

Diagnosing, modeling, and testing a multiplicative stochastic Gent-McWilliams parameterization

Ian Grooms*, William Kleiber

Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, United States



ARTICLE INFO

Keywords:
Stochastic parameterization
Mesoscale parameterization
Gent-McWilliams

ABSTRACT

A depth-independent isotropic Gent-McWilliams (GM) transport parameter κ is diagnosed from an idealized eddy-resolving primitive equation simulation. The optimal depth-independent isotropic GM parameterization is only able to model less than 50% of the diagnosed total tendency of temperature induced by unresolved mesoscale eddies. A spatio-temporal stochastic model of the GM parameter is developed based on the diagnosed values; the graphical lasso is used to estimate the spatial correlation structure. The stochastic model is used as a stochastic parameterization in low-resolution model simulations. The low-resolution stochastic simulation does a poor job of reproducing the temporal mean of large-scale temperature. Deterministic GM parameterizations and multiplicative stochastic GM parameterizations with unrealistic structure result in significantly more-accurate large-scale temperature in the low-resolution simulations. These results suggest that either the depth-independence or the isotropy of the GM parameterization are unrealistic as models of the eddy tracer transport, or that a stochastic GM parameterization should include an additive component.

1. Introduction

Global ocean models (hereinafter GCMs, general circulation models) are used for a variety of purposes including centennial climate forecasts, shorter-term ensemble forecasting or reanalysis scenarios, and paleoclimate research. For these particular purposes the computational cost associated with coupled ensemble simulations, long-term simulations, or simulations with many online tracers (e.g. biogeochemistry) precludes the use of spatial resolution fine enough to represent the ocean mesoscale. The models used for these purposes typically have horizontal spatial resolution on the order of 1° or larger, and will be referred to hereafter as ‘coarse’ GCMs. The inability to resolve mesoscale eddies is one of the foremost obstacles to developing accurate coarse GCMs.

The primary impact of unresolved dynamics in coarse GCMs is the transport of tracers and momentum. In GCMs where mesoscale eddy dynamics are completely unresolved, the primary subgrid-scale transport is of tracers rather than momentum (Grooms et al., 2011). The dominant parameterization paradigm for mesoscale tracer transport is the Gent-McWilliams framework (GM; Gent and McWilliams, 1990; Gent et al., 1995) which essentially codifies the fact that mesoscale eddies typically transport tracers along isopycnals in such a way as to reduce potential energy by flattening isopycnals. Mesoscale eddies also mix tracers along isopycnal directions; this effect is parameterized

separately from GM, usually by some form of Redi parameterization (Redi, 1982). The GM parameterization is deterministic in the sense that subgrid-scale tracer fluxes are modeled as deterministic functions of the resolved model variables. This approach would lead to realistic parameterizations if the ocean mesoscale were much smaller and faster than the dynamics resolved by the GCM. Scale-separation assumptions of this kind are made in multiscale asymptotic analyses of ocean eddy dynamics (Grooms et al., 2011; 2012). But in coarse GCMs, with grid scales on the order of 1° or larger, mesoscale eddies are not significantly smaller than the grid. As a result, the true subgrid-scale tracer fluxes are not completely determined by the resolved-scale variables; the large variability in tracer flux that remains even after averaging over square cells of width approximately 85 km was demonstrated by Grooms (2016). The turbulent mesoscale dynamics generate chaotic fluxes that are spatially and temporally correlated, and are dependent on, but not entirely determined by, the large-scale flow.

The fact that the subgrid-scale tracer fluxes are chaotic implies that parameterizations should be stochastic. The idea of using random numbers in numerical climate models goes back at least to Lorenz (1975). A major development in the practical use of stochastic parameterizations in atmospheric and oceanic GCMs came when a stochastic parameterization was implemented in the operational ECMWF Ensemble Prediction System (Buizza et al., 1999). Since then

* Corresponding author.

E-mail address: ian.grooms@colorado.edu (I. Grooms).

<https://doi.org/10.1016/j.oceomod.2018.10.009>

Received 29 August 2018; Received in revised form 20 October 2018; Accepted 28 October 2018

Available online 29 October 2018

1463-5003/ © 2018 Published by Elsevier Ltd.

stochastic parameterizations have been adopted by operational centers almost exclusively for use in atmospheric ensemble prediction systems because they increase ensemble spread, leading to better calibrated probabilistic forecasts (Berner et al., 2017; Leutbecher et al., 2017).

Many stochastic parameterizations (too numerous to list) have been developed for both atmosphere and ocean applications, but operational use of stochastic parameterizations has been dominated by the style of parameterization advocated by Buizza et al. (1999). Buizza et al. (1999) advocated taking the terms in tracer and momentum budgets that are associated with parameterizations (the ‘parameterization tendencies’) and multiplying them (‘perturbing’ them) by $(1 + e)$ where e is a random field with zero mean. This approach has come to be called ‘Stochastically-Perturbed Parameterization Tendencies’ (SPPT). Christensen et al. (2017) recently explored an alternative whereby different parameterizations are multiplied by different independent random fields in an approach called independent SPPT (iSPPT). SPPT is not conservative of tracers or momentum, and it has recently been found to dry out the atmosphere in long-running simulations (Davini et al., 2017). Conservation within the context of SPPT can be enforced using a method proposed by Leutbecher et al. (2017). As an alternative the parameters within a parameterization can be perturbed, rather than the tendency produced by the parameterization, and this approach is called ‘Stochastically-Perturbed Parameters’ (SPP). (Grooms, 2016 used SPP but incorrectly called it SPPT.)

SPPT and SPP have recently been applied to coarse GCMs in ocean-only and coupled simulations. Brankart et al. (2015) and Andrejczuk et al. (2016) used SPPT in a coarse GCM but only applied the perturbations to mixing parameterizations and not to the GM parameterization. Juricke et al. (2017) and Juricke et al. (2018) used SPP with the GM parameterization. There are a few other stochastic parameterizations for coarse GCMs that are not based on SPP or SPPT. These include the stochastic equation of state developed by Brankart (2013), and the addition of stochastic noise forcing to the temperature equation of a GCM by Williams et al. (2016). Grooms (2016) developed a framework for stochastic GM parameterizations and developed a preliminary non-Gaussian stochastic GM parameterization, albeit in an idealized model.

The goal of the present investigation is to push the SPP approach to GM parameterization to its limits by choosing the multiplicative perturbation structure to be as realistic as possible. Recall that the SPP approach requires specifying the structure of the random field e where $(1 + e)$ is used to perturb the parameter of interest. The random field is usually specified via some sort of pattern generator with tunable time and length scales (see, e.g. Leutbecher et al., 2017, for examples). In the present investigation the parameter κ in the standard isotropic GM parameterization is estimated directly from eddy-resolving simulations, and is allowed to vary in time and horizontally in space. The results are then used to develop a random-field model for κ , which is equivalent to modeling the perturbation e in the SPP approach. This model is then used in an implementation of a stochastic GM parameterization, which is compared to deterministic GM parameterizations and to a less-realistic SPP-GM parameterization.

We begin in Section 2 by diagnosing the GM parameter from an eddy-resolving simulation in an idealized domain. This section includes a description of the results of the diagnosis, while the next section, Section 3, develops the stochastic model for κ . Section 4 briefly describes the configuration of parameterized coarse-model simulations, and then compares the results of various coarse-model simulations with deterministic and stochastic GM parameterizations. Results are discussed in Section 5, and conclusions are offered in Section 6.

2. Diagnosing the GM parameter

2.1. Gent-McWilliams background

Before detailing the method used to diagnose the GM parameter we

begin with a brief review of the deterministic GM framework (Gent and McWilliams, 1990; Gent et al., 1995). In the GM framework the divergence of the subgrid-scale tracer flux is parameterized by a bolus velocity advection

$$\nabla \cdot (\mathbf{u}^{\dagger} \tau') = \mathbf{u}^{\dagger} \cdot \nabla \tau. \quad (1)$$

Throughout this section resolved-scale variables have no primes and subgrid-scale variables are denoted by a prime '. The overbar (τ) denotes a spatial average to the coarse grid scale. In the above expression τ denotes a tracer concentration, and the bolus velocity is

$$\mathbf{u}^{\dagger} = \nabla \times \Psi \quad (2)$$

where $\Psi = (\Psi^x, \Psi^y, 0)^T$ is a vector streamfunction. Eq. (2) guarantees that the bolus velocity is incompressible, which is a key property of the GM framework, namely that the flux is non-mixing, i.e. adiabatic.

