

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

Ian Grooms^{a,*}; William Kleiber^a

^a*Department of Applied Mathematics, University of Colorado, Boulder, CO 80309.*

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.

Keywords: Stochastic parameterization, Mesoscale parameterization, Gent-McWilliams

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

*Corresponding author

Email address: ian.grooms@colorado.edu(Ian Grooms)

21 are not completely determined by the resolved-scale variables; the large variability in tracer flux that remains even
22 after averaging over square cells of width approximately 85 km was demonstrated by Grooms (2016). The turbulent
23 mesoscale dynamics generate chaotic fluxes that are spatially and temporally correlated, and are dependent on, but
24 not entirely determined by, the large-scale flow.

25 The fact that the subgrid-scale tracer fluxes are chaotic implies that parameterizations should be stochastic. The
26 idea of using random numbers in numerical climate models goes back at least to Lorenz (1975). A major develop-
27 ment in the practical use of stochastic parameterizations in atmospheric and oceanic GCMs came when a stochastic
28 parameterization was implemented in the operational ECMWF Ensemble Prediction System (Buizza et al., 1999).
29 Since then stochastic parameterizations have been adopted by operational centers almost exclusively for use in atmo-
30 spheric ensemble prediction systems because they increase ensemble spread, leading to better calibrated probabilistic
31 forecasts (Berner et al., 2017; Leutbecher et al., 2017).

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

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

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

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

66 2. Diagnosing the GM parameter

67 2.1. Gent-McWilliams Background

68 Before detailing the method used to diagnose the GM parameter we begin with a brief review of the deterministic
69 GM framework (Gent and McWilliams, 1990; Gent et al., 1995). In the GM framework the divergence of the subgrid-

70 scale tracer flux is parameterized by a bolus velocity advection

$$\nabla \cdot (\bar{\mathbf{u}}^0 \bar{\tau}^0) = \bar{\mathbf{u}}^{\dagger} \cdot \nabla \bar{\tau}. \quad (1)$$

71 Throughout this section resolved-scale variables have no primes and subgrid-scale variables are denoted by a prime
 72 0 . The overbar ($\bar{\cdot}$) denotes a spatial average to the coarse grid scale. In the above expression $\bar{\tau}$ denotes a tracer
 73 concentration, and the bolus velocity is

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

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

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

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

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

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

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

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

$$\begin{aligned} c^2 \partial_z^2 + \frac{g \partial_z \rho}{\rho_0} \Psi^y &= -\kappa \frac{g}{\rho_0} \partial_x \rho, \\ \# \end{aligned} \quad (4)$$

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

88 with boundary conditions $\Psi^x = \Psi^y = 0$. When c^2 is chosen appropriately the streamfunction approximates (3), and the
 89 solutions smoothly transition through unstratified and convectively unstable regions. In practice the density gradient
 90 $\partial_z \rho$ that appears in these equations is truncated to negative values; here the value of $\bar{N}(z) = -g \partial_z \rho / \rho_0$ is truncated for
 91 values less than 10^{-24} s^{-2} . Ferrari et al. (2010) chose the following specification for c

$$c = \max \left(c_{\min}, \frac{\pi^{-1}}{0} \int_0^H N(z) dz \right) \quad (6)$$

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

93 2.2. Eddy-resolving model configuration

94 The eddy-resolving simulation uses the MITgcm (Marshall et al., 1997) with a linear equation of state based solely
 95 on temperature (no salinity); the expansion coefficient for temperature is $\alpha_T = 1.5 \times 10^{-4} \text{ K}^{-1}$. The model domain is
 96 3,200 km square and 4 km deep. A Cartesian tangent plane approximation is used, with the Coriolis parameter taking
 97 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
 98 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
 99 boundary conditions are no-slip at the sides and bottom, with a rigid lid. The vertical diffusivity is $\kappa_v = 10^{-4} \text{ m}^2/\text{s}$
 100 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
 101 to $\kappa_h = 50 \text{ m}^2/\text{s}$.

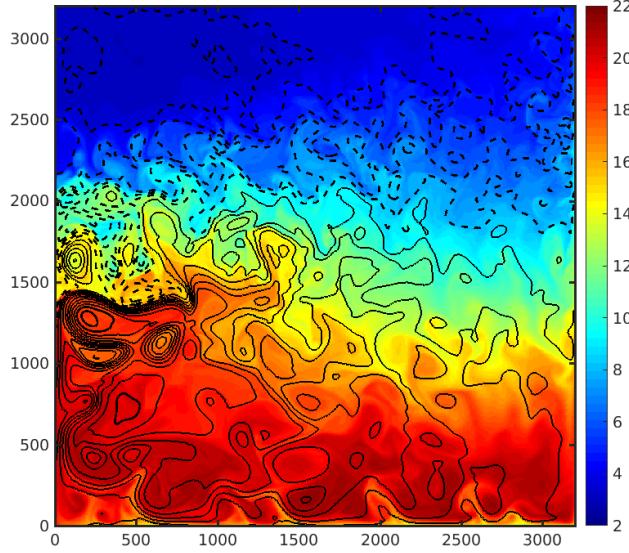


Figure 1: A snapshot of sea surface temperature (color) and sea surface height anomaly (contours) from the high-resolution reference simulation. The contour interval is 10 cm; negative contours are dashed and positive contours are solid.

102 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
 103 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
 104 cosine varying from 22 ° C at the southern boundary to 2 ° C at the northern boundary. The heat flux is converted
 105 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
 106 boundaries.

107 The initial condition was taken from an equilibrated solution of the ‘deterministic GM’ model described in Grooms
 108 (2016), which has grid size of 40 km and 61 vertical levels with depths varying linearly from 25 m at the surface to
 109 106 m at depth. This low-resolution equilibrated state was interpolated to the eddy-resolving resolution of 8 km.
 110 The circulation exhibits a large subtropical gyre and a weaker subpolar gyre, with strong eddies along the western
 111 boundary and into the interior, very similar to the simulations of Henning and Vallis (2004). A snapshot of the sea
 112 surface temperature and the sea surface height anomaly in the statistically steady state are shown in Fig. 1.

