

Deep ocean learning of small scale turbulence

Ali Mashayek¹, Nick Reynard¹, Fangming Zhai¹, Kaushik Srinivasan², Adam
Jelley³, Alberto Naveira Garabato⁴, Colm-cille P. Caulfield⁵

¹Imperial College London, UK

²University of California Los Angeles, USA

³University of Edinburgh, UK

⁴University of Southampton, UK

⁵University of Cambridge, UK

Key Points:

- Machine learning can be used to infer ocean turbulent mixing from basic seawater and geometric properties.
- The machine learning models are trained based on limited available direct turbulence measurements.
- The trained models can be applied to data from global observational programs, which do not sample turbulence directly.

Corresponding author: Ali Mashayek, mashayek@ic.ac.uk

Abstract

Turbulent mixing at the sub-meter scale is an essential component of the ocean’s meridional overturning circulation and its associated global redistribution of heat, carbon, nutrients, pollutants and other tracers. Whereas direct turbulence observations in the ocean interior are limited to a modest collection of field programs, basic information such as temperature, salinity and depth is available globally. Here, we show that supervised machine learning algorithms can be trained on the existing turbulence data to develop skillful predictions of the key properties of turbulence from T, S, Z and topographic data. This constitutes a promising first step toward a hybrid physics-artificial intelligence approach to parameterization of turbulent mixing in climate models.

Plain Language Summary

Ocean turbulence plays an important role in sustaining the general ocean circulation and in the mixing of heat, carbon, nutrients, and other processes within the ocean interior. Turbulent mixing is technically challenging to measure and is often inferred from measurable quantities using parameterizations that are based on numerous simplifying assumptions about the physics of turbulence. In this study, we show that artificial intelligence (more specifically, various machine learning algorithms) can be successfully employed to infer turbulent mixing from quantities measured routinely by global observational programs.

Introduction

Turbulent mixing across density surfaces (i.e. diapycnal mixing) in the ocean interior is key to sustaining the meridional overturning circulation and its global regulation of heat, carbon and nutrient distributions, as well as other climatically and environmentally important tracers (Talley et al., 2016). Such turbulence is primarily excited at the ocean surface by winds, or at the bottom boundary via flow impingement on topography (Garabato & Meredith, 2022). The spatio-temporal variability of turbulence makes its measurement especially challenging. However, turbulence can leave an imprint on vertical temperature (T) and salinity (S) profiles obtained from hydrographic surveys. T, S and depth (Z) are regularly sampled through global international programs, such as ship-based efforts like WOCE (Gouretski & Koltermann, 2004), GO-SHIP (GO-SHIP, 2018), GEOTRACES (GEOTRACERS, 2019) or globally-distributed floats deployed by the Argo Program (Argo, 2000) (see Supplementary Materials for a visual summary, and Davis *et al.* (2019) for a review (Davis et al., 2019)). While turbulence characteristics may be inferred from these T, S, Z data (Polzin et al., 2014; Whalen et al., 2012), such estimates involve many assumptions and uncertainties.

The gold standard in measuring turbulence in the ocean interior is represented by ship-deployed microstructure profiler observations, which include concurrent sampling of T, S and Z , but are limited in number due to their technical complexity and cost (Shroyer et al., 2018). In this study, we train machine learning models on a unique collection of observations from microstructure field programs enabling prediction of turbulence characteristics based on T, S, Z and topographic data, rendering our approach applicable to major global surveys that do not measure turbulence directly. Our aim is to demonstrate that such predictions from microstructure-trained physics-inspired machine-learning models yield better estimates for dynamically-significant quantities than classical finestructure parameterizations.

Physics of Turbulence

A key property of density-stratified ocean turbulence is the significantly enhanced rate (as compared to molecular diffusion) at which it mixes density and tracers in the vertical (Thorpe, 2005). In observations and climate models, such mixing is often encapsulated

63 in a turbulent diffusion coefficient (or diffusivity for short) defined as

$$\kappa = \Gamma \frac{\varepsilon}{N^2}, \quad (1)$$

64 where Γ is a coefficient that determines the fraction of the energy available to turbulence
 65 that contributes to mixing (Peltier & Caulfield, 2003), ε is the rate of dissipation of turbulent
 66 kinetic energy (due to viscosity of seawater)¹, and $N = \sqrt{-[g/\rho_0]\partial\rho/\partial z}$ is the buoyancy
 67 frequency (Osborn, 1980). While Γ is known to be variable (Mashayek & Peltier, 2013;
 68 Mashayek et al., 2017; Gregg et al., 2018), for the purpose of this study it suffices to consider
 69 it a constant, specifically 0.2, in line with operational physical oceanography (Gregg et al.,
 70 2018; Mashayek et al., 2021). N can be directly inferred from measurements of T, S and Z
 71 through the construction of the vertical density gradient, a characteristic density ρ_0 , and the
 72 gravitational acceleration g . On the other hand, the dissipation rate ε , as it is determined
 73 from the strain-rate tensor, cannot be inferred from T, S, Z (which are available from global
 74 observational programs) and is best inferred from microstructure profilers, that measure
 75 spatial gradients of velocity. In this study, we show that machine learning models can be
 76 trained on microstructure data to predict ε (or directly predict κ) based on T, S, Z , and
 77 height above the bottom (Hab).² This allows for global inference of κ from observational
 78 surveys, thereby providing a route for direct application to climate models that assimilate
 79 data from such surveys (e.g. Forget et al. (2015); Verdy and Mazloff (2017)).