In the standard isotropic version of the GM parameterization the vector streamfunction Ψ is linearly related to the local isopycnal slope $\mathbf{s} = -\nabla_h \rho / \partial_z \rho$ where $\nabla_h = (\partial_x, \partial_y)^T$ and ρ is the density :

$$\Psi^x = -\kappa s^y, \quad \Psi^y = \kappa s^x. \quad (3)$$

(When using a nonlinear equation of state \mathbf{s} should be replaced by the slope of neutral surfaces (McDougall, 1987).) When $\kappa > 0$ Eq. (3) guarantees that the parameterization acts to remove potential energy by tilting isopycnal surfaces towards horizontal (Gent et al., 1995).

There are two difficulties associated with this formulation of the GM parameterization: how to deal with unstratified or convectively-unstable regions ($\partial_z \rho \geq 0$), and how to guarantee that the vertical part of the bolus velocity w^{\dagger} is zero at the ocean surface or floor. The vertical bolus velocity is

$$w^{\dagger} = \partial_x \Psi^y - \partial_y \Psi^x.$$

The vertical component of the bolus velocity can be set to zero at a given depth by setting $\Psi = 0$ at that depth. Ferrari et al. (2010) proposed a simple and elegant approach that simultaneously ensures impenetrability at the ocean surface and deals with unstratified or convectively-unstable regions: the vector streamfunction is specified as the solution of the following two-point boundary value problem

$$\left[c^2 \partial_z^2 + \frac{g \partial_z \rho}{\rho_0} \right] \Psi^y = -\kappa \frac{g}{\rho_0} \partial_x \rho, \quad (4)$$

$$\left[c^2 \partial_z^2 + \frac{g \partial_z \rho}{\rho_0} \right] \Psi^x = \kappa \frac{g}{\rho_0} \partial_y \rho \quad (5)$$

with boundary conditions $\Psi^x = \Psi^y = 0$. When c^2 is chosen appropriately the streamfunction approximates (3), and the solutions smoothly transition through unstratified and convectively unstable regions. In practice the density gradient $\partial_z \rho$ that appears in these equations is truncated to negative values; here the value of $N^2(z) = -g \partial_z \rho / \rho_0$ is truncated for values less than 10^{-24} s^{-2} . Ferrari et al. (2010) chose the following specification for c

$$c = \max \left\{ c_{\min}, \pi^{-1} \int_0^H N(z) dz \right\} \quad (6)$$

where the ocean depth is H , and c_{\min} is a tunable parameter, here set to 1 m/s.

2.2. Eddy-resolving model configuration

The eddy-resolving simulation uses the MITgcm (Marshall et al., 1997) with a linear equation of state based solely on temperature (no salinity); the expansion coefficient for temperature is $\alpha_T = 1.5 \times 10^{-4} \text{ K}^{-1}$. The model domain is 3,200 km square and 4 km deep. A Cartesian tangent plane approximation is used, with the Coriolis parameter taking the value $f_0 = 5 \times 10^{-5} \text{ s}^{-1}$ at the southern edge of the domain and increasing to 10^{-4} s^{-1} at the northern edge. The vertical viscosity is

$A_v = 10^{-3} \text{ m}^2/\text{s}$ and the horizontal biharmonic viscosity is $A_{4,h} = 2 \times 10^{10} \text{ m}^4/\text{s}$. The momentum boundary conditions are no-slip at the sides and bottom, with a rigid lid. The vertical diffusivity is $\kappa_v = 10^{-4} \text{ m}^2/\text{s}$ except in regions where $\partial_z T \leq 0$, in which case it increases to $10 \text{ m}^2/\text{s}$. The horizontal diffusivity of temperature is set to $\kappa_h = 50 \text{ m}^2/\text{s}$.

The model is forced by an asymmetric double-gyre zonal wind stress of the form $0.2 \sin(\pi y/L_y - \pi/6) \text{ N/m}^2$, and by an interactive heat flux (out of the ocean) equal to $\gamma(T - T_0)$ where $\gamma = 35 \text{ W/m}^2\text{K}^{-1}$ and T_0 is a half-period cosine varying from 22° C at the southern boundary to 2° C at the northern boundary. The heat flux is converted to a temperature flux using a heat capacity of $4000 \text{ J K}^{-1} \text{ kg}^{-1}$. No heat flux is allowed through the bottom or side boundaries.

The initial condition was taken from an equilibrated solution of the ‘deterministic GM’ model described in Grooms (2016), which has grid size of 40 km and 61 vertical levels with depths varying linearly from 25 m at the surface to 106 m at depth. This low-resolution equilibrated state was interpolated to the eddy-resolving resolution of 8 km . The circulation exhibits a large subtropical gyre and a weaker subpolar gyre, with strong eddies along the western boundary and into the interior, very similar to the simulations of Henning and Vallis (2004). A snapshot of the sea surface temperature and the sea surface height anomaly in the statistically steady state are shown in Fig. 1.

2.3. Diagnosing the GM parameter

The GM parameter κ is diagnosed by minimizing the integral of the square of the error between the true eddy density flux divergence and the modeled divergence

$$\kappa = \arg \min \int (\nabla \cdot (\mathbf{u}\rho) - \nabla \cdot (\bar{\mathbf{u}}\bar{\rho}) - \mathbf{u}^\dagger \cdot \nabla \bar{\rho})^2 dx \quad (7)$$

where the notation ‘arg min’ indicates that κ is the value that minimizes this expression, the integral is over the volume, and κ is allowed to vary horizontally but not vertically. The restriction that κ is depth-independent is perhaps overly severe (cf. Abernathy et al., 2013), but it significantly reduces the size and computational cost of the optimization problem; limitations of the model are discussed further below. Minimizing the error in the flux divergence avoids problems related to gauge freedom that would arise if one instead minimized the error in the flux itself (Mak et al., 2016). Because the bolus velocity \mathbf{u}^\dagger depends linearly on κ , this is a linear least squares problem. Solving the least-squares problem requires computing the true eddy density flux

divergence, which requires defining a low-pass spatial filter $\bar{\cdot}$. This was implemented as a moving average with a Gaussian kernel

$$\bar{u}(x, y, z, t) = \frac{1}{Z} \int e^{-\frac{(x-\xi)^2 + (y-\eta)^2}{2L^2}} u(\xi, \eta, z, t) d\xi d\eta, \quad Z = \int e^{-\frac{\xi^2 + \eta^2}{2L^2}} d\xi d\eta.$$

The kernel width L was 60 km and Z is a normalization constant. Values outside the boundaries were computed using periodic extensions of the interior values of the fields. The density and vertical velocity used even extensions at all boundaries. The zonal velocity used even extensions at the northern and southern boundaries and odd extensions at the eastern and western boundaries, while the meridional velocity used odd extensions at the northern and southern boundaries and even extensions at the eastern and western boundaries.

A kernel length scale of 60 km is arguably close enough to typical eddy length scales that it is possible that it fails to make a clear separation between the eddies and the scales resolvable by the coarse model (though one could argue that the length scale is actually $\sqrt{2} \times 60 \text{ km} \approx 85 \text{ km}$, which is comparable to a typical eddy radius). Repeating the analysis with a variety of filter widths involves prohibitive computational expense, so as an alternate way to test this, the power spectral density of the filtered, large-scale part of temperature was computed over a period from the 98th year of the eddy-resolving simulation to the 208th year, with data taken every tenth day. The power spectral density of sea surface temperature at a point 380 km west of the eastern boundary and 1580 km north of the southern boundary (at the location marked \times (a) in Fig. 5) is shown as a function of the oscillation period in Fig. 2. Most of the power is clearly associated with long time scales on the order of 10 weeks or longer; this slow variability is characteristic of location and depth where the power spectral density was computed. Eddy variability dominates the total variability and is expected to occur on timescales of at most a few weeks, so the fact that the temporal variability of the large-scale part of temperature occurs on much longer time scales implies that the low-pass spatial filter is ineffective at removing eddy length scales.

