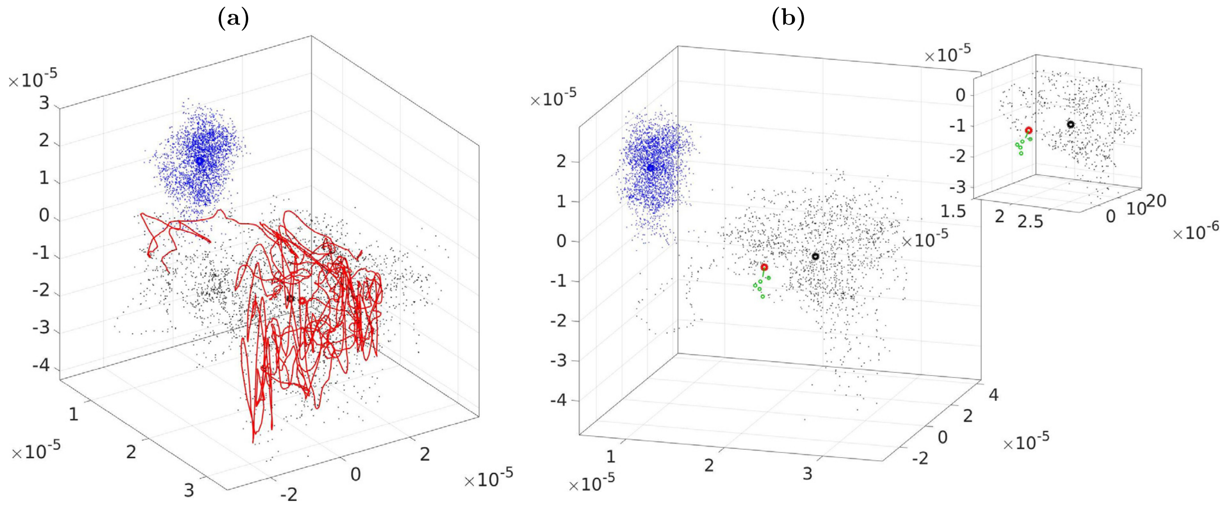

Fig. 3. Shown is **(a)** 3D projection of the reference phase space (black dots), modelled phase space (blue dots), hyper-parameterized solution (red trajectory); in other words, we plot 3 random coordinates from the multi-dimensional phase space of surface relative vorticity (i,e, the axes represent the values of surface relative vorticity at these 3 randomly chosen coordinates), **(b)** build-up effect for the second year of the hyper-parameterized solution (short green line near the red circle) for $N = M = 5$, and $\eta = 0.001$; the green circles are those stalled five points in the neighbourhood. The time-means for every solution are denoted by bold circles of corresponding colours.

where $\mathcal{F}$ is a problem-specific function, and T is the length of the reference solution $\mathbf{x}(t)$. For example, $\mathcal{F}$ can be defined as a norm of the difference between the reference and hyper-parameterized solutions. Our choice of N, M, and η is driven by our measure of goodness.

It is also important to note that the way the neighbourhood $\mathcal{U}(\mathbf{y}(t))$ in Eq. (2) is computed affects the hyper-parameterized solution. One can change the algorithm on how to pickup points from the neighbourhood to tailor the method for a given problem. It is also worth noting that these experiments are very fast and computationally cheap, even relative to a single run of the coarse-grid model. The optimal choice of hyper-parameters and the neighbourhood size, and how they affect the solution is a topic beyond the scope of the present study.

The measure of goodness and evolution in phase space. The measure of goodness (i.e., the proximity of the modelled or parameterized solution to the reference one) in a given metric depends on the specific purpose. Our measure of goodness is how close the hyper-parameterized solution is to the reference phase space. We use this measure to allow the hyper-parameterized solution to evolve in the neighbourhood of the reference phase space, since the failure of the coarse-grid model (blue dots in Fig. 3) to reproduce large- and small-scale features of the flow dynamics is because it steers away from where it should be (black dots in Fig. 3, i.e. the reference phase space). Measuring the proximity of individual trajectories in phase space is of no value, because small initial perturbations will grow exponentially due to the inherently chaotic dynamics.

