

Data-driven Equation Discovery of Ocean Mesoscale Closures

Laure Zanna¹ and Thomas Bolton²

¹New York University

²University of Oxford

November 26, 2022

Abstract

The resolution of climate models is limited by computational cost. Therefore, we must rely on parameterizations to represent processes occurring below the scale resolved by the models. Here, we focus on parameterizations of ocean mesoscale eddies and employ machine learning (ML), namely relevance vector machines (RVM) and convolutional neural networks (CNN), to derive computationally efficient parameterizations from data, which are interpretable and/or encapsulate physics. In particular, we demonstrate the usefulness of the RVM algorithm to reveal closed-form equations for eddy parameterizations with embedded conservation laws. When implemented in an idealized ocean model, all parameterizations improve the statistics of the coarse-resolution simulation. The CNN is more stable than the RVM such that its skill in reproducing the high-resolution simulation is higher than the other schemes; however, the RVM scheme is interpretable. This work shows the potential for new physics-constrained interpretable ML turbulence parameterizations for use in ocean climate models.

Data-driven Equation Discovery of Ocean Mesoscale Closures

Thomas Bolton¹, Laure Zanna^{1,2}

¹Department of Physics, University of Oxford, Oxford OX1 3PU, United Kingdom.

²Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

Key Points:

- We present two machine learning algorithms for ocean mesoscale parameterizations.
- We discover closed-form equations for eddy momentum, temperature and energy parameterizations.
- Deep learning closure is more stable than closed-form equations when implemented in an ocean model.

Abstract

The resolution of climate models is limited by computational cost. Therefore, we must rely on parameterizations to represent processes occurring below the scale resolved by the models. Here, we focus on parameterizations of ocean mesoscale eddies and employ machine learning (ML), namely relevance vector machines (RVM) and convolutional neural networks (CNN), to derive computationally efficient parameterizations from data, which are interpretable and/or encapsulate physics. In particular, we demonstrate the usefulness of the RVM algorithm to reveal closed-form equations for eddy parameterizations with embedded conservation laws. When implemented in an idealized ocean model, all parameterizations improve the statistics of the coarse-resolution simulation. The CNN is more stable than the RVM such that its skill in reproducing the high-resolution simulation is higher than the other schemes; however, the RVM scheme is interpretable. This work shows the potential for new physics-constrained interpretable ML turbulence parameterizations for use in ocean climate models.

Plain Language Summary

The complexity of numerical models used for future climate projections is limited by their computational cost. Many key processes, such as ocean eddies, are not adequately resolved and must be approximated using parameterizations. However, parameterizations are often imperfect and reduce the accuracy of the simulations. Machine learning is now opening new avenues to improve climate simulations by extracting such parameterizations directly from data, rather than using idealized theories as typically done. We show that efficient modern machine learning algorithms can accurately represent the physics of ocean eddies, be constrained by physical laws, and can be interpretable. Our results simultaneously open the door to the discovery of new physics from data and the improvement of climate simulations.

1 Introduction

Turbulent processes are critical components of the climate system and influence the circulation of both the ocean and atmosphere. For example, ocean mesoscale eddies, which are turbulent features of scale 10-100 km, dominate the oceanic kinetic energy reservoir (Ferrari & Wunsch, 2009) and are key for the lateral and vertical transport of tracers, such as heat, carbon, and oxygen. These turbulent processes occur on scales that are below the resolution of typical global climate models, which is roughly 25 km-100 km (IPCC, 2013). Therefore, the effects of these turbulent processes on the large-scale must be approximated.

These approximations, called parameterizations or closures, are often developed using idealized theories of the bulk effect of the subgrid process on the large scale (Warner, 2010). This approach has been used for many decades but is not necessarily optimal as it neglects certain physical effects. Imperfections in current parameterizations and missing physics in climate models introduce significant biases in simulations and considerable uncertainty in anthropogenic climate change projections (IPCC, 2013). For example, current parameterizations of ocean eddies target the effect of i) buoyancy fluxes by removing large-scale available potential energy (Gent & McWilliams, 1990), and ii) momentum fluxes using viscous closures which dissipate momentum (Zanna et al., 2020).

While improving certain properties of the flow (Danabasoglu et al., 1994), eddy parameterizations are missing key energy pathways such as the conversion of available potential energy into subgrid kinetic energy, or the backscatter of kinetic energy to the large-scale flow (Jansen et al., 2015; Zanna et al., 2017; Bachman, 2019). In addition, these parameterizations spuriously dissipate kinetic energy (Jansen & Held, 2014; Kjellsson

61 & Zanna, 2017). These imperfect representations of ocean eddy physics in models can
 62 affect the strength and variability of large-scale ocean currents and ocean heat uptake
 63 (Zanna et al., 2017; Kuhlbrodt & Gregory, 2012). Increasing resolution can reduce some
 64 of these biases; however, due to the computational expense of running turbulence-resolving
 65 simulations, subgrid parameterizations will be in demand for several decades.

