Analyzing Lagrangian Statistics of Eddy-Permitting Models

by

Amy Chen

B.S., University of Colorado Boulder, 2018

M.S., University of Colorado Boulder, 2018

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Master of Science

Department of Applied Mathematics

2018

This thesis entitled:
Analyzing Lagrangian Statistics of Eddy-Permitting Models
written by Amy Chen
has been approved for the Department of Applied Mathematics

_____

Dr. Ian Grooms

_____

Dr. James Curry

_____

Dr. William Kleiber

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Chen, Amy (M.S., Applied Mathematics)

Analyzing Lagrangian Statistics of Eddy-Permitting Models

Thesis directed by Dr. Ian Grooms

Mesoscale eddies are the strongest currents in the world oceans and transport properties such as heat, dissolved nutrients, and carbon. The current inability to effectively diagnose and parameterize mesoscale eddy processes in oceanic turbulence is a critical limitation upon the ability to accurately model large-scale oceanic circulations. This investigation analyzes the Lagrangian statistics for four faster and less computationally expensive eddy-permitting models — Biharmonic, Leith, Jansen & Held Deterministic, and Jansen & Held Stochastic — and compares them against each other and an eddy-resolving quasigeostrophic Reference model. Results from single-particle climatology show that all models exhibit similar behavior in large-scale movement over long times scales and their Lagrangian statistics display Gaussian behavior. However, differences between the models arise in smaller-scale particle-pair climatology. Root-mean-square forecasting errors are found to be uniform across models because the error in the initial condition affects the accuracy of the forecast more than the model chosen.

Dedication

To Anne Dougherty and Ian Grooms —

for believing in me and making sure I always landed on my feet.

Acknowledgements

To Anne Dougherty, for the many times I have cried in your office over the years about things applied math and not: thank you for being more than just my academic advisor. Thank you for caring about each and every one of us, for always being upfront and realistic about what needs to be done, for finding opportunities that help us grow, and for serving as a role model for all women in STEM. Your endless dedication to our department is a critical component of our success.

To my dearest roommate, teammate, and friend: Vickie, you have been my saving grace this semester. Moving in with someone I barely knew turned out to be the best thing I've ever done. Thank you for laughing, eating ice cream, and sharing memes with me. We sure spent a lot of time in the erg room on Friday nights, but not nearly as much time as we spent cleaning the kitchen, dancing in the car to 3OH!3, and complaining to each other about pretty much everything.

Maddie, thank you for always showing me the light and for always having the right words. Jacqui, for our monthly dinners and for being the strongest woman I know. Amy, for being so

endlessly supportive over the years; I hope to be half the Amy you are when I grow up. Grace, for dancing at concerts, sharing music, and encouraging me to be excited and expressive about how rad the world is. Mark (and Puff), for keeping it real and teaching me *that's the way she goes*. When you're a heart surgeon and I'm a secret agent, I hope we're still sharing Spotify playlists, you're still taking beta fish on road trips, and I'm still dropping beers off on your doorstep.

This department has been the greatest place to find my way, even through I still don't know where I'm going. Thank you to the applied math boys — Greg, Felix, Wilder, Ian, Ryan, Evan, Denis — for the witty humor, weird debates, passion for knowledge, and always helping me make it through regardless of how confused I am. You are all far more prepared, perceptive, and intelligent than I could ever hope to be and are sure to succeed at everything and anything you set your mind to. Julia and Quinn, for always trying so hard and for making me a better teacher — I am so privileged to see you grow. James Curry, William Kleiber, Manuel Lladser, Adam Norris, Harvey Segur, and Sujeet Bhat: the applied mathematics professors who have believed in me and my ability despite everything that has happened — thank you for being kinder to me than I am to myself. Thank you for believing that I will still be able to do what I used to do. I am so lucky that you care about my well-being and success not only as a student, but as a human being.

To the athletes and coaches of Colorado Crew: thank you for challenging me to push past my limits every single morning and for being my greatest support system. More than one million meters on the erg and who knows how many laps around the Boulder Reservoir later, you have made me stronger and more resilient than I ever thought I could be; thank you for the laughter, frustration, motivation, sunrises, blood, sweat, and tears. I will always be here for you as you have been here for me, even if it involves jumping into the ocean to rescue the launch we forgot to tie down — you know I'll always take one for the team. To the Boettcher Foundation, thank you for investing in my potential as a sixteen year-old. I hope the rest of my life will exceed your expectations. To Nii Armah: for African dance class, your life lessons, making me feel alive again, and the classmates who have become a family — and especially for bringing me Allie and Nick, the kind of friends who insist on doing the dishes after dinner. Jane, for teaching me that I am more

Contents

Chapter

Tables

Table

Figures

Figure

Chapter 1

Introduction

The circulation of the world oceans affects weather and climate, the oceans' ecosystems, shipping and transport, and oil and natural gas extraction. It exhibits dynamics on timescales from seconds to millenia and on spatial scales from millimeters to tens of thousands of kilometers. Global ocean circulation models, or *general circulation models (GCMs)*, are computational models that simulate the dynamics of the world oceans. The extreme range of oceanic spatial scales cannot be completely resolved with current computing power; the horizontal spatial resolution of ocean GCMs is currently limited to a grid size of five to ten kilometers, at best. Ensemble simulations and decadal — or longer — simulations are more computationally expensive and therefore have a lower resolution and larger grid size.

The strongest currents in the world oceans are associated with mesoscale eddies [5]. These eddies are ubiquitous, transient features with kinetic energy an order of magnitude larger than broad, persistent currents. More than half the kinetic energy of the ocean is contained within mesoscale eddies. Eddies exhibit different properties to their surroundings, transporting properties such as heat, dissolved nutrients, and carbon. They exist on a range of time and space scales, but are, on average, about 150 kilometers in diameter with peak velocities on the order of 20 $\frac{cm}{sec}$ [5]. Thoroughly *resolving* these eddies — or making sure they are well-depicted in ocean GCMs — requires grid sizes on the order of five to ten kilometers. Models with grid sizes on the order of 100 kilometers are completely unable to represent these eddies since the model resolution is too low. Models with grid sizes in between these extremes are called *eddy-permitting* since mesoscale eddies

are revealed in their simulations, but their dynamics are not well resolved. The current inability to effectively diagnose and parameterize mesoscale eddy processes in oceanic turbulence remains a critical limitation upon the ability to accurately model the oceans large-scale circulation.

A wide variety of approaches exists for improving the representation of these mesoscale eddies in eddy-permitting GCMs. The simplest approach is to add a scale-selective biharmonic viscosity, or smoothing, term to the models' momentum equations [20, 4]. This keeps the models' solutions smooth without overly smearing out the partially-resolved eddies. Other approaches achieve the same goal by allowing the coefficient of viscous smoothing to vary in space and time [6, 2].

Eddy-permitting GCMs often have too little kinetic energy. There is a class of methods that inject kinetic energy into the resolved scales of the simulation, mimicking the 'backscatter' of energy from scales too small to be resolved [10, 16]. A modification on these methods, proposed by Jansen & Held [13, 14], injects energy into the resolved scales to replace energy that was spuriously smoothed away by viscous terms in the models' momentum equations. The goal of the present investigation is to compare four methods of improving eddy-permitting ocean models: the biharmonic viscosity method as a baseline, a method based on spatial variation of the viscous coefficient, a method with deterministic energy injection into the resolved scales, and a method with stochastic energy injection into the resolved scales. For the remainder of the investigation, these models will be referred to as *Biharmonic, Leith, Jansen & Held Deterministic (JHD),* and *Jansen & Held Stochastic (JHS)*, respectively. These four eddy-permitting approximate models, as well as the traditional two-equal-layer quasigeostrophic, or *Reference*, eddy-resolving model they are compared against, are described in Chapter Two. Analysis and comparison of these approximate eddy-permitting models is valuable because they are less computationally expensive than eddy-resolving models. Accurate eddy-permitting models can be used in place of eddy-resolving models to reduce the time and computational power necessary for GCMs.

Different methods of improving eddy-permitting GCMs are often compared by analyzing the *Eulerian statistics* of the resolved scales in eddy-permitting models to the same statistics in eddy-resolving models or observations [7, 2]. Eulerian statistics are statistics of the velocity field in

space and time. For example, the time-mean kinetic energy 100 kilometers west of Cape Hatteras is an Eulerian statistic, as is mean kinetic energy and mean energy spectrum. All four eddy-permitting models considered in this investigation can be tuned to achieve reasonably correct Eulerian statistics, as discussed in following chapter.