The eddy-resolving simulation was allowed to run for 200 years after initializing from the low-resolution model before taking the data used to compute κ . The fields in Eq. (7) were computed on the coarse-model grid of 40 km , and the optimal value of κ was computed by solving the normal equations. The results are extremely rough; the left panel of Fig. 3 shows the optimal κ computed at a single time. It may be the case that the extreme values of κ evident in the left panel of Fig. 3 result from locations where the GM model predicts a near-zero flux-

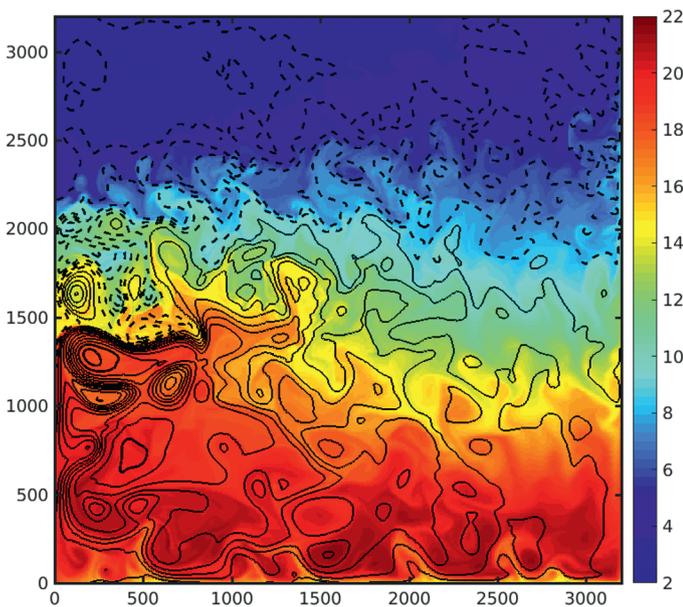


Fig. 1. A snapshot of sea surface temperature (color) and sea surface height anomaly (contours) from the high-resolution reference simulation. The contour interval is 10 cm ; negative contours are dashed and positive contours are solid. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

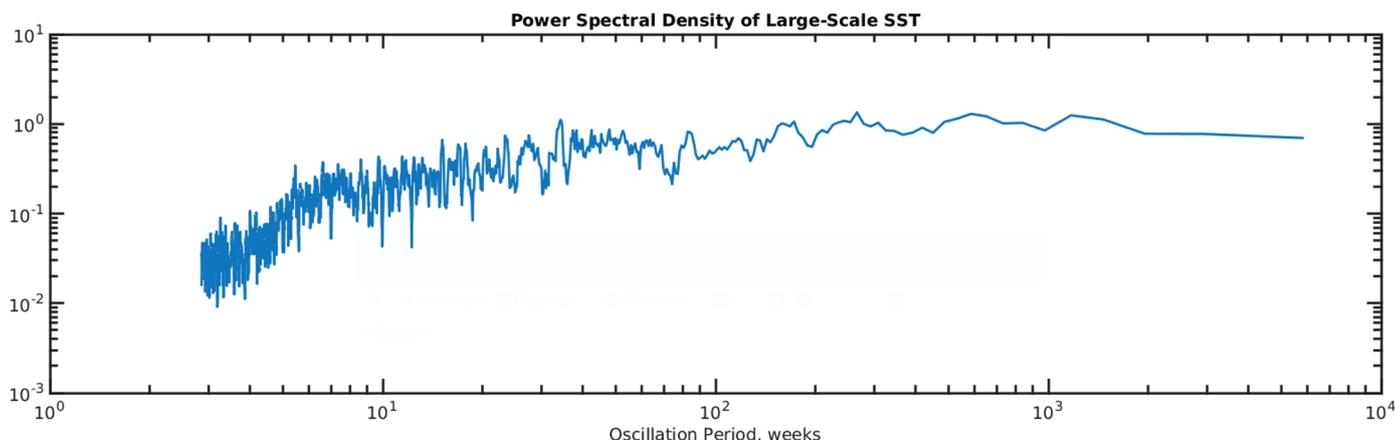


Fig. 2. Power spectral density of the large-scale part of sea surface temperature (SST) at the point marked \times (a) in Fig. 5: 380 km east of the western boundary and 1580 km north of the southern boundary.

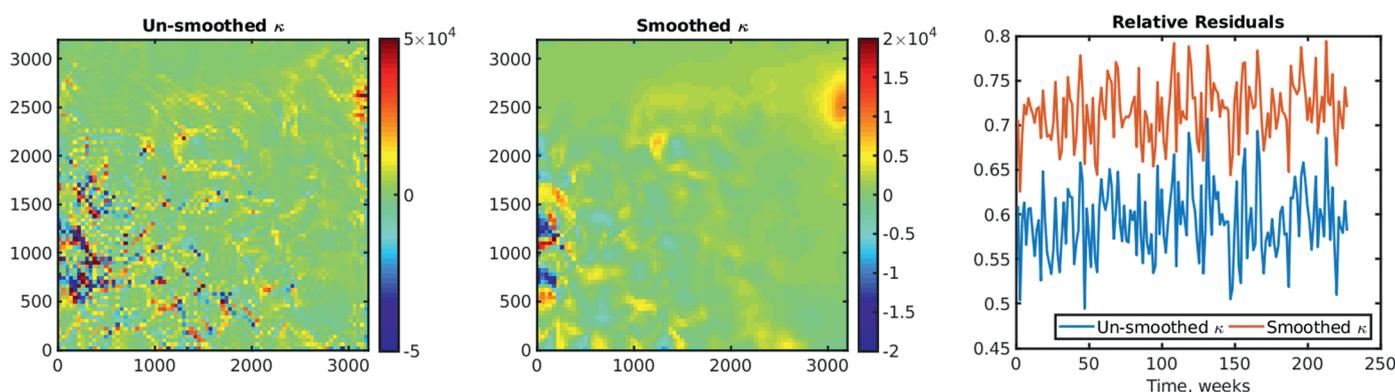


Fig. 3. Left: A snapshot of κ computed using Eq. (7); units are m^2/s ; axis units are kilometers. Center: a snapshot of κ computed using Eq. (8) at the same time as the left panel; note the different color scale compared to the left panel. Right: Relative residuals as a function of time for the κ computed using Eqs. (7) (blue) and (8) (orange). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

divergence (e.g. in regions of small isopycnal slope) while the true flux-divergence is large, and κ becomes large to compensate. Regardless of their origin, extreme values of κ on the order of $10^5 \text{ m}^2/\text{s}$ would necessitate extremely small time steps to avoid numerical instability in a low-resolution model.

In order to avoid such extremely-rough solutions, a term was added to the least squares problem that penalizes the gradient of κ

$$\kappa = \arg \min \int (\nabla \cdot (\mathbf{u}\rho) - \nabla \cdot (\mathbf{u}\bar{\rho}) - \mathbf{u}^T \cdot \nabla \bar{\rho})^2 + \epsilon |\nabla_h \kappa|^2 dx. \quad (8)$$

After some investigation a value of $\epsilon = 5 \times 10^{-17}$ was deemed sufficient to improving the smoothness of κ without overly damaging the accuracy. The center panel of Fig. 3 shows the value obtained from the penalty-smoothed least-squares problem. The optimal value of κ was computed every 10 days for 1590 days (approximately 4.36 years) both with and without the smoothing penalty term.