80 The Training Dataset

81 We employ a global dataset of microstructure profiles compiled by the Climate Process
 82 Team on internal wave-driven ocean mixing (MacKinnon et al., 2017). Fig. 1 shows the lo-
 83 cation of the field measurements, spanning a wide range of geographic locations, depths, and
 84 turbulence-inducing physical processes. A sample microstructure transect from the DIMES
 85 experiment is shown in panel *c* (more specifically, transect T1 in Fig. 5a). Fig. 1 also
 86 provides the list of the field experiments, and the fraction of the total data associated with
 87 each experiment. The data are available at <https://microstructure.ucsd.edu/>, and data
 88 description and relevant references may be found in Waterhouse et al. (2014). The same
 89 dataset was employed by Cael and Mashayek (2021) to show that the data ‘collapses’ on
 90 a seemingly universal log-skew-normal statistical distribution. This finding motivated the
 91 present study by suggesting that such universality might be detectable through data-driven
 92 methods. Together, the experiments in Fig. 1 contain over 700 full-depth microstructure
 93 profiles, binned into 10 m vertical bins (amounting to $\sim 2 \times 10^5$ data points). The concurrent
 94 measurements of ε, T, S, Z in this dataset allow for the construction of the aforementioned
 95 predictor list (i.e. the list of features used in training) used to predict ε and κ . More specifi-
 96 cally, neutral density is calculated from T, S, Z , latitude, and longitude information (Jackett
 97 & McDougall, 1997), the local depth for each profile is looked up from the global bathymet-
 98 ric map of Sandwell et al. (2014), and height above is then calculated by subtracting the
 99 sample depth from the local depth.

¹ More precisely, $\varepsilon = \nu \frac{\partial u'_i}{\partial x_j} \frac{\partial u'_i}{\partial x_j}$, where u'_i represent perturbation velocity components (i.e. departures from the mean flow), x_i represents the three Cartesian dimensions, and the overbar represents an ‘appropriate’ averaging.

² We found that inclusion of both Z and Hab is crucial as they represent the distance from the top and bottom boundaries, both of which are turbulence generation sites. Knowledge of Hab requires topographic data, which has become increasingly more accurate in recent decades thanks to advanced satellite-based gravity measurements and deep-ocean echo-sounding records (Sandwell et al., 2014) (see Supplementary Materials Fig. S1).

Machine Learning Models

Fig. 2 illustrates the overall flowchart for the research presented here: assembling the training datasets (as shown in Fig. 1); training two machine learning models with distinctly different underlying algorithms; assessing the models' skills (as shown in Fig. 3); and independent verification of the models through their application to individual field programs (as shown in Figs 4 and 5). This section describes the construction of the two machine learning algorithms.

Classification And Regression Trees (CART)

We employ CART, one of the most common machine learning predictive models (Wu et al., 2008; Breiman et al., 1984). The method uses a decision tree to connect observations of a parameter of interest (represented in the branches) to predictions about its value (represented in the leaves). When applied to target variables that take continuous values (such as ε or κ in this study), such decision trees are referred to as regression trees. Additionally, we employ an ensemble method, bootstrap aggregating, to improve the stability and accuracy of the decision tree algorithm, reduce variance, and avoid overfitting. Bootstrap aggregated decision trees (hereafter bagging trees) construct multiple trees by repeatedly re-sampling the training data with replacements, and voting the trees for a consensus prediction (Breiman, 1996).

Figs 3a,b show the application of the bagging tree to the training microstructure dataset (from which $\log_{10}(\varepsilon)$, our prediction target, is calculated). The model was trained based on 10 cross-validation k-folds of all data across 13 field experiments. This method involves splitting the dataset into equally sized 'k' number of groups, or 'folds', and taking it in turn to use each group as the test data while the rest of the data is used to train the model, with an average of the results being adopted. A k-fold validation approach is useful when input data is limited, and ensures that every data point is used within the training and test dataset, hence reducing bias when compared to other methods. The fits in Figs 3a,b are satisfactory, with a coefficient of determination (R^2) of 0.83 for $\log_{10}(\varepsilon)$ and 0.84 for $\log_{10}(\kappa)$.³ To analyze further the quality of the agreement between predictions and data, panels *e*–*h* display the cumulative contribution of various predictors to increases in R^2 and decreases in the mean squared error (MSE).

We consider two sets of predictors, with nearly equal skills. First we consider T , S , their gradients, Z , Hab , and latitude. Secondly, we just use $\log_{10}(N^2)$, Z , Hab , and latitude.⁴ While the two sets show similar skills, it is worth noting that the former contains more raw information about the temperature and salinity structures (which get combined into one parameter once N is calculated). It is conceivable that in regions of the ocean where salinity structures play a key role in turbulence generating processes (e.g. double diffusion in the Arctic Ocean; Middleton et al. (2021)), retaining T , S and their derivatives may prove fruitful. We postpone the investigation of application of our methodology to such regions to future work.