The goal of the present investigation is to instead compare the four different eddy-permitting methods introduced above by analyzing their *Lagrangian statistics* against the eddy-resolving Reference model and each other. Lagrangian statistics are statistics associated with particles advected by the oceanic flow, which are called *Lagrangian particles*. The flow fields in question are turbulent such that particles launched from the same place at different times will have markedly different trajectories. Lagrangian statistics directly relate to how the eddying flow transports important quantities like heat, dissolved nutrients, and carbon, and are thus well-suited to compare the flow fields produced by different models. Finding eddy-permitting models that accurately approximate the Lagrangian statistics of eddy-resolving models can greatly reduce the time and computational cost of GCMs that focus on particle transport. In many regions, such as the Southern Ocean, mesoscale eddies are the dominant form of heat transport and are important in understanding the climate in those regions. Global climate models are developed to take these eddies into account for the analysis and simulation of climate change. As mentioned earlier, these models are computationally expensive because the grid size for eddy-resolving models is very small compared to that of the globe, so finding eddy-permitting models with the correct Lagrangian statistics can significantly decrease the computational and time costs of these models. The results of the present investigation aim to shed some light on the strengths of each model and potentially provide guidance as to when certain models should be used in certain situations, as well as inform further work and improvements on eddy-permitting models for eddy-permitting GCMs.

A Lagrangian particle released at time $t = 0$ from initial position $\mathbf{x}(t = 0) = \mathbf{x}_0$ evolves according to the system of ordinary differential equations

$$\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}(t), t)$$

where $\mathbf{v}(\mathbf{x}, t)$ is the velocity field. The models used here to generate velocity fields are highly idealized and the velocity fields produced are spatially homogeneous random fields since the initial conditions are random. In other words, their statistics are invariant under translations in space. However, these fields are anisotropic and differ in the x- and y-directions. The Lagrangian particle trajectory $\mathbf{x}(t)$ is a multivariate stochastic process and the homogeneity of the velocity field implies that the statistics of the trajectory are independent of the initial location of the particle in the velocity field. The statistics of single-particle Lagrangian trajectories can therefore be used to assess how rapidly particles move away from their place of origin and mix with the surrounding flow.

A second Lagrangian quantity of interest is the separation of two particles over time that were initially close together. This multivariate stochastic process is independent of the initial location of the particle pair — but not independent of the initial distance between them — because of the homogeneity of the velocity field. The statistics of the separation between a pair of Lagrangian particles can also be used to assess how rapidly quantities are mixed by the flow, as well as the potential final position of a particle given error in the initial location.

The preceding single-particle and particle-pair statistics measure generic or *climatological* properties of the flow field. Ocean GCMs are used to study climatological properties of the ocean and the Earth; the foregoing Lagrangian statistics are valuable ways to quantify this behavior. Ocean GCMs are also used to make short-term *forecasts* in instances such as search and rescue or to predict where currents will take spilled pollutants. The third Lagrangian quantity of interest in this investigation is used in forecasting. It is the separation between two particles initialized at the same time and location but transported by two different velocity fields: the true, eddy-resolving velocity field, Reference, and the velocity field produced by an eddy-permitting model, which, for this investigation, is one of the four aforementioned approximate models. The statistics of this separation can be used to measure the short-term forecast accuracy of eddy-permitting models.

The thesis is organized as follows: The four eddy-permitting models used to generate the flow fields, along with an eddy-resolving reference model, are described in Chapter Two. The second

chapter also describes the basic Eulerian statistics of the models and how the model parameters are tuned to have the correct Eulerian statistics of the eddy-resolving model. Chapter Three describes the general research methods used and how Lagrangian statistics for the models were obtained. Chapters Four and Five describe the specifics in computing the trajectories of the eddy-resolving model and the eddy-permitting models, as well as the resulting Lagrangian statistics. Chapter Four focuses on climatology, both single-particle and particle-pair, and Chapter Five focuses on forecasting. Analysis of the results and conclusions from the investigation are discussed in Chapter Six.

Chapter 2

Model Description

2.1 Eddy-Resolving Reference Model

Ocean mesoscale eddy dynamics are modeled in this investigation using the traditional two-equal-layer quasigeostrophic model on an f-plane forced by an imposed mean zonal baroclinic shear. The nondimensional governing equations are

$$\partial_t q_1 + \mathbf{u}_1 \cdot \nabla q_1 = -\nu_8 \nabla^8 q_1 \tag{2.1}$$

$$\partial_t q_2 + \mathbf{u}_2 \cdot \nabla q_2 = -c_d \operatorname{curl}\left[|(\mathbf{u}_2 + \hat{\mathbf{x}})|\,(\mathbf{u}_2 + \hat{\mathbf{x}})\right] - \nu_8 \nabla^8 q_2 \tag{2.2}$$

$$q_1 = y + \nabla^2 \psi_1 + \frac{1}{2}(\psi_2 - \psi_1) \tag{2.3}$$

$$q_2 = -y + \nabla^2 \psi_1 + \frac{1}{2}(\psi_1 - \psi_2) \tag{2.4}$$

where q_i and ψ_i are the potential vorticity and streamfunction, respectively, in layer i. The components of ψ and q that are linear in y are related to the imposed uniform zonal baroclinic shear. The velocity is related to the gradient of the streamfunction by the equation

$$\mathbf{u}_i = (-\partial_y \psi_i, \partial_x \psi_i)$$

The term multiplied by c_d in equation (2.2), $\operatorname{curl}\left[|(\mathbf{u}_2 + \hat{\mathbf{x}})|\,(\mathbf{u}_2 + \hat{\mathbf{x}})\right]$, is a standard quadratic drag where the imposed zonal velocity in the lower layer, $-\hat{\mathbf{x}}$, has been subtracted from u_2, which includes the mean flow, before computing the drag. Hyperviscous dissipation is included with the identical hyperviscosity coefficients, ν_8, in each layer. These equations have been nondimensionalized using the imposed zonal velocity as a velocity scale and the deformation radius as a length scale.

The model is configured exactly as in [24, 12] except for the use of eighth order hyperviscosity instead of biharmonic hyperviscosity. Parameter values for the model are $c_d = 0.1$ and $\nu_8 = 5\times10^{-7}$. The domain is square and has a nondimensional width of 32π. Approximate solutions are computed using 256×256 nonzero Fourier modes and a fourth-order semi-implicit Runge-Kutta method as described in [11]. The nondimensional grid size is 0.39 such that there are just more than two grid points per deformation radius, which is sufficient for eddy-resolving computations of this kind [22]. The nondimensional time step is adaptively selected as in [11], or set to a fixed nondimensional value of 0.005. This configuration is used as the eddy-resolving reference model, referred to as Reference for the remainder of the investigation, to which the subsequently described eddy-permitting models are compared.

This model is nondimensional, but it is valuable to choose dimensional units to facilitate comparison with real ocean models. In [24], dimensional units for this model were chosen by comparing the size and strength of the eddies in this nondimensional model to the size and strength of real eddies observed via satellite. Following the observational results of [5], the radius of an eddy was set to 75 kilometers and the maximum current speed was set to 18 $\frac{cm}{sec}$. The result of this comparison is that the nondimensional unit length is 38 kilometers (such that the model's 'deformation radius' is 38 kilometers), the nondimensional unit time is six weeks, and the nondimensional unit velocity is 1 $\frac{cm}{sec}$. Results of this investigation will be reported using these dimensional units, unless stated otherwise. The dimensional grid size of the reference eddy-resolving model is thus 15 kilometers, which is ten times smaller than the average diameter of an eddy, and the dimensional domain size is $256 \times 15 = 3840$ kilometers.

2.2   Eddy-Permitting Models

The defining feature of an eddy-permitting model is not the form of the equations, but rather the grid size. The four models described below are eddy-resolving models which are run as eddy-permitting models by coarsening the grid. Simulations of the eddy-permitting model equations

were run at two eddy-permitting grid sizes: 40 kilometers (a 96×96 grid) and $26.\bar{6}$ kilometers (a 144×144 grid). In the coarser grid there are only about 4 grid points per eddy, and in the finer grid there are between 5 and 6 grid points per eddy. For comparison, Jansen and Held [13] used a configuration with a deformation radius of 12 km. Their primary 'eddy-permitting' resolution of $\Delta_x = 15$ km corresponds to 1.3 grid points per deformation radius, whereas here the grid size of $\Delta_x = 40$ km corresponds to 1.04 grid points per deformation radius.

All four of the eddy-permitting simulations are carried out using the same numerical methods as the eddy-resolving reference simulation with the following exceptions:

(1) The equations are integrated using a third-order, three-stage Runge-Kutta method instead of a fourth-order, semi-implicit Runge Kutta method.

(2) The advection terms $\mathbf{u}_i \cdot \nabla q_i$ are discretized using the fourth-order energy and enstrophy conserving method found in [1].

These changes are made to make the numerical methods of the eddy-resolving models more similar to the numerical methods used in eddy-permitting ocean GCMs.

2.2.1   Biharmonic

The simplest model is the *Biharmonic* model. It has exactly the same form as the original reference model above, except that the hyperviscous terms $-\nu_8 \nabla^8$ in equations (2.1) and (2.2) are replaced by the biharmonic terms $-\nu_4 \nabla^4$. This approach to eddy-permitting and eddy-resolving models is widespread and considered as the baseline for eddy-permitting approximate models [20, 4].

2.2.2   Leith

The next-simplest model is the *Leith* model, based on [17]. The terms $-\nu_8 \nabla^8 q_i$ in equations (2.1) and (2.2) are replaced by $\nabla \cdot (\nu_* \nabla q_i)$. The viscosity coefficient varies in space and time via

$$\nu_* = \left(\frac{\Gamma \Delta_x}{\pi} \right)^3 |\nabla^3 \psi_i|$$

where Γ is a nondimensional constant and Δ_x is the grid spacing. The Leith parameterization for eddy-permitting models is poised to become operational for many eddy-permitting GCMs [6, 2, 18].