Negative values of κ appear in the smoothed results almost everywhere, excepting only in the extreme north end of the domain, and negative values occur on average in various parts of the domain. There are mathematical difficulties with partial differential equations (PDEs) with negative viscosity, and one might want to impose a non-negativity condition on the least-squares problem to avoid this. Numerically solving a PDE with negative viscosity is problematic because the computed solution depends strongly on the discretization and need not converge as the grid is refined (Durran, 2010). On the other hand the negative viscosity exists here because a specific, fixed coarse resolution and associated spatial filter was chosen for the diagnosis of κ and it would not make sense to use that same negative value of κ with increased resolution and decreased filter scale. There are no mathematical problems

with the discretized model with negative viscosity at fixed coarse resolution, so negative values are retained here to make the GM parameterization as accurate as possible.

The optimal value of κ reduces the norm of the error (technically the ‘residual’) on the right hand side of Eq. (7) below the value it would have if $\kappa = 0$. One can measure the goodness-of-fit using the relative residual: the norm of the error in Eq. (7) using the optimal value of κ divided by the norm of the error with $\kappa = 0$. This relative residual is shown as a function of time for both the original and penalty-smoothed versions of κ in the right panel of Fig. 3; in the penalty-smoothed relative residual the penalty term is not included in the error, so that the comparison to the original least-squares problem is more precise and physically relevant. The un-smoothed κ has errors of about 60% on average, while the smoothed κ has errors of about 72% on average. Neither of these is very accurate, which suggests that the isotropic, depth-independent GM parameterization is simply a poor model of the true eddy density flux divergence. This should not be overly surprising since Abernathy et al. (2013) directly diagnosed the GM parameter κ in a differently-configured simulation and observed significant vertical variation. Smith and Marshall (2009) and Bachman and Fox-Kemper (2013) and others have similarly observed depth-dependent isopycnal and diapycnal mixing parameters.

3. A random-field model for κ

This section describes the construction of a random-field model for κ based on the diagnosed values described in the preceding section. Fig. 4 shows the median (upper row) and standard deviation (lower

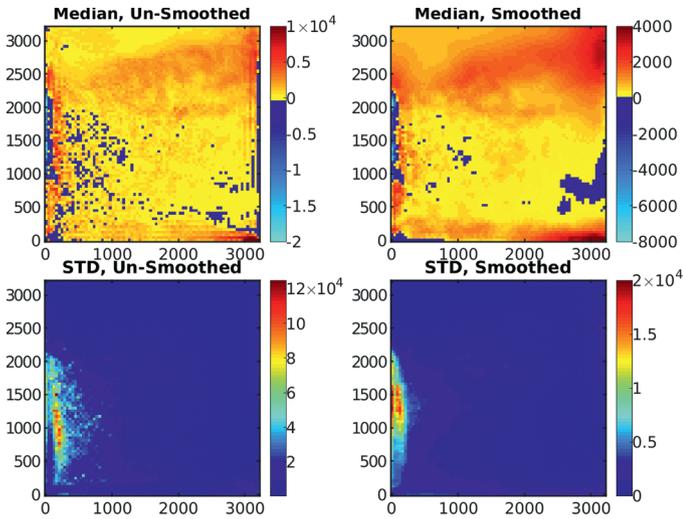


Fig. 4. Top Row: Median value of κ . Bottom row: Standard Deviation (STD) of κ . Left column: κ obtained from Eq. (7), i.e. without a smoothing penalty term. Right column: κ obtained from Eq. (8), i.e. with a smoothing penalty term. Units of median and standard deviation are m^2/s . Axis units are kilometers. The color map in the top column has been chosen to make negative values more obvious. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

row) of the diagnosed values of κ using the un-smoothed results of Eq. (7) (left column) and the smoothed results from Eq. (8) (right column). There is clearly large variability in κ , particularly for the un-smoothed results, which is why the median is used as a central estimator for the field instead of the mean. The median is much less than one standard deviation from the mean, and is robust to outliers; also the median is somewhat smoother spatially than the mean (not shown). The grainy quality of the median and standard deviation of the un-smoothed κ suggest that these statistics are hard to estimate accurately from the time series of 159 points. This consideration, together with the extreme roughness of the un-smoothed fields shown in the left panel of Fig. 3 motivates us to focus henceforth on the smoothed version of κ . In both the smoothed and unsmoothed versions of κ the increased variability shown in the lower row of Fig. 3 is associated with and presumably driven by the increased eddy energy along the western boundary.

We will develop a random field model for the smoothed version of κ . The mean and standard deviation will be given by the values shown in the right column of Fig. 4. The following two sections study the spatial and temporal correlation structure of the smoothed version of κ .

3.1. Temporal structure

The temporal structure in particular is hard to accurately estimate with a time series with spacing of 10 days. In particular, the correlation

drops to 0 in less than 10 days near the western boundary because of the increased eddy energy and variability in that region. So the eddy-resolving simulation was extended by a further 201 days keeping daily snapshots. The left panel of Fig. 5 shows the largest time lag for which the temporal autocorrelation of κ is greater than 0.5 (in this section κ should be understood to be the smoothed version computed from Eq. (8)). The center panel of Fig. 5 shows the empirical autocorrelation function at the location marked \times (a) in the left panel, and the right panel shows the empirical autocorrelation function at the location marked \times (b) in the left panel. The autocorrelation function in the right panel of Fig. 5 was estimated using the original time series of length 1590 days with 10-day gaps. There is clearly huge variability in the temporal correlation structure of κ , with decorrelation times on the order of a few days near the center of the subtropical gyre and decorrelation times longer than 3 weeks in the northeast corner of the domain. It is doubtful that the slow behavior seen in the northeast corner is eddy-driven, since eddies presumably decorrelate far more rapidly. Instead, this low frequency variability in κ may instead be reflecting low-frequency variability in the large-scale structure of the density field itself; such low-frequency variability is not uncommon in this kind of model (Arzel et al., 2006).

3.2. Spatial structure

To investigate the spatial correlation structure of κ (in this section κ

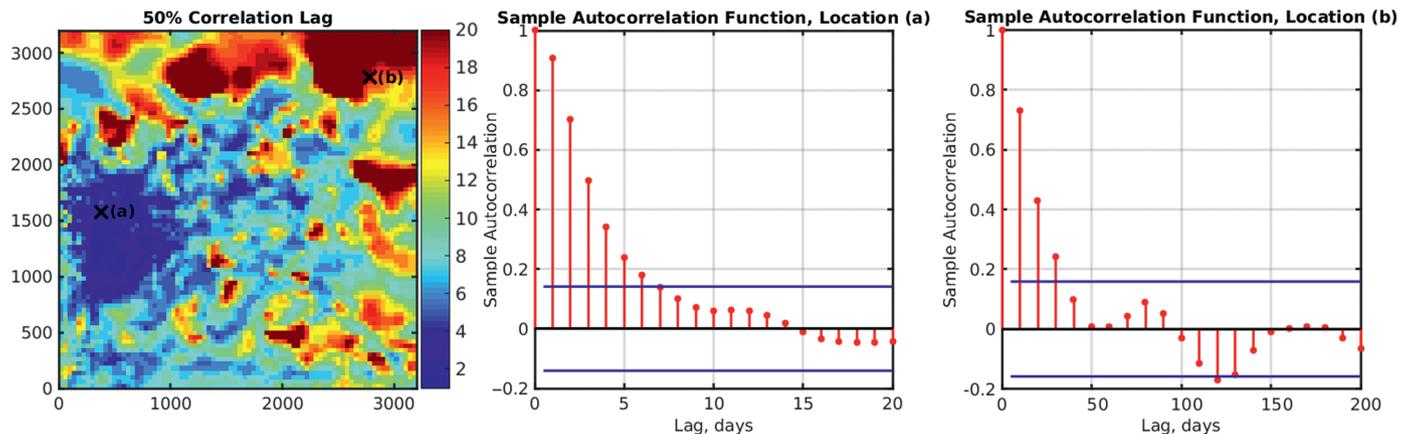


Fig. 5. Left: The longest time lag, in days, for which the autocorrelation of κ remains above 50%. Axis units are kilometers. Center: The empirical autocorrelation function of κ at the point marked (a) in the left panel. The blue line shows the level at which the autocorrelation is not statistically significantly different from zero. Right: The empirical autocorrelation function of κ at the point marked (b) in the left panel. Note the range of x-axis values on the sample autocorrelation functions differ. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

should be understood to be the smoothed version computed from Eq. (8)) we first subtract the pointwise time average and divide by the standard deviation to achieve a centered field with unit variance. This is called the ‘centered, scaled’ version of κ , and it is the subject of analysis in this section.