It is worth noting that we also tried another standard choice, namely the Least Squares Boost (LSBoost) algorithm, as an alternative ensemble learning method. LSBoost is a gradient boosting method in which the mean squared error is chosen as the cost function (Breiman et al., 1984). While we found LSBoost to outperform bagging tree for a smaller number of features (up to 3), bagging tree was superior for the number of features employed herein,

³ R^2 is a statistical metric of how well the regression predictions approximate the real data, and so is a measure of the goodness of fit of a model.

⁴ Since quantities like ε , κ , and N vary over orders of magnitude, employing their logarithms renders the training algorithms more efficient.

144 and thus is our method of choice. Finally, we note that application of a linear regression
 145 model to the dataset proved entirely futile.

146 **Neural Networks**

147 As an entirely different approach, we also train neural networks with the same data.
 148 Specifically, we use a fully-connected feed-forward neural network (FNN), also referred to
 149 as a Multi-layer Perceptron (MLP) in the broader ML community. Standard FNNs consist
 150 of an input layer, an output layer and multiple hidden layers in between (Goodfellow et
 151 al., 2016) but we use a slight modification of this FNN architecture by making the hidden
 152 layers actually *residual* layers. Unlike standard hidden layers, which recursively perform
 153 operations on the previous layer, residual layers are added on to the main input-to-output
 154 information flow. (Such additions are also referred to as *skip connections* (He et al., 2016).)
 155 NNs with predominantly residual layers, or Resnets, have been found to outperform direct
 156 NNs, not just on the class of problems relevant to this study (Gorishniy et al., 2021), but
 157 across almost all modalities in AI/ML in general (Drozdzal et al., 2016; Vaswani et al.,
 158 2017) and hence residual layers are correspondingly ubiquitous features of most modern
 159 NNs. Each hidden layer combines the (learned) features of the previous layer to build up a
 160 non-linear transformation of the input predictors to predict turbulence properties (ε and κ)
 161 in the output layer. Adding additional layers to make the network deeper incorporates more
 162 parameters to be learned, which allows for a more flexible mapping between the easily mea-
 163 surable predictors and the less widely available turbulence properties. Typically, adding
 164 more parameters requires more data to learn an effective generalizable mapping without
 165 overfitting. However, we use a specific training algorithm called stochastic gradient descent
 166 with warm restarts (Loshchilov & Hutter, 2016) that provides a strong implicit regulariza-
 167 tion, essentially eliminating the issue of overfitting, even in low data regimes. A 10-fold
 168 cross-validation algorithm is used to ensure coverage of the entire dataset using the trained
 169 NN models. The Supplementary Materials contain more information on the resnet-FNN
 170 model architecture as well as details of the training and optimization procedure.

171 Figs 3c,d show that the deep neural network is also skillful in predicting both ε and
 172 κ . Deep learning algorithms like neural networks require less human intervention compared
 173 to more traditional machine learning algorithms (e.g. the bagging tree), and so generally
 174 have larger data requirements and their performance increases more strongly with the size
 175 of data. This makes the high R^2 values for NNs in Fig 3b particularly promising, given
 176 the limited nature of the training data compared to data sizes typically employed in deep
 177 learning. Thus, investment in extending the training data through a community effort
 178 appears worthwhile. We note that while CART seems to give a slightly higher R^2 than NN
 179 in Fig. 3, as we will show next, NN proves more skillful in capturing the vertical patterns
 180 of turbulence for individual experiments when they are considered separately.

181 **Application to Individual Datasets**

182 Fig. 4 shows the results of separate analyses for each of the 13 different field programs
 183 listed in Fig 1. Importantly, the data from each experiment are excluded from training
 184 of the models before the models are applied to it. While both NN and CART show skills
 185 in predicting the patterns and, in some cases, the order of magnitude adequately, NN is
 186 clearly superior in both respects. While Fig. 4 shows predictions based on models trained
 187 to infer κ directly, we have also repeated the exercise based on models trained to predict ε ,
 188 and then inferred κ from that prediction using Eq. 1 with $\Gamma = 0.2$. The outcome, shown
 189 in Supplementary Materials Fig. S2, is qualitatively similar, although, importantly, the
 190 direct prediction of κ is more skilled at predicting the turbulence-induced diffusivity in
 191 the vicinity of seafloor. This superiority of ‘direct’ estimation of diffusivity is significant since
 192 such turbulent mixing is key to the upwelling of the deep waters formed and sunk at high

193 latitudes, a process necessary for closure of the oceanic meridional overturning circulation (de
194 Lavergne et al., 2022).

195 Indirect inferences of turbulent mixing from T and S finestructure, practically the only
196 alternative when microstructure data is unavailable, can be inaccurate by as much as two
197 orders of magnitude (Polzin et al., 2014). Furthermore, such parameterizations are based
198 on somewhat restrictive assumptions regarding the nature of the underlying turbulence-
199 generating processes. Thus, the accuracy of NN showcased in Fig. 4, in light of its ag-
200 nosticism towards the underlying physics, is appealing. To further highlight this point, in
201 Fig. 5 we assess the skill of CART and NN for the data sampled along three transects
202 (shown in Fig. 5a) as a part of the DIMES experiment. For these transects, both direct
203 (from microstructure profilers) and indirect finestructure-based parameterizations of ε and
204 κ were reported in Sheen et al. (2013), allowing for testing our models against conventional
205 finestructure parameterizations (Figs 5b-g). Both NN and CART outperform the finestruc-
206 ture parameterization, particularly for κ (which is ultimately the parameter of interest). It
207 is worth noting that the study of Sheen et al. (2013) is one of the more successful appli-
208 cations of finestructure parameterization; examples of much larger disagreements between
209 finestructure and microstructure estimates abound in the literature.