The build-up effect is demonstrated in Fig. 3b; the short green line near the time-mean over the second year (red circle) is the 1-year evolution of the hyper-parameterized solution affected by the bad choice of hyper-parameters, namely $N = M = 5$, and $\eta = 0.001$. Evolution of the Euclidean distance from the reference time mean (the time-mean of the reference solution) to the hyper-parameterized solution (Fig. 4) shows that the hyper-parameterized solution affected by the build-up effect (green line) stops to evolve after one year (the period over which the reference solution is available). In this case, we observe no numerical blow-up as the solution quickly settles to a constant in time field. When the build-up effect is prevented, the hyper-parameterized solution (red line) continues to evolve with the reference phase space.

Coupled fields. Comprehensive ocean models have several prognostic fields (velocity, temperature, etc.), and a good choice of hyper-parameters (N, M, and η) for one field is not necessarily applicable to another. For example, $N = 15$, $M = 5$, and $\eta = 0.001$ choice works well only for the relative vorticity. However, it leads to a build up for the coupled fields (relative vorticity and temperature, not shown). Thus, a set of hyper-parameters for a coupled case has to be found separately. Namely, for the coupled fields (relative vorticity and temperature) we have found that $N = 18$ and $M = 4$ remove the build up effect, hence, the hyper-parameterized solution is of high quality (Fig. 5). The nudging strength in this case remains unchanged from the single relative vorticity case (Fig. 1), $\eta = 0.001$.

This example of the coupled fields demonstrates that one should exercise caution when it comes to setting up the hyper-parameters. On the other hand, it also shows that the hyper-parameterization method works for coupled fields (even with huge differences between the fields, which is 7 orders of magnitude for the relative vorticity and temperature, see colorbars in Fig. 5). The main difficulty here is that the significant difference in the magnitude of these fields can lead to a fast accumulation of the computational error (when computing the l_2-norm) that eventually contaminates the fields and further results in incorrect computations. If the problem comes out it can be fixed by normalizing relative vorticity and temperature, but we chose not to do this because the method performed well for our choice of data.

4. Conclusions and discussion

In this work we have applied the hyper-parameterization method “Advection of the image point” to the Massachusetts Institute of Technology general circulation model in the North Atlantic configuration and, thus, tested the method in a significantly more complicated setting, as compared to the earlier idealized-model tests. Our results show that the hyper-parameterization method significantly outperforms the non-eddy-resolving 1/3°-grid ocean model for both single and coupled fields (even with large difference between the fields, it is 7 orders of magnitude in our case) by reproducing both the large-scale (the Gulf Stream eastward jet extension) and small-scale (vortices) flow features of the reference eddy-resolving solution (on a 1/12°-grid and then projected onto the 1/3°-grid). It is important to note that the hyper-parameterization method reproduces both large- and small-scale flow features not only over the period for which the reference solution is available (1 year in our case), but also over the second year for which there is no reference data. The ability of the hyper-parameterization method to work well beyond the reference data set is explained by the fact that the reference data set is sufficient to capture the structure of the reference phase space. In other words, the reference data is a representative sample of the flow dynamics; if it were not the case, the method would not work. We deliberately chose to demonstrate