The spatial structure is estimated using the original time series of length 159 with 10-day spacing, plus every tenth day of the time series of length 201 with 1-day spacing. This yields a time series of length 180 with 10 day spacing. As noted in the previous section, κ displays low-frequency variability in the northeast corner of the domain that is probably not directly the result of eddy variability. To avoid modeling this low-frequency variability a high-pass time filter is applied to κ before examining the spatial structure. This high-pass filter simply consists in removing a moving average from the data with weights [0.25 0.5 0.25], and the first and last points are removed leaving a time series of length 178. Since the time series has a 10 day spacing and most of the variability in κ has time scales on the order of 10 days or less (left panel of Fig. 5), the high-pass filter only removes time scales significantly longer than 10 days and leaves the primary eddy variability intact.

The centered, scaled field is anisotropic and highly inhomogeneous and cannot be described by any of the standard homogeneous random field models, i.e. it cannot be accurately modeled using the kinds of pattern generators typically used for SPP (Leutbecher et al., 2017). A standard technique for dealing with inhomogeneous fields is to compute empirical orthogonal functions (EOFs). EOFs are approximations to the eigenvectors of the covariance matrix of the field, and are obtained through a singular value decomposition of the data. Each EOF explains a component of the total variance of the field, and EOF analysis is particularly useful when a small number of EOFs describe most of the variance. Unfortunately, as shown in Fig. 6, each EOF accounts for only a small amount of the total variance of the centered, scaled field; the leading EOF only accounts for 6% of the total variance. Moreover, an EOF-based model for κ would have at most 177 degrees of freedom while the field itself (on the 80×80 coarse grid) has 6400 degrees of freedom.

The ‘graphical lasso’ (Friedman et al., 2008) is a method to estimate the precision matrix of a set of variables (the precision matrix is the inverse of the covariance matrix). It generates a sparse approximation to the precision matrix by minimizing a cost function that is the sum of a rank-deficient L^2 term measuring the mismatch between the precision matrix and the data and an L^1 penalty term that regularizes the problem and enhances sparsity of the result. In the statistical literature, it is well known that the precision matrix (inverse covariance matrix) encodes a graphical structure on the random field, and is typically expected to be a sparse matrix that relays information on a conditional dependence structure (Rue and Held, 2005). A penalty parameter $\lambda \geq 0$ determines

the sparsity of the estimate: large λ leads to a very sparse but less-accurate approximation while small λ leads to a more dense estimate ($\lambda = 0$ results in the un-penalized empirical covariance estimate). In particular, we use the graphical lasso to estimate a random field model as it imposes no assumptions of homogeneity.

A key benefit of the graphical lasso is that the precision matrix that it produces is invertible, i.e. full-rank, unlike the sample covariance matrix associated with EOF analysis. While the EOF analysis produces 177 spatial patterns of variability (EOFs) from a time series of length 178, the graphical lasso produces a complete set of 6400 spatial patterns of variability. The spatial patterns associated with the graphical lasso are the eigenvectors of the precision matrix; these are not the same as the EOFs, but they are approximating the same thing: the eigenvectors of the true covariance matrix. Fig. 7 shows that the first two EOFs are essentially the same as the corresponding eigenvectors of the graphical lasso precision matrix with penalty parameter $\lambda = 0.01$.

Like in EOF analysis, each of the graphical lasso spatial patterns explains a component of the total variance: the amount of variance explained is the inverse of the eigenvalue of the precision matrix associated with the spatial pattern. Fig. 6 shows the percentage of the total variance that is explained by each spatial pattern for the EOF analysis and for the graphical lasso analysis with both $\lambda = 0.1$ and $\lambda = 0.01$. The variances associated with the leading patterns of variability are equal in the two approaches and for both values of λ , but while the EOF analysis fits the total variance into 177 spatial patterns the graphical lasso pushes much of the variance into 6400 separate spatial patterns. At $\lambda = 0.1$ the graphical lasso underestimates the variance in several of the leading patterns of variability, and instead puts this variance into the long tail of patterns that is associated with spatially-incoherent noise. As λ decreases the estimation becomes more accurate and the graphical lasso agrees with the EOF analysis for most of the leading patterns of variability. Nevertheless, at $\lambda = 0.01$ some of the total variance is associated with a long tail of patterns that is associated with spatially-incoherent noise. The EOF and graphical lasso agree at large scales, but by forcing all the variance into 177 global spatial patterns the EOF analysis attributes too much long-range spatial coherence to the noise at small scales.

3.3. A stochastic model for κ

This section formulates a stochastic model that attempts to approximate the spatial and temporal structure discussed in the preceding two sections. The approach is similar to that used by Berloff (2005).

We begin with the spatial structure. The graphical lasso yields a spatial precision matrix $\mathbf{P} = \mathbf{C}^{-1}$ where \mathbf{C} is an approximation to the spatial covariance matrix. Let $\mathbf{L}\mathbf{L}^T = \mathbf{P}$ be the Cholesky decomposition

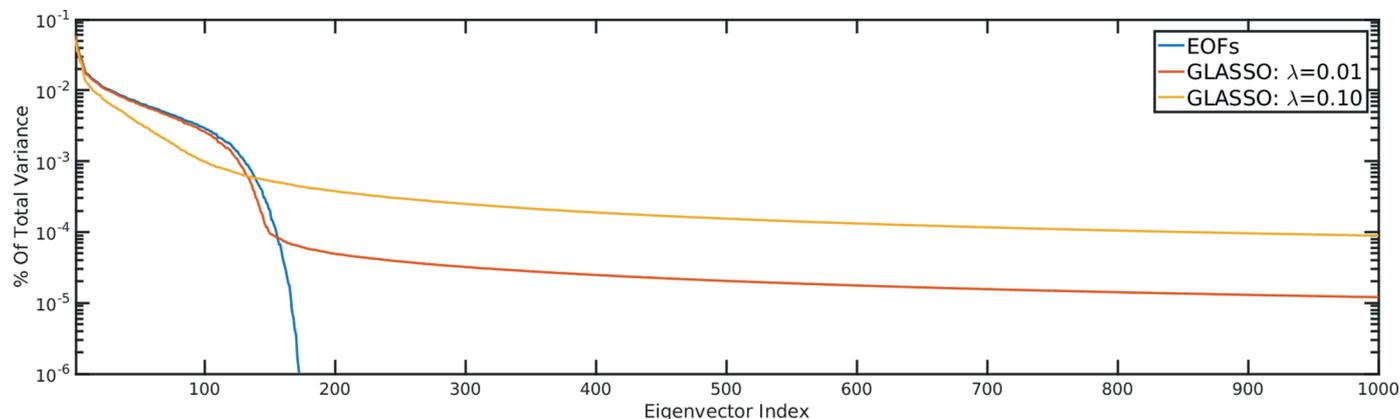


Fig. 6. Percent of total variance explained by (i) EOFs (blue), (ii) eigenvectors of the graphical lasso (GLASSO) precision matrix at (orange) and $\lambda = 0.10$ (yellow). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