210 Discussion & Outlook

211 The primary message of this study is that AI can indeed be successfully employed to
212 use data from global observational programs, which lack direct turbulence measurements, to
213 predict small scale turbulent mixing in the ocean, and in particular, more accurately than
214 conventional finestructure parameterizations. More specifically, this study implies that the
215 knowledge of parameters most basic to turbulence, i.e. finescale density stratification, dis-
216 tance from turbulence-generating boundaries, and latitude, suffice to leading order to obtain
217 an estimate of the turbulence intensity and the associated turbulent (density) diffusivity.

218 There are numerous factors that can contribute to the misfits between the predictions
219 and the data. Three important ones are: (I) the percentage of the training and validation
220 data can vary significantly between the experiments (as shown in Fig. 1); (II) the rele-
221 vance of the underlying physics in each experiment to the rest of the data used for training
222 might be limited; (III) ocean mixing is not entirely ‘local’ in nature, e.g. waves generated
223 thousands of kilometers away can contribute to mixing, and no such information was in-
224 cluded in our training by construction (de Lavergne et al., 2019). Factor (I) can only be
225 addressed through application of AI to larger datasets. In particular, the success of deep
226 learning directly scales with the data size, and what was achieved in this study lies at the
227 lower bound of the data volume required. Our analyses show that while the bagging tree
228 algorithm converges to the optimal performance once a few hundreds of profiles are con-
229 sidered, the NN algorithm does not show such convergence and retains a large standard
230 deviation even when all profiles are included. Thus, further community efforts are required
231 to pull turbulence datasets together and subject them to the consistent high levels of qual-
232 ity control and grid interpolations. Furthermore, adding microstructure sensors to global
233 observational endeavors (such as the Argo float program), while ambitious, is within reach
234 and conceivable in the coming decades (Roemmich et al., 2019). Factor (II) will naturally
235 advance as our physical understanding of ocean turbulence keeps progressing. A conscious
236 effort towards connecting such physical understanding to data-driven parameterizations is
237 required. Addressing factor (III) is more readily achievable in the near future, as it will
238 require inclusion of theoretical estimates of local and non-local energy injected to internal
239 waves from various sources (winds, tides, etc.) in training algorithms. In summary, we have
240 demonstrated here that AI provides a valuable tool to harness our observational, theoretical
241 and statistical knowledge of ocean turbulence to direct the development of a next-generation
242 ‘smart’ turbulence parameterization for climate models.

243 Acknowledgments

244 The authors report no conflict of interests. The authors thank Kathy Sheen for providing
 245 the data from Sheen et al. (2013) for the purpose of constructing Fig. 5, and Lois Baker
 246 and two anonymous reviewers for constructive comments.

247 Data Availability Statement

248 The microstructure data employed for the training of the Machine Learning algorithms
 249 may be obtained from <https://microstructure.ucsd.edu/> by locating the names of the
 250 experiments in Figure 3; also see (Waterhouse et al., 2014) and (MacKinnon et al., 2017)
 251 for further information. The AI algorithms will be shared via an online depository upon
 252 acceptance of the manuscript.