66 Recently, the advent of machine learning (ML) has given rise to a new class of data-
 67 driven parameterizations. Studies rely on ML to optimally tune parameters of existing
 68 closures (Schneider et al., 2017; Ling et al., 2016). This approach, while useful, still ne-
 69 glects the missing physics not encapsulated in the current parameterizations. Instead,
 70 several studies have shown the promise of new ML parameterizations of subgrid processes
 71 in the atmosphere (Gentine et al., 2018; Rasp et al., 2018; O’Gorman & Dwyer, 2018;
 72 Brenowitz & Bretherton, 2018) and ocean (Bolton & Zanna, 2019). However, this new
 73 class of ML parameterizations often uses black-box algorithms (e.g., neural networks)
 74 such that the laws of physics are not necessarily respected unless imposed (Beucler et
 75 al., 2019; Ling et al., 2016), and interpreting the data-driven parameterization becomes
 76 intractable.

77 Here, we propose a complementary or alternative route to both the traditional physics-
 78 driven bulk approach and the ML-black box approach of deep learning. We use ML to
 79 discover closed-form equations for mesoscale eddy parameterizations for coarse-resolution
 80 ocean models using high-resolution model data. Given some spatio-temporal dataset of
 81 the subgrid eddy forcing, we uncover an equation that could have produced that dataset
 82 (Rudy et al., 2017; Zhang & Lin, 2018). This approach has the following advantages over
 83 more complex methods such as convolutional neural networks: the end result is signif-
 84 icantly easier to interpret physically, the computational cost of implementation is lower,
 85 and training time of the algorithm is also lower. Data-driven discovery of equations has
 86 been extensively used to reveal known-equations, such as Burger’s or Navier-Stokes’ equa-
 87 tions (Kutz, 2017). However, unlike in these studies, we use the algorithm to discover
 88 unknown equations for the subgrid eddy forcings.

89 2 Data and Methods

90 2.1 Training Data and Coarse-Graining

91 We use a primitive equation model, MITgcm (J. Marshall et al., 1997), to gener-
 92 ate high-resolution data and construct new eddy momentum, temperature and energy
 93 parameterizations. We run highly-idealized double-gyre eddy-resolving barotropic and
 94 baroclinic simulations in a square-domain. The simulations use a beta-plane approxima-
 95 tion, free-slip boundary conditions on lateral walls and no-slip boundary condition on
 96 the bottom, and a constant surface zonal wind forcing. These simulations are designed
 97 to create highly turbulent flow regimes, with mesoscale eddies shedding from the jet ex-
 98 tension.

99 The barotropic model has a single layer of depth 500 m and length 3840 km, sim-
 100 ilar to Cooper and Zanna (2015). We spin-up the model from rest for 10 years, at a spa-
 101 tial resolution of 3.75 km. The baroclinic model comprises of 15 vertical levels, with a
 102 total depth of 3600 m. Due to the increased computational cost of running the baroclinic
 103 simulation compared to the barotropic model, we decreased the domain size from 3840
 104 km in length to 1920 km, with a spatial resolution of 7.5 km. The meridional temper-
 105 ature gradient is imposed via surface restoring to a linear profile. We spin-up the baro-
 106 clinic model for 100 years and then run for a further 10 years for data collection. Fur-
 107 ther details about the simulations are given in the Supplementary Information (SI, S1).

108 After spin-up, we select 1000 time-slices of model output, with 4 days between each
 109 time-slice. We remove information at small-scales by applying a horizontal Gaussian fil-
 110 ter of width 30 km, and then coarse-grain to a 30 km grid, which is denoted by $(\bar{\cdot})$ (Bolton

111 & Zanna, 2019) (SI, S2). The subgrid eddy momentum and temperature forcing terms,
 112 for each vertical level, are then defined by

$$113 \quad \mathbf{S}_u = \begin{pmatrix} S_x \\ S_y \end{pmatrix} = (\bar{\mathbf{u}} \cdot \bar{\nabla}) \bar{\mathbf{u}} - \overline{(\mathbf{u} \cdot \nabla) \mathbf{u}}, \quad (1)$$

$$114 \quad S_T = (\bar{\mathbf{u}} \cdot \bar{\nabla}) \bar{T} - \overline{(\mathbf{u} \cdot \nabla) T}, \quad (2)$$

115 respectively. Here ∇ is the horizontal 2D gradient operator, T is the temperature, and
 116 the horizontal velocity $\bar{\mathbf{u}} = (\bar{u}, \bar{v})$. These terms reflect the turbulent nonlinear terms
 117 truncated in coarse-resolution models which need to be parameterized (Berloff, 2005; Mana
 118 & Zanna, 2014). At every grid-point for every time-slice, we both i) calculate the tar-
 119 get eddy forcing, i.e, Eqs. (1) and (2), and ii) construct a library of diverse functions
 120 which are necessary for the RVM method described below and are relevant to the pro-
 121 cess being parameterized.