2.2.3   Jansen & Held

The final two models under consideration in this investigation are based on the idea of putting kinetic energy that was spuriously dissipated by biharmonic viscosity back into an eddy-permitting model. One model uses a deterministic method and the other uses a stochastic method to re-inject energy into the eddy-permitting model. Since these models are based on the work of Jansen and Held [13], they are called *Jansen & Held Deterministic (JHD)* and *Jansen & Held Stochastic (JHS)*, respectively.

2.2.3.i   Deterministic

The potential vorticity equations for the deterministic model, JHD, take the form

$$\partial_t q_1 + \mathbf{u}_1 \cdot \nabla q_1 = -\nu_2 \nabla^4 \psi_1 - \nu_4 \nabla^6 \psi_1 \tag{2.5}$$

$$\partial_t q_2 + \mathbf{u}_2 \cdot \nabla q_2 = -c_d \operatorname{curl}\left[\left|(\mathbf{u}_2 + \hat{\mathbf{x}})\right|(\mathbf{u}_2 + \hat{\mathbf{x}})\right] - \nu_2 \nabla^4 \psi_1 - \nu_4 \nabla^6 \psi_2. \tag{2.6}$$

The sign on the Laplacian terms $\nabla^2 q_i$ is anti-diffusive and these terms are responsible for injecting energy into the model. The biharmonic terms $\nabla^4 q_i$ maintain dissipation at small scales and prevent blowup of the solutions. The anti-diffusive coefficient ν_2 varies in time in such a way that it re-injects 90% of the energy dissipated by the biharmonic terms. Specifically,

$$\nu_2 = -0.9\nu_4 \frac{\int \psi_1 \nabla^6 \psi_1 + \psi_2 \nabla^6 \psi_2}{\int \psi_1 \nabla^4 \psi_1 + \psi_2 \nabla^4 \psi_2}$$

where the integrals are over the spatial domain.

2.2.3.ii   Stochastic

The potential vorticity equations for the stochastic model, JHS, take the form

$$\partial_t q_1 + \mathbf{u}_1 \cdot \nabla q_1 = F_1 - \nu_4 \nabla^6 \psi_1 \tag{2.7}$$

$$\partial_t q_2 + \mathbf{u}_2 \cdot \nabla q_2 = -c_d \operatorname{curl}\left[\left|(\mathbf{u}_2 + \hat{\mathbf{x}})\right|(\mathbf{u}_2 + \hat{\mathbf{x}})\right] + F_2 - \nu_4 \nabla^6 \psi_2. \tag{2.8}$$

The stochastic forcing terms in this model are Gaussian and uncorrelated in time. In Jansen and Held's study, the stochastic forcing terms were also uncorrelated in space. Following similar arguments for three-dimensional turbulence, it was argued in [9] that the stochastic forcing should instead have a spatial energy spectrum that is concentrated at small scales, but that noise should not be injected directly at the grid scale to avoid numerical errors. To accomplish this, the noise terms F_1 and F_2 were obtained from unit-variance, spatiotemporal white noises $\hat{F}_1^{(0)}$ and $\hat{F}_2^{(0)}$ in a four-stage process.

(1) First, the white noise was smoothed by applying a local moving average as in [9]

$$\hat{F}_1^{(1)} = \mathcal{S}\hat{F}_1^{(0)}, \quad \hat{F}_2^{(1)} = \mathcal{S}\hat{F}_2^{(0)}$$

where $\mathcal{S}$ denotes the smoothing operator.

(2) Next, the noises were scaled to have amplitudes proportional to the vertical structure of the solution ψ:

$$\hat{F}_1^{(2)} = \frac{\|\psi_1\|}{\|\psi_1\| + \|\psi_2\|} \frac{\hat{F}_1^{(1)}}{\|\hat{F}_1^{(1)}\|}, \quad \hat{F}_2^{(2)} = \frac{\|\psi_2\|}{\|\psi_1\| + \|\psi_2\|} \frac{\hat{F}_2^{(1)}}{\|\hat{F}_2^{(1)}\|}.$$

The norms $\|\cdot\|$ are the L^2 norm of the function over the spatial domain. Dividing by $\|\hat{F}_i^{(1)}\|$ ensures that the forcing has exactly unit amplitude before scaling it to match the vertical structure of ψ.

(3) At this point, $\hat{F}_i^{(2)}$ is the forcing that would appear in equations for ψ_i, up to normalization. To obtain the forcing in the q_i equations, perform

$$\hat{F}_1^{(3)} = \nabla^2 \hat{F}_1^{(2)} + \frac{1}{2}(\hat{F}_2^{(2)} - \hat{F}_1^{(2)})$$

$$\hat{F}_2^{(3)} = \nabla^2 \hat{F}_2^{(2)} + \frac{1}{2}(\hat{F}_1^{(2)} - \hat{F}_2^{(2)})$$

(4) All that remains is to scale the forcing so that it produces a rate of energy injection equal to 90% of the biharmonic energy dissipation. This is accomplished by setting

$$F_1 = A\hat{F}_1^{(3)}, \quad F_2 = A\hat{F}_2^{(3)} \text{ where}$$

$$A = -0.9\nu_4 \frac{\int \psi_1 \nabla^6 \psi_1 + \psi_2 \nabla^6 \psi_2}{\int \hat{F}_1^{(2)} \hat{F}_1^{(3)} + \hat{F}_2^{(2)} \hat{F}_2^{(3)}}.$$

The spatial smoothing operator $\mathcal{S}$ is defined on the computational grid as a moving average with weights 0.25, 0.5, and 0.25 performed in both the x and y directions.

2.3   Tuning Parameters for Eulerian Statistics

Each of the eddy-permitting models has a single tunable parameter: Γ for Leith and ν_4 for Biharmonic, JHD, and JHS. At each resolution, these parameters are tuned to produce accurate Eulerian statistics: correct upper-layer kinetic energy spectrum and correct upper-layer total kinetic energy. During this tuning process, simulations were run from random initial conditions until the total kinetic energy reached a statistically steady state. Table 2.1 reports the optimal parameter values and the square root of the average (in space and time) of $|\mathbf{u}_1|^2$ for each model. Figure 2.1 shows the kinetic energy spectra for the upper layer for the each model, with two different eddy-permitting resolutions.

Table 2.1: Optimal parameter values (nondimensional) for the four eddy-permitting models, along with the square root of the average (in space and time) of $|\mathbf{u}_1|^2$. The reported parameter values are nondimensional, while the mean current speeds are reported in cm/s. The mean current speed in the reference model is 10.9 cm/s.

	Biharmonic	Leith	JHD	JHS
$\Delta_x = 26.\bar{6}$ km: ν_4 or Γ	5×10^{-3}	1.08	10^{-2}	10^{-2}
$\Delta_x = 26.\bar{6}$ km: Mean current speed	10.4	10.6	10.8	10.5
$\Delta_x = 40$ km: ν_4 or Γ	10^{-2}	0.96	5×10^{-2}	1.7×10^{-2}
$\Delta_x = 40$ km: Mean current speed	10.6	10.5	10.8	10.4