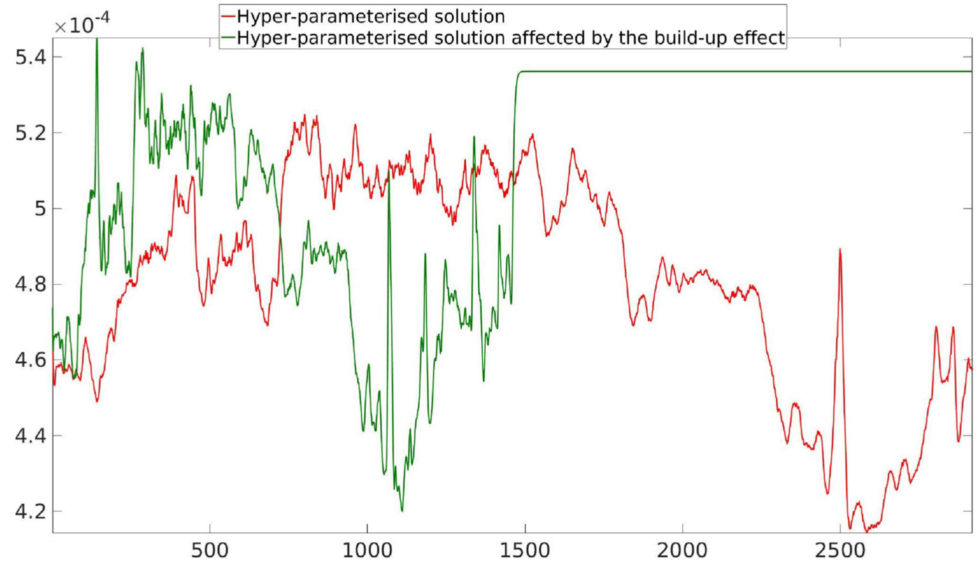

Fig. 4. Shown is evolution of the Euclidean distance (vertical axis) from the reference two-year time-mean to the hyper-parameterized solution (red line) and to the hyper-parameterized solution affected by the build-up effect (green line). The horizontal axis shows the simulation time in time steps (6 h in our case).

that the method works decently even with a short reference solution. As a matter of fact, the whole method is intrinsically stable and can provide a hyper-parameterized solution over any time interval. The quality of the hyper-parameterized solution still depends on the amount of data fed into the data-training, that is, into increasingly more detailed reconstruction of the phase space. The more data is fed in, especially in terms of longer time spans, the better will be quality of the hyper-parameterized solution.

We have also explained the build-up effect and showed that a bad choice of hyper-parameters leads to the build-up effect, and eventually to a significant degradation of the hyper-parameterized solution. In the worst-case scenario, the build-up can lead to a numerical blow up (which we did not observe in our experiments though). The build-up can be avoided by a proper setup of hyper-parameters which we have found through a series of numerical experiments. In addition, the hyper-parameters have to be found separately for hyper-parameterizing single and coupled fields, as well as for different single fields (for instance, the hyper-parameters that work well for relative vorticity may not be optimal for temperature, and vice versa).

It is important to keep in mind that the proposed method is data-driven and, therefore, can suffer from lack of data as any data-driven method. The hyper-parameterization approach has other methods (Shevchenko and Berloff, 2022b,a) that can be used to reproduce effects of mesoscale oceanic eddies on the large-scale ocean circulation, but demonstration of their skills on the level of comprehensive models is left for the future. In other words, staying complimentary to the mainstream physics-based parameterization approach, we propose to work in phase space of the corresponding dynamical system and to interpret the lack of eddy effects as persistent tendency of phase space trajectories (representing the modelled solution) to escape the correct reference phase-space region.

The proposed hyper-parameterization method is purely data-driven, hence, it cannot accommodate for changing physics. This only implies that the method should be used for specific problems; for example, as representation of the ocean circulation in atmosphere-only model runs or in ocean transport process studies. It will not be superfluous to also note that the data-driven nature of the method arrests its ability to operate in phases space regions that are not presented (or sparsely presented) in reference data. For instance, studying decadal variability can require tens of years of reference data to ensure that the reference data set is a representative sample of the studied phenomenon. However, this weakness is inherent in the data-driven paradigm. On the other note, within the hyper-parameterization approach there are other methods (Shevchenko and Berloff, 2022a,b), which are hybrids between pure data-driven and physics-driven dynamics — they can handle physical feedbacks.

The proposed hyper-parameterization method offers many advantages: (1) it is computationally very cheap (it is much faster than even a single run of the coarse-grid ocean model) and is easy to implement; (2) it enables research situations involving only part/component of the ocean circulation subject to hyper-parameterization: some geographical part (e.g. ocean surface or sub-basin), some specific field (e.g. temperature or velocity), or even some specific observational data (satellite images, drifters, etc.); if needed, the reference data can be always filtered to get rid of its fluctuations and only then be used as the reference phase space; (3) it can be straightforwardly used with comprehensive ocean models (e.g. MITgcm, NEMO, etc.) without modifying the model dynamic core; (4) it is well-suited for generating ensembles of solutions that can be used to sample the multitude of possible nonlinear, chaotic ocean/climate states, and to separate the response to external forcing from the internal, unforced nonlinear variability; an ensemble can be generated by changing initial positions in the phase space; (5) it can take as input data not only the reference solution but also real measurements from different sources (drifters, weather stations, etc.), or combination of both.