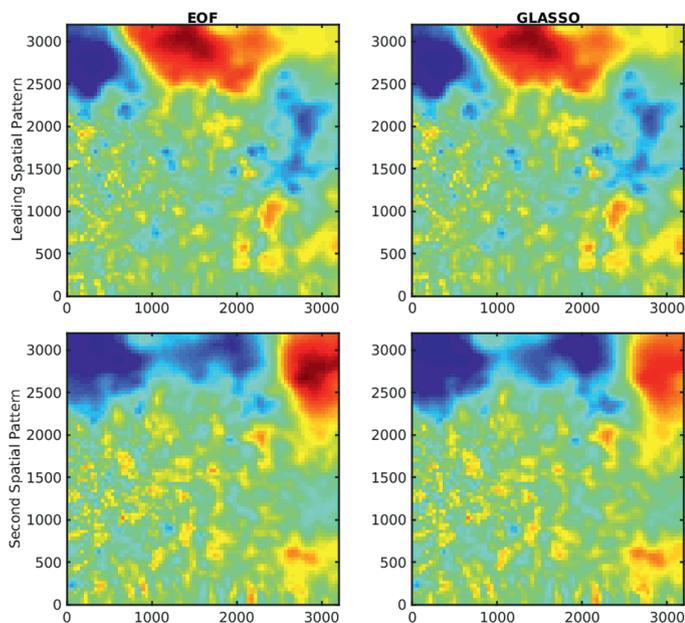


Fig. 7. Left column: Empirical orthogonal functions (EOFs). Right column: Eigenvectors of the graphical lasso (GLASSO) precision matrix. Upper row: leading spatial pattern of variability. Lower row: secondary spatial pattern of variability. Axis units are kilometers. The spatial patterns are dimensionless and their amplitudes are arbitrary, though they have been scaled consistently for ease of visual comparison.

of \mathbf{P} and let \mathbf{x} be a vector of zero-mean, unit-variance independent random variables. Then $\mathbf{z} = \mathbf{L}^{-T}\mathbf{x}$ is a vector with covariance matrix \mathbf{C} . This shows how to construct a random field with a spatial correlation structure that mimics κ : compute the Cholesky factorization of the precision matrix computed by the graphical lasso, and use it to give spatial correlation to a field of independent random variables.

A simple approach to modeling the temporal correlation structure is to construct an AR-1 (first order autoregressive) process at each spatial location on the coarse grid. AR-1 processes are unable to model negative correlations, but the autocorrelation functions diagnosed across the domain are all similar to those shown in the center and left panels of Fig. 5, i.e. they are monotonically decreasing until they become statistically insignificant. In a more general setting AR-2 or higher processes could be used. The decorrelation time of this process can be tuned to match the observed decorrelation time of κ . It is not straightforward to combine the two foregoing ideas (Cholesky plus AR-1 processes) into a model that accurately reflects the full spatiotemporal correlation structure of the true field κ . Our model takes the following form: At each point of the coarse grid the centered, scaled version of κ is modeled as an AR-1 process in the form

$$\hat{\kappa}_{n+1} = r\hat{\kappa}_n + \sqrt{1 - r^2}\chi_n \quad (9)$$

where χ_n are unit-variance, zero-mean random variables that are independent from one time to the next (the subscript n indicates time), and r governs the local decorrelation time. Specifically, r is chosen so that the AR-1 process reaches a correlation of 50% at a lag time given by the 50% lag time shown in the left panel of Fig. 5, with one caveat: The times shown in Fig. 5 are in one-day increments which leads to sharp jumps in the decorrelation time. The field shown in Fig. 5 is therefore smoothed by applying a moving average with weights [0.25 0.5 0.25] in each direction before being used to set r for the AR-1 processes.

Spatial structure is given to the centered, scaled $\hat{\kappa}$ by endowing the innovations χ_n with spatial correlation. This is achieved using the Cholesky method described above, with the precision matrix produced by the graphical lasso with $\lambda = 0.01$. The final value of κ is obtained by multiplying the scaled-centered version $\hat{\kappa}$ by the standard deviation shown in the lower right panel of Fig. 4 and then adding the median shown in the upper right panel of Fig. 4.

4. Coarse model simulations

Coarse model simulations are configured as in Grooms (2016) and using parameter values from Section 2.2, with the following exceptions. The horizontal grid size is 40 km and diffusion of temperature is accomplished by a third-order upwind scheme rather than by an explicit diffusivity. The horizontal viscosity is Laplacian rather than bi-harmonic; the value $8000 \text{ m}^2/\text{s}$ is applied to the baroclinic part of the flow and the value $1000 \text{ m}^2/\text{s}$ is applied to the barotropic part. Different values are used in an effort to make the barotropic and baroclinic viscous boundary layers both have widths on the order of one grid cell. The time-mean state of the large-scale part of the eddy-resolving simulation described in Section 2.2 was used as an initial condition for the coarse model simulations.

We run simulations with a constant value of $\kappa = 6000 \text{ m}^2/\text{s}$ for comparison to the stochastic GM scheme described in the preceding section; this simulation is called the ‘deterministic GM’ simulation. In addition, we run simulations where the GM parameter κ is perturbed by multiplying by $1 + e$ where e is a unit-variance, zero-mean Gaussian random field. The temporal structure of e is an AR-1 process with decorrelation time of 10 days, and the spatial structure results from applying a moving average with weights [0.25 0.5 0.25] in each direction to an uncorrelated field (see Grooms, 2016 for more details). This simulation is called the ‘SPP-GM’ simulation. The simulation using the model described in Section 3 is simply called the ‘stochastic GM’ simulation.

The temperature field in the coarse model simulations is compared to the large-scale part of the temperature field from the eddy-resolving simulation described in Section 2.2. The large-scale part of the temperature is obtained by the Gaussian kernel smoother described in Section 2.3; it was computed every 10 days starting in year 98 of the eddy-resolving simulation and ending in year 208 – a 110 year time series.

Fig. 8 shows the time-mean sea surface temperature (SST) from the eddy-resolving simulation (top left), from the deterministic GM simulation (top right), from the SPP-GM simulation (lower left), and from the stochastic GM simulation (lower right). The deterministic and SPP-GM schemes provide reasonably accurate predictions for the mean SST, while the tuned stochastic GM simulation is badly incorrect. The true SST has a steady decrease of SST with latitude, with very little zonal variability. The stochastic GM simulation has a plume of warm water originating near the southwest corner of the domain and extending

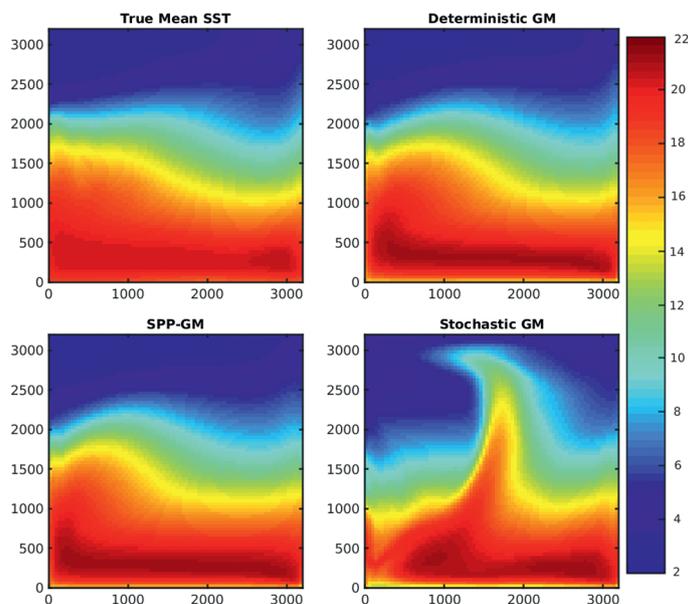


Fig. 8. Time-mean sea surface temperature for the eddy-resolving simulation (upper left), the deterministic GM simulation (upper right), the SPP-GM simulation (lower left) and the stochastic GM simulation (lower right). Temperature is in degrees Celsius and axis units are kilometers.

north and east up through the center of the domain. The mean temperature structure in the SPP-GM simulation is similarly inaccurate throughout the main thermocline (not shown). In view of the very poor accuracy of the optimized stochastic GM parameterization that is the primary focus of the investigation, there is little point in proceeding to test whether the scheme generates accurate variability, or in tuning other parameters of the deterministic and SPP-GM simulations.