Figure 2.1: Nondimensional kinetic energy spectra in the upper layer for the reference model and the four eddy-permitting models at the coarser resolution ($\Delta_x = 40$ km; left) and finer resolution ($\Delta_x = 26.\bar{6}$ km; right).
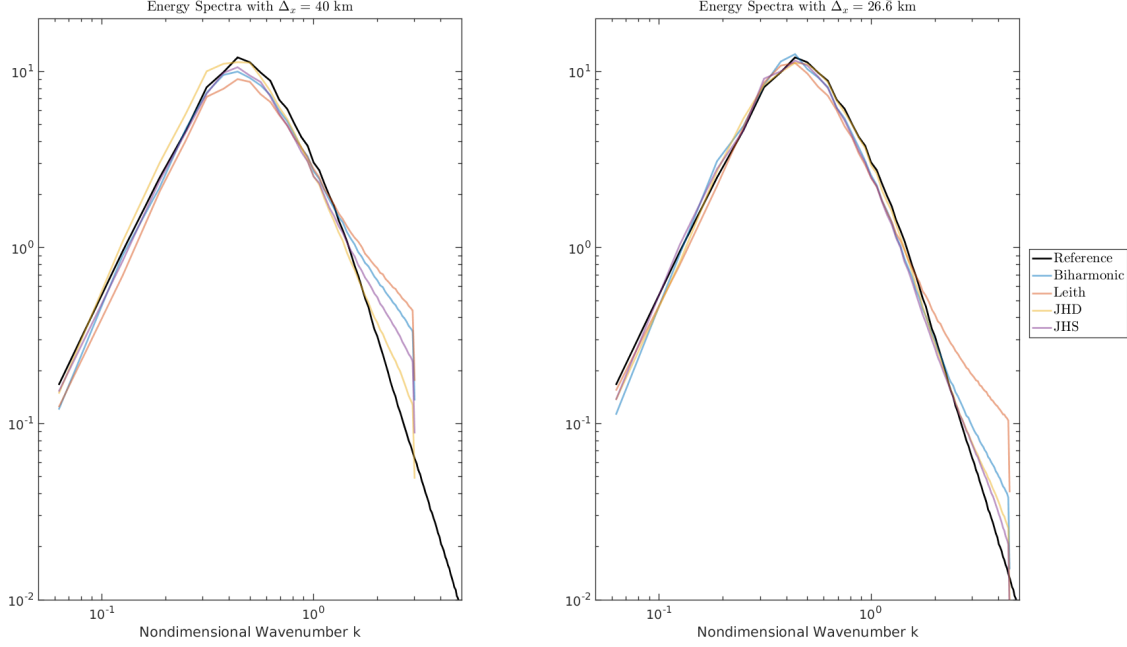
All the models are able to achieve mean current speeds close to the reference model. However, compromises have to be made to achieve this at the level of the energy spectrum. In particular, the Leith model has far too much energy at small scales (scales near the grid scale). This error in the energy spectrum can be alleviated by increasing Γ, but this comes at the cost of lowering the total energy of the model. In this regard, the Leith model is less accurate than the standard biharmonic model, whereas both of the Jansen & Held approaches yield more accurate spectra. This is presumably because the Jansen & Held models allow for strong biharmonic damping to keep the energy levels near the grid scale appropriately small. In the other two models this strong damping would decrease the total energy, but the Jansen & Held models re-inject energy to the large scales. The two Jansen & Held models are quite similar at the finer grid resolution ($\Delta_x = 26.\bar{6}$ km), but at the coarser resolution ($\Delta_x = 40$ km) the deterministic version yields better results. At the coarser resolution the optimal biharmonic viscosity parameter in the deterministic model is 5×10^{-2}, whereas in the stochastic model it is 1.7×10^{-2}. Increasing the biharmonic viscosity

parameter in the stochastic model does improve the energy spectrum at small scales but it also reduces the total energy; this suggests that the stochastic approach re-injects energy less efficiently than the deterministic approach.

Chapter 3

Research Methods

Once the parameter values of the approximate models have been tuned to achieve the correct *Eulerian statistics* as described in Chapter Two, the models are run to compare their *Lagrangian statistics* to the Reference model and each other. The simulations that generate the velocity fields in this investigation were run from a randomized initialization (for climatological statistics) or prescribed initial condition (for forecasting statistics) until they reached a *Eulerian steady state*, that is, until the velocity field reached and maintained the correct energy spectrum. To be precise, the climatological statistics in each model were taken starting at the first time step that the model's upper-layer velocity reached an average amplitude of 10.7 cm/s. Once steady state was reached, Lagrangian particles were placed in the velocity fields to track Lagrangian trajectories and Lagrangian velocities.

To obtain Lagrangian statistics, fifty different trials of each model were run to create fifty velocity fields for each model. As mentioned above, climatology trials were all randomly initialized and forecast trials shared the same fifty initial conditions across the five models. In each trial, 64 Lagrangian particles are placed in an equally spaced 8×8 grid on the periodic, nondimensional $32\pi \times 32\pi$ square domain. The initial spacing between the points in each trial is large enough to keep the resulting trajectories uncorrelated, so 50 trials of each model gives $50 \times 64 = 3200$ independent trajectories per model. The periodicity of the model allowed for the continued tracking of the trajectory even if the particle moved outside the initial grid.

Velocity fields of the approximate models were run at two different resolutions: a $26.\bar{6}$ kilo-

meter grid size (144×144) and a coarser 40 kilometer grid size (96×96). In both cases, the velocity field from the coarser grid was interpolated to a 256×256 grid using Fourier interpolation. Velocities off the grid are determined using cubic spline interpolation. Then, trajectories are computed by solving the system of ordinary differential equations

$$\dot{\mathbf{x}} = \mathbf{v}(\mathbf{x}(t), t)$$

using standard fourth-order Runge-Kutta. The nondimensional time step is 0.01, which is verified to be small enough to avoid numerical errors.

From these trajectories, Lagrangian statistics such as mean, standard deviation, absolute deviation, variance, skew, and kurtosis for both the Lagrangian trajectories and Lagrangian velocities are calculated. As mentioned before, the eddy-resolving Reference model is anisotropic, or not identical in the x- and y- directions. Thus, the Lagrangian statistics were also computed in the x- and y-directions for all the eddy-permitting models. Minor changes to this method are made for each of the three Lagrangian statistics of focus — single-particle climatology, particle-pair climatology, and forecasting — and are described in their respective sections in the following chapters.

Chapter 4

Climatology

4.1 Single-Particle Lagrangian Statistics

4.1.1 Method

The simplest of the three Lagrangian statistics to simulate is the *single-particle Lagrangian statistic*, or the statistics that assess how rapidly particles move away from their initial position and mix with the surrounding flow. These trials were run exactly as described in Chapter Three, with a few adjustments:

(1) The velocity fields for each of the four models were randomized. That is, that fifty random trials for one model were not the same fifty random trials used to obtain trajectories for another model.

(2) These simulations were calculated out to 336 days (48 weeks).

The trajectories are expected to exhibit *Gaussian* behavior after a long-enough time as a result of the Central Limit Theorem: the distribution of averages should be Gaussian. Since the trajectory is an integral of the velocity, it should eventually appear Gaussian. The test for whether or not the Reference model's trajectory statistics are Gaussian is inconclusive, but the Lagrangian statistics of the Reference model are expected to exhibit Gaussian properties — namely, that the skew is 0 and kurtosis is 3 — and this investigation considers whether or not the approximate models exhibit Gaussian Lagrangian statistics as well.

4.1.2   Results

All of the eddy-permitting approximate models performed well in the long term, single-particle scale: no model exhibited significant differences from the Reference model or the other approximate models.

4.1.3.i   Eddy-Permitting 144×144 Models

The four eddy-permitting models behave similarly to each other and the Reference model in single-particle dispersion. Absolute dispersion (where absolute is used in the sense of historical time) is the main calculation of comparison, which is computed by

$$\text{absolute dispersion} = \sqrt{E[(x - x_0)^2 + (y - y_0)^2]}$$

where x_0 and y_0 are the initial $x-$ and $y-$locations of the Lagrangian particle at time 0, respectively. Figure 4.1 shows that the absolute dispersion of all four approximate models closely corresponds to the Reference model at this finer grid. In addition, the models are noticeably more accurate in the ballistic regime than in the diffusive regime.

Figure 4.1: The absolute dispersion of all five eddy-permitting models at the finer resolution.
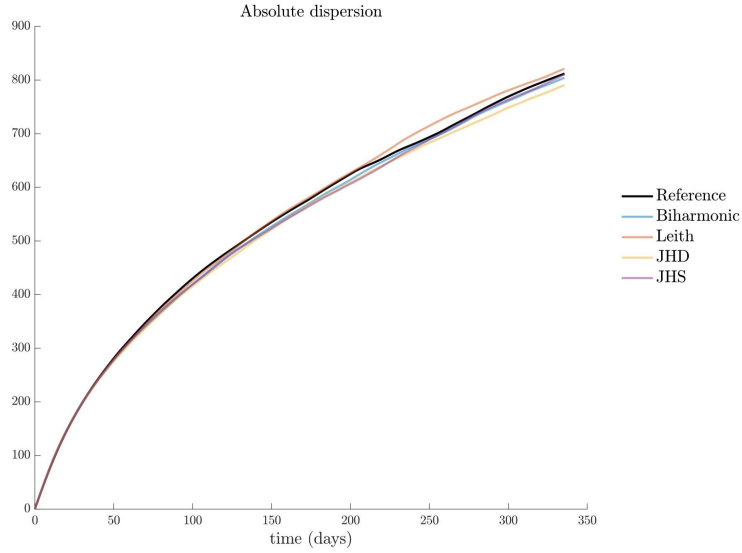
The four eddy-permitting approximate models, like the eddy-resolving Reference model, also

exhibited Gaussian Lagrangian statistics of trajectory as expected — with the exception of the kurtosis of trajectory in the ballistic regime. All of the models, including the Reference model, have kurtosis values significantly above 3 — the expected Gaussian kurtosis — in the ballistic regime. Figure 4.2 shows the kurtosis of trajectory in the x-direction and the y-direction; it is evident that kurtosis in the x-direction is even higher than in the y-direction. Over the entire time scale, however, kurtosis of trajectory averages out to the expected Gaussian value of 3 across all models.