122 2.2 Data-Driven Algorithms

123 **Relevance Vector Machine.** Here, we employ the sparse Bayesian regression method
 124 used in Zhang and Lin (2018) based on relevance vector machines (RVM) (Tipping, 2001)
 125 to reveal new eddy parameterizations. RVM is a regression technique that assumes Gaus-
 126 sian prior distributions for each regression weight (Bishop, 2006). The width of the Gaus-
 127 sian prior of each regression weight provides a measure of uncertainty of that regression
 128 weight. The method relies on a library of functions, which can comprise of any function
 129 such as products or derivatives of relevant quantities defined as basis functions (e.g., ve-
 130 locity shears, temperature shears). The sparse regression is applied iteratively to the li-
 131 brary of functions, and then a pruning of the library of functions is carried out by dis-
 132 carding the functions with an uncertainty higher than a pre-specified threshold (Zhang
 133 & Lin, 2018). This uncertainty threshold, δ , is the only parameter that requires setting
 134 in the Zhang and Lin (2018) method. The algorithm finishes when the uncertainty mea-
 135 sures of each regression weight stop changing from iteration to iteration. We found the
 136 Zhang and Lin (2018) method to be more robust than the sequential threshold ridge re-
 137 gression (STRidge) of Rudy et al. (2017). For example, using data to discover the known
 138 2D advection-diffusion equations, we found that STRidge required substantially more
 139 data for training than the RVM method, STRidge has a large number of tunable hyper-
 140 parameters which substantially influenced the discovered equation compared to the RVM
 141 method which has only one hyperparameter. In addition, unlike STRidge, Zhang and
 142 Lin (2018) method provided an error associated with the weights discovered. Given these
 143 tests were performed on known equations in which we knew the answers, we opted for
 144 the use of Zhang and Lin (2018) RVM method to discover unknown parameterizations.

145 At every grid-point for every time-slice from the MITgcm coarse-grained output
 146 (described above) we construct a library of diverse functions, ϕ_i , which are derived from
 147 a set of basis functions relevant to the process being parameterized. We build the library
 148 from the filtered velocities \bar{u} , \bar{v} , and \bar{T} using up to second-order for both spatial deriva-
 149 tives and polynomial products, mainly due to memory limitations. The basis of func-
 150 tions used for the momentum and temperature eddy parameterizations differ and will
 151 be discussed in the next section. We normalized each function individually such that they
 152 have zero mean and unit variance. We use 50% of the 1000 time-slices for training and
 153 the other 50% for validation. For both the eddy momentum and temperature forcing,
 154 we impose a physical constraint for global conservation. To do so, we only specify library
 155 functions that can be written as the divergence of a flux (or as the divergence of a ten-
 156 sor \mathbf{T} for the eddy momentum forcing, i.e. $\bar{\nabla} \cdot \mathbf{T}$), such that with the appropriate bound-
 157 ary conditions there is no net input of momentum or temperature.

158 We then apply the iterative RVM algorithm to prune the library of functions and
 159 construct the final equation for the subgrid forcing (independently for S_x , S_y and S_T)
 160 as a linear sum of the functions, ϕ_i , each weighted by the regression coefficient, w_i . We

Figure 1. A) Illustration of the RVM procedure; B) Schematic of the architecture of the physics-constrained fully-convolutional neural network (FCNN); C) Online validation of the sub-grid momentum forcing from the barotropic simulations for three parameterizations, denoted as \hat{S} { the physics-driven \hat{S}^{AZ} , \hat{S}^{BT} revealed by the RVM algorithm (Eq. 5), and the FCNN { against the diagnosed forcing from high-resolution data, S . Top Row shows the mean [ms^{-2}], Middle Row the Standard Deviation [ms^{-2}], and the Bottom Row the Pearson correlation of the zonal component of the eddy momentum forcing, S_x and \hat{S}_x (the meridional component is shown in SI). The x- and y-axis are longitude and latitude, respectively; the extent is 3840 km in each direction.