113 2.3. Diagnosing the GM parameter

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

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

116 where the notation ‘ $\arg \min$ ’ indicates that κ is the value that minimizes this expression, the integral is over the volume,
 117 and κ is allowed to vary horizontally but not vertically. The restriction that κ is depth-independent is perhaps overly
 118 severe (cf. Abernathey et al., 2013), but it significantly reduces the size and computational cost of the optimization
 119 problem; limitations of the model are discussed further below. Minimizing the error in the flux divergence avoids
 120 problems related to gauge freedom that would arise if one instead minimized the error in the flux itself (Mak et al.,
 121 2016). Because the bolus velocity $\bar{\mathbf{u}}^\dagger$ depends linearly on κ , this is a linear least squares problem. Solving the least-
 122 squares problem requires computing the true eddy density flux divergence, which requires defining a low-pass spatial
 123 filter $\bar{\cdot}$. This was implemented as a moving average with a Gaussian kernel

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

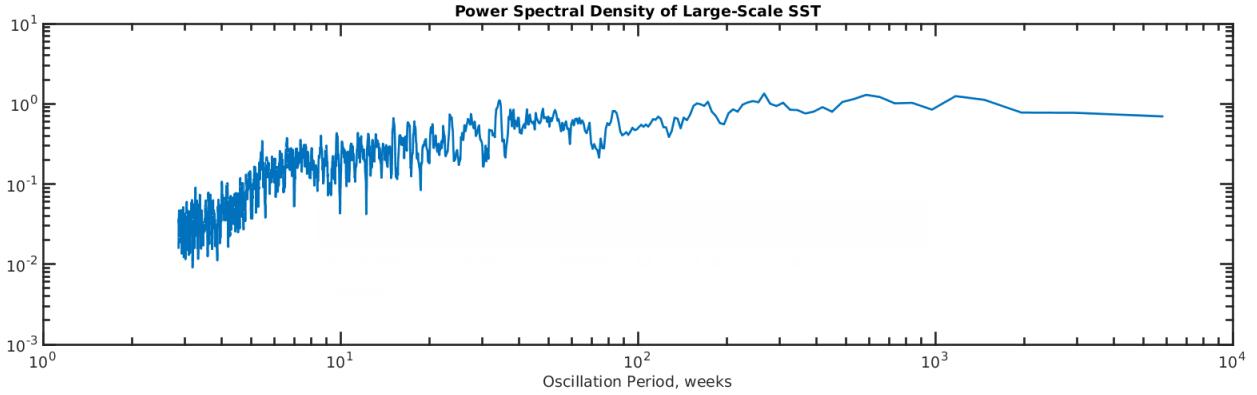


Figure 2: Power spectral density of the large-scale part of sea surface temperature (SST) at the point marked \times (a) in Fig. 5: 380 km east of the western boundary and 1580 km north of the southern boundary.

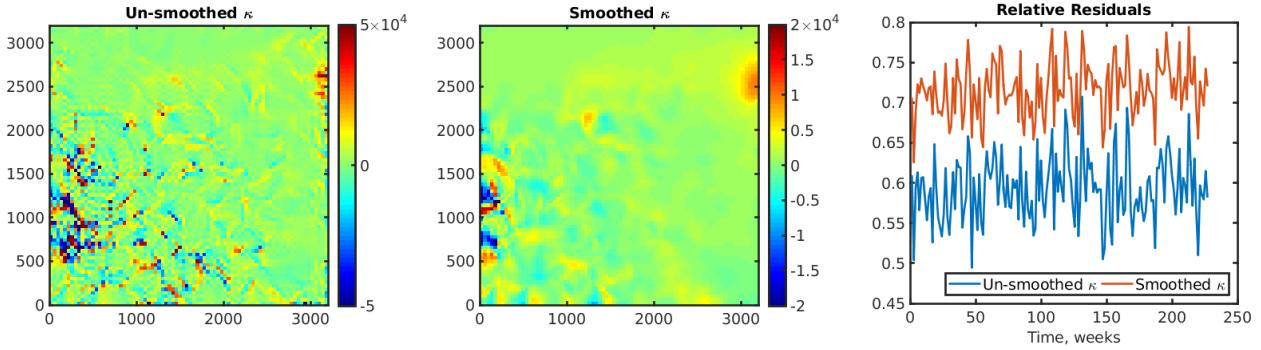


Figure 3: Left: A snapshot of κ computed using Eq. 7; units are m^2/s ; axis units are kilometers. Center: A snapshot of κ computed using Eq. 8 at the same time as the left panel; note the different color scale compared to the left panel. Right: Relative residuals as a function of time for the κ computed using Eq. 7 (blue) and Eq. 8 (orange).

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

129 A kernel length scale of 60 km is arguably close enough to typical eddy length scales that it is possible that it fails
 130 to make a clear separation between the eddies and the scales resolvable by the coarse model (though one could argue
 131 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
 132 analysis with a variety of filter widths involves prohibitive computational expense, so as an alternate way to test this,
 133 the power spectral density of the filtered, large-scale part of temperature was computed over a period from the 98th
 134 year of the eddy-resolving simulation to the 208th year, with data taken every tenth day. The power spectral density of
 135 sea surface temperature at a point 380 km west of the eastern boundary and 1580 km north of the southern boundary
 136 (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
 137 is clearly associated with long time scales on the order of 10 weeks or longer; this slow variability is characteristic of
 138 location and depth where the power spectral density was computed. Eddy variability dominates the total variability
 139 and is expected to occur on timescales of at most a few weeks, so the fact that the temporal variability of the large-scale
 140 part of temperature occurs on much longer time scales implies that the low-pass spatial filter is effective at removing
 141 eddy length scales.