Figure 4.2: x- and y-kurtosis of Lagrangian trajectory of the five models at the finer resolution.
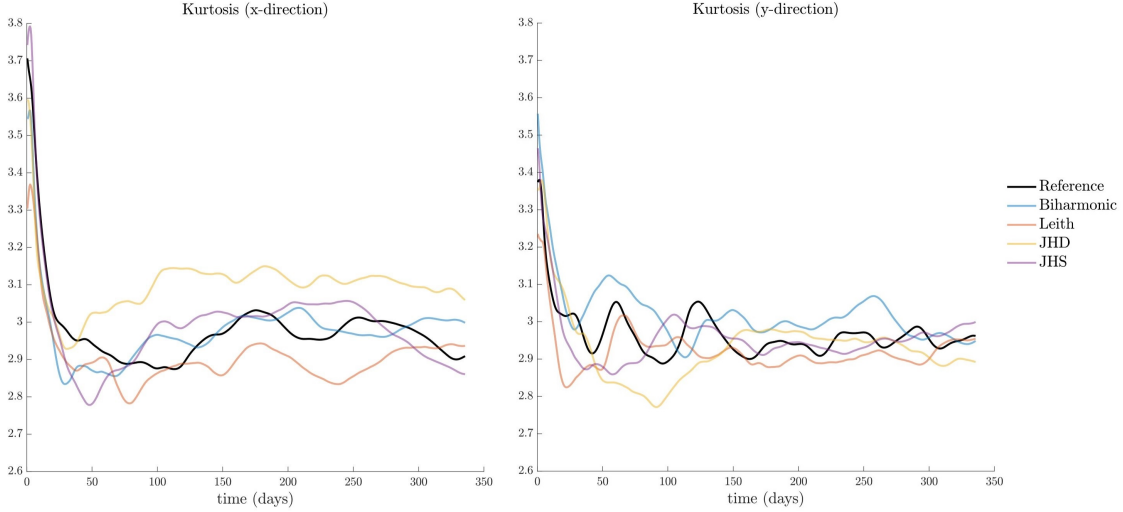
Trajectory statistics are not stationary; they change in time. At different times, the trajectory has different distributions and this result reveals that the ballistic regime is less Gaussian than the dispersive regime for single-particle dispersion. This makes sense because the ballistic regime is highly correlated and the Central Limit Theorem does not act strongly upon smaller sample sizes (since there are less time steps). The Lagrangian statistics for velocity were also similarly Gaussian, except that kurtosis of velocity in both the $x-$ and $y-$directions were slightly higher than Gaussian across the entire time scale. Table 4.1 lists the Lagrangian statistics of velocity, where the values of kurtosis can be seen to be higher than expected. This result suggests that velocity is less Gaussian than trajectory for single-particle dispersion.

Table 4.1: Table of values of Lagrangian velocity for the 144×144 models

	x-standard deviation (km)	y-standard deviation (km)	x-skew	y-skew	x-kurtosis	y-kurtosis
Reference	7.5871	7.6707	0.0688	-0.0018	3.5343	3.4733
Biharmonic	7.5406	7.5976	0.0645	0.0046	3.5022	3.5002
Leith	7.5284	7.5590	0.0892	-0.0082	3.3152	3.2172
JHD	7.5643	7.5731	0.0641	-0.0084	3.4982	3.5425
JHS	7.5032	7.4953	0.0906	0.0081	3.6053	3.5341

4.1.3.ii   Eddy-Permitting 96×96 Models

Since there were minimal differences between all four approximate eddy-permitting models and the Reference model at the finer grid size of 144×144, the four approximate models were further coarsened down to a grid size of 96×96 before interpolating back up to the 256×256 grid to determine if any larger discrepancies would arise. The results at this coarser grid size are similar to before — all four approximate models closely represent the Reference model. Figure 4.3 exhibits the accuracy in absolute dispersion of all the eddy-permitting models. Again, the eddy-permitting models are more precise in the ballistic region than the diffusive region.

Figure 4.3: The absolute dispersion of all five eddy-permitting models at the coarser resolution.
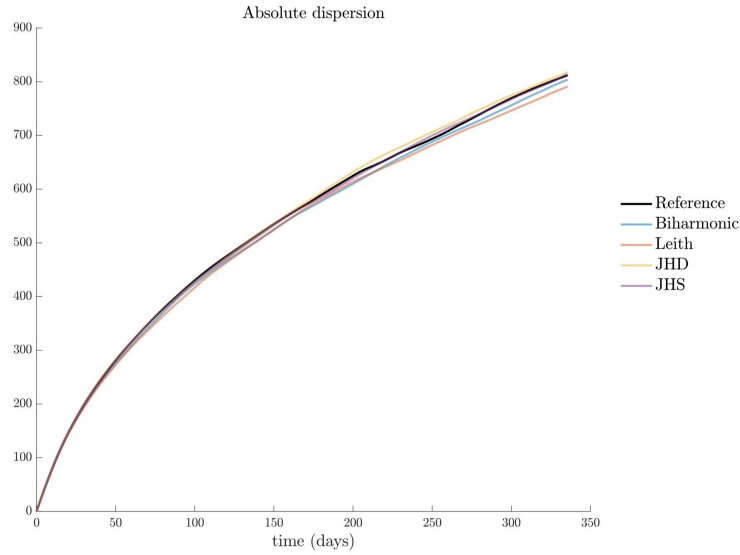
As in the finer resolution model, the Lagrangian statistics of the approximate eddy-permitting

models are Gaussian with the exception of the kurtosis of trajectory. Once again, the kurtosis is much higher than the expected Gaussian kurtosis of 3 in the ballistic region with kurtosis in the x-direction higher than kurtosis in the y-direction, as demonstrated in Figure 4.4. At a coarser resolution, the approximate models also exhibit higher values overall for kurtosis in the x-direction than in the finer resolution.

Figure 4.4: x- and y-kurtosis of Lagrangian trajectory of the five models at the coarser resolution.
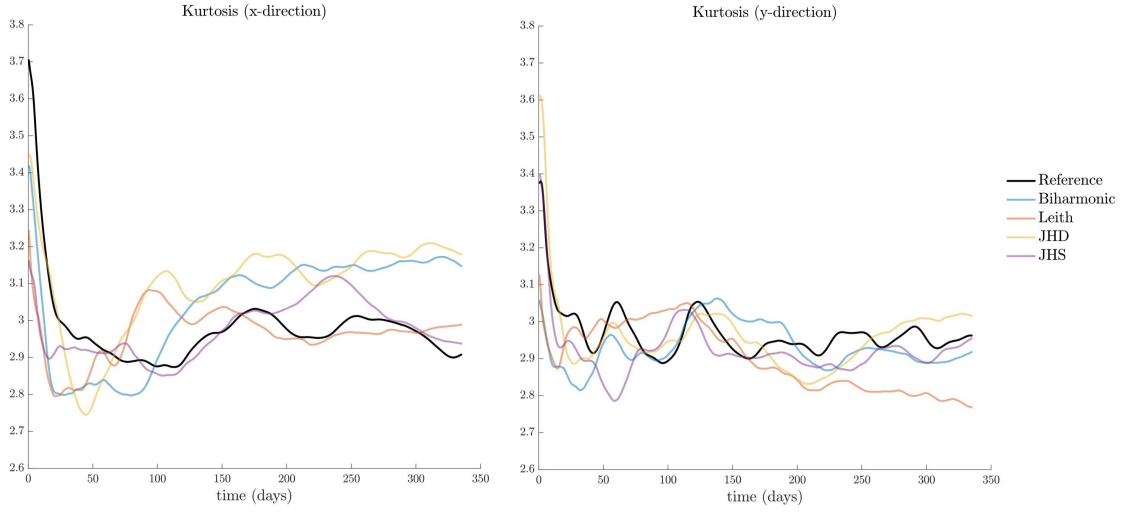
As in the finer grid size, the Lagrangian statistics for velocity were also similarly Gaussian for the coarser grid size. The exception, again, is that kurtosis of velocity in both the $x-$ and $y-$directions are higher than Gaussian across the entire time scale. Table 4.2 lists the Lagrangian statistics of velocity at the coarser grid size, where the values of kurtosis can be seen to be higher than expected. This result supports the previous conclusion that velocity is less Gaussian than trajectory for single-particle dispersion.

Table 4.2: Table of values of Lagrangian velocity for the 96 × 96 models

	x-standard deviation (km)	y-standard deviation (km)	x-skew	y-skew	x-kurtosis	y-kurtosis
Reference	7.5871	7.6707	0.0688	-0.0018	3.5343	3.4733
Biharmonic	7.5589	7.6188	0.0886	0.0080	3.2521	3.2736
Leith	7.5449	7.6066	0.1160	0.0049	3.1902	3.1692
JHD	7.6098	7.6995	0.0871	0.0033	3.4681	3.4378
JHS	7.4940	7.5600	0.0792	0.0057	3.2743	3.3135

4.1.3   Regimes of Absolute Deviation

All five models exhibit two regimes of absolute deviation typical of GCMs: the ballistic regime and the diffusive regime. In the beginning, right after the particles are placed, particles go through a *ballistic phase*. Here, absolute deviation is nearly linear since the absolute deviation of a particle is still highly correlated to its initial position and the initial velocities acting upon it when time is small. As time goes on, the particles enter the *diffusive regime*. After enough time has passed, the particle has moved away from its initial position enough that the initial conditions no longer affect it and it enters a random walk. In this investigation, the end of the ballistic regime was found to be around 250 nondimensional time steps, which is 52.5 days (7.5 weeks). The boundary can be seen in the plots of absolute deviation in both Figure 4.1 and Figure 4.3: the initial linear nature of the absolute dispersion ends across all the eddy-permitting models and begins to resemble a random walk. All of the models are similar over the diffusive regime and tend to perform similarly as time extends past the ballistic regime; for this reason, the remaining Lagrangian statistics will only be considered in the ballistic regime.