253 References

- 254 Argo, G. (2000). Argo float data and metadata from global data assembly centre (argo
 255 gdac). *SEANOE*.
- 256 Breiman, L. (1996). Bagging predictors. *Machine learning*, *24*(2), 123–140.
- 257 Breiman, L., Friedman, J., Olshen, R., & Stone, C. (1984). Classification and regression
 258 trees. *wadsworth int. Group*, *37*(15), 237–251.
- 259 Cael, B., & Mashayek, A. (2021). Log-skew-normality of ocean turbulence. *Physical Review
 260 Letters*, *126*(22), 224502.
- 261 Davis, R. E., Talley, L. D., Roemmich, D., Owens, W. B., Rudnick, D. L., Toole, J., . . .
 262 Barth, J. A. (2019). 100 years of progress in ocean observing systems. *Meteorological
 263 Monographs*, *59*, 3–1.
- 264 de Lavergne, C., Falahat, S., Madec, G., Roquet, F., Nycander, J., & Vic, C. (2019, 5). To-
 265 ward global maps of internal tide energy sinks. *Ocean Modelling*, *137*, 52–75. Retrieved
 266 from <https://linkinghub.elsevier.com/retrieve/pii/S1463500318302890> doi:
 267 10.1016/j.ocemod.2019.03.010
- 268 de Lavergne, C., Groeskamp, S., Zika, J., & Johnson, H. L. (2022). The role of mixing in
 269 the large-scale ocean circulation. *Ocean Mixing*, 35–63.
- 270 Drozdal, M., Vorontsov, E., Chartrand, G., Kadoury, S., & Pal, C. (2016). The importance
 271 of skip connections in biomedical image segmentation. In *Deep learning and data
 272 labeling for medical applications* (pp. 179–187). Springer.
- 273 Forget, G., Campin, J.-M., Heimbach, P., Hill, C. N., Ponte, R. M., & Wunsch, C. (2015).
 274 ECCO version 4: an integrated framework for non-linear inverse modeling and global
 275 ocean state estimation.
- 276 Garabato, A. N., & Meredith, M. (2022). Ocean mixing: oceanography at a watershed. In
 277 *Ocean mixing* (pp. 1–4). Elsevier.
- 278 GEOTRACERS. (2019). Geotraces. <https://www.geotraces.org/>. Retrieved from [https://](https://www.geotraces.org/)
 279 www.geotraces.org/
- 280 Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. ([http://](http://www.deeplearningbook.org)
 281 www.deeplearningbook.org)
- 282 Gorishniy, Y., Rubachev, I., Khrulkov, V., & Babenko, A. (2021). Revisiting deep learning
 283 models for tabular data. *Advances in Neural Information Processing Systems*, *34*.
- 284 GO-SHIP. (2018). Go-ship. <http://www.go-ship.org/>. Retrieved from [http://www.go-ship](http://www.go-ship.org/)
 285 [.org/](http://www.go-ship.org/)
- 286 Gouretski, V., & Koltermann, K. P. (2004). {WOCE} global hydrographic climatology.
 287 *Berichte des BSH*, *35*, 1–52.
- 288 Gregg, M., D’Asaro, E., Riley, J., & Kunze, E. (2018, 1). Mixing Efficiency in the
 289 Ocean. *Annual Review of Marine Science*, *10*(1), 443–473. Retrieved from [http://](http://www.annualreviews.org/doi/10.1146/annurev-marine-121916-063643)
 290 www.annualreviews.org/doi/10.1146/annurev-marine-121916-063643 doi: 10
 291 .1146/annurev-marine-121916-063643
- 292 He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition.
 293 In *Proceedings of the ieee conference on computer vision and pattern recognition* (pp.

- 770–778).
- 294 Jackett, D. R., & McDougall, T. J. (1997, 2). A Neutral Density Variable for the
 295 World’s Oceans. *Journal of Physical Oceanography*, *27*(2), 237–263. doi: 10.1175/
 296 1520-0485(1997)027(0237:ANDVFT)2.0.CO;2
- 297 Loshchilov, I., & Hutter, F. (2016). SGDR: Stochastic gradient descent with warm restarts.
 298 *arXiv preprint arXiv:1608.03983*.
- 299 MacKinnon, J. A., Zhao, Z., Whalen, C. B., Waterhouse, A. F., Trossman, D. S., Sun,
 300 O. M., . . . others (2017). Climate process team on internal wave–driven ocean mixing.
 301 *Bulletin of the American Meteorological Society*, *98*(11), 2429–2454.
- 302 Mashayek, A., Caulfield, C. P., & Alford, M. H. (2021). Goldilocks mixing in oceanic
 303 shear-induced turbulent overturns. *Journal of Fluid Mechanics*, *928*, A1.
- 304 Mashayek, A., & Peltier, W. R. (2013). Shear-induced mixing in geophysical flows: does
 305 the route to turbulence matter to its efficiency? *Journal of Fluid Mechanics*, *725*,
 306 216–261.
- 307 Mashayek, A., Salehipour, H., Bouffard, D., Caulfield, C. P., Ferrari, R., Nikurashin, M., . . .
 308 Smyth, W. D. (2017). Efficiency of turbulent mixing in the abyssal ocean circulation.
 309 *Geophysical Research Letters*, *44*(12), 6296–6306.
- 310 Middleton, L., Fine, E., MacKinnon, J., Alford, M., & Taylor, J. (2021). Estimating dissipa-
 311 tion rates associated with double diffusion. *Geophysical Research Letters*, *48*(15),
 312 e2021GL092779.
- 313 Osborn, T. R. (1980). Estimates of the local rate of vertical diffusion from dissipation
 314 measurements. *Journal of Physical Oceanography*, *10*, 83–89.
- 315 Peltier, W. R., & Caulfield, C. P. (2003). Mixing efficiency in stratified shear flows. *Annual*
 316 *Rev. Fluid Mech.*, *35*, 135–167.
- 317 Polzin, K. L., Garabato, A. C. N., Huussen, T. N., Sloyan, B. M., & Waterman, S. (2014).
 318 Finescale parameterizations of turbulent dissipation. *Journal of Geophysical Research:*
 319 *Oceans*, *119*(2), 1383–1419.
- 320 Roemmich, D., Alford, M., Claustre, H., Johnson, K., King, B., Moum, J., et al. (2019). On
 321 the future of argo: an enhanced global array of physical and biogeochemical sensing
 322 floats. *front. Mar. Sci*, *6*, 439.
- 323 Sandwell, D. T., Müller, R. D., Smith, W. H., Garcia, E., & Francis, R. (2014). New global
 324 marine gravity model from cryosat-2 and jason-1 reveals buried tectonic structure.
 325 *Science*, *346*(6205), 65–67.
- 326 Sheen, K. L., Brearley, J. A., Naveira Garabato, A. C., Smeed, D. A., Waterman, S., Ledwell,
 327 J. R., . . . others (2013). Rates and mechanisms of turbulent dissipation and mixing
 328 in the Southern Ocean: Results from the Diapycnal and Isopycnal Mixing Experiment
 329 in the Southern Ocean (DIMES). *Journal of Geophysical Research: Oceans*, *118*(6),
 330 2774–2792.
- 331 Shroyer, E. L., Nash, J. D., Waterhouse, A. F., & Moum, J. N. (2018). Measuring
 332 Ocean Turbulence. In (pp. 99–122). Springer, Cham. Retrieved from [https://](https://link.springer.com/chapter/10.1007/978-3-319-66493-4_6)
 333 link.springer.com/chapter/10.1007/978-3-319-66493-4_6 doi: 10.1007/978-3-
 334 -319-66493-4{_}6
- 335 Talley, L. D., Feely, R. A., Sloyan, B. M., Wanninkhof, R., Baringer, M. O., Bullister,
 336 J. L., . . . Zhang, J. Z. (2016, 1). Changes in Ocean Heat, Carbon Content, and
 337 Ventilation: A Review of the First Decade of GO-SHIP Global Repeat Hydrography.
 338 *Annual Review of Marine Science*, *8*, 185–215. Retrieved from www.go-ship.org doi:
 339 10.1146/annurev-marine-052915-100829
- 340 Thorpe, S. A. (2005). The turbulent ocean. *Cambridge University Press*.
- 341 Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., . . . Polosukhin,
 342 I. (2017). Attention is all you need. *Advances in neural information processing systems*,
 343 *30*.
- 344 Verdy, A., & Mazloff, M. R. (2017). A data assimilating model for estimating South-
 345 ern Ocean biogeochemistry. Retrieved from <http://hycom.org> doi: 10.1002/
 346 2016JC012650
- 347 Waterhouse, A. F., MacKinnon, J. A., Nash, J. D., Alford, M. H., Kunze, E., Simmons,
- 348