142 The eddy-resolving simulation was allowed to run for 200 years after initializing from the low-resolution model
 143 before taking the data used to compute κ . The fields in Eq. (7) were computed on the coarse-model grid of 40 km, and

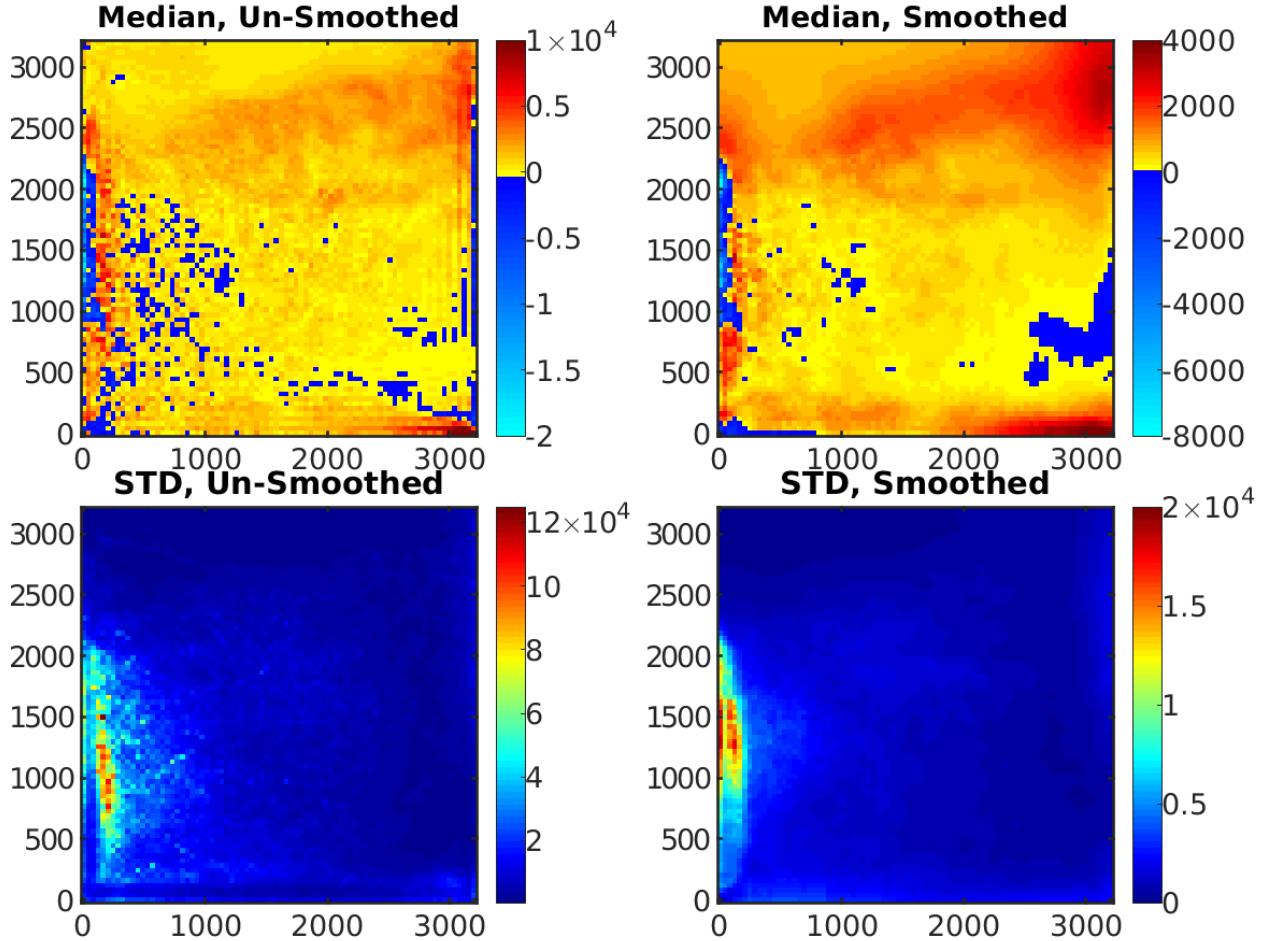


Figure 4: Top Row: Median value of κ . Bottom Row: Standard Deviation (STD) of κ . Left Column: κ obtained from Eq. (7), i.e. without a smoothing penalty term. Right Column: κ obtained from Eq. (8), i.e. with a smoothing penalty term. Units of median and standard deviation are m^2/s . Axis units are kilometers. The color map in the top column has been chosen to make negative values more obvious.

144 the optimal value of κ was computed by solving the normal equations. The results are extremely rough; the left panel
 145 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
 146 left panel of Fig. 3 result from locations where the GM model predicts a near-zero flux-divergence (e.g. in regions of
 147 small isopycnal slope) while the true flux-divergence is large, and κ becomes large to compensate. Regardless of their
 148 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
 149 instability in a low-resolution model.

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

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

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

156 Negative values of κ appear in the smoothed results almost everywhere, excepting only in the extreme north
 157 end of the domain, and negative values occur on average in various parts of the domain. There are mathematical
 158 difficulties with partial differential equations (PDEs) with negative viscosity, and one might want to impose a non-

159 negativity condition on the least-squares problem to avoid this. Numerically solving a PDE with negative viscosity
160 is problematic because the computed solution depends strongly on the discretization and need not converge as the
161 grid is refined (Durran, 2010). On the other hand the negative viscosity exists here because a specific, fixed coarse
162 resolution and associated spatial filter was chosen for the diagnosis of κ and it would not make sense to use that same
163 negative value of κ with increased resolution and decreased filter scale. There are no mathematical problems with the
164 discretized model with negative viscosity at fixed coarse resolution, so negative values are retained here to make the
165 GM parameterization as accurate as possible.

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

178 3. A random-field model for κ

179 This section describes the construction of a random-field model for κ based on the diagnosed values described in
180 the preceeding section. Figure 4 shows the median (upper row) and standard deviation (lower row) of the diagnosed
181 values of κ using the un-smoothed results of Eq. (7) (left column) and the smoothed results from Eq. (8) (right column).
182 There is clearly large variability in κ , particularly for the un-smoothed results, which is why the median is used as
183 a central estimator for the field instead of the mean. The median is much less than one standard deviation from the
184 mean, and is robust to outliers; also the median is somewhat smoother spatially than the mean (not shown). The
185 grainy quality of the median and standard deviation of the un-smoothed κ suggest that these statistics are hard to
186 estimate accurately from the time series of 159 points. This consideration, together with the extreme roughness of
187 the un-smoothed fields shown in the left panel of Fig. 3 motivates us to focus henceforth on the smoothed version of
188 κ . In both the smoothed and unsmoothed versions of κ the increased variability shown in the lower row of Fig. 3 is
189 associated with and presumably driven by the increased eddy energy along the western boundary.