## 4.2   Particle-Pair Lagrangian Statistics

4.2.1   Method

Particle-pair dispersion, the separation of two particles over time that were initially close together, is the second Lagrangian statistic of interest in this investigation. As mentioned before,

this multivariate stochastic process is independent of the initial location of the particle pair — but not independent of the initial distance between them — because of the homogeneity of the velocity field. Like single-particle dispersion, statistics of the separation between a pair of Lagrangian particles that start North-South, East-West, or Northeast from each other can be used to assess how rapidly quantities are mixed by the flow. These trials were run exactly as described in Chapter Three, with a few adjustments:

(1) The velocity fields for each of the four models were randomized. That is, that fifty random trials for one model were not the same fifty random trials used to obtain trajectories for another model.

(2) These simulations were only calculated in the ballistic region, which is approximately calculated to be 52.5 days (7.5 weeks) since the results of single-particle dispersion indicate that all models perform similarly over the diffusive regime.

(3) Instead of placing a grid of 8×8 single Lagrangian particles in the domain, a grid of 8×8 clusters of four Lagrangian particles were placed in the domain. These clusters consist of four Lagrangian points in a square configuration with a side length of 15 kilometers. Figure 4.5 shows which points were used to the North-South, East-West, and Northeast dispersions of each cluster. The initial distance between the points used for the horizontal and vertical dispersions is 15 kilometers and the initial distance between the points used for the diagonal distance is $15\sqrt{2}$ kilometers.

Figure 4.5: The setup of the particle-pair dispersion model.
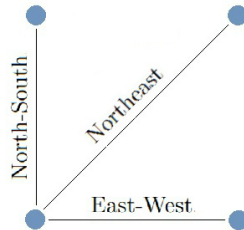
Pair-dispersion is calculated as the distance between the two closely-initialized particles over time according to the function

$$\text{pairdispersion} = \sqrt{E[(x_2 - x_1)^2 - (y_2 - y_1)^2]}$$

The particles are initialized on neighboring grid spaces, so $\Delta x = \Delta y = 15$ kilometers for the 256×256 grid size, $\Delta x = \Delta y = 26.\bar{6}$ kilometers for the 144×144 grid size, and $\Delta x = \Delta y = 40$ kilometers for the 96×96 grid size at $t = 0$. For the North-South initialization, $x_2(t = 0) = x_1(t = 0)$ and $y_2(t = 0) = y_1(t = 0) + \Delta y$. East-West initialization has $x_2(t = 0) = x_1(t = 0) + \Delta x$ and $y_2(t = 0) = y_1(t = 0)$, and Northeast initialization is $x_2(t = 0) = x_1(t = 0) + \Delta x$ and $y_2(t = 0) = y_1(t = 0) + \Delta y$.

4.2.2   Results

The particle-pair Lagrangian statistics exhibit some significant differences between the approximate eddy-permitting models. It is also interesting to note that as the models are coarsened down to lower resolutions, the Reference model consistently has less dispersion (and less energy) than the four eddy-permitting approximate models. This occurs because the original resolution of the eddy-resolving Reference model is 256×256, which is higher than the original resolution of 144×144 or 96×96 of the four eddy-permitting approximate models. As the approximate models are interpolated up to a higher resolution, small-scale movement is not as defined. Similarly to how particles behave in shear flow, the particle pairs at the interpolated resolutions do not experience as detailed small-scale movement as the Reference model and, on average, move away from each other faster in the approximate models.

4.2.2.i   Eddy-Permitting 144×144 Models

First, the four approximate models were run at the 144×144 grid and then interpolated up to the 256×256 grid. Although the dispersion of Reference, Biharmonic, JHD, and JHS seem relatively similar across the three dimensions, Leith stands out as having a larger dispersion from all three initializations. These differences across the approximate models can be related back to the

original tuning of the energy spectrum and Figure 2.1: the Leith model has far too much energy at small scales (scales near the grid scale) and therefore performs less accurately on this small-scale Lagrangian statistic. The higher energy in the small scales for the Leith model lead to it having more dispersion than the other models. This can be seen in Figure 4.6.

Figure 4.6: The North-South, East-West, and Northeast dispersion between pairs of particles of all five eddy-permitting models at the finer resolution.
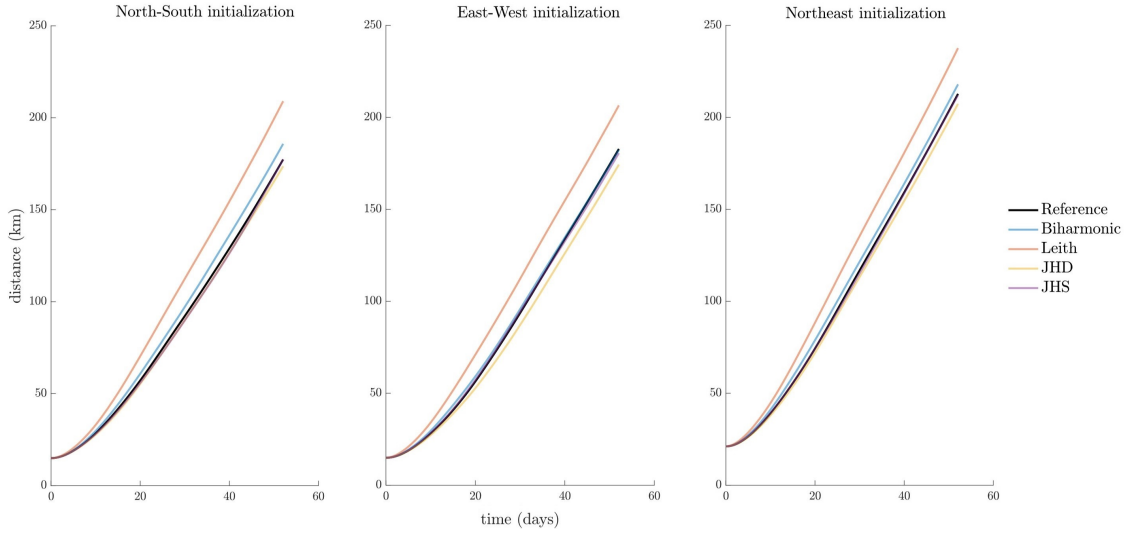
One of the real-world applications of particle-pair dispersion is in search and rescue. These particle-pair Lagrangian statistics provide a radius for a circular search area given a one-grid space (15 kilometer) error in initial position, or location last seen. The dispersion value of the last time step was used as the radius to calculate the circular search areas in a search and rescue situation in 4.7. For the circles below, the center of the circle is the location of the particle given an accurate initial condition. Although the velocity fields are anisotropic, the dispersion of the particles is symmetric and allows for the use of circles rather than ellipses in this visualization. If the initial condition has an error of less than 15 kilometers, the actual location of the particle after 52 days is somewhere inside the circle plotted in the figure. Clearly, the Leith model has a significantly larger search area than the other four eddy-permitting models, and does not model particle-pair dispersion behavior as well at the 144×144 grid size.
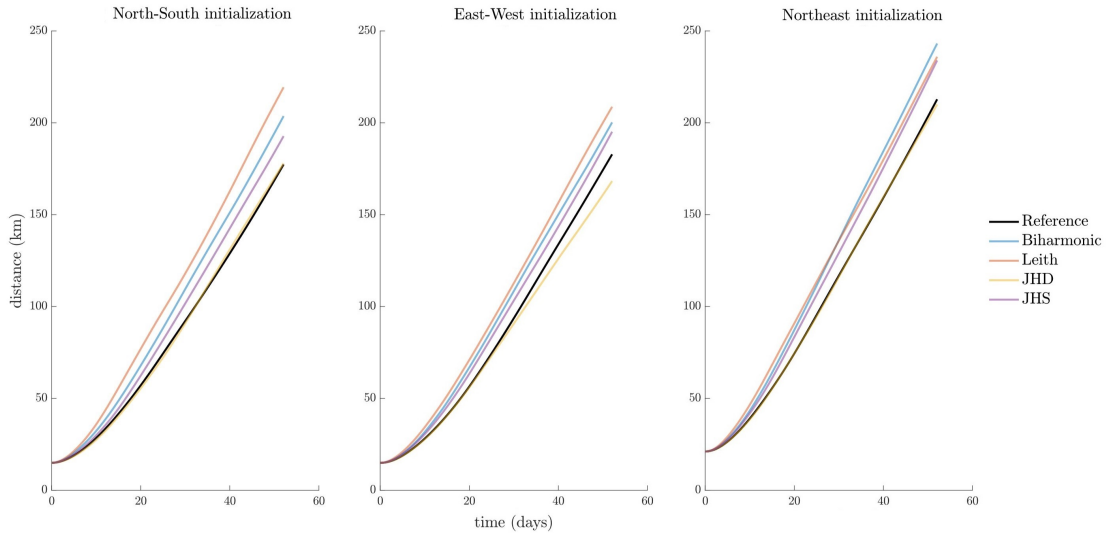
Figure 4.7: Circular search areas generated using the pair-dispersion value at the final time step as the radius
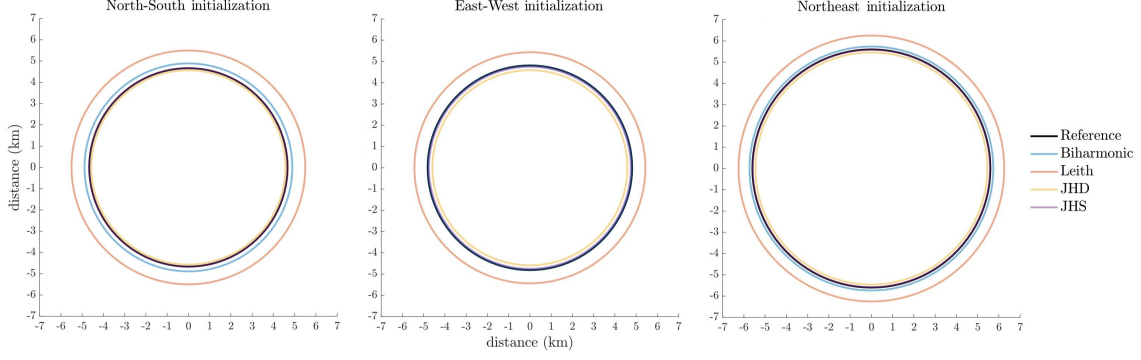
4.2.2.ii   Eddy-Permitting 96 × 96 Models