- 349 H. L., ... others (2014). Global patterns of diapycnal mixing from measurements of
350 the turbulent dissipation rate. *Journal of Physical Oceanography*, *44*(7), 1854–1872.
- 351 Whalen, C., Talley, L., & MacKinnon, J. (2012). Spatial and temporal variability of global
352 ocean mixing inferred from argo profiles. *Geophysical Research Letters*, *39*(18).
- 353 Wu, X., Kumar, V., Quinlan, J. R., Ghosh, J., Yang, Q., Motoda, H., ... others (2008).
354 Top 10 algorithms in data mining. *Knowledge and information systems*, *14*(1), 1–37.

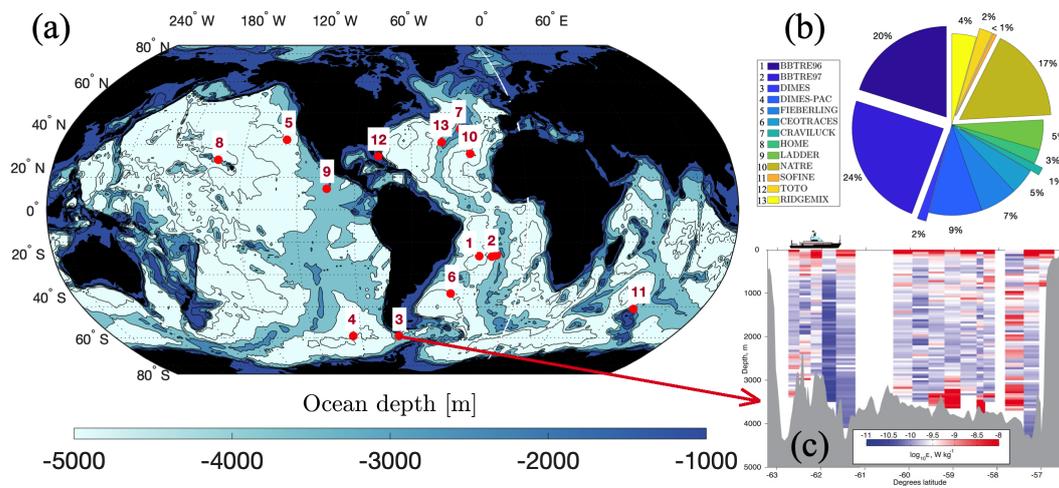


Figure 1. Direct turbulence measurements can be used to train machine learning algorithms to predict turbulent mixing where direct measurements are not available. (a) Location of the field programs that include direct measurements of turbulence (specifically, turbulent kinetic energy dissipation rate ϵ from microstructure profilers) along with co-located temperature, salinity and depth sampling. (b) The experiments' name and associated contributions to the total data. More details about the data sources are available at <https://microstructure.ucsd.edu/> (Waterhouse et al., 2014) and in the Supplementary Materials. The data contains a total of ~ 700 profiles, with ϵ binned into 10 m vertical bins. (c) A sample transect of microstructure data from the DIMES experiment (transect T1 in Fig. 5); from Sheen et al. (2013).