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

193 3.1. Temporal structure

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

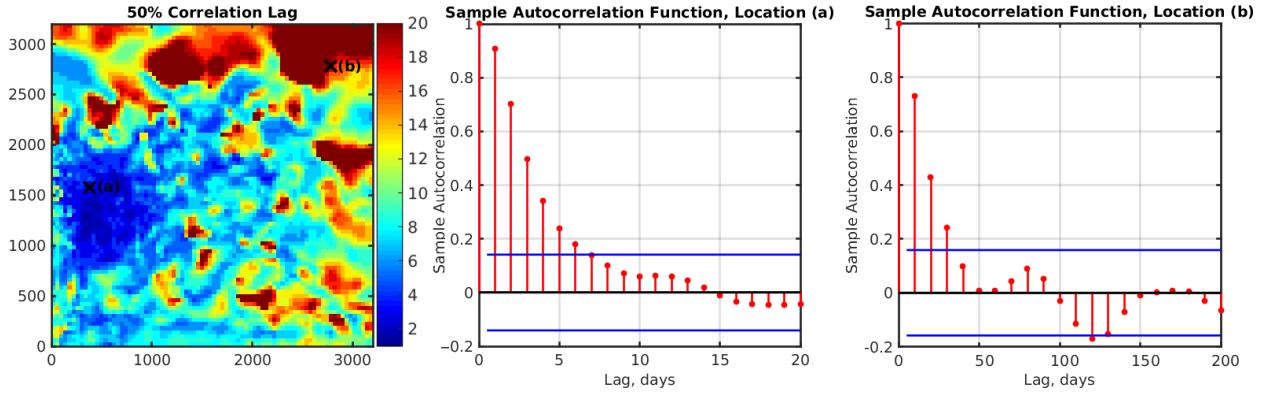


Figure 5: Left: The longest time lag, in days, for which the autocorrelation of \mathbf{K} remains above 50%. Axis units are kilometers. Center: The empirical autocorrelation function of \mathbf{K} at the point marked (a) in the left panel. The blue line shows the level at which the autocorrelation is not statistically significantly different from zero. Right: The empirical autocorrelation function of \mathbf{K} at the point marked (b) in the left panel. Note the range of x-axis values on the sample autocorrelation functions differ.

208 3.2. Spatial structure

209 To investigate the spatial correlation structure of \mathbf{K} (in this section \mathbf{K} should be understood to be the smoothed
 210 version computed from Eq. (8)) we first subtract the pointwise time average and divide by the standard deviation to
 211 achieve a centered field with unit variance. This is called the ‘centered, scaled’ version of \mathbf{K} , and it is the subject of
 212 analysis in this section.

213 The spatial structure is estimated using the original time series of length 159 with 10-day spacing, plus every tenth
 214 day of the time series of length 201 with 1-day spacing. This yields a time series of length 180 with 10 day spacing.
 215 As noted in the previous section, \mathbf{K} displays low-frequency variability in the northeast corner of the domain that is
 216 probably not directly the result of eddy variability. To avoid modeling this low-frequency variability a high-pass time
 217 filter is applied to \mathbf{K} before examining the spatial structure. This high-pass filter simply consists in removing a moving
 218 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
 219 length 178. Since the time series has a 10 day spacing and most of the variability in \mathbf{K} has time scales on the order of
 220 10 days or less (left panel of Fig. 5), the high-pass filter only removes time scales significantly longer than 10 days
 221 and leaves the primary eddy variability intact.

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

232 The ‘graphical lasso’ (Friedman et al., 2008) is a method to estimate the precision matrix of a set of variables
 233 (the precision matrix is the inverse of the covariance matrix). It generates a sparse approximation to the precision
 234 matrix by minimizing a cost function that is the sum of a rank-deficient L^2 term measuring the mismatch between the
 235 precision matrix and the data and an L^1 penalty term that regularizes the problem and enhances sparsity of the result.
 236 In the statistical literature, it is well known that the precision matrix (inverse covariance matrix) encodes a graphical
 237 structure on the random field, and is typically expected to be a sparse matrix that relays information on a conditional
 238 dependence structure (Rue and Held, 2005). A penalty parameter $\lambda \geq 0$ determines the sparsity of the estimate: large
 239 λ leads to a very sparse but less-accurate approximation while small λ leads to a more dense estimate ($\lambda = 0$ results in
 240 the un-penalized empirical covariance estimate). In particular, we use the graphical lasso to estimate a random field
 241 model as it imposes no assumptions of homogeneity.

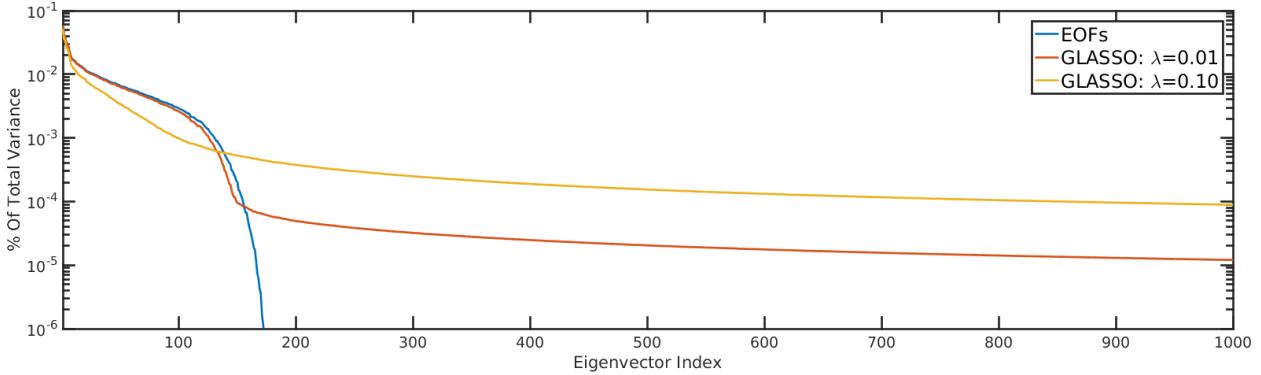


Figure 6: Percent of total variance explained by (i) EOFs (blue), (ii) eigenvectors of the graphical lasso (GLASSO) precision matrix at $\lambda = 0.01$ (orange) and $\lambda = 0.10$ (yellow).

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

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

262 3.3. A stochastic model for \mathbf{K}

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

265 We begin with the spatial structure. The graphical lasso yields a spatial precision matrix $\mathbf{P} = \mathbf{C}^{-1}$ where \mathbf{C} is
 266 an approximation to the spatial covariance matrix. Let $\mathbf{L}\mathbf{L}^T = \mathbf{P}$ be the Cholesky decomposition of \mathbf{P} and let \mathbf{x} be
 267 a vector of zero-mean, unit-variance independent random variables. Then $\mathbf{z} = \mathbf{L}^{-T}\mathbf{x}$ is a vector with covariance
 268 matrix \mathbf{C} . This shows how to construct a random field with a spatial correlation structure that mimics \mathbf{K} : compute the
 269 Cholesky factorization of the precision matrix computed by the graphical lasso, and use it to give spatial correlation
 270 to a field of independent random variables.