Further coarsening of all five models to the 96 × 96 grid size was carried out to see if any more significant differences would manifest. Figure 4.8 shows that at this lower resolution, JHD is most accurate to the Reference model — especially in the North-South and Northeast directions — and Leith stands out less than it did at the higher resolution. Again, this relates back to the tuning of the models and Figure 2.1: JHD has the best spectrum at the small-scales and performs well at this small-scale Lagrangian statistic.

Figure 4.8: The North-South, East-West, and Northeast dispersion between pairs of particles of all five eddy-permitting models at the coarser resolution.

These differences between the models are more evident in the circular search areas in Figure 4.9, which are calculated as in the previous section. The Leith model no longer stands out as much as it did in the finer grid and JHD performs exceptionally similarly to the Reference model.

Figure 4.9: Circular search areas generated using the pair-dispersion value at the final time step as the radius
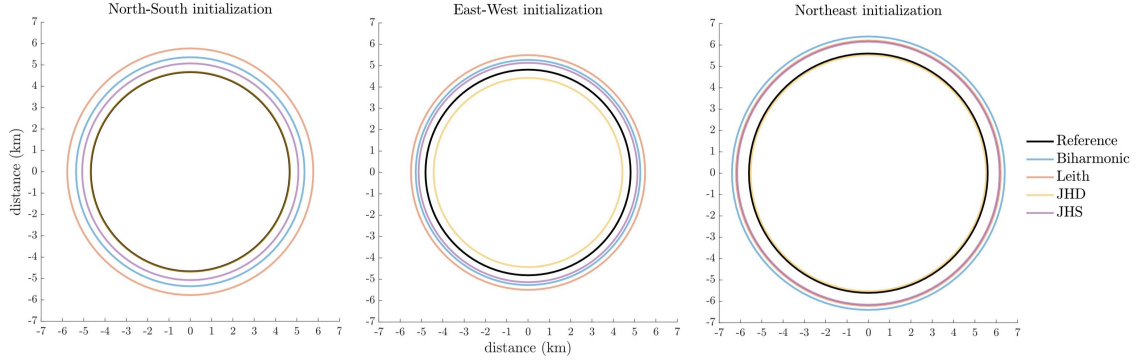
Chapter 5

Forecasting

5.1   Method

Single-particle forecast dispersion is the final Lagrangian quantity of interest in this investigation. It is the separation between two particles initialized at the same time and location but transported by two different velocity fields: the true, eddy-resolving velocity field, Reference, course-grained down to an equivalent grid size, and the velocity field produced by one of the four approximate eddy-permitting models. The statistics of this separation can be used to measure the short-term forecast accuracy of eddy-permitting models. These trials were run exactly as described in Chapter Three, with a few adjustments:

(1) The same initial conditions were used to generate fifty velocity fields across all the models. That is, the fifty velocity fields used for one model were the same fifty velocity fields used to obtain trajectories for another model.

(2) These simulations were only calculated in the ballistic region, which is approximately calculated to be 52.5 days, or 7.5 weeks, since forecasts become useless around the time the ballistic regime ends.

(3) The Reference model was calculated at its eddy-resolving 256×256 grid, then coarsened down to 144×144 or 96×96 before interpolating back up to 256×256 with the other approximate models to make the comparison for root-mean-squared error. In this calculation for root-mean-squared error, the Reference model was used as the expected value.

5.2   Results

The forecasting errors across all four approximate models are very similar. Absolute deviation of single-particle forecasting is nearly identical between models, so the investigation uses the root-mean-squared (RMS) error to analyze the difference between forecasts predicted by the approximate eddy-permitting models and the Reference model at the same grid size.

5.2.1   Eddy-permitting 144×144 Models

To calculate the RMS forecast error for the approximate models run at the finer grid size, the Reference model was also coarsened down to a 144×144 grid. There are no significant differences between the forecast accuracies of the models, as shown in Figure 5.1.

Figure 5.1: Root-mean-squared (RMS) forecast error between the approximate eddy-permitting models and the Reference model at the finer resolution.
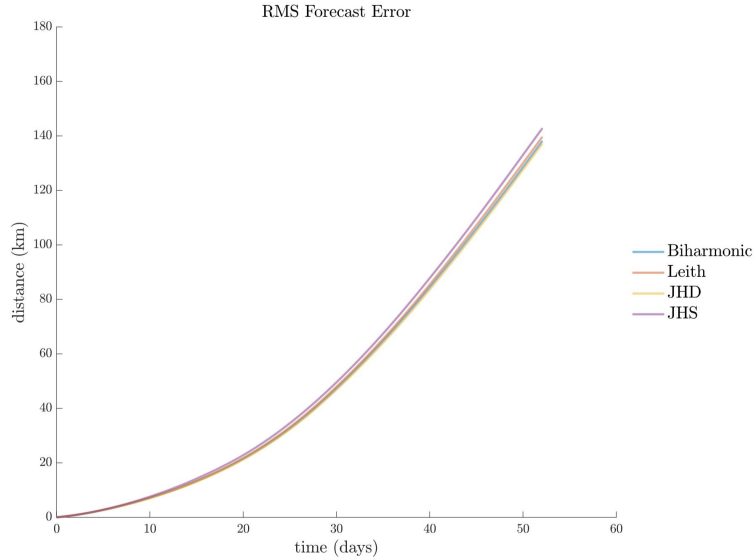
5.2.2   Eddy-permitting 96×96 Models

All five models were further coarsened down to the 96×96 grid size to see if any significant differences would be found at a lower resolution. Again, the four eddy-permitting approximate models exhibit very similar forecasting errors, with none of the models standing out as being significantly better or worse at single-particle forecasting, which is evident in Figure 5.2.
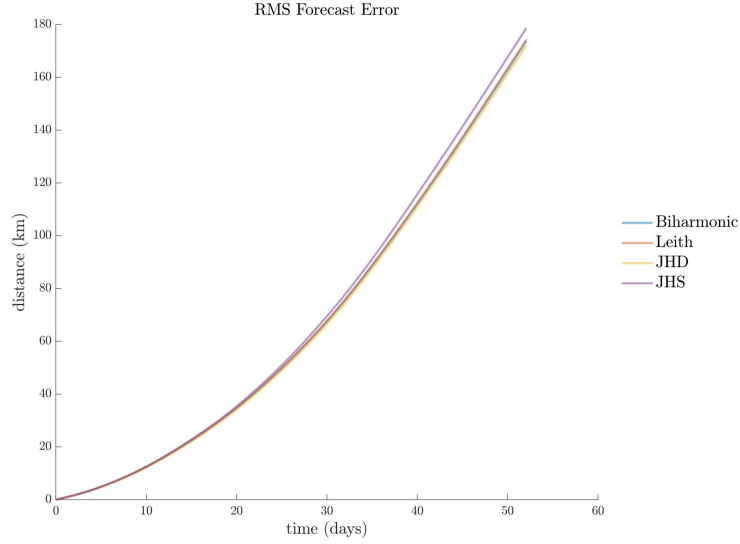
Figure 5.2: Root-mean-squared (RMS) forecast error between the approximate eddy-permitting models and the Reference model at the coarser resolution.

The lack of differences between models can be attributed to the accurate initialization of the velocity fields and the fact that the errors in the approximate models do not have enough time to accumulate in a forecast. These approximate models are different in their modeling errors and because they do not accumulate enough of these errors in the ballistic regime to have significantly different forecasts, they perform very similarly to each other. This result suggests that what matters most in accurate forecasting is having an accurate initial condition, not the model used in the forecast.