In order to use real measurements together with a numerical solution, one needs to compute the variables of interest and corresponding directional vectors (from the observational data) in the modelled region (say, relative vorticity and temperature in the North Atlantic as in this study) at the same resolution as the numerical solution and then use both the numerical solution and the variables (and the directional vectors) computed from observations as new reference data. For example, if the numerical solution and the observational data consists of N_{num} and N_{obs} records of relative vorticity and temperature, respectively; then the phase space consists of $\mathbf{x}_i$ points and corresponding directional vectors $\mathbf{F}(\mathbf{x}_i)$, $i = 1, \dots, N_{num} + N_{obs}$. We recommend using data assimilation, first, and then using reanalysis for data-driven reconstruction of the attractor, as observational data might be inconsistent in time and/or in space, thus requiring extra interpolation (or other) procedures to ensure its proper representation at a given resolution. However, other approaches are likely to exist. All this offers a great flexibility not only to ocean modellers working with mathematical models but also to those working with measurements.

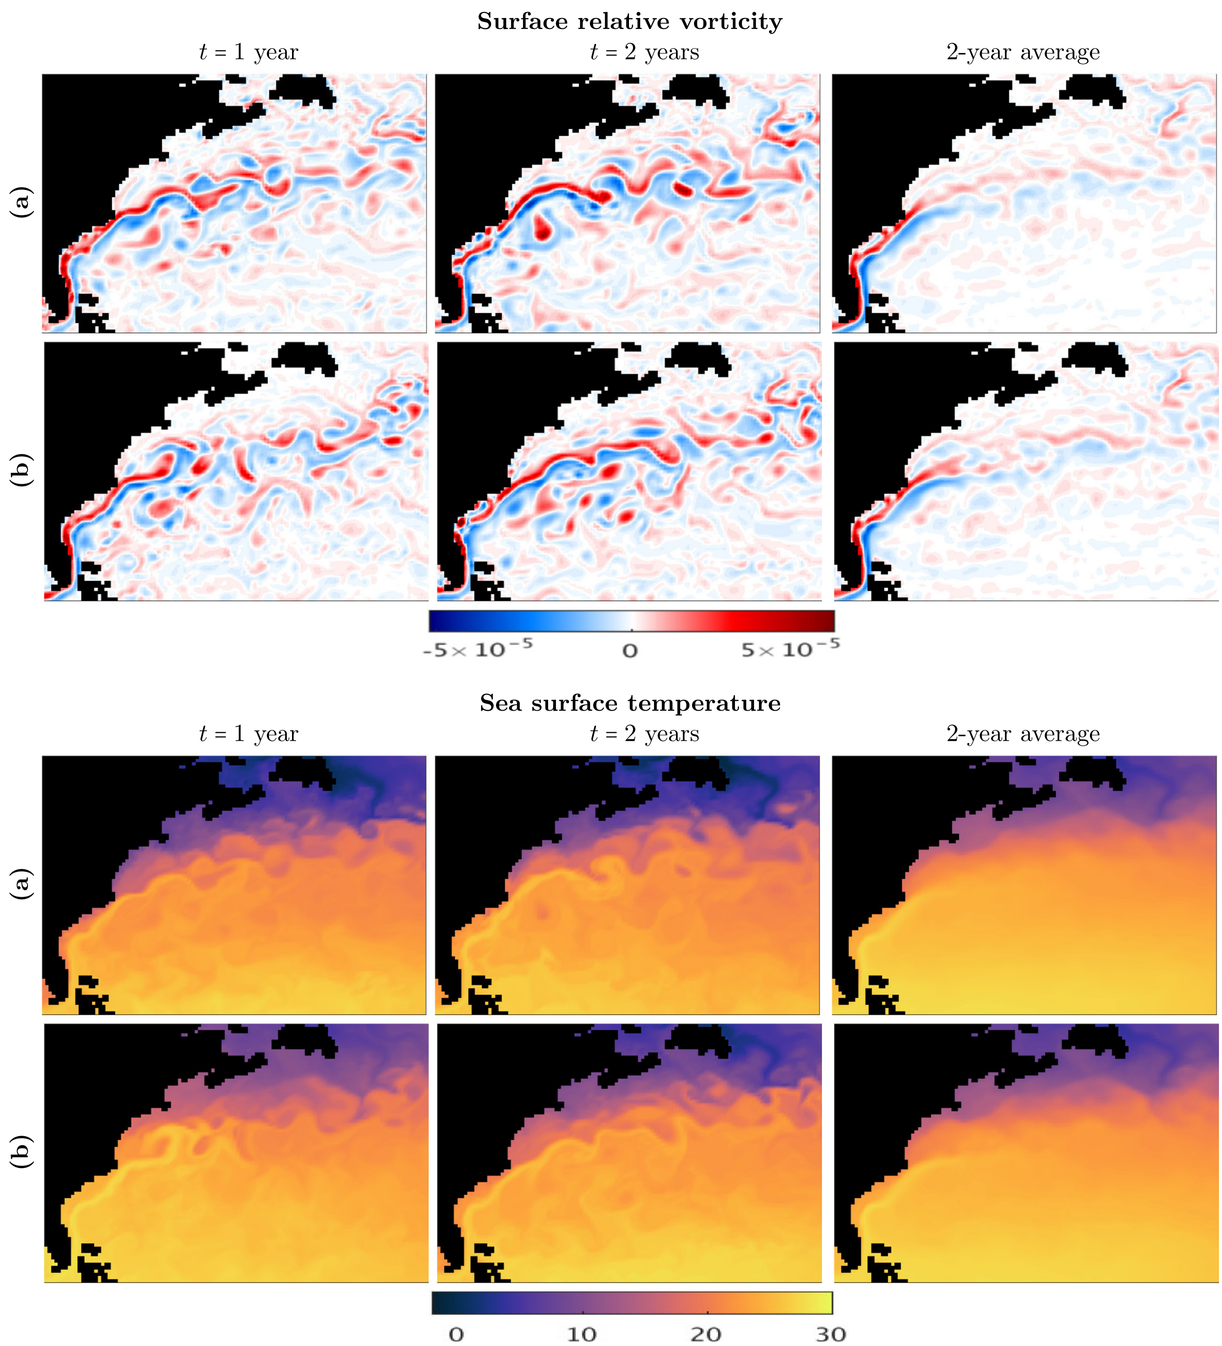