271 A simple approach to modeling the temporal correlation structure is to construct an AR-1 (first order autoregres-
 272 sive) process at each spatial location on the coarse grid. AR-1 processes are unable to model negative correlations, but
 273 the autocorrelation functions diagnosed across the domain are all similar to those shown in the center and left panels
 274 of Fig. 5, i.e. they are monotonically decreasing until they become statistically insignificant. In a more general setting
 275 AR-2 or higher processes could be used. The decorrelation time of this process can be tuned to match the observed
 276 decorrelation time of \mathbf{K} . It is not straightforward to combine the two foregoing ideas (Cholesky plus AR-1 processes)

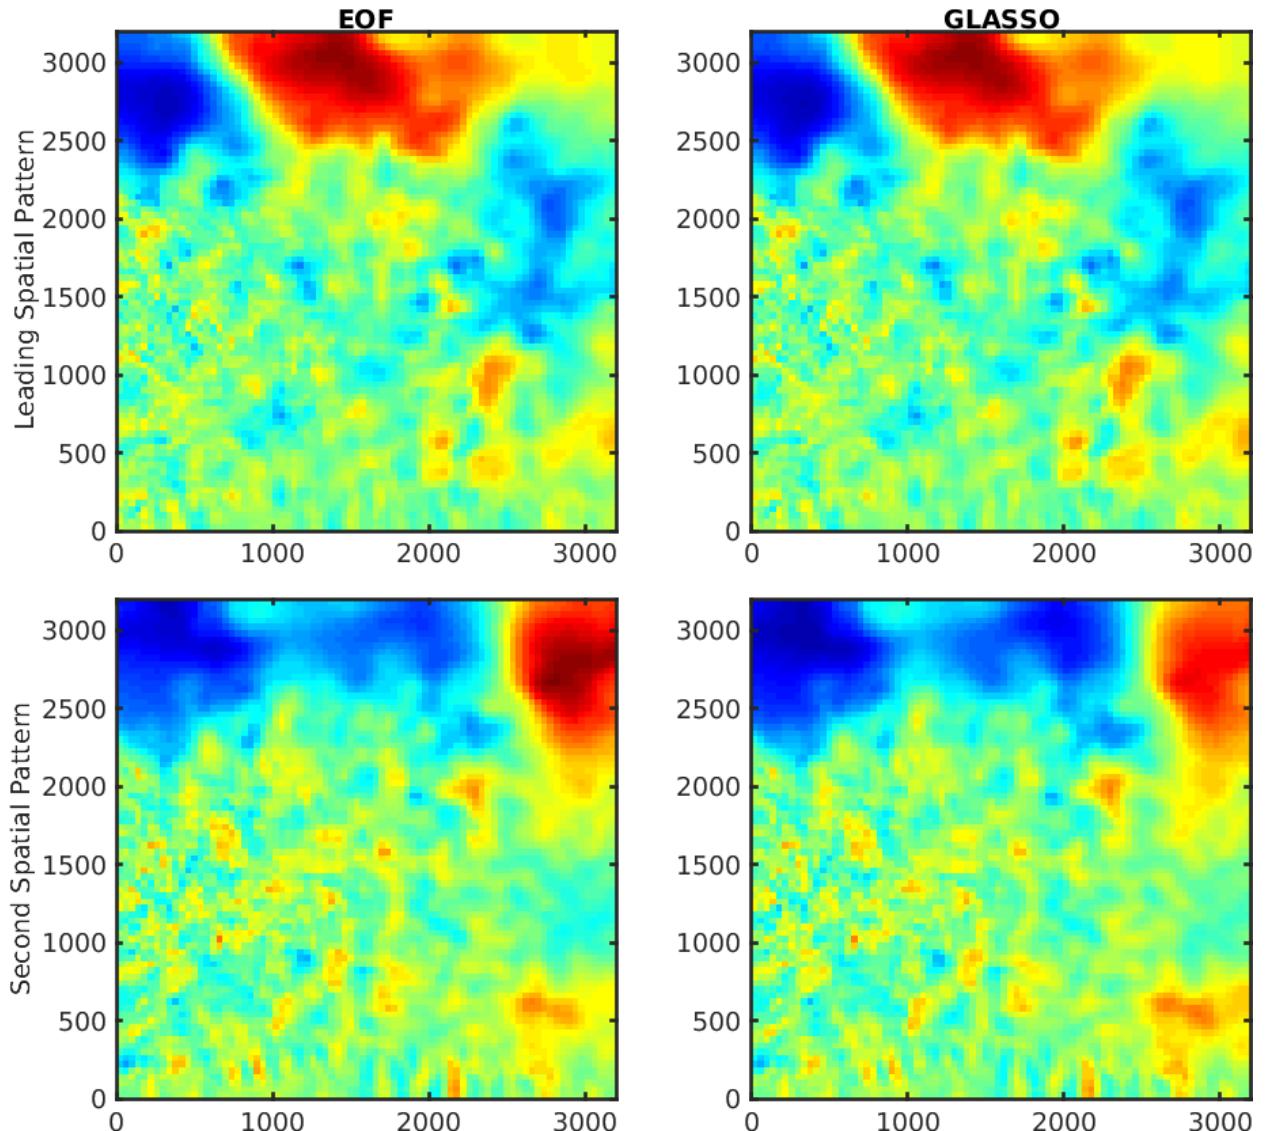


Figure 7: Left column: Empirical orthogonal functions (EOFs). Right column: Eigenvectors of the graphical lasso (GLASSO) precision matrix. Upper row: Leading spatial pattern of variability. Lower row: Secondary spatial pattern of variability. Axis units are kilometers. The spatial patterns are dimensionless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

277 into a model that accurately reflects the full spatiotemporal correlation structure of the true field κ . Our model takes
 278 the following form: At each point of the coarse grid the centered, scaled version of κ is modeled as an AR-1 process
 279 in the form

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

280 where χ_n are unit-variance, zero-mean random variables that are independent from one time to the next (the subscript
 281 n indicates time), and r governs the local decorrelation time. Specifically, r is chosen so that the AR-1 process reaches
 282 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
 283 times shown in Fig. 5 are in one-day increments which leads to sharp jumps in the decorrelation time. The field shown
 284 in Fig. 5 is therefore smoothed by applying a moving average with weights [0 .25 0.5 0.25] in each direction before
 285 being used to set r for the AR-1 processes.