5. Discussion

The ultimate goal of this research is to formulate a GM parameterization that produces an accurate coarse model simulation, and there are two routes to achieve this. One route, not taken here, involves tuning the parameterization by running repeated coarse model simulations and comparing the results to known reference data; this process is costly because of the infinite number of possible ways to tune the parameterization. The other route, taken here, involves studying the process or quantity that is being parameterized and attempting to build an accurate model thereof. This route is based on the assumption that a parameterization that accurately models the true eddy flux divergence should result in an accurate coarse model simulation. The primary result of the foregoing section is the fact that a simulation using a stochastic GM parameterization that is configured to match as closely as possible the GM parameter diagnosed from a high resolution reference experiment yields significantly worse results than a basic simulation with a constant value for κ . If the underlying assumption is true then the contrapositive must also be true: a parameterization that yields an inaccurate coarse model simulation must be an inaccurate model of the true eddy flux divergence. The stochastic GM parameterization derived in Section 3 must therefore be an inaccurate model of the true eddy flux divergence. (The converse, i.e. if a parameterization yields an accurate simulation then it must also be an accurate model of the true eddy flux divergence, is clearly false since Section 2 shows that constant κ gives a less-accurate model of the eddy flux divergence while Section 4 shows that constant κ yields a more-accurate coarse model simulation.)

There are many ways in which the stochastic GM model developed in Section 3 might be inaccurate. For example, as with all SPP stochastic parameterizations the modeled variations in κ are independent of the large-scale temperature, whereas the real variations in κ might depend on the large-scale temperature. Alternatively it might be the case that despite the care taken in developing a model with realistic spatial and

temporal variability, some crucial aspect of the true variability of κ has been missed in the stochastic model of κ . But the results of Section 2 indicate a more fundamental problem: the isotropic GM parameterization with depth-independent κ yields a very poor fit to the true eddy flux divergence, so that no matter how well the diagnosed κ is represented with a stochastic model it will never be an accurate model of the true eddy flux divergence. (Recall that the optimal κ computed using this configuration reduced the error between the parameterized eddy tracer flux divergence and the true eddy tracer flux divergence by only 30 to 40% as compared to $\kappa = 0$ (see Fig. 3).) It is possible that an anisotropic formulation (Smith and Gent, 2004) or a depth-dependent κ (Abernathey et al., 2013) might yield a better fit to the flux-divergence. In particular allowing arbitrary depth dependence would significantly increase the number of degrees of freedom available for fitting the true eddy flux divergence. But at a more general level it may be the case that it is simply not possible to model the true eddy flux divergence purely by varying κ (or the flux tensor in anisotropic GM) in a GM parameterization.

Recall that the advective formulation of the GM parameterization implies that the GM parameterization models the eddy tracer flux divergence as $\nabla \cdot (\mathbf{u} \tau) = \mathbf{u} \cdot \nabla \tau$ where $\mathbf{u} \tau$ is the GM advective flux. Griffies (1998) demonstrated that one can alternatively parameterize the eddy tracer flux divergence by the divergence of ‘skew’ flux that equals the advective flux $\mathbf{u} \tau$ plus a non-divergent gauge; the GM skew-flux has the form

$$\mathbf{F}_{\text{skew}} = \Psi \times \nabla \tau. \tag{10}$$

For a standard isotropic configuration of GM the horizontal part of the skew flux of density has the form

$$F_{\text{skew}}^x = -\kappa \partial_x \rho, \quad F_{\text{skew}}^y = -\kappa \partial_y \rho. \tag{11}$$

Grooms (2016) showed that stochastic GM parameterizations could be built by producing stochastic parameterizations of the skew flux of density. Denoting a stochastic parameterization of the horizontal density flux by $\mathbf{F}_{\text{stoch}}$, one can invert Eq. (10) to obtain the corresponding stochastic parameterization of the GM vector streamfunction:

$$\Psi_{\text{stoch}}^y = \frac{F_{\text{stoch}}^x}{\partial_z \rho}, \quad \Psi_{\text{stoch}}^x = -\frac{F_{\text{stoch}}^y}{\partial_z \rho}. \tag{12}$$

The stochastic GM model developed in Section 3, as well as the SPP-GM approach of Juricke et al. (2017, 2018), fits within this framework; in

these parameterizations the stochastic parameterization of the skew flux takes the form of Eq. (11) with a stochastic κ . As shown in Section 2, a stochastic GM model of this form does not fit the diagnosed eddy flux divergences well.

A more successful approach to building a stochastic GM parameterization that accurately models the eddy flux divergence might need to break out of this paradigm of randomizing κ . For example, one might find more success in both fitting the true eddy flux divergence and in coarse model simulations by using the following more general form for the eddy skew flux of density

$$F_{\text{stoch}}^x = -\kappa \partial_x \rho + \chi^x, \quad F_{\text{stoch}}^y = -\kappa \partial_y \rho + \chi^y \quad (13)$$

where χ^x and χ^y are components of a random eddy flux. This type of parameterization has both a multiplicative component associated with κ and an additive component associated with the random eddy flux. Such a parameterization would, for example, easily generate an eddy flux in regions with small isopycnal slope, whereas the multiplicative approach requires large κ to generate eddy fluxes in regions of small isopycnal slopes. The additive terms are qualitatively similar to the nonlocal transport term in the KPP parameterization (Large et al., 1994). Williams et al. (2016) developed an additive stochastic parameterization in the temperature tendency of an ocean GCM; the additive terms above would have a similar effect, but would guarantee that the parameterization remains adiabatic by remaining within the GM framework. The development of the details of such a parameterization is beyond the scope of this article, though preliminary strides in that direction were made by Grooms (2016).

6. Conclusions

This paper addresses the question of parameterizing the tracer flux divergence associated with unresolved mesoscale eddies in non-eddy ocean models. We begin by diagnosing the parameter κ in an isotropic Gent-McWilliams parameterization (Gent and McWilliams, 1990; Gent et al., 1995) by comparing to eddy-resolving simulations. Rather than estimate κ by comparing the diagnosed fluxes to the parameterized fluxes (cf. Eden et al., 2007; Bachman and Fox-Kemper, 2013), we use a least-squares approach to fit the parameterized eddy flux divergence to the diagnosed eddy flux divergence (cf. Mak et al., 2016). The diagnosed κ displays huge variability with values as large as $\pm 30,000 \text{ m}^2/\text{s}$. Crucially, the optimal value of κ leads to a parameterization that only accounts for less than 50% of the diagnosed eddy flux divergence. A stochastic model is developed that carefully models the spatial and temporal structure of the diagnosed κ . Coarse model simulations are then run using the new stochastic model for κ , constant κ , and constant κ perturbed by simple stochastic noise. The coarse model simulations with constant κ and with the simple stochastic model are far more accurate than simulations with the new stochastic model. We argue that the poor agreement of the optimal diagnosed κ with the diagnosed eddy flux divergence explains the failure of the new stochastic model to produce accurate coarse model simulations, and suggest that future GM parameterizations might be more successful, both in terms of modeling the true eddy flux divergence and in terms of producing realistic coarse model simulations, if they include an additive component to the skew flux rather than relying solely on variations in κ . Of course the κ diagnosed here was constrained to be depth-independent, and significant improvements might still be obtained by allowing depth dependence or anisotropy (Smith and Gent, 2004; Abernathy et al., 2013).

These results demonstrate that accurate coarse model simulations can be produced using parameterizations that are not accurate models of the true eddy flux divergence. It may therefore be possible to improve coarse model simulations with a stochastically perturbed GM parameter κ by tuning the structure of the perturbations of κ , as in Juricke et al. (2017, 2018). But these improvements to the coarse model simulations will not come as a result of making the modeled eddy flux

divergence a more accurate model of the true eddy flux divergence.