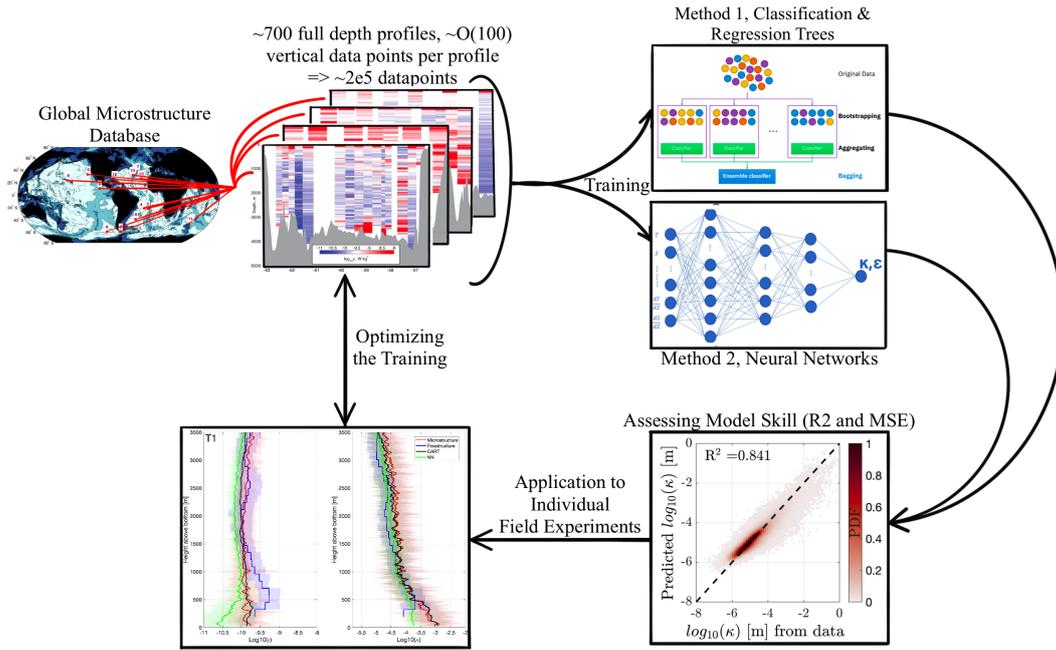


Figure 2. Two distinctly different machine learning algorithms can successfully reproduce turbulent mixing estimates in agreement with microstructure data.

A flowchart, illustrating the sequence of data assembly, training, model skill assessment, application to original data, verification, and fine-tuning. The source for the sample microstructure profiles shown in the top row is Sheen et al. (2013)– see Fig. 5 for details. The source for the CART diagram in the top row is https://en.wikipedia.org/wiki/Bootstrap_aggregating. Note that CART and NN are not applied sequentially, but are independent algorithms.

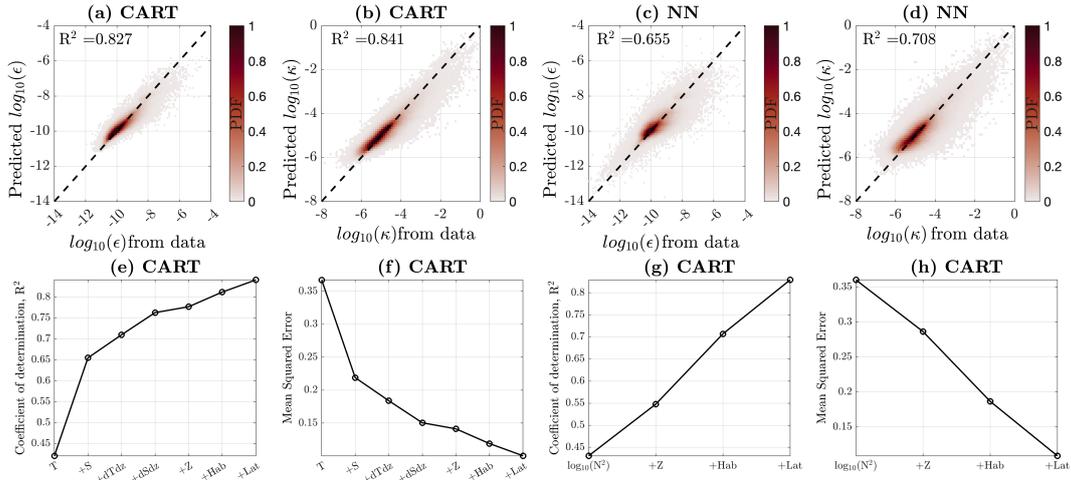


Figure 3. Machine learning can successfully fit the global microstructure data based on few predicting features.

(a-d) Bivariate histograms (in the form of a probability density function, PDF) of predicted rate of dissipation of turbulent kinetic energy (ϵ ; [m^2/s^3]) and turbulent diffusivity (κ ; [m^2/s]) based on use of the Classification And Regression Tree algorithm (CART; panels *a,b*), and Neural Networks (NN; panels *c,d*), versus the actual data. Both CART and NN models are validated using k-fold validation with 10 folds to avoid overfitting (see main text). (e-f) Cumulative contributions from each of the training features to the increase in the coefficient of determination (R^2) and the decrease in the mean squared error (MSE). (g-h) same as panels *e, f* but for a smaller number of features. All the datasets shown in Fig. 1 are employed in this figure.