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

290 4. Coarse Model Simulations

291 Coarse model simulations are configured as in Grooms (2016) and using parameter values from section 2.2,
 292 with the following exceptions. The horizontal grid size is 40 km and diffusion of temperature is accomplished by
 293 a third-order upwind scheme rather than by an explicit diffusion. The horizontal viscosity is Laplacian rather than
 294 biharmonic; 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
 295 barotropic part. Different values are used in an effort to make the barotropic and baroclinic viscous boundary layers
 296 both have widths on the order of one grid cell. The time-mean state of the large-scale part of the eddy-resolving
 297 simulation described in section 2.2 was used as an initial condition for the coarse model simulations.

298 We run simulations with a constant value of $\kappa = 6000 \text{ m}^2/\text{s}$ for comparison to the stochastic GM scheme described
 299 in the preceding section; this simulation is called the ‘deterministic GM’ simulation. In addition, we run simulations
 300 where the GM parameter κ is perturbed by multiplying by $1 + e$ where e is a unit-variance, zero-mean Gaussian
 301 random field. The temporal structure of e is an AR-1 process with decorrelation time of 10 days, and the spatial
 302 structure results from applying a moving average with weights [0 .25 0.5 0.25] in each direction to an uncorrelated
 303 field (see Grooms (2016) for more details). This simulation is called the ‘SPP-GM’ simulation. The simulation using
 304 the model described in section 3 is simply called the ‘stochastic GM’ simulation.

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

309 Figure 8 shows the time-mean sea surface temperature (SST) from the eddy-resolving simulation (top left), from
 310 the deterministic GM simulation (top right), from the SPP-GM simulation (lower left), and from the stochastic GM
 311 simulation (lower right). The deterministic and SPP-GM schemes provide reasonably accurate predictions for the
 312 mean SST, while the tuned stochastic GM simulation is badly incorrect. The true SST has a steady decrease of SST
 313 with latitude, with very little zonal variability. The stochastic GM simulation has a plume of warm water originating
 314 near the southwest corner of the domain and extending north and east up through the center of the domain. The mean
 315 temperature structure in the SPP-GM simulation is similarly inaccurate throughout the main thermocline (not shown).
 316 In view of the very poor accuracy of the optimized stochastic GM parameterization that is the primary focus of the
 317 investigation, there is little point in proceeding to test whether the scheme generates accurate variability, or in tuning
 318 other parameters of the deterministic and SPP-GM simulations.

319 5. Discussion

320 The ultimate goal of this research is to formulate a GM parameterization that produces an accurate coarse model
 321 simulation, and there are two routes to achieve this. One route, not taken here, involves tuning the parameterization by
 322 running repeated coarse model simulations and comparing the results to known reference data; this process is costly