Fig. 5. Shown are snapshots of the coupled fields (sea surface relative vorticity $\zeta = v_x - u_y$ [1/s] and surface temperature [°C]) for **(a)** the reference solution (computed at horizontal resolution 1/12° and then projected on the 1/3°-grid), **(b)** hyper-parameterized solution computed at horizontal resolution 1/3° for $N = 18$, $M = 4$, $\eta = 0.001$, and the 2-year time-average (last column). Snapshots are taken after 1 year (left column) and 2 years (middle column) of simulations. Note that the hyper-parameterized solution **(b)** reproduces both large- (the Gulf Stream) and small-scale (vortices) features of the flow dynamics in both fields.

CRediT authorship contribution statement

Igor Shevchenko: Conceptualization, Methodology, Software, Validation, Formal analysis, Writing – original draft. **P. Berloff:** Writing – review & editing, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

The authors thank The Leverhulme Trust, United Kingdom for the support of this work through the grant RPG-2019-024, and the anonymous referees for their constructive comments, suggestions, and efforts which helped us improve the paper. Pavel Berloff was supported by the NERC grant NE/T002220/1 and by the Moscow Centre for Fundamental and Applied Mathematics (supported by the Agreement

075-15-2019-1624 with the Ministry of Education and Science of the Russian Federation, Russia).

References

Berloff, P., 2015. Dynamically consistent parameterization of mesoscale eddies. Part I: simple model. Ocean Model. 87, 1–19.