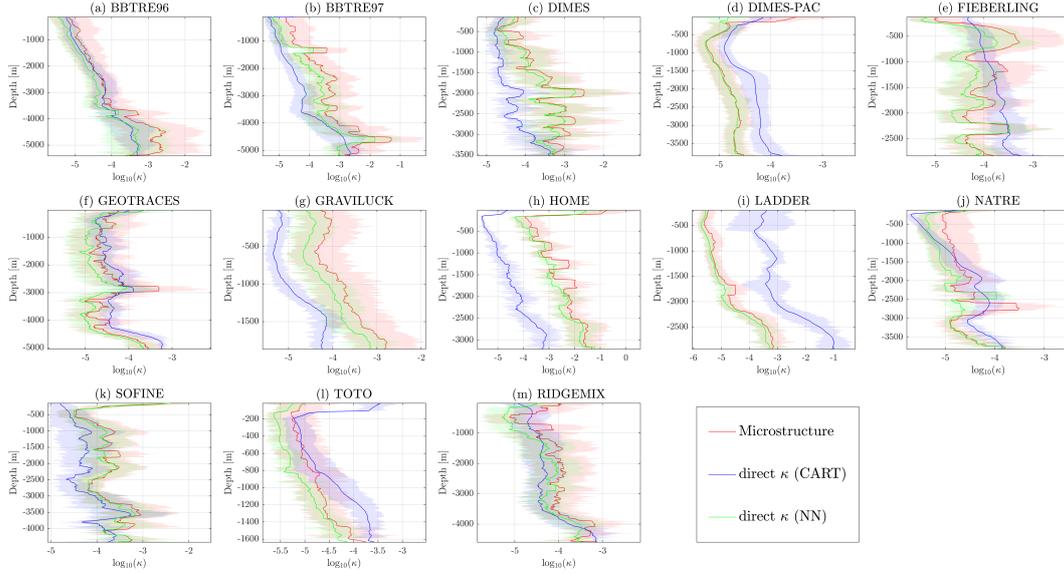


Figure 4. Predictions for individual field programs also show promise.

Comparison of the predictions of the machine learning algorithms (CART and NN) to each of the 13 field programs introduced in Fig 1. For each case, the solid lines represent the mean over all the profiles in that experiment and the corresponding shadings represent standard deviation. Note that individual predictions are made for each profile of each experiment, before averaging. For each experiment, the models were trained based on the data from all other 12 experiments, **excluding** the data from the given experiment itself, to avoid overfitting. This figure shows results from models trained to predict the turbulent diffusivity κ directly. A similar plot, showing qualitatively the same level of success, is included in the Supplementary Materials in which the models predict ϵ and κ is constructed using (1). Direct inference of κ seems to be better for predicting turbulence near the seafloor.

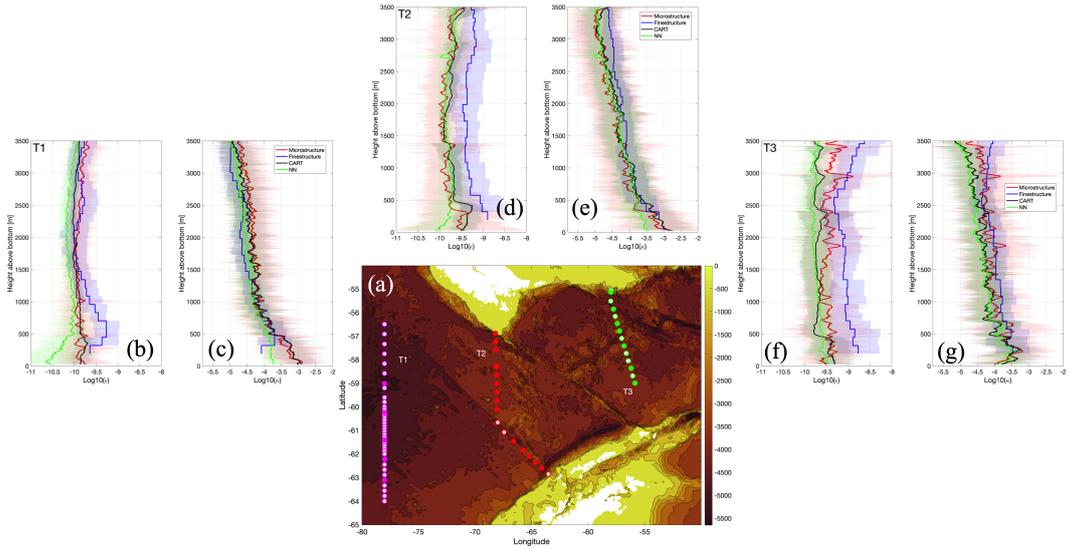


Figure 5. Machine learning competes well against physics-based parameterizations. Comparison of the predictions from machine learning against finestructure parameterization for three transects of the DIMES experiment. (a) The Drake Passage of the Southern Ocean, with the three cruise transects marked. Each circle represents a sampling. Transect T1 is the most western, plotted in magenta. Filled circles mark locations where microstructure data was taken, along with T, S, Z (from which finestructure based estimates of ϵ and κ are inferred). The circles with a white filling do not include microstructure sampling. The means over all profiles for each transect are calculated for microstructure-based, finestructure-based, and machine learning-based (both CART and NN) estimates of ϵ and κ for transect T1 in (b,c), transect T2 in (d,e), and transect T3 in (f,g). The plots are in height-above-bottom (Hab) coordinate, due to strong bottom-enhanced topographically-induced turbulence in the Drake Passage. Microstructure and finestructure estimates are from Sheen et al. (2013).