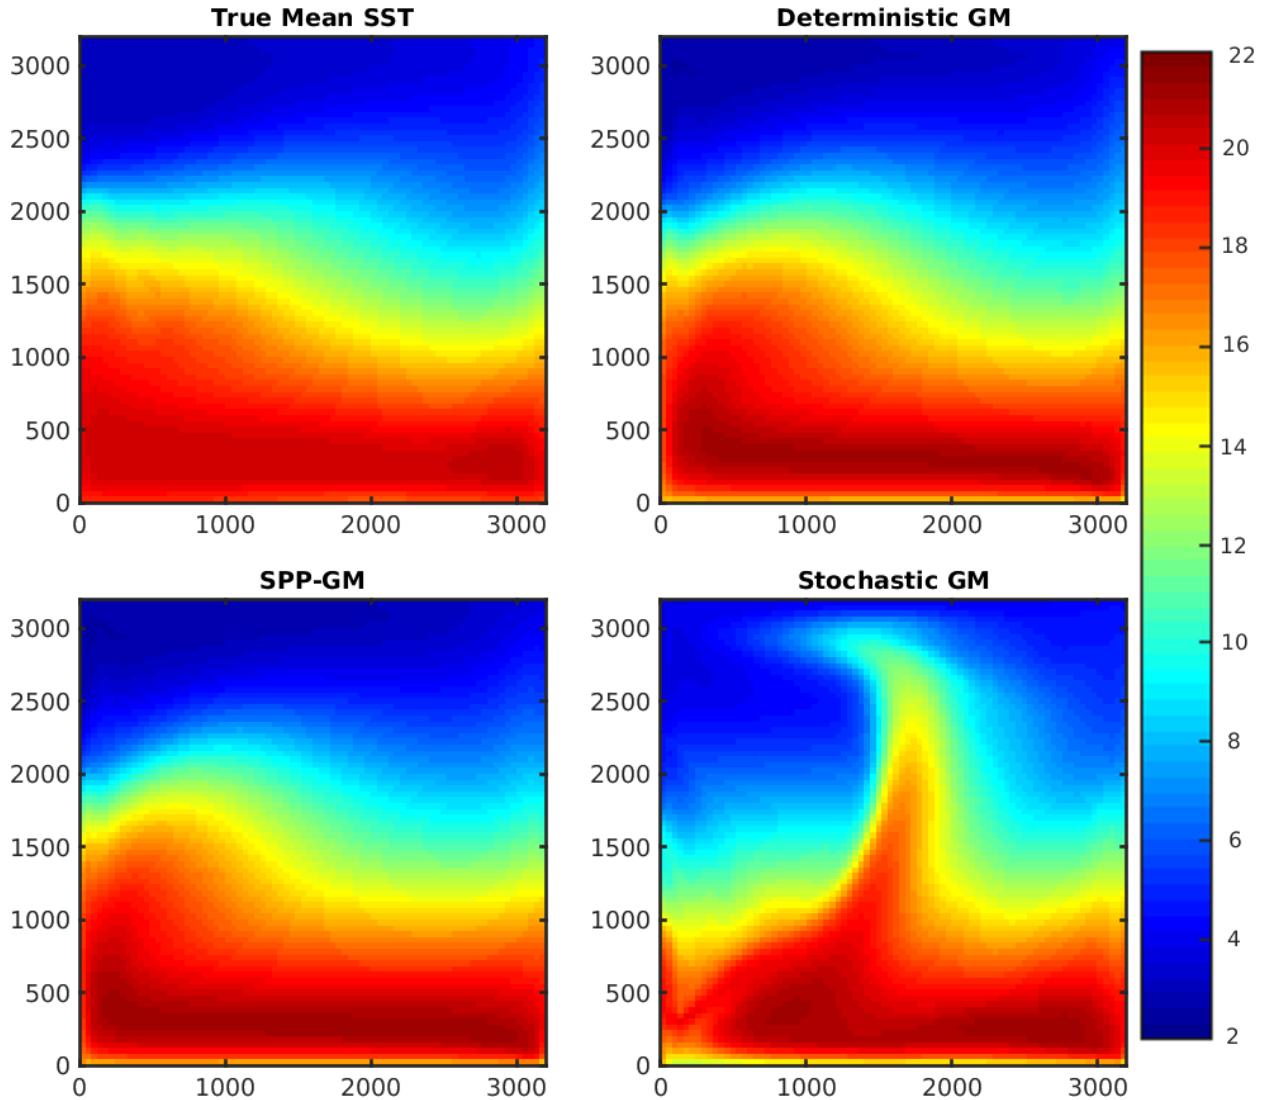


Figure 8: Time-mean sea surface temperature for the eddy-resolving simulation (upper left), the deterministic GM simulation (upper right), the SPP-GM simulation (lower left) and the stochastic GM simulation (lower right). Temperature is in degrees Celsius and axis units are kilometers.

323 because of the infinite number of possible ways to tune the parameterization. The other route, taken here, involves
 324 studying the process or quantity that is being parameterized and attempting to build an accurate model thereof. This
 325 route is based on the assumption that a parameterization that accurately models the true eddy flux divergence should
 326 result in an accurate coarse model simulation. The primary result of the foregoing section is the fact that a simulation
 327 using a stochastic GM parameterization that is configured to match as closely as possible the GM parameter diagnosed
 328 from a high resolution reference experiment yields significantly worse results than a basic simulation with a constant
 329 value for κ . If the underlying assumption is true then the contrapositive must also be true: a parameterization that yields
 330 an inaccurate coarse model simulation must be an inaccurate model of the true eddy flux divergence. The stochastic
 331 GM parameterization derived in section 3 must therefore be an inaccurate model of the true eddy flux divergence.
 332 (The converse, i.e. if a parameterization yields an accurate simulation then it must also be an accurate model of the
 333 true eddy flux divergence, is clearly false since section 2 shows that $\text{constant}\kappa$ gives a less-accurate model of the eddy
 334 flux divergence while section 4 shows that $\text{constant}\kappa$ yields a more-accurate coarse model simulation.)

335 There are many ways in which the stochastic GM model developed in section 3 might be inaccurate. For example,

336 as with all SPP stochastic parameterizations the modeled variations in κ are independent of the large-scale temperature,
 337 whereas the real variations in κ might depend on the large-scale temperature. Alternatively it might be the case that
 338 despite the care taken in developing a model with realistic spatial and temporal variability, some crucial aspect of
 339 the true variability of κ has been missed in the stochastic model of κ . But the results of section 2 indicate a more
 340 fundamental problem: the isotropic GM parameterization with depth-independent κ yields a very poor fit to the true
 341 eddy flux divergence, so that no matter how well the diagnosed κ is represented with a stochastic model it will never
 342 be an accurate model of the true eddy flux divergence. (Recall that the optimal κ computed using this configuration
 343 reduced the error between the parameterized eddy tracer flux divergence and the true eddy tracer flux divergence by
 344 only 30 to 40% as compared to $\kappa = 0$ (see Fig. 3).) It is possible that an anisotropic formulation (Smith and Gent, 2004)
 345 or a depth-dependent κ (Abernathay et al., 2013) might yield a better fit to the flux-divergence. In particular allowing
 346 arbitrary depth dependence would significantly increase the number of degrees of freedom available for fitting the true
 347 eddy flux divergence. But at a more general level it may be the case that it is simply not possible to model the true
 348 eddy flux divergence purely by varying κ (or the flux tensor in anisotropic GM) in a GM parameterization.

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

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

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

$$F_{\text{skew}}^x = -\kappa \partial_z \rho, \quad F_{\text{skew}}^y = -\kappa \partial_x \rho. \quad (11)$$

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

$$\Psi_{\text{stoch}} = \frac{F_{\text{stoch}}^x}{\partial_z \rho}, \quad \Psi_{\text{stoch}}^x = \frac{F_{\text{stoch}}^y}{\partial_x \rho}. \quad (12)$$

357 The stochastic GM model developed in section 3, as well as the SPP-GM approach of Juricke et al. (2017) and Juricke
 358 et al. (2018), fits within this framework; in these parameterizations the stochastic parameterization of the skew flux
 359 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
 360 the diagnosed eddy flux divergences well.