Berloff, P., 2016. Dynamically consistent parameterization of mesoscale eddies. Part II: eddy fluxes and diffusivity from transient impulses. Fluids 1, 1–19.

Berloff, P., 2018. Dynamically consistent parameterization of mesoscale eddies. Part III: Deterministic approach. Ocean Model. 127, 1–15.

Cooper, F., Zanna, L., 2015. Optimization of an idealised ocean model, stochastic parameterisation of sub-grid eddies. Ocean Model. 88, 38–53.

Cotter, C., Crisan, D., Holm, D., Pan, W., Shevchenko, I., 2019. Numerically modelling stochastic Lie transport in fluid dynamics. Multiscale Model. Simul. 17, 192–232.

Cotter, C., Crisan, D., Holm, D., Pan, W., Shevchenko, I., 2020a. A Particle Filter for Stochastic Advection by Lie Transport (SALT): A case study for the damped and forced incompressible 2D Euler equation. SIAM/ASA J. Uncertain. Quant. 8, 1446–1492.

Cotter, C., Crisan, D., Holm, D., Pan, W., Shevchenko, I., 2020b. Data assimilation for a quasi-geostrophic model with circulation-preserving stochastic transport noise. J. Stat. Phys. 179, 1186–1221.

Cotter, C., Crisan, D., Holm, D., Pan, W., Shevchenko, I., 2020c. Modelling uncertainty using stochastic transport noise in a 2-layer quasi-geostrophic model. Found. Data Sci. 2, 173–205.

Deremble, B., Wienders, N., Dewar, W., 2013. CheapAML: A simple, atmospheric boundary layer model for use in ocean-only model calculations. Mon. Wea. Rev. 141, 809–821.

Duan, J., Nadiga, B., 2007. Stochastic parameterization for large eddy simulation of geophysical flows. Proc. Amer. Math. Soc. 135, 1187–1196.

Frederiksen, J., O'Kane, T., Zidikheri, M., 2012. Stochastic subgrid parameterizations for atmospheric and oceanic flows. Phys. Scr. 85, 068202.

Gent, P., Mcwilliams, J., 1990. Isopycnal mixing in ocean circulation models. J. Phys. Oceanogr. 20, 150–155.

Grooms, I., Majda, A., Smith, K., 2015. Stochastic superparametrization in a quasi-geostrophic model of the Antarctic Circumpolar Current. Ocean Model. 85, 1–15.

Jamet, Q., Dewar, W., Wienders, N., Deremble, B., 2019. Spatiotemporal patterns of chaos in the Atlantic overturning circulation. Geophys. Res. Lett. 46, 7509–7517.

Mana, P., Zanna, L., 2014. Toward a stochastic parameterization of ocean mesoscale eddies. Ocean Model. 79, 1–20.

Marshall, J., Adcroft, A., Hill, C., Perelman, L., Heisey, C., 1997. A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. J. Geophys. Res. 102, 5753–5766.

Ryzhov, E., Kondrashov, D., Agarwal, N., Berloff, P., 2019. On data-driven augmentation of low-resolution ocean model dynamics. Ocean Model. 142, 101464.

Ryzhov, E., Kondrashov, D., Agarwal, N., McWilliams, J., Berloff, P., 2020. On data-driven induction of the low-frequency variability in a coarse-resolution ocean model. Ocean Model. 153, 101664.

Shevchenko, I., Berloff, P., 2021a. A method for preserving large-scale flow patterns in low-resolution ocean simulations. Ocean Model. 161, 101795.

Shevchenko, I., Berloff, P., 2021b. On a minimum set of equations for parameterisations in comprehensive ocean circulation models. Ocean Model. 168, 101913.

Shevchenko, I., Berloff, P., 2022a. A method for preserving nominally-resolved flow patterns in low-resolution ocean simulations: Constrained dynamics. Ocean Model. 178, 102098.

Shevchenko, I., Berloff, P., 2022b. A method for preserving nominally-resolved flow patterns in low-resolution ocean simulations: Dynamical system reconstruction. Ocean Model. 161, 101795.