Acknowledgments

The authors gratefully acknowledge discussion of SPPT and SPP with J. Berner and H. M. Christensen, and the efforts of two reviewers. This research was supported by the NSF OCE grant 1736708. This work utilized the RMACC Summit supercomputer, which is supported by the National Science Foundation (awards ACI-1532235 and ACI-1532236), the University of Colorado Boulder, and Colorado State University. The Summit supercomputer is a joint effort of the University of Colorado Boulder and Colorado State University.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.ocemod.2018.10.009](https://doi.org/10.1016/j.ocemod.2018.10.009)

References

- Abernathy, R., Ferreira, D., Klocker, A., 2013. Diagnostics of isopycnal mixing in a circumpolar channel. *Ocean Model.* 72, 1–16.
- Andrejczuk, M., Cooper, F., Juricke, S., Palmer, T., Weisheimer, A., Zanna, L., 2016. Oceanic stochastic parameterizations in a seasonal forecast system. *Mon. Weather Rev.* 144 (5), 1867–1875.
- Arzel, O., Huck, T., Colin de Verdière, A., 2006. The different nature of the interdecadal variability of the thermohaline circulation under mixed and flux boundary conditions. *J. Phys. Ocean.* 36 (9), 1703–1718.
- Bachman, S., Fox-Kemper, B., 2013. Eddy parameterization challenge suite I: eady spindown. *Ocean Model.* 64, 12–28.
- Berloff, P.S., 2005. Random-forcing model of the mesoscale oceanic eddies. *J. Fluid Mech.* 529, 71–95.
- Berner, J., Achatz, U., Batté, L., Bengtsson, L., de la Cámara, A., Christensen, H.M., Colangeli, M., Coleman, D.R.B., Crommelin, D., Dolaptchiev, S.I., Franzke, C.L.E., Friederichs, P., Imkeller, P., Järvinen, H., Juricke, S., Kitsios, V., Lott, F., Lucarini, V., Mahajan, S., Palmer, T.N., Penland, C., Sakradzija, M., von Storch, J.-S., Weisheimer, A., Weniger, M., Williams, P.D., Yano, J.-I., 2017. Stochastic parameterization: toward a new view of weather and climate models. *B. Am. Meteorol. Soc.* 98 (3), 565–588. <https://doi.org/10.1175/BAMS-D-15-00268.1>.
- Brankart, J.-M., 2013. Impact of uncertainties in the horizontal density gradient upon low resolution global ocean modelling. *Ocean Model.* 66, 64–76.
- Brankart, J.-M., Candille, G., Garnier, F., Calone, C., Melet, A., Bouttier, P.-A., Brasseur, P., Verron, J., 2015. A generic approach to explicit simulation of uncertainty in the NEMO ocean model. *Geosci. Model Dev.* 8 (5), 1285–1297.
- Buizza, R., Milleer, M., Palmer, T., 1999. Stochastic representation of model uncertainties in the ECMWF ensemble prediction system. *Q. J. R. Meteor. Soc.* 125 (560), 2887–2908.
- Christensen, H., Lock, S.-J., Moroz, I., Palmer, T., 2017. Introducing independent patterns into the stochastically perturbed parameterization tendencies (SPPT) scheme. *Q. J. R. Meteor. Soc.* 143 (706), 2168–2181.
- Davini, P., Hardenberg, J.v., Corti, S., Christensen, H.M., Juricke, S., Subramanian, A., Watson, P.A., Weisheimer, A., Palmer, T.N., 2017. Climate sphinx: evaluating the impact of resolution and stochastic physics parameterizations in the ec-earth global climate model. *Geosci. Model Dev.* 10 (3), 1383–1402.
- Durran, D.R., 2010. *Numerical Methods for fluid Dynamics with Applications to Geophysics*, second ed. Springer Dordrecht.
- Eden, C., Greatbatch, R.J., Willebrand, J., 2007. A diagnosis of thickness fluxes in an eddy-resolving model. *J. Phys. Ocean.* 37 (3), 727–742.
- Ferrari, I., Griffies, S.M., Nurser, A.G., Vallis, G.K., 2010. A boundary-value problem for the parameterized mesoscale eddy transport. *Ocean Model.* 32 (3), 143–156.
- Friedman, J., Hastie, T., Tibshirani, R., 2008. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 9 (3), 432–441.
- Gent, P.R., McWilliams, J.C., 1990. Isopycnal mixing in ocean circulation models. *J. Phys. Ocean.* 20 (1), 150–155.
- Gent, P.R., Willebrand, J., McDougall, T.J., McWilliams, J.C., 1995. Parameterizing eddy-induced tracer transports in ocean circulation models. *J. Phys. Ocean.* 25 (4), 463–474.
- Griffies, S.M., 1998. The Gent-McWilliams skew flux. *J. Phys. Ocean.* 28 (5), 831–841.
- Grooms, I., 2016. A Gaussian-product stochastic Gent-McWilliams parameterization. *Ocean Model.* 106, 27–43.
- Grooms, I., Julien, K., Fox-Kemper, B., 2011. On the interactions between planetary geostrophy and mesoscale eddies. *Dyn. Atmos. Oceans* 51. <https://doi.org/10.1016/j.dynatmoce.2011.02.002>.
- Grooms, I., Smith, K.S., Majda, A.J., 2012. Multiscale models for synoptic-mesoscale interactions in the oceans. *Dyn. Atmos. Oceans* 58.
- Henning, C.C., Vallis, G.K., 2004. The effects of mesoscale eddies on the main subtropical thermocline. *J. Phys. Ocean.* 34 (11), 2428–2443.
- Juricke, S., MacLeod, D., Weisheimer, A., Zanna, L., Palmer, T., 2018. Seasonal to annual ocean forecasting skill and the role of model and observational uncertainty. *Q. J. R. Meteor. Soc.* <https://doi.org/10.1002/qj.3394>. Accepted manuscript

- Juricke, S., Palmer, T.N., Zanna, L., 2017. Stochastic subgrid-scale ocean mixing: impacts on low-frequency variability. *J. Climate* 30 (13), 4997–5019.
- Large, W.G., McWilliams, J.C., Doney, S.C., 1994. Oceanic vertical mixing: a review and a model with a nonlocal boundary layer parameterization. *Rev. Geophys.* 32 (4), 363–403.
- Leutbecher, M., Lock, S.-J., Ollinaho, P., Lang, S., Balsamo, G., Bechtold, P., Bonavita, M., Christensen, H., Diamantakis, M., Dutra, E., English, S., Fisher, M., Forbes, R., Goddard, J., Haiden, T., Hogan, R., Juricke, S., Lawrence, H., MacLeod, D., Magnusson, L., Malardel, S., Massart, S., Sandu, I., Smolarkiewicz, P., Subramanian, A., Vitart, F., Wedi, N., Weisheimer, A., 2017. Stochastic representations of model uncertainties at ECMWF: state of the art and future vision. *Q. J. R. Meteor. Soc.* 143 (707), 2315–2339.
- Lorenz, E., 1975. Climatic predictability. *WMO The Phys. Basis of Climate and Climate Modelling* p 132–136(SEE N 76–19675 10–47).
- Mak, J., Maddison, J.R., Marshall, D.P., 2016. A new gauge-invariant method for diagnosing eddy diffusivities. *Ocean Model.* 104, 252–268.
- 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.-Oceans* 102 (C3), 5753–5766.
- McDougall, T.J., 1987. Neutral surfaces. *J. Phys. Ocean.* 17 (11), 1950–1964.
- Redi, M.H., 1982. Oceanic isopycnal mixing by coordinate rotation. *J. Phys. Ocean.* 12 (10), 1154–1158.
- Rue, H., Held, L., 2005. *Gaussian Markov Random Fields Theory and Applications*. Chapman & Hall/CRC.
- Smith, K.S., Marshall, J., 2009. Evidence for enhanced eddy mixing at middepth in the Southern Ocean. *J. Phys. Ocean.* 39 (1), 50–69.
- Smith, R.D., Gent, P.R., 2004. Anisotropic Gent–McWilliams parameterization for ocean models. *J. Phys. Ocean.* 34 (11).
- Williams, P., Howe, N., Gregory, J., Smith, R., Joshi, M., 2016. Improved climate simulations through a stochastic parameterization of ocean eddies. *J. Climate* 29, 8763–8781.