Chapter 6

Conclusions

The goal of the present investigation was to determine if there are any significant differences in the Lagrangian statistics of four approximate eddy-permitting models used in eddy-permitting GCMs: Biharmonic, Leith, Jansen & Held Deterministic (JHD), and Jansen & Held Stochastic (JHS). Comparison of these eddy-permitting models to a Reference model and to each other over two different eddy-permitting resolutions resulted in no significant differences between models for single-particle dispersion or forecasting. However, there were differences in pair-particle dispersion.

6.1   Single-Particle Climatology

Single-particle dispersion had the longest time scale of the three Lagrangian statistics of interest in this investigation. The results show that all four approximate eddy-permitting models behave similarly to the coarsened eddy-resolving Reference model at both resolutions throughout time. All five models exhibit Lagrangian statistics that suggest Gaussian model behavior — as they should, since the Central Limit Theorem states that the distribution of averages should be Gaussian. The only exception is the kurtosis of trajectory, which is significantly higher than 3, the expected Gaussian kurtosis, in the ballistic regime across all five models. In addition, velocity seems to be less Gaussian than trajectory as velocity of kurtosis is significantly higher than 3 across the entire time scale.

These results indicate that the four eddy-permitting approximate models, once tuned to have the correct Eulerian statistics, accurately approximate the Reference model in the ballistic

and diffusive regimes. All four approximate models perform similarly in this long-term behavior and there are no significant outliers that perform better or worst than the others.

6.2 Particle-Pair Climatology

Particle-pair dispersion was analyzed on a smaller time scale in the ballistic regime in this investigation. The results showed that the four models do perform differently in this shorter-term, smaller-scale behavior: Leith has a significantly larger dispersion than the other models at the finer resolution of 144×144 and JHD is clearly the most accurate to the Reference model at the coarser resolution of 96×96. This relates back to tuning the energy of the models to achieve the correct Eulerian statistics: the Leith model has far too much energy at the small scales and therefore performs poorly at this small-scale Lagrangian statistic while JHD achieves relatively similar energies at the small scales to the Reference model and therefore performs well in particle-pair dispersion.

Understanding how these models perform for particle-pair dispersion is important because this type of behavior is of interest in search and rescue situations. When something disappears in the ocean, its last known position may not be completely accurate. Eddy-permitting approximate models that perform pair-dispersive behavior better would be favorable in this situation since they are more accurate in determining where the missing item is even when there is error in its last known location. In search and rescue, narrowing the search radius by even a small amount can contribute significantly to a successful retrieval.

6.3 Forecasting

The Lagrangian statistics of forecasting were also run in the ballistic regime because forecasts become useless around the time that the diffusive regime begins and particle movement becomes a random walk. Results from comparing the trajectories of the four eddy-permitting approximate models to the coarsened Reference model at both resolutions found no significant differences. The four approximate models perform forecasting behavior similarly, with no one model performing

significantly better or worse than the others. This can be attributed to the fact that velocity fields were initialized accurately and the model errors — the main differences between the models — do not have enough time to accumulate in the forecasting time scale to exhibit significant differences. Forecasting accuracy depends more on the accuracy of the initial conditions than the specific model used.

Evaluating forecast performance is valuable in knowing if models can predict where a particle will be given its initial location and, if so, how accurate that prediction is. Forecasting, like particle-pair dispersion, is also important in situations like search and rescue, but it is also important in tracking the movement of spilled pollutants such as oil or nuclear waste. Knowing where something in the ocean is going and being able to predict its position accurately is critical in these aforementioned situations and in larger climate-change forecasts.

The conclusions of this investigation aim to show that the four approximate eddy-permitting models — Biharmonic, Leith, Jansen & Held Deterministic, and Jansen & Held Stochastic — perform quite similarly and quite accurately in the Lagrangian statistics of single-particle dispersion, particle-pair dispersion, and forecasting. No single model stands out as significantly more accurate across the board, although Leith does exhibit the largest particle-pair dispersion at lower resolutions and Jansen & Held Deterministic is nearly-identical to the Reference model for particle-pair dispersion at lower resolutions. Besides those two cases, no one eddy-permitting model is better suited to approximating the standard quasigeostrophic eddy-resolving model.

Ultimately, these results aim to inform that there is no clearly preferable model out of the four considered here in approximating all three of these Lagrangian behaviors. This investigation hopes to convey that further work and improvements on eddy-permitting GCMs can be done to create distinctly preferable models that are advantageous in scenarios involving different types of Lagrangian behavior.

Bibliography

[1] A Arakawa. Computational design for long-term numerical integration of the equations of fluid motion: Two-dimensional incompressible flow. Part I. J. Comp. Phys., 1(1):119–143, 1966.

[2] Scott D Bachman, Baylor Fox-Kemper, and Brodie Pearson. A scale-aware subgrid model for quasi-geostrophic turbulence. J. Geophysical Research: Oceans, 122(2):1529–1554, 2017.

[3] GK Batchelor. Diffusion in a field of homogeneous turbulence: II. The relative motion of particles. Proc. Cambridge Philos. Soc., 48(2):345–362, 1952.

[4] Claus W Böning and Reinhard G Budich. Eddy dynamics in a primitive equation model: Sensitivity to horizontal resolution and friction. J. Phys. Ocean., 22(4):361–381, 1992.

[5] Dudley B Chelton, Michael G Schlax, and Roger M Samelson. Global observations of nonlinear mesoscale eddies. Prog. Oceanogr., 91(2):167–216, 2011.

[6] B. Fox-Kemper and D. Menemenlis. Can large eddy simulation techniques improve mesoscale rich ocean models? In Ocean modeling in an eddying regime, volume 177 of Geophysical Monograph Series, pages 319–337. American Geophysical Union, 2008.

[7] Jonathan Pietarila Graham and Todd Ringler. A framework for the evaluation of turbulence closures used in mesoscale ocean large-eddy simulations. Ocean Model., 65:25–39, 2013.

[8] Ian Grooms, Yoonsang Lee, and Andrew J Majda. Ensemble filtering and low-resolution model error: covariance inflation, stochastic parameterization, and model numerics. Monthly Weather Review, 143(10):3912–3924, 2015.

[9] Ian Grooms, Yoonsang Lee, and Andrew J Majda. Numerical schemes for stochastic backscatter in the inverse cascade of quasigeostrophic turbulence. Multiscale Modeling & Simulation, 13(3):1001–1021, 2015.

[10] Ian Grooms and Andrew J. Majda. Efficient stochastic superparameterization for geophysical turbulence. Proc. Natl. Acad. Sci. USA, 110, 2013.

[11] Ian Grooms and Andrew J Majda. Stochastic superparameterization in quasigeostrophic turbulence. J. Comp. Phys., 271, 2014.

[12] Ian Grooms and Laure Zanna. A note on toward a stochastic parameterization of ocean mesoscale eddies. Ocean Model., 113:30–33, 2017.

[13] Malte F Jansen and Isaac M Held. Parameterizing subgrid-scale eddy effects using energetically consistent backscatter. Ocean Model., 80:36–48, 2014.

[14] Malte F Jansen, Isaac M Held, Alistair Adcroft, and Robert Hallberg. Energy budget-based backscatter in an eddy permitting primitive equation model. Ocean Model., 94:15–26, 2015.

[15] Shane R Keating, K Shafer Smith, and Peter R Kramer. Diagnosing lateral mixing in the upper ocean with virtual tracers: Spatial and temporal resolution dependence. J. Phys. Ocean., 41(8):1512–1534, 2011.

[16] V Kitsios, JS Frederiksen, and MJ Zidikheri. Scaling laws for parameterisations of subgrid eddy–eddy interactions in simulations of oceanic circulations. Ocean Model., 68:88–105, 2013.

[17] CE Leith. Stochastic models of chaotic systems. Physica D, 98(2):481–491, 1996.

[18] Brodie Pearson, Baylor Fox-Kemper, Scott Bachman, and Frank Bryan. Evaluation of scale-aware subgrid mesoscale eddy models in a global eddy-rich model. Ocean Model., 115:42–58, 2017.

[19] Robert Sadourny and Claude Basdevant. Parameterization of subgrid scale barotropic and baroclinic eddies in quasi-geostrophic models: Anticipated potential vorticity method. J. Atmos. Sci., 42(13):1353–1363, 1985.

[20] AJ Semtner and Y Mintz. Numerical simulation of the gulf stream and mid-ocean eddies. J. Phys. Ocean., 7(2):208–230, 1977.

[21] Geoffrey I Taylor. Diffusion by continuous movements. Proc. London Math. Soc., 2(1):196–212, 1921.

[22] Andrew F Thompson and William R Young. Scaling baroclinic eddy fluxes: Vortices and energy balance. J. Phys. Ocean., 36(4):720–738, 2006.

[23] Geoffrey K Vallis and Bach-lien Hua. Eddy viscosity of the anticipated potential vorticity method. J. Atmos. Sci., 45(4):617–627, 1988.

[24] Jeffrey B Weiss and Ian Grooms. Assimilation of ocean sea-surface height observations of mesoscale eddies. Chaos: An Interdisciplinary Journal of Nonlinear Science, 27(12):126803, 2017.