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

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

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

374 **6. Conclusions**

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

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

397 **Acknowledgements**

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

403 Abernathey, R., Ferreira, D., Klocker, A., 2013. Diagnostics of isopycnal mixing in a circumpolar channel. *Ocean Model.* 72, 1–16.
404 Andreczuk, M., Cooper, F., Juricke, S., Palmer, T., Weisheimer, A., Zanna, L., 2016. Oceanic stochastic parameterizations in a seasonal forecast
405 system. *Mon. Weather Rev.* 144, 1867–1875.
406 Arzel, O., Huck, T., Colin de Verdière, A., 2006. The different nature of the interdecadal variability of the thermohaline circulation under mixed
407 and flux boundary conditions. *J. Phys. Ocean.* 36, 1703–1718.
408 Bachman, S., Fox-Kemper, B., 2013. Eddy parameterization challenge suite I: Eady spindown. *Ocean Model.* 64, 12–28.
409 Berloff, P.S., 2005. Random-forcing model of the mesoscale oceanic eddies. *J. Fluid Mech.* 529, 71–95.
410 Berner, J., Achatz, U., Batt, L., Bengtsson, L., de la Cmara, A., Christensen, H.M., Colangeli, M., Coleman, D.R.B., Crommelin, D., Dolaptchiev,
411 S.I., Franzke, C.L.E., Friederichs, P., Imkeller, P., Jrvinen, H., Juricke, S., Kitsios, V., Lott, F., Lucarini, V., Mahajan, S., Palmer, T.N., Penland,
412 C., Sakradzija, M., von Storch, J.S., Weisheimer, A., Weniger, M., Williams, P.D., Yano, J.I., 2017. Stochastic parameterization: Toward a new
413 view of weather and climate models. *B. Am. Meteorol. Soc.* 98, 565–588. doi:10.1175/BAMS-D-15-00268.1
414 Brankart, J.M., 2013. Impact of uncertainties in the horizontal density gradient upon low resolution global ocean modelling. *Ocean Model.* 66,
415 64–76.
416 Brankart, J.M., Candille, G., Garnier, F., Calone, C., Melet, A., Bouttier, P.A., Brasseur, P., Verron, J., 2015. A generic approach to explicit
417 simulation of uncertainty in the NEMO ocean model. *Geosci. Model Dev.* 8, 1285–1297.
418 Buizza, R., Milleer, M., Palmer, T., 1999. Stochastic representation of model uncertainties in the ECMWF ensemble prediction system.
419 *Q. J. Roy. Meteor. Soc.* 125, 2887–2908.
420 Christensen, H., Lock, S.J., Moroz, I., Palmer, T., 2017. Introducing independent patterns into the stochastically perturbed parametrization ten-
421 dencies (SPPT) scheme. *Q. J. Roy. Meteor. Soc.* 143, 2168–2181.
422 Davini, P., Hardenberg, J.v., Corti, S., Christensen, H.M., Juricke, S., Subramanian, A., Watson, P.A., Weisheimer, A., Palmer, T.N., 2017. Climate
423 sphinx: evaluating the impact of resolution and stochastic physics parameterisations in the ec-earth global climate model. *Geoscientific Model
424 Development* 10, 1383–1402.
425 Durran, D.R., 2010. Numerical methods for fluid dynamics with applications to geophysics. Second ed., Springer Dordrecht.
426 Eden, C., Greatbatch, R.J., Willebrand, J., 2007. A diagnosis of thickness fluxes in an eddy-resolving model. *J. Phys. Ocean.* 37, 727–742.

427 Ferrari, R., Griffies, S.M., Nurser, A.G., Vallis, G.K., 2010. A boundary-value problem for the parameterized mesoscale eddy transport. *Ocean*
 428 *Model.* 32, 143–156.

429 Friedman, J., Hastie, T., Tibshirani, R., 2008. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* 9, 432–441.

430 Gent, P.R., McWilliams, J.C., 1990. Isopycnal mixing in ocean circulation models. *J. Phys. Ocean.* 20, 150–155.

431 Gent, P.R., Willebrand, J., McDougall, T.J., McWilliams, J.C., 1995. Parameterizing eddy-induced tracer transports in ocean circulation models.
 432 *J. Phys. Ocean.* 25, 463–474.

433 Griffies, S.M., 1998. The Gent-McWilliams skew flux. *J. Phys. Ocean.* 28, 831–841.

434 Grooms, I., 2016. A Gaussian-product stochastic Gent-McWilliams parameterization. *Ocean Model.* 106, 27–43.

435 Grooms, I., Julien, K., Fox-Kemper, B., 2011. On the interactions between planetary geostrophy and mesoscale eddies. *Dyn. Atmos. Oceans* 51.
 436 doi:10.1016/j.dynatmoce.2011.02.002

437 Grooms, I., Smith, K.S., Majda, A.J., 2012. Multiscale models for synoptic-mesoscale interactions in the oceans. *Dyn. Atmos. Oceans* 58.

438 Henning, C.C., Vallis, G.K., 2004. The effects of mesoscale eddies on the main subtropical thermocline. *J. Phys. Ocean.* 34, 2428–2443.

439 Juricke, S., MacLeod, D., Weisheimer, A., Zanna, L., Palmer, T., 2018. Seasonal to annual ocean forecasting skill and the role of model and
 440 observational uncertainty. *Q. J. Roy. Meteor. Soc.* doi:10.1002/qj.3394 accepted manuscript.

441 Juricke, S., Palmer, T.N., Zanna, L., 2017. Stochastic subgrid-scale ocean mixing: Impacts on low-frequency variability. *J. Climate* 30, 4997–5019.

442 Large, W.G., McWilliams, J.C., Doney, S.C., 1994. Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameteriza-
 443 tion. *Rev. Geophys.* 32, 363–403.

444 Leutbecher, M., Lock, S.J., Ollinaho, P., Lang, S., Balsamo, G., Bechtold, P., Bonavita, M., Christensen, H., Diamantakis, M., Dutra, E., English,
 445 S., Fisher, M., Forbes, R., Goddard, J., Haiden, T., Hogan, R., Juricke, S., Lawrence, H., MacLeod, D., Magnusson, L., Malardel, S., Massart,
 446 S., Sandu, I., Smolarkiewicz, P., Subramanian, A., Vitart, F., Wedi, N., Weisheimer, A., 2017. Stochastic representations of model uncertainties
 447 at ECMWF: State of the art and future vision. *Q. J. Roy. Meteor. Soc.* 143, 2315–2339.

448 Lorenz, E., 1975. Climatic predictability. WMO The Phys. Basis of Climate and Climate Modelling p 132-136(SEE N 76-19675 10-47) .

449 Mak, J., Maddison, J.R., Marshall, D.P., 2016. A new gauge-invariant method for diagnosing eddy diffusivities. *Ocean Model.* 104, 252–268.

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

452 McDougall, T.J., 1987. Neutral surfaces. *J. Phys. Ocean.* 17, 1950–1964.

453 Redi, M.H., 1982. Oceanic isopycnal mixing by coordinate rotation. *J. Phys. Ocean.* 12, 1154–1158.

454 Rue, H., Held, L., 2005. Gaussian Markov random fields theory and applications. Chapman & Hall/CRC.

455 Smith, K.S., Marshall, J., 2009. Evidence for enhanced eddy mixing at middepth in the Southern Ocean. *J. Phys. Ocean.* 39, 50–69.

456 Smith, R.D., Gent, P.R., 2004. Anisotropic Gent-McWilliams parameterization for ocean models. *J. Phys. Ocean.* 34.

457 Williams, P., Howe, N., Gregory, J., Smith, R., Joshi, M., 2016. Improved climate simulations through a stochastic parameterization of ocean
 458 eddies. *J. Climate* 29, 8763–8781.